

# Context-I2W: Mapping Images to Context-Dependent Words for Accurate Zero-Shot Composed Image Retrieval

Yuanmin Tang<sup>1,2</sup>, Jing Yu<sup>1,2\*</sup>, Keke Gai<sup>3</sup>, Jiamin Zhuang<sup>1,2</sup>, Gang Xiong<sup>1,2</sup>, Yue Hu<sup>1,2</sup>, Qi Wu<sup>4</sup>

<sup>1</sup>Institute of Information Engineering, Chinese Academy of Sciences

<sup>2</sup>School of Cyber Security, University of Chinese Academy of Sciences

<sup>3</sup>School of Cyberspace Science and Technology, Beijing Institute of Technology

<sup>4</sup>Australia Institute of Machine Learning, University of Adelaide

{tangyuanmin, yujing02, zhuangjiamin, xionggang, huyue}@iie.ac.cn, gaikeke@bit.edu.cn, qi.wu01@adelaide.edu.au

## Abstract

Different from Composed Image Retrieval task that requires expensive labels for training task-specific models, Zero-Shot Composed Image Retrieval (ZS-CIR) involves diverse tasks with a broad range of visual content manipulation intent that could be related to domain, scene, object, and attribute. The key challenge for ZS-CIR tasks is to learn a more accurate image representation that has adaptive attention to the reference image for various manipulation descriptions. In this paper, we propose a novel context-dependent mapping network, named Context-I2W, for adaptively converting description-relevant Image information into a pseudo-word token composed of the description for accurate ZS-CIR. Specifically, an Intent View Selector first dynamically learns a rotation rule to map the identical image to a task-specific manipulation view. Then a Visual Target Extractor further captures local information covering the main targets in ZS-CIR tasks under the guidance of multiple learnable queries. The two complementary modules work together to map an image to a context-dependent pseudo-word token without extra supervision. Our model shows strong generalization ability on four ZS-CIR tasks, including domain conversion, object composition, object manipulation, and attribute manipulation. It obtains consistent and significant performance boosts ranging from 1.88% to 3.60% over the best methods and achieves new state-of-the-art results on ZS-CIR. Our code is available at <https://github.com/Pter61/context-i2w>.

## Introduction

Given a reference image and a text description, Composed Image Retrieval (CIR) (Vo et al. 2019b) aims to retrieve an image that is visually similar to a reference image while having visual modification according to the description. Unlike traditional content-based image retrieval (Datta et al. 2008), CIR is more flexible and accurate for users to express their search intent by incorporating both visual and language content, which brings emerging attention to internet search and e-commerce. Several CIR tasks have been proposed, including composing and manipulating objects for scene image search, attribute changes for fashion image search, and converting the image style for creative image search. There exist two core challenges of CIR: (1) accurately grounding the

\*Corresponding author

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

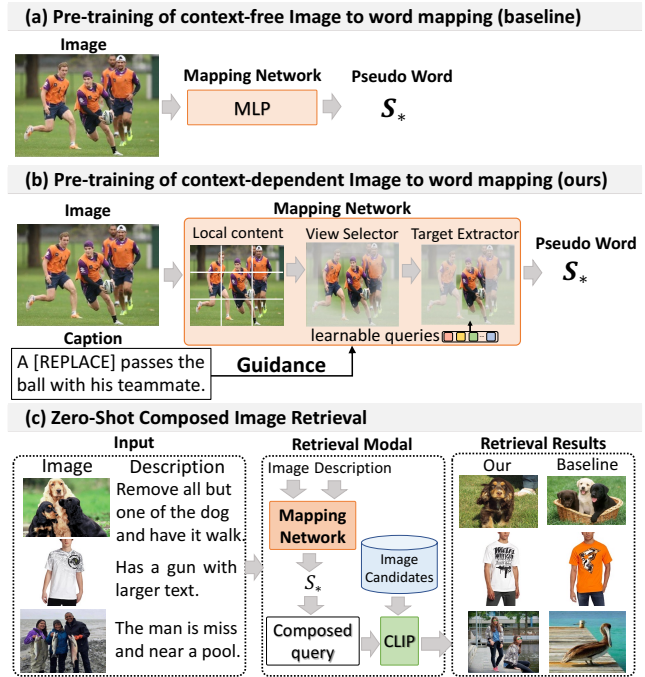


Figure 1: An illustration of our motivation. (a) Global visual mapping. (b) Our context-dependent visual mapping. (c) ZS-CIR process with the results of different mapping strategies.

visual content in the reference image according to the description and (2) composing the relevant visual and textual content from different modalities for image retrieval.

Several supervised approaches have been proposed to solve CIR problems (Chen, Gong, and Bazzani 2020; Liu et al. 2021; Goenka et al. 2022; Baldrati et al. 2022), which require a large amount of annotated triplets consisting of a reference image, a description, and a retrieved image. The supervised approaches trained on a specific task are also hard to generalize. To tackle these problems, the latest works (Saito et al. 2023; Baldrati et al. 2023) introduce the Zero-Shot Composed Image Retrieval (ZS-CIR) task. ZS-CIR aims to perform various CIR tasks without supervised training on task-specific triplet data. Current solutions treat the

ZS-CIR task as a traditional text-based image retrieval problem, as illustrated in Figure 1(c). They first learn a mapping network to convert the reference image embedding into a pseudo-word token, which is concatenated with the description to form a composed text query. Then, a pre-trained CLIP is leveraged to encode the query and the image candidates for text-to-image retrieval in a zero-shot mode. However, these models convert the whole visual information of an image into the same pseudo-word for different descriptions, as shown in Figure 1(a). This limits the model’s flexibility of adaptively selecting and mapping the visual information. In fact, only part of the visual content is relevant to the manipulation intent. Considering the queries in Figure 1(c), existing approaches are vulnerable when the description manipulates a certain object from multiple ones (*e.g.*, one of all the dogs), partial attributes of the whole object (*e.g.*, text size on the T-shirt) or the background/foreground of the global image (*e.g.*, the sea in the background).

In this work, we regard a description as the context of a reference image and propose a *Context-dependent mapping network to adaptively convert description-relevant Image information into a pseudo Word (Context-I2W)* for accurate ZS-CIR. To be generic for diverse CIR tasks, Context-I2W adaptively selects description-relevant information from the image in a hierarchical mode as illustrated in Figure 1(b): the Intent View Selector module first learns various mapping rules to dynamically map the visual embeddings to different views in a context-dependent way. Thus, each view captures the visual information from the task-specific manipulation intent (*e.g.*, domain, scene, and attribute) for even identical input. Then, the Visual Target Extractor module further collects local information from distinct aspects (*e.g.*, foreground, background, objects, and details) under the guidance of multiple learnable queries. The two modules work together to map an image to a context-dependent pseudo-word token. Independent of expensive description-region labels, Context-I2W is trained with a contrastive loss between the retrieved image and the pseudo-word enhanced composed query.

The main contributions are summarized as follows: (1) We propose a novel image to context-dependent word mapping network augmented by view selection and target extraction. We consider manipulation descriptions and learnable queries as multi-level constraints for visual information filtering, which sheds new light on the vision-to-language alignment mechanism. (2) The proposed mapping network of Context-I2W is incorporated into the ZS-CIR framework and demonstrated beneficial to tackle the existing challenges of foreground/background differentiation, multiple object combination, and fine-grained image editing. Context-I2W outperforms context-free mapping approaches and most supervised solutions. (3) Our Context-I2W mapping approach is consistently effective for diverse ZS-CIR tasks. It significantly improves CIR from 1.88% to 3.60% across four CIR tasks, including domain conversion, object composition, object/scene manipulation, and attribute manipulation. It establishes new state-of-the-art results and further impacts a broader range of vision and language applications.

## Related Work

**Composed Image Retrieval.** Composed Image Retrieval (CIR) integrates image and text for retrieval (Vo et al. 2019a). Current models use late-fusion techniques based on separately extracted visual and language features to merge them for retrieval (Baldrati et al. 2022; Liu et al. 2021; Chen, Gong, and Bazzani 2020; Shi et al. 2023). Another kind of approach (Goenka et al. 2022; Han et al. 2022) incorporates image, text, and tag features through pre-trained models. CLIP4Cir (Baldrati et al. 2022) achieves state-of-the-art by fine-tuning a CLIP (Radford et al. 2021) text encoder and employing the Combiner late-fusion module but still relies on annotated triplets. In contrast, Pic2Word (Saito et al. 2023) and SEARLE (Baldrati et al. 2023) are zero-shot CIR models trained on common image-text pairs without costly annotated CIR datasets. Pic2Word aligns the whole image features with text features, while SEARLE further integrates a pseudo-word token with a GPT-based (Brown et al. 2020) caption for data generation. However, these methods ignore the context clues for mapping description-desired visual content, thus introducing much irrelevant noise for query composition and resulting in incorrect results. To address this, we propose a context-dependent word mapping strategy, enabling the text encoder to access essential image features for accurate retrieval.

**Vision and Language Pre-training Models.** Typical Vision and Language Pre-training (VLP) models like CLIP (Radford et al. 2021) and ALIGN (Li et al. 2021) are trained on large-scale image-text pairs, which obtain pre-trained knowledge of implicit alignment between images and texts. Recent advance in VLP (Zhou et al. 2022; Song et al. 2022) explores frozen VLP models to map the encoded image features and text features to the aligned semantic space for various zero-shot tasks (Li et al. 2022; Song et al. 2022). In this work, we utilize the pre-trained CLIP as the mapping network backbone with improvement by adaptively selecting description-relevant visual information for image-text alignment. Besides, recent works (Alayrac et al. 2022; Li et al. 2023) inspired by DETR (Carion et al. 2020) utilize learnable queries for image information selection, leading to performance boost by capturing targeted information before prediction. Our work defines multiple learnable queries as guidance to capture specific parts of visual content, which provides explainable clues and fine-grained visual features for more accurate ZS-CIR.

**Mapping Image as One Word.** Several methods (Li et al. 2020; Zhang et al. 2021) aim to represent image regions as word tokens via VLP models, which rely on the effectiveness of object detectors. However, ZR-CIR tasks extend the alignment ability beyond objects to scenes, styles, attributes, *ect.* Our Context-I2W addresses this issue by using patch-based image features instead of object features to align description-relevant visual content to the word token space dynamically. PALAVRA (Cohen et al. 2022) addresses personalized image retrieval through cycle contrastive loss, which needs class-wise and caption annotations. Our proposed model achieves fine-grained image-to-word mapping without the need of any extra annotations. Some other meth-

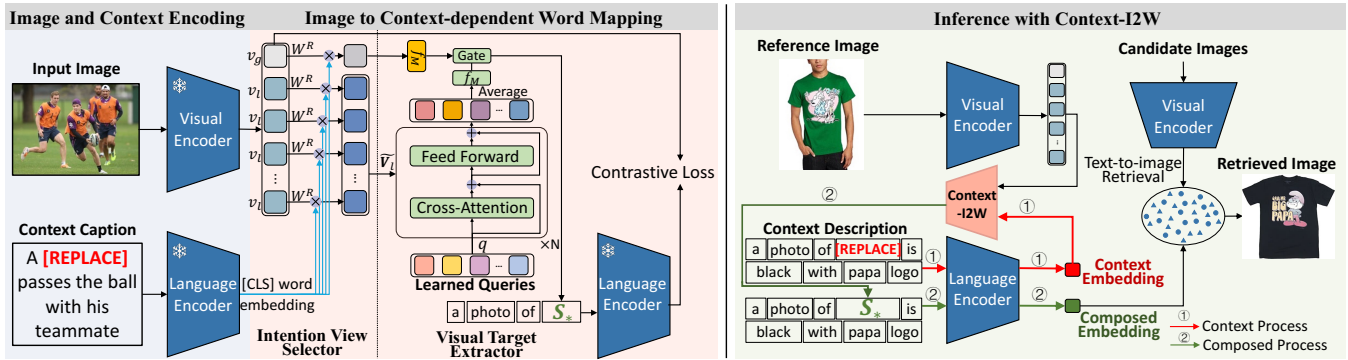


Figure 2: An overview of our Context-I2W model. Pre-training (left): Image to Context-dependent Word Mapping aims to extract caption-relevant visual content from the view level to the target level and map it to a pseudo-word token  $S_*$ . Inference (right): Map the inference image to  $S_*$  and form the composed query in a unified language space for ZS-CIR.

ods (Kumari et al. 2023; Mokady, Hertz, and Bermano 2021; Zhu et al. 2023; Tam, Raffel, and Bansal 2023) use one word token to represent multiple images of the same object for text-to-image generation. Our model doesn’t require expensive image grouping or supervised training.

## Methodology

Given a reference image  $I$  and a manipulation description  $T$ , ZS-CIR aims to retrieve images from an image database that are visually similar to  $I$  while having visual modification required in  $T$ . Figure 2 gives a detailed illustration of our model. We first introduce a new approach of encoding context based on pre-trained CLIP (Radford et al. 2021), which contains vision-language alignment knowledge to provide contextual constraints for visual information selection. This process is consistent for both contextual caption in the pre-training stage and contextual description in the reference stage. Then we learn a mapping network of Context-I2W to convert the reference image  $I$  into a pseudo-word token  $S_*$ , which is not an actual word but a representation of the image in the word token embedding space. In this work,  $S_*$  accurately depicts the visual content of context-guided manipulation intention (Intent View Selector module) and target-attended manipulation details (Visual Target Extractor module) specified in  $T$ . To effectively compose  $I$  and  $T$  across different modalities for zero-shot image retrieval, we construct a composed query in the form of a sentence  $P$  “a photo of  $S_*$  that  $T$ ” and embed it using the frozen text encoder of CLIP. Given the composed query embedding, we embed each candidate image  $I_c$  by the frozen image encoder of CLIP and regard ZS-CIR as a traditional text-to-image retrieval task by measuring the similarity between  $P$  and  $I_c$ .

### Image and Context Encoding

Since the pre-trained vision-language models are strong at modeling the cross-modal implicit alignment, we first utilize the pre-trained CLIP model to encode the context and image for fine-grained image-to-word mapping. We apply the visual encoder of frozen CLIP to represent the reference image  $I$  by a set of visual feature vectors  $V = \{v_i\}_{i=1}^m \subseteq \mathbb{R}^{d \times m}$

( $m = 257, d = 1024$ ), where  $v_1$  denotes the global image feature  $v_g$  while other ones denote local patch features  $V_l = \{v_i\}_{i=2}^m$ . In this work, we have a dataset of image-caption pairs for learning the Context-I2W network and a set of image-description pairs for ZS-CIR. We denote both captions and descriptions identically as the context of the corresponding images. Since the context generally introduces a visual target (e.g., rugby player or T-shirt in Figure 2) by describing its relevant information (e.g., surroundings or attributes). We represent the context by target-relevant information except for the target for two goals. First, we need to ensure that the context captures the described intent view of the visual content. In other words, the context features serve as mapping rules that convert the identical image representation to a text-specific view, such as the visual aspects of the surrounding scene, image domain, visual attributes, etc. Second, we need to ensure that such context features can integrate with the visual target feature in a complementary way for accurate pseudo-token mapping. To this end, we extract the first subject term in the context by a part-of-speech tagger *spacy* (Honribal et al. 2020) and replace it with a learnable token [REPLACE]. We then feed the rewritten sentence to the language encoder of frozen CLIP and obtain the [CLS] token embedding  $t = \{t_i\}_{i=1}^d \in \mathbb{R}^{d \times 1}$  as visual feature extraction guidance.

### Image to Context-Dependent Word Mapping

Since each context within an image-context pair contains a view-different and detail-specific description, we propose two modules to progressively learn the context-relevant visual features to map it to a context-dependent pseudo-token for accurate ZS-CIR. The two modules constrain the visual representation from complementary aspects: the *Intent View Selector* (IVS for short) preserves the context-desired view of visual features from the identical image features by context-guidance feature rotation. The *Visual Target Extractor* (VTE for short) further aggregates visual features of context-specific targets from the desired view features. Both the target and its relevant information are adaptively combined and mapped into the word token space by cross-modal contrastive learning.

**Intent View Selector.** Based on the encoded image and context embeddings, this module aims to represent the visual content from the view of intent described in the context without interfering with the visual content. From a mathematical point of view, the context serves by rotating the image embedding space according to the context constraint, thereby selecting the view representation of  $V$  without changing its content. Specifically, we feed each visual embedding  $v_i$  to a fully-connected layer with learnable weight matrix  $\mathbf{W}^R = \{\mathbf{w}_k^R\}_{k=1}^d \subseteq \mathbb{R}^{d \times d}$ . Each column of  $\mathbf{W}^R$  is normalization with  $\|\mathbf{w}_k^R\| = 1$ . The operation of the IVS is formulated as:

$$\begin{aligned} \tilde{v}_i &= ((\mathbf{W}^R)^\top v_i) \odot t \\ &= [t_1 \|v_i\| \|\mathbf{w}_1^R\| \cos \theta_1, \dots, t_d \|v_i\| \|\mathbf{w}_d^R\| \cos \theta_d]^\top \quad (1) \\ &= [t_1 \|\mathbf{w}_1^R\| \cos \theta_1, \dots, t_d \|\mathbf{w}_d^R\| \cos \theta_d]^\top \|v_i\| \\ &= [t_1 \cos \theta_1, \dots, t_d \cos \theta_d]^\top \|v_i\| \end{aligned}$$

where  $\odot$  represents element-wise multiplication and  $\theta_k$  is the angle between  $\mathbf{w}_k^R$  and  $\tilde{v}_i$ . Given the visual feature  $v_i$  and the fixed  $\mathbf{W}^R$ , the selected visual feature  $\tilde{v}_i$  is only affected by the context  $t$ . Furthermore, we normalize  $t$  by  $\sqrt{\sum_{k=1}^d (t_k \cos \theta_k)^2}$  in Equation 1 and thus “rotate”  $v_i$  in the feature space without changing the information content. That is, we represent the visual feature from the intent view required in context while keeping the information content unchanged, which is essential for providing comprehensive information to extract the visual target further.

**Visual Target Extractor.** Given the visual features from a specific intent view, the AI agent needs to focus further on diverse partial content in the image with respect to the context. The contextual descriptions mostly attend to several typical visual targets, including global information in domain conversion, foreground or background in scene manipulation, specific objects in object composition, visual details in attribute manipulation, *etc.* To adaptively capture the desired visual content, we define a set of learnable query embeddings  $\mathbf{X} = \{\mathbf{x}_k\}_{k=1}^n \in \mathbb{R}^{d \times n}$ , where  $d$  is the embedding dimension and  $n$  is the number of learnable queries. Each learnable query  $\mathbf{x}_k$  represents a kind of context-mentioned visual target. As illustrated in Figure 2(left), we apply cross-attention to gather target-relevant visual information from all the local patch features  $\tilde{V}_l = \{\tilde{v}_i\}_{i=2}^m$  for the learnable queries  $\mathbf{X}$  as follows. First, we compute the query, key and value through linear projections, *i.e.*,  $\mathbf{Q} = \mathbf{X}\mathbf{W}^Q$ ,  $\mathbf{K} = [\mathbf{X}, \tilde{V}_l]\mathbf{W}^K$ ,  $\mathbf{V} = [\mathbf{X}, \tilde{V}_l]\mathbf{W}^V$ .  $[\mathbf{X}, \tilde{V}_l]$  denotes concatenating the two matrices, which enhances the interaction between learnable queries and local patches with better performance. Then, the learnable queries from the current cross-attention block  $\mathbf{X}^i$  is calculated as:

$$\mathbf{X}_{att}^i = \text{Att}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d}}\right) \mathbf{V} \quad (2)$$

$$\mathbf{X}^i = \text{FFW}(\mathbf{X}_{att}^i + \mathbf{X}^{i-1}) + \mathbf{X}_{att}^i \quad (3)$$

where  $\mathbf{X}^{i-1}$  are learnable queries from the previous block and  $\text{FFW}(\cdot)$  denotes 2-layer feed-forward networks. When mapping the visual content to a pseudo-word token, both the target and relevant visual content are complementary to form the complete context information. We design a learnable scalar *gate* to decide the contribution of the target content and integrate the two parts to form the final pseudo-word token embedding  $S_*$  as follows:

$$S_* = f_{M_l}(\text{gate} \cdot \text{Avg}(\mathbf{X}_{output})) + f_{M_g}(\tilde{v}_g) \quad (4)$$

where  $\mathbf{X}_{output}$  is the output query embeddings from  $N$  transformer blocks,  $\text{Avg}(\cdot)$  denotes average pooling  $f_{M_l}(\cdot)$  and  $f_{M_g}(\cdot)$  respectively denote local and global mapping of 3-layer feed-forward networks.

**Cross-modal Contrastive Loss.** To map the pseudo token  $S_*$  to the word token space, we first replace the special token [REPLACE] in a prompt sentence “a photo of [REPLACE]” with  $S_*$  and feed it to the language encoder of CLIP to obtain the sentence embedding  $\hat{t}_s$ . We aim to match an image to its paired context-dependent prompt sentence while separating unpaired ones. We minimize the symmetric contrastive loss between the global visual embedding  $v_g$  and the prompt sentence embedding  $\hat{t}_s$  as follows:

$$\mathcal{L} = \mathcal{L}_{t2i}(\hat{t}_s, \hat{v}_g) + \mathcal{L}_{i2t}(\hat{t}_s, \hat{v}_g) \quad (5)$$

The two contrastive loss terms with a temperature hyperparameter  $\tau$  that controls the strength of penalties on hard negative samples are defined as:

$$\mathcal{L}_{t2i}(\hat{t}_s, \hat{v}_g) = -\frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \log \frac{\exp(\tau \hat{t}_{s_i}^T \hat{v}_{g_i})}{\sum_{j \in \mathcal{B}} \exp(\tau \hat{t}_{s_i}^T \hat{v}_{g_j})} \quad (6)$$

$$\mathcal{L}_{i2t}(\hat{t}_s, \hat{v}_g) = -\frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \log \frac{\exp(\tau \hat{v}_{g_i}^T \hat{t}_{s_i})}{\sum_{j \in \mathcal{B}} \exp(\tau \hat{v}_{g_i}^T \hat{t}_{s_j})} \quad (7)$$

where  $\hat{t}_{s_i} = \frac{\hat{t}_{s_i}}{\|\hat{t}_{s_i}\|}$  and  $\hat{v}_{g_j} = \frac{v_{g_j}}{\|v_{g_j}\|}$  are the normalized features of  $i$ -th prompt sentence embedding  $\hat{t}_{s_i}$  and the  $j$ -th image global embedding  $v_{g_j}$  in a batch  $\mathcal{B}$ .

### Inference with Context-I2W

In the inference stage, we compose the reference image with the paired context description and compare the composed query with candidate images for retrieval. As shown in Figure 2 (right), we first modify the context description with the learnt [REPLACE] token to form a prompt sentence and feed it to the pre-trained language encoder of CLIP and subsequently the Context-I2W network. Then we obtain the mapped pseudo-word token embedding  $S_*$  with the context-dependent visual information. Since the pseudo-word token is in the same space as actual words, we replace the [REPLACE] token with  $S_*$  in the prompt sentence to form a composed query. The composed query as well as each candidate image, is encoded by CLIP for similarity measurement. We rank the candidates based on their similarity scores.

Since our focus is on studying the context-dependent word mapping for ZS-CIR, we utilize the same prompt in



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 (CVPR 2023)	8.0	21.9	13.5	25.6	8.7	21.6	10.0	23.8	10.1	23.2
	<b>Context-I2W</b>	<b>10.2</b>	<b>26.1</b>	<b>17.5</b>	<b>28.7</b>	<b>11.6</b>	<b>27.4</b>	<b>12.1</b>	<b>28.2</b>	<b>12.9</b>	<b>27.6</b>
CIRR	Combiner (CVPR 2022)	6.1	14.8	10.5	21.3	7.0	17.7	8.5	20.4	8.0	18.5
Fashion-IQ	Combiner (CVPR 2022)	6.0	16.9	7.6	20.2	2.7	10.9	8.0	21.6	6.0	17.4

Table 1: Results on ImageNet for domain conversion.

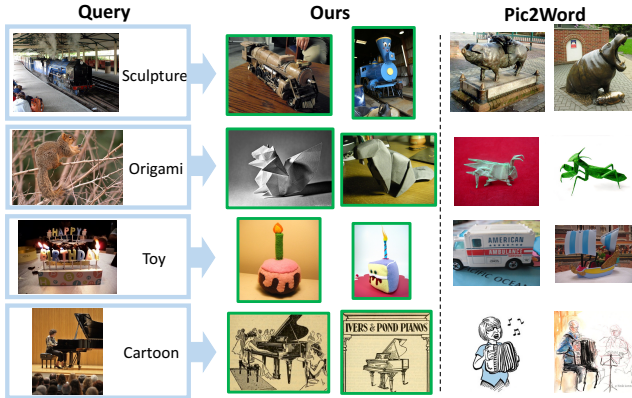


Figure 3: Retrieved results on the domain conversion task.

the most recent work (Saito et al. 2023) for a fair comparison. We show prompt examples for different ZS-CIR tasks: **(a) Domain conversion** aims to modify the domain of the reference image. The prompt is defined as a [domain tag] of [REPLACE]; **(b) Object composition** retrieves an image that contains an object in the reference image and other object tags. The prompt is in the format of a photo of [REPLACE], [obj<sub>1</sub> tag] and [obj<sub>2</sub> tag], ..., and [obj<sub>n</sub> tag]; **(c) Sentence manipulation** modifies the reference image based on a sentence. We simply append the sentence with the special token as a photo of [REPLACE], [sentence].

## Experiments

**Datasets.** We evaluate our model on four ZS-CIR datasets, *i.e.*, COCO (Lin et al. 2014) for object composition, ImageNet (Deng et al. 2009; Hendrycks et al. 2021) for domain conversion, CIRR (Liu et al. 2021) for object/scene manipulation, and Fashion-IQ (Wu et al. 2021) for attribute manipulation. All the dataset settings follow the recent works (Saito et al. 2023; Baldrati et al. 2023) for a fair comparison. (1) Domain conversion. This dataset comprises 16,983 images of 200 classes from four domains, *i.e.*, cartoon, origami, toy, and sculpture. We use the prompt (a) in inference. (2) Object composition. The dataset contains images with corresponding lists of object labels and instance masks of query images. We randomly crop one object and mask its back-

ground using its instance mask to create a reference image. We use the prompt (b) in inference. (3) Object/scene manipulation. A reference image is an instruction for manipulating an object or the background scene. We apply the prompt (c) in inference. (4) Attribute manipulation. This dataset includes various description sentences for manipulating image attributes. We utilize the prompt (c) in inference.

**Implementation Details.** We adopt ViT-L/14 CLIP (Radford et al. 2021) pre-trained on 400M image-text paired data. For training Context-I2W, we utilize the Conceptual Caption dataset (Sharma et al. 2018), which comprises 3M images. The number of cross-attention blocks is 6. The number of learnable queries is 4. To improve training stability, we initialize the learnable scalar of tanh-gating to 0 (Bachlechner et al. 2021). We employ AdamW (Loshchilov and Hutter 2018) with a learning rate of  $5 \times 10^{-5}$ , weight decay of 0.1, and a linear warmup of 10000 steps. The batch size for contrastive learning is 1024. Our model is trained on 4 Tesla V100 (32G) GPUs for 24 hours. To ensure reliable results, we report the performance averaged over three trials.

## Quantitative and Qualitative Results

We compare Context-I2W with several ZS-CIR methods, including: 1) **Text-only**: the similarity is computed based on the CLIP features of the descriptions and the candidate images; 2) **Image-only**: retrieves the most similar images to the reference image via CLIP visual features; 3) **Image + Text**: the summation of CLIP features of the reference image and the description; 4) **Pic2Word** (Saito et al. 2023): maps the entire visual features of the reference image into a pseudo-word token within the CLIP token embedding space; 5) **SEARLE-XL** (Baldrati et al. 2023): Similar to Pic2Word, it further integrates the pseudo-word token with the relative caption generated by GPT (Brown et al. 2020) and distill for efficiency. We show the reported results of SEARLE-XL on CIRR and Fashion-IQ. Also, we compare Context-I2W with the published results of widely compared supervised models, including Combiner (Baldrati et al. 2022) (Combiner\* in tables indicates using ResNet50x4 as a backbone), TIRG (Vo et al. 2019b), ARTEMIS (Delmas et al. 2022), CIR-PLANT (Liu et al. 2021) and MAAF (Dodds et al. 2020).

Table 1-4 present the quantitative results. Figure 3-6 show the corresponding qualitative results of our model and the most recent work Pic2Word. In the domain conversion results (Table 1), Context-I2W consistently outperforms exist-

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
	<b>Context-I2W</b>	<b>13.5</b>	<b>28.5</b>	<b>38.1</b>
CIRR	Combiner	9.9	22.8	32.2
Fashion-IQ	Combiner	13.2	27.1	35.2

Table 2: Results of the object composition task using COCO.

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
	SEARLE-XL	24.2	52.4	66.3	88.6
	<b>Context-I2W</b>	<b>25.6</b>	<b>55.1</b>	<b>68.5</b>	<b>89.8</b>
CIRR	Combiner	30.3	60.4	73.2	92.6
Fashion-IQ	Combiner	20.1	47.7	61.6	85.9
CIRR	Combiner*	33.6	65.4	77.4	95.2
CIRR	TIRG	14.6	48.4	64.1	90.0
CIRR	ARTEMIS	17.0	46.1	61.3	87.7
CIRR	CIRPLANT	19.6	52.6	68.4	92.4

Table 3: Results on CIRR for object composition.

ing approaches, including the supervised ones, and remarkably outperforms the State-of-the-Art (SoTA) Pic2Word by 3.60% on average. As exemplified in Figure 3, Context-I2W accurately selects the targets (*e.g.*, train, cake, squirrel, and piano) for domain conversion. Pic2Word fails to focus on the correct visual parts because it coarsely maps the global image features without object selection ability, which is compensated by our specially designed VTE.

In the experiments of object composition (Table 2), Context-I2W also outperforms the SoTA ZS-CIR model by 3.47% on average while surpassing the supervised methods on all the metrics. This further proves that Context-I2W enables the model to accurately convert the target visual content in the language token space to achieve cross-modal object combination, as depicted in Figure 4.

We also evaluate the ability of foreground/background differentiation and fine-grained image editing through the object/scene manipulation task (Table 3). Context-I2W consistently outperforms existing ZS-CIR models and achieves a boost of 1.88% on average over the best model. This is contributed to Context-I2W of selecting context-relevant visual information before mapping, effectively mitigating the impact of dataset-specific bias on ZS-CIR models (Saito et al. 2023). In Figure 5, Context-I2W adaptively and accurately removes the specific object (row 1), changes the background (row 2), and even replaces the entire content of the image (row 3) based on the different context of descriptions.

The attribute manipulation task requires accurately localizing specific attributes in the entire image. As depicted in Table 4, Context-I2W achieves an obvious average improve-

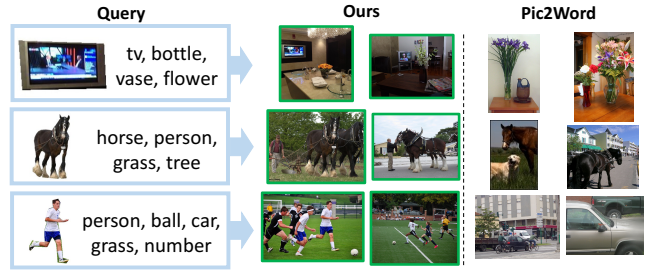


Figure 4: Retrieved results on the object composition task.

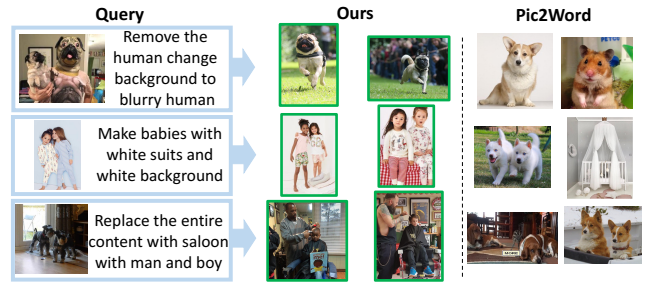


Figure 5: Retrieved results on the object manipulation task.

ment of 2.35% over the SoTA. Figure 6 provides further proves Context-I2W adeptly selects view representations of attributes and intricate visual details for word mapping.

## Ablation Study

In Table 5, we evaluate the contribution of the core components in Context-I2W on CIRR and FashionIQ. **(1) In models ‘2-3’, we evaluate the importance of context encoding approach.** Removing the context embedding  $t$  (model ‘2’) results in a significant drop of 4.86% on average. When subject replacement is not applied (model ‘3’), the performance declined by an average of 2.04%, indicating that [REPLACE] is beneficial for extracting target-relevant visual information. **(2) In models ‘4-6’, we assess the importance of key modules in image-to-word mapping process.** Removing either IVS (model ‘4’) or VTE (model ‘5’) causes obvious performance decrease. Notably, model ‘5’ brings a bigger decrease since the two tasks require manipulation of detailed visual parts, which is effectively captured by our extractor module. By directly summing the local and global features instead of using the gating strategy (model ‘6’), the performance drops by 1.48%. It indicates the necessity to capture complementary information from the two sources adaptively. **(3) Models ‘7-8’ evaluate the effect of alternative solutions for key modules.** In model ‘7’, we replace the Context-I2W with a typical cross-attention network ( $Q = \text{description}, K = V = \text{image}$ ). The results drop significantly by 10.30% on average, confirming the effectiveness of the Context-I2W mapping strategy. In model ‘8’, we replaced all nouns in the captions with [REPLACE], which disrupts the relevant information about the target and results in an average drop of 1.46%.

Supervision	Methods	Dress		Shrit		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 (CVPR 2023)	20.0	40.2	26.2	43.6	27.9	47.4	24.7	43.7
	SEARLE-XL (ICCV 2023)	20.3	43.2	27.4	45.7	29.3	50.2	25.7	46.3
	<b>Context-I2W</b>	<b>23.1</b>	<b>45.3</b>	<b>29.7</b>	<b>48.6</b>	<b>30.6</b>	<b>52.9</b>	<b>27.8</b>	<b>48.9</b>
CIRR	Combiner (CVPR 2022)	17.2	37.9	23.7	39.4	24.1	43.9	21.7	40.4
Fashion-IQ	Combiner (CVPR 2022)	30.3	54.5	37.2	55.8	39.2	61.3	35.6	57.2
Fashion-IQ	Combiner* (CVPR 2022)	31.6	56.7	36.4	58.0	38.2	62.4	35.4	59.0
Fashion-IQ	CIRPLANT (ICCV 2021)	17.5	40.4	17.5	38.8	21.6	45.4	18.9	41.5
Fashion-IQ	ARTEMIS (ICLR 2022)	27.2	52.4	21.8	43.6	29.2	54.8	26.1	50.3
Fashion-IQ	MAAF (arXiv 2020)	23.8	48.6	21.3	44.2	27.9	53.6	24.3	48.8

Table 4: Results on Fashion-IQ for attribute manipulation.



Figure 6: Retrieved results of attribute manipulation.

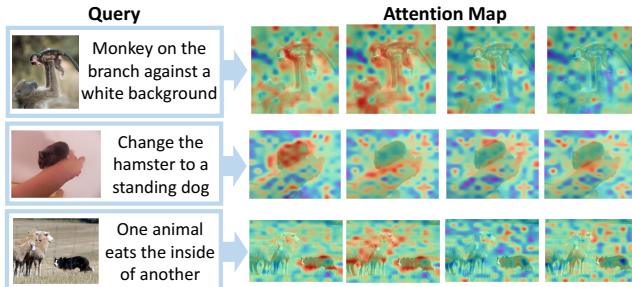


Figure 7: Attention map visualization of learnable queries.

### Interpretability of Learnable Query

In Figure 7, we visualize the attention maps of each learnable query from the last block. The four queries obviously attend to different targets in the whole image: the first two queries mainly focus on the object attributes and foreground information, while the last two queries mostly consider the background and scene information. This pattern is usually consistent (Please refer to the supplementary materials for more samples). These attention maps provide insight that our model is interpretable in extracting specific visual features of different targets, which supports fine-grained visual information selection for the context.

Methods	CIRR			Fashion-IQ	
	R1	R5	R10	R10	R50
1. full model	25.6	55.1	68.5	27.8	48.9
2. w/o context	21.4	50.3	63.6	23.8	42.5
3. w/o [REPLACE]	24.4	53.3	66.3	26.4	45.3
4. w/o View Selector	24.0	53.5	67.2	25.7	46.1
5. w/o Target Extractor	23.3	51.5	65.7	26.1	44.8
6. w/o gate	24.5	54.3	68.4	25.9	45.4
7. cross-attention	17.0	41.7	55.0	20.8	39.9
8. replace all noun	24.7	53.7	67.1	26.6	46.5

Table 5: Ablation study on CIRR and FashionIQ.

### Effectiveness and Efficiency Analysis

Our approach obtains significant improvement over four widely compared ZR-CIR tasks from 1.88% to 3.60% compared to the SoTA models. Since we design dedicated modules for accurate mapping, our model size (65.3M) is bigger than the simple 3-layer MLP mapping (0.9M) of Pic2Word. Accordingly, our training time (24h) increased by 8 hours compared to Pic2Word in the same settings. It's worth noting that our model using only 50% of the pre-training data achieves comparable performance to SoTA models<sup>1</sup>, which makes up for training time. Our inference time (0.026s) is slightly 0.007s slower than Pic2Word.

### Conclusion

In this paper, we propose a novel image-to-word mapping network that leverages manipulation descriptions and learnable queries for context-dependent visual information selection and accurate mapping. Context-I2W shows strong generalization ability and remarkably improves the best performance of existing approaches on four diverse ZS-CIR tasks. It inspires the vision-to-language alignment mechanism and impacts diverse vision and language applications. How to design more lightweight models with high retrieval performance will be the future work.

<sup>1</sup>Our Appendix at <https://arxiv.org/abs/2309.16137>

## Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No. 62006222).

## References

- Alayrac, J.-B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millican, K.; Reynolds, M.; Ring, R.; Rutherford, E.; Cabi, S.; Han, T.; Gong, Z.; Samingo, S.; Monteiro, M.; Menick, J. L.; Borgeaud, S.; Brock, A.; Nematzadeh, A.; Sharifzadeh, S.; Bińkowski, M. a.; Barreira, R.; Vinyals, O.; Zisserman, A.; and Simonyan, K. 2022. Flamingo: a Visual Language Model for Few-Shot Learning. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 23716–23736.
- Bachlechner, T.; Majumder, B. P.; Mao, H.; Cottrell, G.; and McAuley, J. 2021. Rezero is all you need: Fast convergence at large depth. In *Uncertainty in Artificial Intelligence*, 1352–1361.
- Baldrati, A.; Agnolucci, L.; Bertini, M.; and Del Bimbo, A. 2023. Zero-Shot Composed Image Retrieval with Textual Inversion. *arXiv:2303.15247*.
- Baldrati, A.; Bertini, M.; Uricchio, T.; and Del Bimbo, A. 2022. Effective Conditioned and Composed Image Retrieval Combining CLIP-Based Features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21466–21474.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020. Language Models are Few-Shot Learners. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 1877–1901. Curran Associates, Inc.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, 213–229.
- Chen, Y.; Gong, S.; and Bazzani, L. 2020. Image Search With Text Feedback by Visiolinguistic Attention Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3001–3011.
- Cohen, N.; Gal, R.; Meir, E. A.; Chechik, G.; and Atzmon, Y. 2022. “This Is My Unicorn, Fluffy”: Personalizing Frozen Vision-Language Representations. In *European conference on computer vision*, 558–577.
- Datta, R.; Joshi, D.; Li, J.; and Wang, J. Z. 2008. Image retrieval: Ideas, influences, and trends of the new age. *ACM Computing Surveys*, 40(2): 1–60.
- Delmas, G.; Rezende, R. S.; Csurka, G.; and Larlus, D. 2022. ARTEMIS: Attention-based Retrieval with Text-Explicit Matching and Implicit Similarity. In *International Conference on Learning Representations*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition*, 248–255.
- Dodds, E.; Culpepper, J.; Herdade, S.; Zhang, Y.; and Boakye, K. 2020. Modality-agnostic attention fusion for visual search with text feedback. *arXiv:2007.00145*.
- Goenka, S.; Zheng, Z.; Jaiswal, A.; Chada, R.; Wu, Y.; Hedau, V.; and Natarajan, P. 2022. FashionVLP: Vision Language Transformer for Fashion Retrieval With Feedback. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14105–14115.
- Han, X.; Yu, L.; Zhu, X.; Zhang, L.; Song, Y.-Z.; and Xiang, T. 2022. FashionViL: Fashion-Focused Vision-and-Language Representation Learning. In *European conference on computer vision*, 634–651.
- Hendrycks, D.; Basart, S.; Mu, N.; Kadavath, S.; Wang, F.; Dorundo, E.; Desai, R.; Zhu, T.; Parajuli, S.; Guo, M.; Song, D.; Steinhardt, J.; and Gilmer, J. 2021. The Many Faces of Robustness: A Critical Analysis of Out-of-Distribution Generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8340–8349.
- Honnibal, M.; Montani, I.; Van Landeghem, S.; Boyd, A.; et al. 2020. spaCy: Industrial-strength natural language processing in python.
- Kumari, N.; Zhang, B.; Zhang, R.; Shechtman, E.; and Zhu, J.-Y. 2023. Multi-Concept Customization of Text-to-Image Diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1931–1941.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. BLIP-2: Bootstrapping Language-Image Pre-training with Frozen Image Encoders and Large Language Models. *arXiv:2301.12597*.
- Li, J.; Li, D.; Xiong, C.; and Hoi, S. 2022. BLIP: Bootstrapping Language-Image Pre-training for Unified Vision-Language Understanding and Generation. In *Proceedings of the 39th International Conference on Machine Learning*, 12888–12900.
- Li, J.; Selvaraju, R.; Gotmare, A.; Joty, S.; Xiong, C.; and Hoi, S. C. H. 2021. Align before Fuse: Vision and Language Representation Learning with Momentum Distillation. In *Advances in Neural Information Processing Systems*, 9694–9705.
- Li, X.; Yin, X.; Li, C.; Zhang, P.; Hu, X.; Zhang, L.; Wang, L.; Hu, H.; Dong, L.; Wei, F.; et al. 2020. Oscar: Object-semantic aligned pre-training for vision-language tasks. In *European Conference on Computer Vision*, 121–137.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft COCO: Common Objects in Context. In Fleet, D.; Pajdla, T.; Schiele, B.; and Tuytelaars, T., eds., *European Conference on Computer Vision*, 740–755.
- Liu, Z.; Rodriguez-Opazo, C.; Teney, D.; and Gould, S. 2021. Image Retrieval on Real-Life Images With Pre-Trained Vision-and-Language Models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2125–2134.



- Loshchilov, I.; and Hutter, F. 2018. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- Mokady, R.; Hertz, A.; and Bermano, A. H. 2021. ClipCap: CLIP Prefix for Image Captioning. arXiv:2111.09734.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the International Conference on Machine Learning*, 8748–8763.
- Saito, K.; Sohn, K.; Zhang, X.; Li, C.-L.; Lee, C.-Y.; Saenko, K.; and Pfister, T. 2023. Pic2Word: Mapping Pictures to Words for Zero-Shot Composed Image Retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19305–19314.
- Sharma, P.; Ding, N.; Goodman, S.; and Soricut, R. 2018. Conceptual Captions: A Cleaned, Hypernymed, Image Alt-text Dataset For Automatic Image Captioning. In *Annual Meeting of the Association for Computational Linguistics*, 2556–2565.
- Shi, J.; Zhang, Y.; Yin, X.; Xie, Y.; Zhang, Z.; Fan, J.; Shi, Z.; and Qu, Y. 2023. Dual Pseudo-Labels Interactive Self-Training for Semi-Supervised Visible-Infrared Person Re-Identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 11218–11228.
- Song, H.; Dong, L.; Zhang, W.-N.; Liu, T.; and Wei, F. 2022. CLIP Models are Few-shot Learners: Empirical Studies on VQA and Visual Entailment. arXiv:2203.07190.
- Tam, D.; Raffel, C.; and Bansal, M. 2023. Simple Weakly-Supervised Image Captioning via CLIP’s Multimodal Embeddings. In *The AAAI-23 Workshop on Creative AI Across Modalities*.
- Vo, N.; Jiang, L.; Sun, C.; Murphy, K.; Li, L.-J.; Fei-Fei, L.; and Hays, J. 2019a. Composing Text and Image for Image Retrieval - an Empirical Odyssey. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6439–6448.
- Vo, N.; Jiang, L.; Sun, C.; Murphy, K.; Li, L.-J.; Fei-Fei, L.; and Hays, J. 2019b. Composing text and image for image retrieval-an empirical odyssey. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6439–6448.
- Wu, H.; Gao, Y.; Guo, X.; Al-Halah, Z.; Rennie, S.; Grauman, K.; and Feris, R. 2021. Fashion IQ: A New Dataset Towards Retrieving Images by Natural Language Feedback. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11307–11317.
- Zhang, P.; Li, X.; Hu, X.; Yang, J.; Zhang, L.; Wang, L.; Choi, Y.; and Gao, J. 2021. Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5579–5588.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022. Conditional Prompt Learning for Vision-Language Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16816–16825.
- Zhu, W.; Yan, A.; Lu, Y.; Xu, W.; Wang, X. E.; Eckstein, M.; and Wang, W. Y. 2023. Visualize Before You Write: Imagination-Guided Open-Ended Text Generation. arXiv:2210.03765.