

# Missing Target-Relevant Information Prediction with World Model for Accurate Zero-Shot Composed Image Retrieval

Yuanmin Tang<sup>1,2</sup> Jing Yu<sup>3\*</sup> Keke Gai<sup>4</sup> Jiamin Zhuang<sup>1,2</sup> Gang Xiong<sup>1</sup> Gaopeng Gou<sup>1</sup> Qi Wu<sup>5</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 Information Engineering, Minzu University of China

<sup>4</sup>Beijing Institute of Technology, <sup>5</sup>University of Adelaide

{tangyuanmin, zhuangjiamin, gougaopeng, xionggang}@iie.ac.cn, jing.yu@muc.edu.cn,  
gaikeke@bit.edu.cn, qi.wu01@adelaide.edu.au

## Abstract

*Zero-Shot Composed Image Retrieval (ZS-CIR) involves diverse tasks with a broad range of visual content manipulation intent across domain, scene, object, and attribute. The key challenge for ZS-CIR tasks is to modify a reference image according to manipulation text to accurately retrieve a target image, especially when the reference image is missing essential target content. In this paper, we propose a novel prediction-based mapping network, named PrediCIR, to adaptively predict the missing target visual content in reference images in the latent space before mapping for accurate ZS-CIR. Specifically, a world view generation module first constructs a source view by omitting certain visual content of a target view, coupled with an action that includes the manipulation intent derived from existing image-caption pairs. Then, a target content prediction module trains a world model as a predictor to adaptively predict the missing visual information guided by user intention in manipulating text at the latent space. The two modules map an image with the predicted relevant information to a pseudo-word token without extra supervision. Our model shows strong generalization ability on six ZS-CIR tasks. It obtains consistent and significant performance boosts ranging from 1.73% to 4.45% over the best methods and achieves new state-of-the-art results on ZS-CIR. Our code is available at <https://github.com/Pter61/predicir>.*

## 1. Introduction

Given a reference image and a human manipulation text, Composed Image Retrieval (CIR) [41] aims to retrieve a target image visually similar to the reference image while

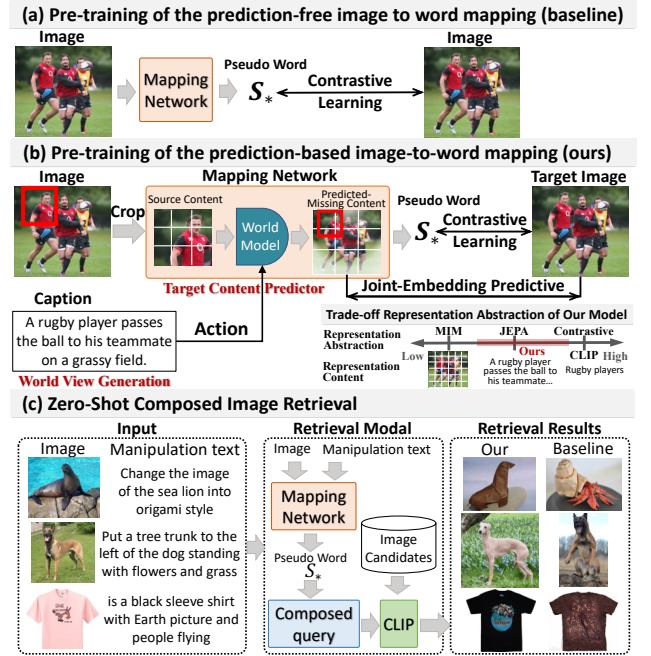


Figure 1. Illustration of our motivation. (a) Prediction-free visual mapping. (b) Our prediction-based visual mapping. (c) ZS-CIR process results from different strategies.

incorporating content modifications specified by the manipulation text. Distinct from traditional content-based image retrieval [10], in the CIR task, the content of the manipulation text often does not appear in the reference images, as illustrated in Figure 1(c). CIR enhances flexibility and improves the accuracy of intent expression by allowing users to integrate visual and textual information into their search queries. This task has gained emerging attention in internet searches and e-commerce applications [9, 33]. As shown in Figure 1(c), CIR tasks include image domain transforma-

\*Corresponding author

tion for creative searches, object composition and manipulation for natural image searches, and attribute modifications for fashion image searches.

There exist two core challenges of CIR: (1) accurately modify the reference image guided by the manipulation text for retrieving the target image, particularly when the necessary visual content is missing in the reference images, and (2) adaptively compose visual and textual content guided by manipulation text for image retrieval. Various supervised [4, 29] and semi-supervised methods [14, 22] have been proposed to address CIR problem, which requires an extensive amount of annotated triplets, *i.e.*, a reference image, a manipulated text, and a target image, for training task-specific retrieval models. These supervised methods are labor-intensive or require large models (*e.g.*, diffusion [32]) for data annotation, which limits their generalizability. To overcome this issue, recent studies have introduced Zero-Shot Composed Image Retrieval (ZS-CIR) [5, 33] by utilizing a pre-trained CLIP model [31] to treat ZS-CIR as a traditional text-based image retrieval challenge. As depicted in Figure 1(c), these methods map the reference image into a pseudo-token of CLIP’s language space, combining it with manipulation text to form a query. This query retrieves target images from a shared semantic space in a zero-shot mode by calculating semantic similarity. Despite these advancements, the mapping networks are inadequate for ZS-CIR tasks for the following reasons:

(1) The CLIP embedding, learned through contrastive methods, is coarse-grained [13]. This leads to losing crucial visual details in the pre-trained pseudo-token, which is essential for CIR tasks. For instance, in the bottom right of Figure 1(b), the CLIP embedding captures the main object (*e.g.*, rugby players) while missing fine-grained details of relation (*e.g.*, pass ball) and sense (*e.g.*, grassy field).

(2) Existing ZS-CIR methods train a network to map reference images into language space, ignoring its missing target content, as shown in Figure 1(a). This limits the model’s ability to generate target image information for retrieval adaptively. In fact, key elements of the target images are often missing from the reference images. Considering the queries in Figure 1(c), existing methods struggle to handle manipulations where relevant elements are missing in the reference images, such as image domains (*e.g.*, origami style), objects/scene (*e.g.*, dog among the flowers with tree), and attributes (*e.g.*, black with Earth logo).

In this paper, we propose a novel approach to *Predict target image feature before retrieval for zero-shot Composed Image Retrieval (PrediCIR)*. Unlike existing ZS-CIR approaches, PrediCIR trains a world model, an effective target content predictor [13, 44], to adaptively predict key visual elements (*e.g.*, objects, senses, attributes, domain) missing from reference images, guided by manipulation text, as illustrated in Figure 1(b): the *World View Generation* mod-

ule first generates both source and target views from existing image-caption pairs without extra supervision for training. Specifically, we corrupt an image’s content via random cropping, regarding these as the source view and the original image as the target view, with the corresponding description as the action. Subsequently, the *Target Content Predictor* module trains a world model to adaptively predict the key element of the target view guided by action, which is missing in the source view. The two modules map an image with the predicted target content to a pseudo-word, providing a high-quality feature with fine-grained visual details.

The main contributions are summarized as follows: (1) We introduce a novel prediction-based image-to-word mapping network, leveraging the world model to simulate user behavior. This facilitates the prediction of potential target image features relevant to manipulation intent for retrieval, offering new insights into the vision-to-language alignment mechanism. (2) The proposed mapping network of PrediCIR adaptively predicts key elements (*e.g.*, objects, senses, attributes, and various details) of target visual content missing in reference images, proving advantageous for challenges such as object combination, foreground/background editing, attribute adjustment, and domain conversion. (3) Our PrediCIR is consistently effective for diverse ZS-CIR tasks. It significantly improves CIR from 1.73% to 4.45% across six CIR tasks with comparable inference times. It establishes new state-of-the-art results and further impacts a broader range of vision and language applications.

## 2. Related works

**Composed Image Retrieval.** Composed Image Retrieval (CIR) combines image and text features for retrieval [41], typically using late fusion to integrate visual and language features separately while requiring extensively annotated triplets CIR datasets. [4, 22, 29, 47]. Zero-shot CIR models [5, 12, 15, 21, 23, 33, 38, 39, 45], trained on image-text pairs, avoid the need for extensive CIR datasets by mapping reference images to text space for query formation. However, they often miss visual content specified by manipulation text, resulting in less accurate queries. To address this, we introduce a prediction-based word mapping, allowing the text encoder to access potential target image features. Unlike CompoDiff [14], which requires multi-step diffusion model training with synthesizing triplets, our model predicts target content in latent space on a single step, enhancing performance without additional supervision. We create pseudo triplets by cropping visual elements to preserve full contextual embeddings under a frozen CLIP, avoiding the mask-based CIR methods [8, 19, 46] that require CLIP fine-tuning. Additionally, unlike diffusion [14], LLMs [23], or external databases [27, 37] methods, which introduce non-negligible computational overhead, our model remains lightweight with comparable inference times.

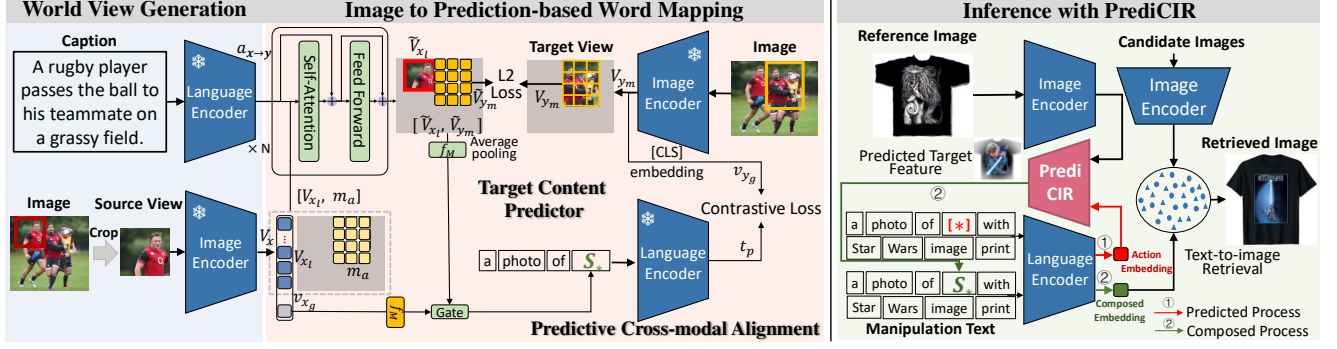


Figure 2. An overview of our PrediCIR model. Pre-training (left): Image to prediction-based word mapping aims to predict target-relevant missing visual content in latent space and map it with reference image content to a pseudo-word token  $S_*$ . Inference (right): Map the inference image to  $S_*$  and form the composed query in a unified language space for ZS-CIR.

**World Model in Vision Representation Learning.** World modeling has proven its effectiveness across several domains [16, 44], showing clear benefits in prediction-based representation learning. Masked Image Modeling (MIM) approaches [17, 43] learn representations by predicting masked image areas in pixel space, aligning their decoders with world models. Similarly, leveraging the Joint Embedding Predictive Architecture (JEPA) [24], I-JEPA [2] predicts masked parts of the image in the latent space. Recent work [13] introduces an image world model that predicts cropped image region in latent space, focusing on fine-grained visual details at the patch level instead of whole-image diffusion [32]. Building on these advances, our work first introduces a novel world model to predict target image features guided by text for vision-language retrieval tasks.

**Vision and Language Pre-training Models.** Vision and Language Pre-training (VLP) models, such as CLIP [31], employ extensive image-text pair training to achieve implicit alignment between visual and textual data. Recent advancements in VLP [36, 48] have utilized static models to merge encoded image and text features, facilitating a range of zero-shot tasks [1, 25, 26, 35, 36]. In our work, we adapt VLP models for CIR tasks in a zero-shot manner. Moreover, we uniquely leverage the pre-trained knowledge of VLP models to encode world views and actions to train a vision-language world model for prediction-based ZS-CIR.

### 3. Methodology

Given a reference image  $I$  and a manipulation text  $T$ , Zero-Shot Composed Image Retrieval (ZS-CIR) aims to retrieve images from an image database that are visually similar to  $I$  while incorporating the visual modifications specified in  $T$ . A detailed illustration of our model is provided in Figure 2. We first introduce a new approach for generating world view from image-caption pairs. This process allows us to provide source and target views with corresponding actions for training a predictor based on the Joint Embedding Predictive Architecture (JEPA) [24]. In this frame-

work, the predictor is the instantiation of the world model [2, 13]. Then, we learn a prediction-based mapping network of PrediCIR to predict visual elements of the target image missing in the reference image guided by manipulation text  $T$  and convert them into a pseudo-word token  $S_*$  in the token embedding space. In this work,  $S_*$  depicts the potential visual content of the target image, combining the existing content of the reference image with the predicted visual elements specified by  $T$ . 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  $S_*, T$ ” and embed it using the frozen text encoder of CLIP. Given the composed query embedding, we embed each candidate image  $I_i$  by the frozen image encoder of CLIP and approach ZS-CIR as a traditional text-to-image retrieval task by measuring the similarity between  $P$  and  $I_i$ .

#### 3.1. World View Generation

Since images and captions in existing ZS-CIR training datasets share the contextual elements, training a world model to predict the absent visual elements in reference images poses a challenge. To address this, we construct images with missing visual elements by modifying existing image-caption datasets, specifically by randomly cropping existing images. This approach aligns better with the frozen CLIP model than the masked-region CIR method [8], which requires fine-tuning to interpret masked inputs. Given an image  $I$  with width  $W$  and height  $H$  in an image-caption pair, we obtain a cropped image  $I_c$  as follows:

$$\begin{aligned} I_c &= \text{Crop}(I, (x, y, W_c, H_c)) \\ &= \text{Crop}(I, (x, y, x + \sqrt{s_c r_c W H}, y + \sqrt{s_c W H / r_c})) \end{aligned} \quad (1)$$

where  $\text{Crop}(\cdot)$  denotes the cropping operation, and  $(x, y)$  denotes the top-left corner of the cropping region. These coordinates are dynamically calculated to focus on areas of interest within the image.  $W_c$  and  $H_c$  are the width and height of the cropped image, respectively. Crop size  $s_c$  and

aspect ratios  $r_c$  ensure the cropping operation dynamically adjusts to various image sizes. Subsequently, we utilize the cropped image as the source view  $\mathbf{x}$ , the original image as the target view  $\mathbf{y}$ , and the caption as the action  $a_{\mathbf{x} \rightarrow \mathbf{y}}$  for training a world model for predicting the visual content of a target view that missing in a source view.

As CLIP shows its strong capabilities, we employ the CLIP model to encode the source/target views and the actions for prediction-based mapping. We utilize the visual encoder of the frozen CLIP model to represent the cropped image  $I_c$  as the source view  $\mathbf{x}$  by a set of visual feature vectors  $\mathbf{V}_x = \{\mathbf{v}_{x_i}\}_{i=1}^m$  where  $m = 257$  and  $d = 1024$ . Here, the  $\mathbf{v}_{x_1}$  denotes the global source feature  $\mathbf{v}_{x_g}$ , and the subsequent vectors  $\{\mathbf{v}_{x_i}\}_{i=2}^m$  represent the local patch features  $\mathbf{V}_{x_l}$ . Similarly, the original image  $I$  is encoded as the target view  $\mathbf{y}$  using a set of visual feature vectors  $\mathbf{V}_y = \{\mathbf{v}_{y_i}\}_{i=1}^m$ , where  $\mathbf{v}_{y_1}$  acting as the global target feature  $\mathbf{v}_{y_g}$ .

In this work, we construct a dataset of triplets  $\langle \text{source view, action, target view} \rangle$  comprising  $\langle \text{cropped image, caption, original image} \rangle$  for training the PrediCIR network. Subsequently, we utilize triplets  $\langle \text{reference image, manipulation text, target image} \rangle$  triplets for ZS-CIR. We treat both captions and manipulation texts as actions  $a_{\mathbf{x} \rightarrow \mathbf{y}}$ , which express the user’s intent to modify the source view (e.g., a rugby player) into the target view (e.g., a rugby player pass the ball to his teammate on a grassy field) as illustrated in Figure 2. Our PrediCIR has two goals: First, it predicts the visual content of the target view missing in the source, guided by the action. Second, it adaptively combines the predicted content with the source’s content for mapping. To this end, we feed the caption to the frozen CLIP language encoder, obtaining the [CLS] token embedding  $\mathbf{t} = \{t_i\}_{i=1}^d \in \mathbb{R}^{d \times 1}$  as the action for predicting.

### 3.2. Image to Prediction-based Word Mapping

Given the constructed triplets  $\langle \text{source view, action, target view} \rangle$ , where the source view  $\mathbf{x} = \mathbf{V}_{x_l}$  comprises the local patch-level features of the cropped image, the target view  $\mathbf{y} = \mathbf{V}_y$  includes the patch-level features of the original image, and the action  $a_{\mathbf{x} \rightarrow \mathbf{y}} = \mathbf{t}$  encapsulates the user manipulation intent. We introduce two modules to progressively predict the target visual content that missing in the source view, and map to a pseudo-token for accurate ZS-CIR: the *Target Content Predictor* (TCP for short) first learns a world model functioning as a user simulator to predict the target visual content guided by the manipulation intent through the JEPA framework. Subsequently, the *Predictive Cross-Modal Alignment* (PMA for short) adaptively combines the predicted and source visual contents, mapping them into the word token space using cross-modal contrastive learning.

**Target Content Predictor.** Given the visual patch features from the triplets  $\langle \text{source view, action, target view} \rangle$ . This module aims to predict the target visual content missing in

the source view, guided by the action. Specifically, we fed with the target visual content in the form of mask tokens as well as  $a_{\mathbf{x} \rightarrow \mathbf{y}}$ . We denote these mask tokens as  $m_a$ , parameterized by a shared learnable vector with an added positional embedding, representing a randomly sampled block  $B$  from the target patch features  $\mathbf{V}_y$ . These mask tokens  $m_a$  corresponding of the position of  $B$  by the target patch features  $\mathbf{V}_{y_m} = \{\mathbf{V}_{y_j}\}_{j \in B}$ . Specifically, we sample the block using the same crop size and aspect ratios described in Eq.1. Subsequently, as illustrated in Figure 2 (left), we apply self-attention and combine the action  $a_{\mathbf{x} \rightarrow \mathbf{y}}$ , the embedded source patches  $\mathbf{V}_{x_l}$ , and mask tokens  $m_a$  as input  $\mathbf{X} = [a_{\mathbf{x} \rightarrow \mathbf{y}}, \mathbf{V}_{x_l}, m_a]$ . First, we compute the query, key and value through linear projections, i.e.,  $\mathbf{Q} = \mathbf{XW}^Q$ ,  $\mathbf{K} = \mathbf{XW}^K$ ,  $\mathbf{V} = \mathbf{XW}^V$ .  $\mathbf{X}$  denotes concatenating the three matrices, which enhances the interaction between mask tokens and source local patches guided by the manipulation intent to achieve a high-quality representation with fine-gained visual details crucial for CIR tasks. Then, the mask token and source local patches from the current self-attention block  $\mathbf{X}^i$  are calculated as:

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

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

where  $\mathbf{X}^{i-1}$  are mask tokens with source local patch features from the previous block and  $\text{FFW}(\cdot)$  denotes 2-layer feed-forward networks. Finally, we calculate the squared  $\mathcal{L}_2$  loss to minimize the distance between the patch prediction  $\tilde{\mathbf{V}}_{y_m} = \mathbf{X}_{out}[\text{select}(m_a)]$  as follow:

$$\mathcal{L}_{pred} = \mathcal{L}_2(\mathbf{x}, \mathbf{y}) = \sum_{i \in B} \|\tilde{\mathbf{V}}_{y_m}^i - \mathbf{V}_{y_m}^i\|_2^2 \quad (4)$$

where  $\mathbf{X}_{out}$  denotes the output embeddings from  $N$  transformer blocks.  $\text{select}(\cdot)$  is used to select the corresponding indexes of features within  $\mathbf{X}_{out}$ ,  $\tilde{\mathbf{V}}_{y_m}^i$  and  $\mathbf{V}_{y_m}^i$  are the  $i^{\text{th}}$  target patch prediction and the corresponding target patch feature, respectively, and  $B$  is the target image block.

**Predictive Cross-Modal Alignment.** Given the target patch prediction  $\tilde{\mathbf{V}}_{y_m}$ , enhanced source patch features  $\tilde{\mathbf{V}}_{x_l} = \mathbf{X}_{out}[\text{select}(x_l)]$  that is high-quality representation, and global source feature  $\mathbf{v}_{x_g}$ , the AI agent aims to form a target embedding optimized for retrieval. When mapping the visual content to a pseudo-word token, both the predict and source content are complementary to form the complete target information. To align the feature of JEPA and CLIP and adaptively weight predicted information on the retrieval process, we introduce a learnable scalar *gate* that decides the contribution of the predicted information  $[\tilde{\mathbf{V}}_{x_l}, \tilde{\mathbf{V}}_{y_m}]$  and integrates the global source information  $\mathbf{v}_{x_g}$  to form the final target embedding  $\mathbf{S}_*$  as follows:

$$\mathbf{S}_* = f_{M_p}(\text{gate} \cdot \text{Avg}([\tilde{\mathbf{V}}_{x_l}, \tilde{\mathbf{V}}_{y_m}])) + f_{M_s}(\mathbf{v}_{x_g}) \quad (5)$$



Backbones	Methods	Conferences	Dress		Shrit		TopTee		Average	
			R10	R50	R10	R50	R10	R50	R10	R50
ViT-L/14	Pic2Word <sup>†</sup>	CVPR 2023	20.0	40.2	26.2	43.6	27.9	47.4	24.7	43.7
	SEARLE-XL <sup>†</sup>	ICCV 2023	20.3	43.2	27.4	45.7	29.3	50.2	25.7	46.3
	LinCIR <sup>†</sup>	CVPR 2024	20.9	42.4	29.1	46.8	28.8	50.2	26.3	46.5
	Context-I2W <sup>†</sup>	AAAI 2024	23.1	45.3	29.7	48.6	30.6	52.9	27.8	48.9
	<b>PrediCIR</b>	–	<b>25.4</b>	<b>49.5</b>	<b>31.8</b>	<b>52.0</b>	<b>33.1</b>	<b>55.4</b>	<b>30.1</b>	<b>52.3</b>
ViT-G/14	CompoDiff <sup>†</sup>	TMLR 2024	37.8	49.1	41.3	55.2	44.3	56.4	39.0	51.7
	LinCIR <sup>†</sup>	CVPR 2024	38.1	60.9	46.8	65.1	50.5	71.1	45.1	65.7
	<b>PrediCIR</b>	–	<b>39.7</b>	<b>62.4</b>	<b>48.2</b>	<b>67.4</b>	<b>53.7</b>	<b>73.6</b>	<b>47.2</b>	<b>67.8</b>

Table 1. Results on Fashion-IQ for attribute manipulation. <sup>†</sup> indicates results from the original paper.

where  $\text{Avg}(\cdot)$  denotes average pooling,  $f_{M_p}(\cdot)$  and  $f_{M_s}(\cdot)$  respectively denote predict and source mapping of 3-layer feed-forward networks. To map the pseudo token  $S_*$  to the word token space, we append  $S_*$  to the end of the token embeddings of the prompted sentence, “a photo of”, and feed it to the language encoder of CLIP to obtain the sentence embedding  $\mathbf{t}_p$ . We aim to match an image to its paired prediction-based prompt sentence while separating unpaired ones. We minimize the symmetric contrastive loss between the global target visual feature  $\mathbf{v}_{y_g}$  and the prompt sentence embedding  $\mathbf{t}_p$  as follows:

$$\mathcal{L}_{align} = \mathcal{L}_{t2i}(\mathbf{t}_p, \mathbf{v}_{y_g}) + \mathcal{L}_{i2t}(\mathbf{t}_p, \mathbf{v}_{y_g}) \quad (6)$$

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

$$\mathcal{L}_{t2i}(\mathbf{t}_p, \mathbf{v}_{y_g}) = -\frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \log \frac{e^{\tau(\mathbf{t}_p^i)^T \mathbf{v}_{y_g}^i}}{\sum_{j \in \mathcal{B}} e^{\tau(\mathbf{t}_p^i)^T \mathbf{v}_{y_g}^j}} \quad (7)$$

$$\mathcal{L}_{i2t}(\mathbf{t}_p, \mathbf{v}_{y_g}) = -\frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \log \frac{e^{\tau(\mathbf{v}_{y_g}^i)^T \mathbf{t}_p^i}}{\sum_{j \in \mathcal{B}} e^{\tau(\mathbf{v}_{y_g}^i)^T \mathbf{t}_p^j}} \quad (8)$$

The final loss to optimize PrediCIR:

$$\mathcal{L} = \mathcal{L}_{pred} + \mathcal{L}_{align} \quad (9)$$

### 3.3. Inference with PrediCIR

In the inference stage, we compose the reference image with the paired manipulation text and compare the composed query with candidate images for retrieval. As shown in Figure 2 (right), we first form a prompt sentence that includes special token  $[*]$  and manipulation text, which is fed to the language encoder of CLIP to obtain an action embedding, followed by predicting through the PrediCIR network. Then, we obtain the mapped pseudo-token embedding  $S_*$  containing the predicted target-relevant information and replace the  $[*]$  token with  $S_*$  to form a composed query. The result is embedded by the language encoder and compared to the visual features of candidate images.

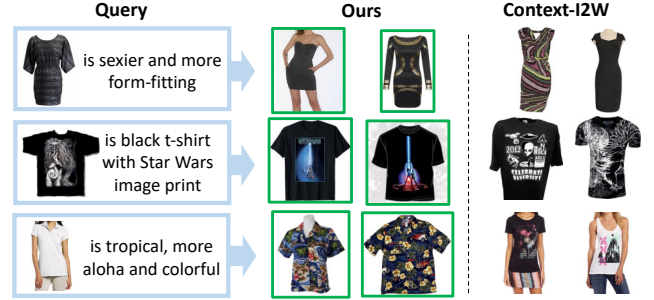


Figure 3. Results of attribute manipulation on FashionIQ.

Since our focus is on studying the prediction-based word mapping for ZS-CIR, we utilize the same prompt in the most recent works [33, 39] for a fair comparison. We show prompt examples for different ZS-CIR tasks. In all examples,  $[*]$  indicates the pseudo token from the mapping network: (a) **Domain conversion** aims to modify the domain of the reference image. The prompt is defined as a [domain tag] of  $[*]$ ; (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  $[*]$ ,  $[\text{obj}_1 \text{ tag}]$  and  $[\text{obj}_2 \text{ tag}]$ , ..., and  $[\text{obj}_n \text{ tag}]$ ; (c) **Sentence manipulation** modifies the reference image based on a sentence. We simply append the sentence with the special token as a photo of  $[*]$ , [sentence].

## 4. Experiments

**Datasets.** We evaluate our model on four ZS-CIR datasets, *i.e.*, COCO [28] and GeneCIS [40] for object/attribute composition, ImageNet [11, 18] for domain conversion, CIRR [29] and CIRCO [5] for object/scene manipulation, and Fashion-IQ [42] for attribute manipulation. All the dataset settings and evaluation metrics (Recall@K and mAP@R) follow the recent works [5, 15, 33] 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/attribute composition. The COCO dataset contains images with corresponding lists of object

Backbones	Methods	R1	R5	R10
ViT-L/14	Pic2Word <sup>†</sup>	23.9	51.7	65.3
	SEARLE-XL <sup>†</sup>	24.2	52.4	66.3
	LinCIR <sup>†</sup>	25.0	53.3	66.7
	Context-I2W <sup>†</sup>	25.6	55.1	68.5
	<b>PrediCIR</b>	<b>27.2</b>	<b>57.0</b>	<b>70.2</b>
ViT-G/14	CompoDiff <sup>†</sup>	26.7	55.1	74.5
	LinCIR <sup>†</sup>	35.3	64.7	76.1
	<b>PrediCIR</b>	<b>37.0</b>	<b>66.1</b>	<b>77.9</b>

Table 2. Results on CIRR for object manipulation.

Backbones	Methods	mAP@5	mAP@10	mAP@25	mAP@50
ViT-L/14	Pic2Word	8.7	9.5	10.6	11.3
	SEARLE-XL <sup>†</sup>	11.7	12.7	14.3	15.1
	LinCIR <sup>†</sup>	12.6	13.6	15.0	15.9
	Context-I2W	13.0	14.6	16.1	17.2
	<b>PrediCIR</b>	<b>15.7</b>	<b>17.1</b>	<b>18.6</b>	<b>19.3</b>
ViT-G/14	CompoDiff <sup>†</sup>	15.3	17.7	19.5	21.0
	LinCIR <sup>†</sup>	19.7	21.0	23.1	24.2
	<b>PrediCIR</b>	<b>23.7</b>	<b>24.6</b>	<b>25.4</b>	<b>26.0</b>

Table 3. Results on CIRCO for object manipulation.

Backbones	Methods	R1	R5	R10
ViT-L/14	Pic2Word <sup>†</sup>	11.5	24.8	33.4
	SEARLE-XL	13.3	28.3	37.6
	LinCIR	11.7	24.9	34.2
	Context-I2W <sup>†</sup>	13.5	28.5	38.1
	<b>PrediCIR</b>	<b>15.1</b>	<b>33.0</b>	<b>42.8</b>
ViT-G/14	LinCIR	14.8	30.6	40.5
	<b>PrediCIR</b>	<b>17.2</b>	<b>34.8</b>	<b>45.9</b>

Table 4. Results on COCO for object composition.

labels and instance masks of query images. Similarly, the GeneCIS dataset introduces four task variations, such as changing a specific attribute or object. We use the prompt (b) in inference. (3) Object/scene manipulation. A reference image is an instruction for manipulating an object or the background. We apply the prompt (c) in inference. (4) Attribute manipulation. This dataset includes various descriptions for manipulating image attributes. We utilize the prompt (c) in inference. Details are in the Supplementary.

**Implementation Details.** We adopt ViT-L/14 CLIP [31] from OpenAI and ViT-G/14 CLIP [20] from OpenCLIP. The crop sizes of random cropped images and blocked target images are the same in the (0.2, 0.25), and the aspect ratios are (0.75, 1.5), respectively. For training PrediCIR, we utilize the Conceptual Caption dataset [34], which comprises 3M images. The number of self-attention blocks is 12 with 384 dimensional embeddings. To improve training stability, we initialize the learnable scalar of tanh-gating to 0 [3]. We employ AdamW [30] with a learning rate of  $1 \times 10^{-5}$ , weight decay of 0.1, and a linear warmup of 10000 steps. The batch size is 1024. All models are trained on 4 NVIDIA A100 (80G) GPUs. To ensure reliable results, we report the performance averaged over three trials.



Figure 4. Results of the object manipulation on CIRR.

#### 4.1. Quantitative and Qualitative Results

We compare PrediCIR with several commonly benchmarked ZS-CIR methods, including: 1) **Pic2Word** [33]: maps a reference image into a pseudo-word token within the CLIP token embedding space. 2) **SEARLE** [5]: Integrates the pseudo-word token with the GPT caption [7]. 3) **Context-I2W** [39]: Selectively extracts text description-relevant visual information before mapping. 4) **LinCIR** [15]: Masks subjects in captions for efficiency training. For a fair comparison, we present the results of methods relying on the ViT-L/14 and ViT-G/14 CLIP models without LLMs [23] or external databases [27, 37]. We also compare with the semi-supervised 5) **CompoDiff** [14]: Training a diffusion model using 18M synthetic data for multi-step entire target image prediction. We report results for CompoDiff on ViT-G/14 CLIP, given its comparable inference times. Since most baselines reported their results on ViT-L/14, we primarily compare results on this backbone and explore the generalization ability of our model on ViT-G/14.

PrediCIR surpasses existing ZS-CIR models on the ViT-L/14 and ViT-G/14 backbones. Tables 1 to 6 present the quantitative results, while Figures 3 to 5 illustrate the corresponding qualitative results of our model and the most recent works, Context-I2W. The attribute manipulation task requires accurately localizing specific attributes within the fashion image. As indicated in Table 1, PrediCIR achieves an average improvement of 2.85% on ViT-L/14, over the State-of-the-Art (SoTA) model, Context-I2W. Context-I2W struggles to retrieve a target image with accurately manipulated fashion attributes, which are missing in the reference images. PrediCIR tackles this challenge by effectively predicting fashion-relevant visual details guided by manipulation text for CIR retrieval. As exemplified in Figure 3, PrediCIR accurately predicts the missing fashion-relevant attribute of sexier and form-fitting (row 1), Star Wars print (row 2), and tropical with more aloha and colorful (row 3).

We further evaluate PrediCIRs' capability in foreground/background differentiation and fine-grained image editing through the object/scene manipulation task (Table 2 and Table 3). PrediCIR consistently surpasses existing ZS-CIR models, achieving an average performance improvement of 1.73% over the best model on CIRR and 2.45%

Backbones	Methods	Conferences	Cartoon		Origami		Toy		Sculpture		Average	
			R10	R50	R10	R50	R10	R50	R10	R50	R10	R50
ViT-L/14	Pic2Word <sup>†</sup>	CVPR 2023	8.0	21.9	13.5	25.6	8.7	21.6	10.0	23.8	10.1	23.2
	SEARLE-XL	ICCV 2023	9.6	24.9	16.1	27.3	7.6	25.4	11.3	26.4	11.2	26.0
	LinCIR	CVPR 2024	9.4	24.2	15.7	26.9	10.8	27.0	11.7	27.9	11.9	26.5
	Context-I2W <sup>†</sup>	AAAI 2024	10.2	26.1	17.5	28.7	11.6	27.4	12.1	28.2	12.9	27.6
	<b>PrediCIR</b>	–	<b>14.2</b>	<b>31.9</b>	<b>20.4</b>	<b>34.3</b>	<b>14.7</b>	<b>30.8</b>	<b>16.3</b>	<b>34.9</b>	<b>16.4</b>	<b>33.0</b>
ViT-G/14	LinCIR	CVPR 2024	13.7	30.2	19.5	32.9	13.8	30.2	15.2	34.0	15.5	31.8
	<b>PrediCIR</b>	–	<b>15.6</b>	<b>34.6</b>	<b>23.7</b>	<b>37.2</b>	<b>17.2</b>	<b>37.5</b>	<b>19.3</b>	<b>37.8</b>	<b>19.0</b>	<b>36.8</b>

Table 5. Results on ImageNet for domain conversion. <sup>†</sup> indicates results from the original paper.

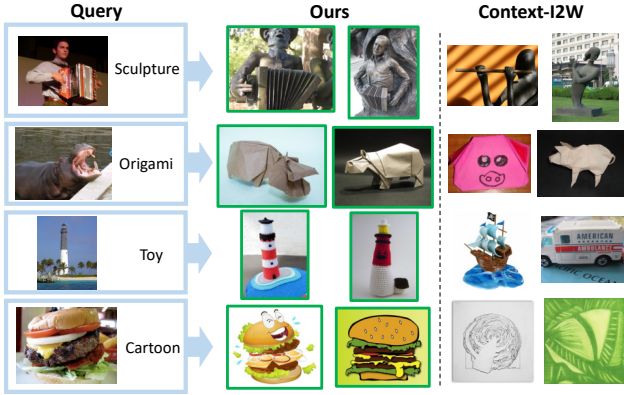


Figure 5. Retrieved results on the domain conversion task.

Backbones	Methods	R1	R2	R3
ViT-L/14	Pic2Word <sup>†</sup>	11.2	21.5	30.4
	SEARLE-XL	12.3	22.1	31.3
	LinCIR	12.2	22.8	32.4
	Context-I2W	12.5	23.2	33.1
	<b>PrediCIR</b>	<b>16.6</b>	<b>26.7</b>	<b>35.8</b>
ViT-G/14	LinCIR <sup>†</sup>	13.7	24.6	34.1
	CompoDiff <sup>†</sup>	15.5	26.6	35.4
	<b>PrediCIR</b>	<b>17.7</b>	<b>28.9</b>	<b>38.6</b>

Table 6. Results on GeneCLS. The average R@1, R@2, R@3 for “Focus Attribute”, “Change Attribute”, “Focus Object”, and “Change Object” are shown. The full table is in the supplementary.

on CIRCO. This improvement is attributed to PrediCIR’s approach of predicting target elements that are missing in reference images guided by manipulation intention before searching and mapping into a pseudo token with fine-grained visual details, enhancing the ability of the CLIP language encoder to compose target image information accurately. In Figure 4, PrediCIR accurately predicts the absent fine-grained visual content of a written notebook (row 1), a bedsheets background (row 2), and a handing hand (row 3).

In the object/attribute composition experiments (Table 4 and 6), PrediCIR significantly outperforms the current SoTA model by an average of 3.60% on COCO and 3.43% on GeneCLS. These results underscore the remarkable effectiveness of our TCP module in predicting missing objects relevant to manipulation text.

Moreover, in the domain conversion experiments (Table 5), PrediCIR consistently outperforms existing approaches

Methods	CIRR			Fashion-IQ	
	R1	R5	R10	R10	R50
1. full model	27.2	57.0	70.2	30.1	52.3
2. w/o cropped images	23.5	53.6	66.0	25.1	45.5
3. w/o action	22.4	52.7	64.9	24.5	43.2
4. w/o Target Predictor	20.2	44.5	56.3	22.5	41.9
5. w/o $\mathcal{L}_2$ loss	22.0	51.5	65.7	24.2	42.8
6. w/o gate	25.9	55.4	67.8	27.5	49.8
7. self-attention	18.2	42.4	55.8	21.3	40.5
8. mask images	22.3	52.2	64.3	24.2	42.8
9. predict entire images	25.1	54.3	66.8	26.5	49.0

Table 7. Ablation study on CIRR and FashionIQ.

and notably surpasses the SoTA Context-I2W by an average of 4.45%. As illustrated in Figure 5, PrediCIR accurately converts image domains guided by manipulation text while maintaining fidelity to the visual content of the reference image (e.g., man playing accordion, a hippo with mouth open, the lighthouse on the island, and juicy burger). In contrast, Context-I2W struggles to map images to other domains as specified by manipulation texts while missing fine-grained details in the contrastive representation.

## 4.2. Ablation Study

Following [5, 33, 39], we evaluate the contribution of the core components in PrediCIR with ViT-L/14 backbone on CIRR and Fashion-IQ in Table 7. (1) **In models ‘2-3’, we evaluate the importance of the world view generation approach.** Using the entire target image as the source view without cropping images (model ‘2’), the performance significantly declined by an average of 4.62%, indicating that a corrupt original image as the source image is essential for learning the ability to predict the target visual content that is missing in the reference image. When removing the action embedding  $a_{x \rightarrow y}$  (model ‘3’) results in a significant drop of 5.82% on average. (2) **In models ‘4-6’, we assess the importance of key modules in the prediction-based image-to-word mapping process.** Removing PTC (model ‘4’) or JEPA framework (model ‘5’) causes obvious performance decrease of 10.28% and 6.12% on average, respectively. By directly summing the predicted and original image features instead of using the gating strategy in PMA (model ‘6’), the performance drops by 2.08%. It indicates the necessity to capture complementary information from the two sources




Figure 6. Analysis of the crop size for source/target view.

adaptively. **(3) Models ‘7-10’ evaluate the effect of alternative solutions for key modules.** In model ‘7’, we replace the PrediCIR with a typical self-attention network with the same input. The results drop significantly by 11.76% on average, confirming the effectiveness of the PrediCIR mapping strategy. In model ‘8’, we employ random masking for constructing source views. The results drop by 6.20% on average, likely due to the frozen CLIP encoder struggling with masked images, whereas our cropped images retain coherent regional context. In model ‘9’, we predict the entire target image, resulting in an average drop of 3.02%, indicating that partial prediction reduces computation and mitigates overfitting. Due to space constraints, please refer to the supplementary (Section A) for more ablation studies.

### 4.3. Analysis

In this subsection, we provide detailed analyses of our design choices, efficiency, and common failure cases.

**Analysis of the Crop Size of World View.** We analyze the influence of crop size for source and target views, as illustrated in Figure 6. A crop size range of (0.15, 0.2) fails to learn sufficient target-relevant missing visual elements due to inadequate context in the source view. Increasing the crop size to (0.45, 0.5) proves redundant, leading to excessive context overlap with the caption. We choose the crop size in the range (0.2, 0.25), which gives the best result among different settings.

**Visualization of Predictor Representations.** Following I-JEPA [2], we freeze our model and train a decoder following the RCDM framework [6] to map the average pool of the predictor outputs back to pixel space. In Figure 7, we show decoder outputs for various random seeds. The PrediCIR correctly predicts the target visual content missing in reference images guided by manipulation texts (e.g., Stars Wars print, open eyes, and origami style). For more details and samples, please refer to our supplementary.

**Effectiveness and Efficiency Analysis.** Our approach achieves significant improvements across six ZS-CIR tasks, with performance gains ranging from 1.73% to 4.45% Over SoTA models. Due to our predictor design for prediction-




Figure 7. Visualization of our predictor representations. Green bounding boxes contain samples from a generative model decoding the output of our pretrained predictor.




Figure 8. Visualization of common failure cases.

based mapping, our model size is larger than the simple MLP mapping of Pic2Word. As a result, in the same setting, our training time (28 hours) is 6 hours longer than Pic2Word and 18 hours longer than SEARLE. Moreover, PrediCIR completes training 203 hours faster than the diffusion-based semi-supervised CompoDiff, achieving significant performance gains. Our inference time(0.03s) is only 0.01s slower than LinCIR and four times faster than CompoDiff (0.12s).

**Discussion on Common Failure Case.** Figure 8 depicts PrediCIR’s common failure cases, particularly with complex and redundant manipulation texts. Challenges include handling multiple objects and attributes (row 1), manipulating objects while converting image domains (row 2), and extensive concurrent manipulation of attributes and scenes (row 3). We believe these difficulties arise from the limitations of the CLIP language encoder in interpreting abstract or redundant intentions for retrieval.

## 5. Conclusion

In this paper, we propose a novel predicted-based image-to-word mapping method that leverages existing image-caption pairs to train a world model for predicting target visual content at latent space that is missing in reference guided by manipulation intention for accurate ZS-CIR. PrediCIR shows strong generalization ability and remarkably improves the best performance of existing approaches on six ZS-CIR tasks. It inspires the vision-to-language alignment mechanism and impacts diverse word modal applications. How to design more lightweight and efficient models with high performance will be the future work.



## Acknowledgements

This work was supported by the Central Guidance for Local Special Project (Grant No. Z231100005923044).

## References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob L Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karén Simonyan. Flamingo: a visual language model for few-shot learning. In *Advances in Neural Information Processing Systems*, pages 23716–23736, 2022. 3
- [2] Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15619–15629, 2023. 3, 8
- [3] Thomas Bachlechner, Bodhisattwa Prasad Majumder, Henry Mao, Gary Cottrell, and Julian McAuley. Rezero is all you need: Fast convergence at large depth. In *Uncertainty in Artificial Intelligence*, pages 1352–1361, 2021. 6
- [4] Alberto Baldrati, Marco Bertini, Tiberio Uricchio, and Alberto Del Bimbo. Effective conditioned and composed image retrieval combining clip-based features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21466–21474, 2022. 2
- [5] Alberto Baldrati, Lorenzo Agnolucci, Marco Bertini, and Alberto Del Bimbo. Zero-shot composed image retrieval with textual inversion. *arXiv:2303.15247*, 2023. 2, 5, 6, 7
- [6] Florian Bordes, Randall Balestriero, and Pascal Vincent. High fidelity visualization of what your self-supervised representation knows about. *arXiv preprint arXiv:2112.09164*, 2021. 8
- [7] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, pages 1877–1901. Curran Associates, Inc., 2020. 6
- [8] Junyang Chen and Hanjiang Lai. Pretrain like you inference: Masked tuning improves zero-shot composed image retrieval. *arXiv preprint arXiv:2311.07622*, 2023. 2, 3
- [9] Yanbei Chen, Shaogang Gong, and Loris Bazzani. Image search with text feedback by visiolinguistic attention learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3001–3011, 2020. 1
- [10] Ritendra Datta, Dhiraj Joshi, Jia Li, and James Z Wang. Image retrieval: Ideas, influences, and trends of the new age. *ACM Computing Surveys*, 40(2):1–60, 2008. 1
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition*, pages 248–255, 2009. 5
- [12] Yongchao Du, Min Wang, Wengang Zhou, Shuping Hui, and Houqiang Li. Image2sentence based asymmetrical zero-shot composed image retrieval. *arXiv preprint arXiv:2403.01431*, 2024. 2
- [13] Quentin Garrido, Mahmoud Assran, Nicolas Ballas, Adrien Bardes, Laurent Najman, and Yann LeCun. Learning and leveraging world models in visual representation learning, 2024. 2, 3
- [14] Geonmo Gu, Sanghyuk Chun, Wonjae Kim, HeeJae Jun, Yoohoon Kang, and Sangdoo Yun. Compodiff: Versatile composed image retrieval with latent diffusion. *arXiv preprint arXiv:2303.11916*, 2023. 2, 6
- [15] Geonmo Gu, Sanghyuk Chun, Wonjae Kim, Yoohoon Kang, and Sangdoo Yun. Language-only efficient training of zero-shot composed image retrieval. In *Conference on Computer Vision and Pattern Recognition*, 2024. 2, 5, 6
- [16] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019. 3
- [17] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 3
- [18] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. 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
- [19] Bohan Hou, Haoqiang Lin, Haokun Wen, Meng Liu, Mingzhu Xu, and Xuemeng Song. Pseudo-triplet guided few-shot composed image retrieval. *arXiv preprint arXiv:2407.06001*, 2024. 2
- [20] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hananeh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. Openclip. 6
- [21] Young Kyun Jang, Dat Huynh, Ashish Shah, Wen-Kai Chen, and Ser-Nam Lim. Spherical linear interpolation and text-anchoring for zero-shot composed image retrieval. *arXiv preprint arXiv:2405.00571*, 2024. 2
- [22] Young Kyun Jang, Donghyun Kim, Zihang Meng, Dat Huynh, and Ser-Nam Lim. Visual delta generator with large multi-modal models for semi-supervised composed image retrieval. In *Proceedings of the IEEE/CVF Conference on*

- Computer Vision and Pattern Recognition*, pages 16805–16814, 2024. 2
- [23] Shyamgopal Karthik, Karsten Roth, Massimiliano Mancini, and Zeynep Akata. Vision-by-language for training-free compositional image retrieval. In *The Twelfth International Conference on Learning Representations*, 2024. 2, 6
- [24] Yann LeCun. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review*, 62(1):1–62, 2022. 3
- [25] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *Proceedings of the 39th International Conference on Machine Learning*, pages 12888–12900, 2022. 3
- [26] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models, 2023. 3
- [27] Haoqiang Lin, Haokun Wen, Xuemeng Song, Meng Liu, Yupeng Hu, and Liqiang Nie. Fine-grained textual inversion network for zero-shot composed image retrieval. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 240–250. ACM, 2024. 2, 6
- [28] 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, 2014. 5
- [29] Zheyuan Liu, Cristian Rodriguez-Opazo, Damien Teney, and Stephen Gould. Image retrieval on real-life images with pre-trained vision-and-language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2125–2134, 2021. 2, 5
- [30] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018. 6
- [31] 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 International Conference on Machine Learning*, pages 8748–8763, 2021. 2, 3, 6
- [32] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2, 3
- [33] Kuniaki Saito, Kihyuk Sohn, Xiang Zhang, Chun-Liang Li, Chen-Yu Lee, Kate Saenko, and Tomas Pfister. Pic2word: Mapping pictures to words for zero-shot composed image retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19305–19314, 2023. 1, 2, 5, 6, 7
- [34] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Annual Meeting of the Association for Computational Linguistics*, pages 2556–2565, 2018. 6
- [35] Jiangming Shi, Yachao Zhang, Xiangbo Yin, Yuan Xie, Zhizhong Zhang, Jianping Fan, Zhongchao Shi, and Yanyun Qu. 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*, pages 11218–11228, 2023. 3
- [36] 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, 2022. 3
- [37] Yucheng Suo, Fan Ma, Linchao Zhu, and Yi Yang. Knowledge-enhanced dual-stream zero-shot composed image retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26951–26962, 2024. 2, 6
- [38] Yuanmin Tang, Xiaoting Qin, Jue Zhang, Jing Yu, Gaopeng Gou, Gang Xiong, Qingwei Ling, Saravan Rajmohan, Dongmei Zhang, and Qi Wu. Reason-before-retrieve: One-stage reflective chain-of-thoughts for training-free zero-shot composed image retrieval, 2024. 2
- [39] Yuanmin Tang, Jing Yu, Keke Gai, Jiamin Zhuang, Gang Xiong, Yue Hu, and Qi Wu. Context-i2w: Mapping images to context-dependent words for accurate zero-shot composed image retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5180–5188, 2024. 2, 5, 6, 7
- [40] Sagar Vaze, Nicolas Carion, and Ishan Misra. Genecis: A benchmark for general conditional image similarity. In *CVPR*, 2023. 5
- [41] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6439–6448, 2019. 1, 2
- [42] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11307–11317, 2021. 5
- [43] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9653–9663, 2022. 3
- [44] Mengjiao Yang, Yilun Du, Kamyar Ghasemipour, Jonathan Tompson, Dale Schuurmans, and Pieter Abbeel. Learning interactive real-world simulators. *arXiv preprint arXiv:2310.06114*, 2023. 2, 3
- [45] Zhenyu Yang, Shengsheng Qian, Dizhan Xue, Jiahong Wu, Fan Yang, Weiming Dong, and Changsheng Xu. Semantic editing increment benefits zero-shot composed image retrieval. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 1245–1254, 2024. 2
- [46] Huaying Zhang, Rintaro Yanagi, Ren Togo, Takahiro Ogawa, and Miki Haseyama. Zero-shot composed image retrieval considering query-target relationship leveraging masked

- image-text pairs. In *2024 IEEE International Conference on Image Processing (ICIP)*, pages 2431–2437. IEEE, 2024. [2](#)
- [47] Kai Zhang, Yi Luan, Hexiang Hu, Kenton Lee, Siyuan Qiao, Wenhui Chen, Yu Su, and Ming-Wei Chang. Magiclens: Self-supervised image retrieval with open-ended instructions, 2024. [2](#)
- [48] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16816–16825, 2022. [3](#)