

Visual Delta Generator with Large Multi-modal Models for Semi-supervised Composed Image Retrieval

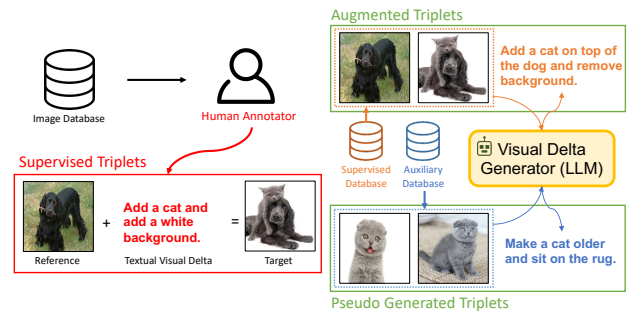
Young Kyun Jang^{*1}, Donghyun Kim^{*2}, Zihang Meng¹, Dat Huynh¹, and Ser-Nam Lim³
 Meta AI¹, Korea University², University of Central Florida³

Abstract

Composed Image Retrieval (CIR) is a task that retrieves images similar to a query, based on a provided textual modification. Current techniques rely on supervised learning for CIR models using labeled triplets of the \langle reference image, text, target image \rangle . These specific triplets are not as commonly available as simple image-text pairs, limiting the widespread use of CIR and its scalability. On the other hand, zero-shot CIR can be relatively easily trained with image-caption pairs without considering the image-to-image relation, but this approach tends to yield lower accuracy. We propose a new semi-supervised CIR approach where we search for a reference and its related target images in auxiliary data and learn our large language model-based Visual Delta Generator (VDG) to generate text describing the visual difference (i.e., visual delta) between the two. VDG, equipped with fluent language knowledge and being model agnostic, can generate pseudo triplets to boost the performance of CIR models. Our approach significantly improves the existing supervised learning approaches and achieves state-of-the-art results on the CIR benchmarks.

1. Introduction

Image-to-image or text-to-image retrieval, where a query image/text is used to retrieve similar ones from a gallery, has grown into a pivotal research field with many practical applications [38]. However, relying solely on image queries is limiting, as they primarily retrieve similar images, making it challenging to understand the user’s intent for modifications in the results. On the other hand, relying solely on text queries can also be restrictive, as it may not effectively convey the user’s desired detailed visual contents. To address this, Composed Image Retrieval (CIR) was introduced [2, 36, 42, 52]. CIR seeks to retrieve images using a query that combines both an image and a textual description of the user’s intent (referred to as the *visual delta*), which allows more flexible retrieval. Due to the convenience and diverse



(a) CIR triplet generation with human supervision (expensive). (b) Visual Delta Generator for generating pseudo triplets.

Figure 1. An illustration of the data preparation process of (a) conventional supervised Composed Image Retrieval (CIR) vs. (b) our proposed semi-supervised CIR. While supervised CIR struggles to scale up due to high annotation costs, our semi-supervised method offers a cost-effective and scalable solution. It augments training samples efficiently by generating pseudo triplets through our Large Language Model (LLM)-based Visual Delta Generator.

applicability of CIR, it has attracted increased attention recently for a variety of real-world applications.

Existing research on CIR has been developed under two major settings: (1) Supervised CIR: Learning with supervised triplets (i.e. \langle reference image, visual delta, target image \rangle) [2, 11, 28, 35, 50] as shown in Fig. 1 (a), and (2) Zero-shot CIR: Learning with massive noisy \langle image, textual caption \rangle pairs [3, 42, 49], without any CIR supervision. Supervised CIR would obviously yield much higher accuracy in retrieval but requires expensive two-stage data collection processes - collecting pairs of related reference and target images, and then annotating them with visual delta that depicts the difference between them. On the other hand, zero-shot CIR does not incur additional labeling costs and utilizes web-collected noisy image text caption pairs directly. However, it has a much lower performance bar compared to supervised approaches and lacks the ability to specialize in specific CIR domain tasks.

In this paper, we investigate a class of CIR called *semi-supervised CIR*, blending supervised and unsupervised samples to enhance generalization (Fig. 1 (b)). This

* Authors contributed equally.

method focuses on boosting CIR performance in specific retrieval domains by creating new triplets from unsupervised data. Building on this concept, we introduce a novel technique, Visual Delta Generator (VDG), designed to tap into the extensive natural language capabilities of Large Language Models (LLMs) [5, 21, 47, 48]. Our approach involves projecting reference and target images from supervised CIR triplets into the language embedding space, making them suitable inputs for the LLM. We then fine-tune the model by using prompts such as ‘Describe the differences between images.’ to induce it to yield human-like visual delta as illustrated in Fig. 2. Furthermore, we employ a parameter-efficient fine-tuning technique, LoRA [19], on the LLM. This choice of design effectively enhances the quality of visual deltas while also preserving the LLM’s original capabilities, without harming its inherent knowledge.

After the VDG is trained, it knows how to distinguish between a given reference and target image and produce visual delta as a textual response. If we forward two similar images with different compositions, we can thus obtain the corresponding visual delta easily with the VDG. This allows us to achieve two purposes. First, we can now augment existing CIR triplets by adding generated visual deltas to pairs of reference and target images from the training set. Second, we can also harvest new reference-target pairs from an unlabeled database based on visual similarity, after which we forward these images to VDG to configure new pseudo triplets for CIR training. Note that our VDG is model agnostic – it simply increases the number of triplet candidates for training any given supervised CIR baselines. This strategy strikes a balance between maintaining the integrity of supervisory concepts derived from a supervised dataset and the capacity for effortless expansion using new, unlabeled image samples. It’s a cost-effective and scalable solution, ensuring uniformity in annotations across extensive datasets. By generating pseudo triplets with VDG, we significantly reduce annotation costs and enhance the performance of CIR models trained solely on supervised learning, as well as those trained without supervised triplets. Our approach leads to state-of-the-art results in CIR benchmarks.

The key advantages of our semi-supervised CIR include:

- To the best of our knowledge, we are the first to transfer knowledge from Large Language Models (LLMs) and connect it with semi-supervised Composed Image Retrieval (CIR).
- We propose a novel Visual Delta Generator (VDG) that generates synthetic visual deltas for augmenting the supervised dataset and allowing the integration of an auxiliary image gallery for CIR model training.
- Comprehensive experimental results confirm the effectiveness of our method, demonstrating state-of-the-art retrieval rankings and showcasing the potential of our work.

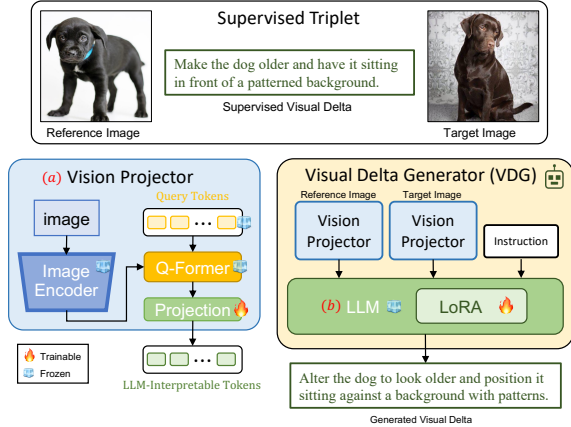


Figure 2. An overview of the VDG tuning process. It includes (a) a vision projector and (b) a Large Language Model (LLM). The VDG is trained to produce visual delta that accurately describes the difference between a reference image and its corresponding target image.

2. Related Work

Composed Image Retrieval. The field of image retrieval has captured the interest of many researchers in our community [6, 38]. One notable area that has seen much progress recently is Composed Image Retrieval (CIR), a problem that focuses on retrieving images that best match a given pair of a query image and textual intent. Supervised CIR methods [2, 11, 35, 50] are trained on human-annotated triplets, consisting of a reference image, a target image, and their textual difference. On the other hand, zero-shot CIR [3, 9, 16, 42] operates without relying on human-guided descriptions of the differences between the two images. Instead, it uses noisy image-text pairs, aiming to find a function that can translate images into words. This approach, designed to discriminate subtle differences between images based solely on text captions, poses challenges, but also scales well, making it easy to add more data for CIR training. Addressing the limitations of both approaches, we explore the field of semi-supervised learning-based retrieval in this work, leveraging the strengths of both labeled and unlabeled data to enhance retrieval performance.

Semi-supervised Learning. Semi-supervised learning has been an active research topic in visual recognition for a long time [4, 18, 20, 25–27, 39, 44, 46, 51]. Semi-supervised learning can be roughly categorized into two groups, consistency learning [25, 39, 46] and pseudo-labeling based learning [27, 44]. Consistency-based methods, as the name implies, encourage consistency in the output of the model by adding noise to model weights or using Exponential Moving Average (EMA) from a teacher model to a student model. In pseudo-labeling based meth-

ods, hard/soft pseudo-labels obtained from a pretrained model are assigned to unlabeled images [15, 44]. These pseudo-labels can be filtered using confidence thresholding or multi-view consistency. However, neither consistency-based nor pseudo labels methods are directly applicable to CIR. This is because the relative textual descriptions of the corresponding visual differences (visual delta) needed in CIR are not the intended outputs of these methods. In this work, we propose the Visual Delta Generator (VDG), a multi-modal pseudo-label generator. VDG processes two input images and generates text that describes their visual differences, making it an ideal candidate for constructing pseudo triplets for CIR.

Multi-modal Models for Image-Text Retrieval. Models like CLIP [41] and BLIP [29] have showcased the advantages of training models on extensive image-text pairs, enabling precise alignment between language and vision representations that is crucial for image-text retrieval. Building on this, there have been significant advancements in utilizing Large Language Models (LLMs) for enhanced vision-language understanding [1, 10, 31, 33, 53]. Especially for CIR, Fromage [23] utilizes LLMs to directly produce embeddings for retrieval, allowing cross-modal compositional search with a single image and textual intent. CoVR [49] and SEARLE [3] apply LLMs to generate visual deltas using image captions without incorporating visual data, which limits the generation of accurate CIR training samples. Our proposed method overcomes these challenges by fine-tuning LLMs with the integration of a pretrained vision-language alignment module. This integration empowers LLMs to perceive and comprehend images, making them applicable even in scenarios with only image datasets. The method excels in generating accurate visual deltas, a key factor in training efficient external CIR models. These enhancements optimize the use of LLMs while ensuring computational efficiency, thereby expanding the versatility of LLMs in image-related tasks.

3. Method

Overview. Our goal is to establish a semi-supervised Composed Image Retrieval (CIR) system that merges image reference features with user textual descriptions to retrieve images from a large-scale database. We face a challenge in the limited availability of supervised triplets necessary for robust CIR model training. To overcome this limitation, we introduce a novel semi-supervised approach for CIR by leveraging an instruction-tuned Large Language Model (LLM), which we call Visual Delta Generator (VDG). The VDG learns to discriminate differences between two images and produces a textual response. This capability allows us to generate additional CIR training triplets, which in turn contributes to the development of more robust CIR models.

Sec. 3.1 provides detailed insights into the construction of the VDG. Sec. 3.2 describes the pseudo triplet generation process. Training of CIR models with pseudo triplets and our additional objective function for better optimization are described in Sec. 3.3.

3.1. Visual Delta Generator Training

While semi-supervised learning methods have been actively developed for standard visual recognition tasks, these cannot be directly applied to CIR. In CIR, pseudo-label generation requires a detailed semantic understanding of two separate images such that their difference can be automatically expressed in the form of text. We leverage vision-language pretraining models and LLMs to achieve the requirements. With the huge success of LLMs, there are approaches that aim to utilize their understanding of the language domain for improving vision tasks. Particularly, LLaVA [33] and InstructBLIP [10] which are trained on top of the chat-bot style instruction tuned LLM, Vicuna [7], have shown interesting results on vision-language tasks. Inspired by these, we propose VDG, which allows the LLM to take two images (reference, target) with similar contents and discriminate their difference in the form of text response (visual delta) as shown in Fig. 2.

First, to enable the LLM to interpret images, we use the Querying Transformer (Q-Former) motivated by InstructBLIP [10] to prepare images for the LLM input (*i.e.*, *Vision Projector* (VP) in Fig. 2(a)). Q-Former, a transformer encoder [12, 21], processes a fixed set of 32 learnable query tokens (embeddings). These tokens are modified through self-attention layers and interact with image features of Vision Transformer (ViT) [12] through cross-attention layers. As a result, the query tokens are infused with the visual information from the provided image, making them suitable for LLM processing.

Stage 1: Alignment. The outputs of the VP are inherently not aligned with the tokenized word embeddings of the inputs to the LLM, which we chose to be LLaMA2 [48] in this work. Alignment between the VP and LLM needs to be attained by fine-tuning trainable projection of VP in Fig. 2(a) as was done in LLaVA [33]. Specifically, we employ large-scale image-caption pairs to foster alignment between the image representations understood by the VP and the LLM. This step includes minimizing standard next token prediction loss which is generally used to train the decoder-based LLM as:

$$L_{LLM} = - \sum_{t=1}^T \log P(w_t | x, I_{inst}, w_{(1, \dots, t-1)}; \theta_{proj}), \quad (1)$$

where w_t denotes a token at time step t , T is the length of the sequence, $P(\cdot)$ is the probability assigned by the model

to the actual token w_t given with the image x , instruction I_{inst} and previous tokens $w_{(1,\dots,t-1)}$, and θ_{proj} is parameters of projection layer. Following standard practice, the prediction for each token is computed using a softmax over the vocabulary of the LLM, and the cross-entropy loss is computed between the predicted probabilities and the textual token labels as a classification.

Stage 2: Instruction Tuning. In this stage, we conduct instruction tuning to equip our LLM with the ability to understand image pairs and produce visual delta, as shown in Fig. 2(b). Referring to Fig. 3, we train the LLM with an instruction (i.e., “Request”) to generate the visual delta in textual format (i.e., “Response”). We forward two distinct images (i.e., “Reference” and “Target”) into the LLM via the VP. We employ a parameter-efficient fine-tuning technique, LoRA [19], directly on the LLM. This fine-tuning process is designed to provide “extra room” for the LLM to undertake new tasks, all without compromising its foundational capabilities. In this stage, we introduce a specific prompt structure to the LLM as below:

Request: Analyze given reference and target images and provide a description that transforms the reference to match the target.
Reference: <Reference image query tokens>
Target: <Target image query tokens>
Response: <Visual delta (text label)>

Figure 3. A template prompt for VDG instruction tuning.

which guides the LLM to generate the corresponding visual delta. We leverage the same training loss from Eqn. 1 to train LoRA parameters.

3.2. Pseudo Triplet Generation for CIR

After training VDG, we can generate visual delta of two images, which can be used to form pseudo triplets for CIR. However, it is important to choose two images that not only share certain attributes and similarities but also present other distinct attributes (i.e., visual delta). To gather suitable pairs of reference and target images, we utilize an image encoder to select them from an auxiliary image gallery, which we denote as G' . Note that, we notate *upper strophe* ($'$) on samples and embeddings that are obtained from G' , in the following.

Following the strategy in CIR [36], we start with an anchor image x_a and retrieve the top 20 images from G' using cosine similarity scores between ResNet 152 [17] embeddings, pretrained on ImageNet [24]. We exclude images with scores above 0.94 and sequentially add images to a subgroup of size 6, skipping those within a 0.002 score of the previously included image. As depicted in Fig. 4, we establish dense connections between all pairs (both consecutive, represented by the outer circle, and non-consecutive,




Figure 4. The process of pseudo triplet generation. First, an image subgroup is constructed based on visual similarity (left). Then, paired reference and target images are fed into the VDG to generate the visual delta, completing the triplet formation (right).

represented by the dotted inner connections), while ensuring no overlaps. The arrow’s starting point denotes the *reference*, while its endpoint indicates the *target*. Given images x_i^r and x_j^t , VDG produces $\delta'_{i,j}$. This allows us to formulate the pseudo triplet as $\{x_i^r, x_j^t, \delta'_{i,j}\}$.

3.3. Semi-supervised CIR Training

Preliminaries. Suppose we have access to a CIR dataset of triplets: $D = \{(x_i^r, x_j^t, \delta_{n(i,j)})\}_{n=1}^N$ where x_i^r, x_j^t denote reference and target images, respectively, $\delta_{n(i,j)}$ represents their visual delta, and N denotes the total triplet counts.

In the pursuit of enhancing CIR through a semi-supervised approach, our method’s strength lies in its model-agnosticism, allowing for seamless integration with a variety of CIR models. Following recent trends, we opt to use the encoders from vision-language pretraining models [10, 29, 31, 41], notably CLIP and BLIP, as our baseline backbones. These encoders are naturally equipped to understand and convert both visual and textual elements into a joint embedding space, making them suitable for CIR tasks.

The image encoder takes patchified image token embeddings $x = [x_1, \dots, x_{K_{img}}]$ with the learnable image *cls* token embedding $[x_{cls}]$, and outputs the visual feature embeddings $E_{img}(x_{cls}, x) = [\hat{x}_{cls}, \hat{x}_1, \dots, \hat{x}_{K_{img}}]$ where $\hat{x} \in R^{d_i}$ of d_i dimension, and K_{img} denotes the number of generated tokens from each image. The text encoder processes tokenizer output embedding of visual delta $\delta = [z_1, \dots, z_{K_{txt}}]$ with the learnable text *cls* token embedding $[z_{cls}]$ to produce its textual feature embeddings $E_{txt}(z_{cls}, \delta) = [\hat{z}_{cls}, \hat{z}_1, \dots, \hat{z}_{K_{txt}}]$, where $\hat{z} \in R^{d_t}$ of d_t dimension, and K_{txt} denotes the number of generated text tokens from each visual delta.

Model Architecture. To carry out CIR, we establish a fusion function f that takes a reference image x^r and a visual delta δ , producing a composed embedding c as: $f(x^r, \delta) = c$, and notates its trainable components as f_θ . We utilize two well-known backbones, CLIP and BLIP, to configure f . In the case of CLIP, we employ the text encoder $E_{txt;\theta}$ and an

additional Combiner module [2] $\mathcal{C}_\theta(\hat{x}_{cls}^T, \hat{z}_{cls})$ that outputs \mathbf{c} to be the components of f_θ (i.e., $f_\theta^{\text{CLIP}} = \{E_{\text{txt}:\theta}, \mathcal{C}_\theta\}$). The Combiner is designed to optimally blend \hat{x}_{cls} and \hat{z}_{cls} carefully weighing their individual impacts while adeptly mixing them. In the BLIP case, we exclusively use BLIP’s text encoder, grounded in the image, and designate it as the trainable $E_{\text{txt}:\theta}$ (i.e., $f_\theta^{\text{BLIP}} = E_{\text{txt}:\theta}$, and $\mathbf{c} = \hat{z}_{cls}$), without incorporating additional modules. BLIP’s text encoder inherently fuses image and text signals in its cross-attention layer, eliminating the need for a separate combining module. Notably, we freeze all vision encoders in our setup to ensure compatibility with existing image retrieval galleries and to enhance training efficiency.

Supervised / Pseudo Separated Contrastive Loss. The training objective of CIR is to achieve strong alignment between the target image’s embedding \mathbf{x} (where $\mathbf{x} = \hat{x}_{cls}^t$ for simplicity), and the composed embedding \mathbf{c} . On this purpose, we utilize HN-NCE [40] loss for a given training batch $\mathcal{B} \sim D$ and $\mathcal{B}' \sim D'$, where D' contains pseudo triplets. Additionally, to mitigate the impact of noise in pseudo triplets and ensure consistent contributions from supervised triplets, we compute a *target-composed contrastive loss* (tcc) as:

$$\mathcal{L}_{tcc}(\mathcal{B}, \mathcal{B}') = \mathcal{L}_c(\mathcal{B}; \tau) + \mathcal{L}_c(\mathcal{B} \oplus \mathcal{B}'; \tau) \quad (2)$$

where \mathcal{L}_c is defined as:

$$\mathcal{L}_c(\mathcal{B}) = - \sum_{i=1}^n \log \left[\frac{e^{\mathbf{x}_i^T \mathbf{c}_i / \tau}}{\alpha \cdot e^{\mathbf{x}_i^T \mathbf{c}_i / \tau} + \sum_{j \neq i}^{\mathcal{B}} e^{\mathbf{x}_i^T \mathbf{c}_j / \tau} \cdot w_{\mathbf{x}_i, \mathbf{c}_j}} \right] - \sum_{i=1}^n \log \left[\frac{e^{\mathbf{x}_i^T \mathbf{c}_i / \tau}}{\alpha \cdot e^{\mathbf{x}_i^T \mathbf{c}_i / \tau} + \sum_{j \neq i}^{\mathcal{B}} e^{\mathbf{x}_i^T \mathbf{c}_j / \tau} \cdot w_{\mathbf{x}_j, \mathbf{c}_i}} \right].$$

where τ, α denote hyper-parameters, and \oplus denotes concatenation along the batch axis, and $w_{\mathbf{x}_i, \mathbf{c}_j}, w_{\mathbf{x}_j, \mathbf{c}_i}$ are set as in [40]. This design facilitates the independent yet concurrent investigation of both supervised and semi-supervised CIR embedding spaces in a contrastive manner.

Target-Delta Matching Loss. In the case of BLIP’s image-grounded text encoder structure, we introduce a new *target-delta matching loss* (tdm). Our insight stems from the observation that, while all reference image patch tokens are considered in the cross-attention layer of the BLIP text encoder, target image patch tokens are overlooked when training solely with contrastive learning between cls token embeddings. As an example shown in Fig. 5, the visual delta can be seen as a *weakly correlated caption* to the target image. Thus, we aim to align the target image and visual delta to enable the BLIP text encoder to identify the image tokens related to textual input. The tdm loss is applied as:

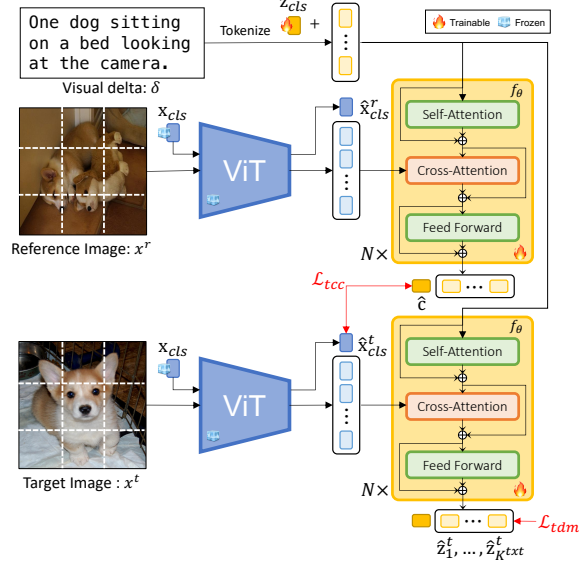


Figure 5. Illustration of our proposed adaptation of the BLIP image-grounded text encoder for CIR. Both reference (x^r) and target image (x^t) patch tokens are processed by the text encoder (f_θ).

$$\mathcal{L}_{tdm} = \mathbb{E}_{(x^t, \delta) \sim (D, D')} \mathbf{H}(\mathbf{y}^{tdm}, \mathbf{p}^{tdm}(x^t, \delta)) \quad (3)$$

where \mathbf{H} is cross entropy, \mathbf{y}^{tdm} is a 2-dimensional one-hot vector label obtained through the hard negative mining process proposed in [30], where a pair of image and text is configured as matched or not. $\mathbf{p}^{tdm} \in \mathbb{R}^2$ generates both positive and negative target-delta matching scores and is computed as:

$$\mathbf{p}^{tdm}(x^t, \delta) = \sum_{i=1}^{K^{t:t}} p_\theta(\hat{z}_i^t) / K^{t:t},$$

where \hat{z}_i^t denotes the i -th token embedding in $f_\theta(x^t, \delta)$ and p_θ is a FC layer that outputs a 2-D vector. This loss ensures the text encoder fully processes the target image features, improving its understanding of the visual delta.

4. Experiments

Sec. 4.1 details the setup and VDG generation. Sec. 4.2 covers quality checks of visual deltas. Evaluations and baseline comparisons are in Sec. 4.3, and further analyses in Sec. 4.4.

4.1. Setup

Implementation Details. We utilize InstructBLIP’s pre-trained weights [10], which have been trained with both a ViT-G/14 [13] and the Q-Former based on the Vicuna-13B model [7], without using prompts for Q-Former. For the stage 1 training described in Sec. 3.1, we employ 595K filtered image-text pairs from CC3M [43] provided by [33] to

find alignment with our baseline LLM, LLaMA2-13B [48]. In Stage 2, we implement instruction tuning with LoRA parameters [19], following the fixed prompt as outlined in Fig. 3. Visual deltas are generated in an autoregressive manner, predicted based on the LLaMA2 vocabulary. For the CLIP-based CIR training, we select the ViT-L/14 model combined with a Combiner [2]. For the BLIP-based model [29], we use a dedicated BLIP text encoder for image-text matching, paired with the ViT-L/16 model. Additional details can be found in the appendix.

Datasets for CIR Evaluation. There are two standard benchmarks in CIR, one is CIRR [36] which deals with natural images, and the other is FashionIQ [52] which focuses on fashion domain images. Each presents unique challenges and datasets that help researchers push the boundaries of what’s possible in CIR. Following the protocols utilized in benchmarks [36, 52], we report the CIR results with recall scores at top K retrieval results (R@K), or results under collected subset (R_s@K). Specifically, CIRR is configured with 4,351 subgroups (subsets), each containing six similar images, sourced from NLVR2 [45]. For experimental purposes these groups are distributed into train (3,345 subgroups of 16,742 images), validation (503 subgroups of 2,265 images), and test (503 subgroups of 2,178 images) sets. FashionIQ is divided into three categories of Dress, Shirt, and Toptee (Tops and Tees). The reference and target images are paired based on their category similarities. The 18,000 CIR triplets for training are pooled from 45,429 images of the training set, and 6,016 CIR triplets for the test are chosen from 15,415 images of the validation set.

Visual Delta Generation. To produce pseudo triplets, we expand our dataset sources to configure an *auxiliary gallery* (G') which is built upon the grouping strategy introduced in Sec. 3.2, while *excluding images that overlap* with the benchmark sets. Once the subgroups are constructed, we further filter them to avoid heavy overlap. In total, we draw upon 42,390 unique subgroups from NLVR2 [45], and 79,427 from COCO [32]. Similarly, we build 27,957 individual subgroups from FashionIQ [52], and 30,880 from DeepFashion [34]. We mark *upper strophe* (\prime) to these datasets. For the semi-supervised settings, we randomly select 3,345 groups from NLVR2 \prime and COCO \prime for the CIRR case, as well as 3,600 groups from FashionIQ \prime and DeepFashion \prime . We denote these datasets used for semi-supervised CIR with the subscript *se*, (e.g., NLVR2 \prime_{se} , COCO \prime_{se}). Fig. 6 shows the comparison between human-annotated visual deltas and those generated by the VDG in both natural and fashion domains. We observe that the VDG is effective in generating high-quality visual deltas – additional results can be found in the appendix.

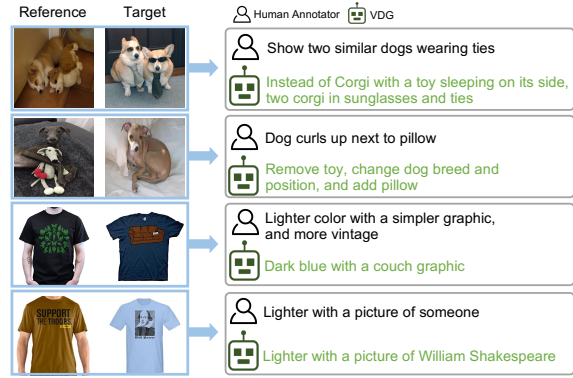


Figure 6. Qualitative comparison on visual deltas, human vs. VDG on CIRR and FashionIQ datasets. For both natural and fashion domain images, VDG can produce informative visual deltas.

Table 1. Retrieval results on CIRR validation set. Human + VDG represents utilizing both human annotated and VDG-generated visual deltas when training the CIR model. BLIP_{tdm} represents a model with target-delta matching loss. The best scores for each Val. set supervision are highlighted in bold.

Training Set Supervision	Val. Set Supervision	Baseline	R@1	R@5	R@10	R@50
(a) Human	Human	Combiner	37.98	71.49	82.52	95.29
		BLIP _{tdm}	53.17	82.09	89.81	97.54
(b) VDG	Human	Combiner	35.47	68.29	80.29	94.31
		BLIP _{tdm}	50.16	80.03	87.78	96.75
(c) Human + VDG	Human	Combiner	39.11	73.02	84.41	95.96
		BLIP _{tdm}	53.67	82.99	89.97	97.84
(d) Human	VDG	Combiner	38.96	73.31	84.36	96.22
		BLIP _{tdm}	51.64	83.33	91.39	97.90
(e) VDG	VDG	Combiner	41.59	77.57	87.35	97.11
		BLIP _{tdm}	52.69	85.17	92.37	98.59
(f) Human + VDG	VDG	Combiner	41.47	77.90	87.42	97.30
		BLIP _{tdm}	52.95	85.27	92.54	98.61

4.2. Quality Check of VDG Responses

To assess the quality of VDG-generated visual deltas, we executed a series of experiments as outlined in Table 1. We use two backbones, Combiner, and BLIP with the tdm loss (denoted as BLIP_{tdm}), as our baselines for evaluation.

Initially, we compare the visual deltas generated by VDG with those annotated by humans. This is done by replacing the deltas in the training set with VDG-generated deltas for the same reference-target pairs. When comparing (a) and (b), we observe only a marginal drop in performance upon switching to VDG-generated deltas. This indicates that the deltas generated by VDG are as effective as those created by humans. More importantly, the improved performance observed when comparing (a) with (c) — which combines human-annotated and VDG-generated deltas — highlights the effectiveness of this hybrid approach as a robust data augmentation strategy to enhance CIR models.

To evaluate VDG’s performance on new image pairs, we conduct experiments by substituting the human-annotated

Table 2. Retrieval results on *CIRR test set*. * denotes our baselines, † denotes VDG generated visual deltas are applied to augmented original training set. We categorize into two distinct groups: one is *Seen: Supervised / Supervised + External / Supervised + Aux. with VDG* which utilize human-annotated visual delta for CIR model training, and the other is *Unseen: Zero-shot / Aux. with VDG* which does not utilize human-annotated visual delta for CIR model training. Within each category, best viewed with **bold**.

Method	R@1	R@5	R@10	R@50	R _s @1	R _s @2	R _s @3
<i>(a) Supervised</i>							
ARTEMIS [11]	16.96	46.10	61.31	87.73	39.99	62.20	75.67
CIRPLANT [35]	19.55	52.55	68.39	92.38	39.20	63.03	79.49
Combiner [2]	33.59	65.35	77.35	95.21	62.39	81.81	92.02
CASE [28]	48.00	79.11	87.25	97.57	75.88	90.58	96.00
CoVR [49]	48.84	78.05	86.10	94.19	75.78	88.22	92.80
Combiner	34.39	66.22	76.58	91.04	68.55	86.36	93.98
Combiner†	36.91	69.21	79.54	92.04	70.00	87.45	94.39
BLIP _{tdm}	48.94	77.83	86.15	94.17	75.71	89.71	95.81
BLIP _{tdm} †	49.08	78.98	86.89	94.24	76.18	90.62	95.86
<i>(b) Supervised + External Dataset for Pretraining</i>							
CASE + LasCo.Ca. [28]	49.35	80.02	88.75	97.47	76.48	90.37	95.71
CoVR + WebVid [49]	49.69	78.60	86.77	94.31	75.01	88.12	93.16
<i>(c) Supervised + Auxiliary Gallery with VDG</i>							
Combiner† + COCO _{se}	38.77	69.25	79.21	91.52	71.25	87.49	94.34
Combiner† + NLVR2 _{se}	38.89	69.84	79.41	91.18	71.92	87.89	94.29
BLIP _{tdm} † + COCO _{se}	49.37	78.12	85.52	93.74	76.68	90.46	96.05
BLIP _{tdm} † + NLVR2 _{se}	50.96	80.15	86.86	94.46	77.45	90.65	96.10
<i>(d) Zero-shot</i>							
Pic2Word by CC3M [42]	23.90	51.70	65.30	87.80	-	-	-
SEARLE by ImageNet [3]	24.22	52.41	66.29	88.63	53.71	74.63	87.61
CASE by LasCo.Ca. [28]	35.40	65.78	78.53	94.63	64.29	82.66	91.61
CoVR by WebVid [49]	38.48	66.70	77.25	91.47	69.28	83.76	91.11
<i>(e) Auxiliary Gallery with VDG</i>							
Combiner by COCO'	26.80	54.05	65.30	83.88	67.28	85.42	92.99
Combiner by NLVR2'	31.57	61.37	72.10	88.94	67.98	86.18	93.18
BLIP _{tdm} by COCO'	43.49	72.07	81.59	93.21	72.36	87.98	94.72
BLIP _{tdm} by NLVR2'	45.74	75.01	82.52	93.13	72.60	87.90	94.77

deltas in the validation set with deltas generated by VDG. The comparison of (a) and (d) reveals comparable performance levels, suggesting that deltas generated by VDG are as effective as human annotations. When we incorporate VDG-generated deltas into the training set (cases (e) and (f)), we achieve a notable performance improvement over the human-annotated validation set across all metrics, with the exception of a small reduction in R@1. For example, there is a 5.14%p increase in R@5 between (b) and (e). The improvement observed may stem from the varied text descriptions by different human annotators, leading to inconsistencies in style and detail that potentially affect the distribution of visual differences between the annotated triplets. VDG offers a more consistent solution in generating visual deltas, thereby reducing errors associated with differences in human annotation. Our comprehensive testing demonstrates VDG's reliability as an annotation tool for CIR.

4.3. Comparison with Other Methods

Tables 2 and 3 present comparative evaluations between our method and other CIR approaches on CIRR and FashionIQ evaluation protocols.

Table 3. Retrieval results on *FashionIQ validation set*. We utilize the same notations and same categorizations as Table 2.

Method	Dress		Shirt		Toptee	
	R@10	R@50	R@10	R@50	R@10	R@50
<i>(a) Supervised</i>						
ARTEMIS [11]	27.16	52.40	21.78	43.64	29.20	53.83
DCNet [22]	28.95	56.07	23.95	47.30	30.44	58.29
FashionVLP [14]	32.42	60.29	31.89	58.44	38.51	68.79
Combiner [2]	31.63	56.67	36.36	58.00	38.19	62.42
CoVR [49]	43.51	67.94	48.28	66.68	51.53	73.60
*Combiner	31.95	55.05	39.21	56.82	38.55	62.16
*Combiner†	35.40	59.99	42.30	61.63	43.09	66.96
*BLIP _{tdm}	44.87	66.83	49.61	66.93	50.54	72.26
*BLIP _{tdm} †	46.90	68.86	50.28	68.04	52.73	74.45
<i>(b) Supervised + External Dataset for Pretraining</i>						
CASE + LasCo.Ca. [28]	47.44	69.36	48.48	70.23	50.18	72.24
CoVR + WebVid [49]	44.55	69.03	48.43	67.42	52.60	74.31
<i>(c) Supervised + Auxiliary Gallery with VDG</i>						
Combiner† + DeepFashion _{se}	35.50	60.09	42.54	62.86	43.27	67.86
Combiner† + FashionIQ _{se}	36.30	60.19	43.98	62.27	44.33	68.06
BLIP _{tdm} † + DeepFashion _{se}	47.10	69.10	49.95	69.96	53.90	74.35
BLIP _{tdm} † + FashionIQ _{se}	47.89	69.81	51.36	71.08	53.29	74.65
<i>(d) Zero-shot</i>						
Pic2Word by CC3M [42]	20.00	40.20	26.20	43.60	27.90	47.40
SEARLE by ImageNet [3]	20.32	43.18	27.43	45.68	29.32	50.17
CoVR by WebVid [49]	21.95	39.05	30.37	46.12	30.78	48.73
<i>(e) Auxiliary Gallery with VDG</i>						
Combiner by DeepFashion'	23.30	46.36	30.86	49.02	31.87	51.96
Combiner by FashionIQ'	28.26	51.46	32.58	51.28	34.88	55.79
BLIP _{tdm} † by DeepFashion'	32.67	54.39	35.48	55.05	39.47	59.92
BLIP _{tdm} † by FashionIQ'	37.48	58.70	37.29	57.11	42.12	62.32

We categorize the experiments into two distinct groups. The first group, *Seen*, includes scenarios where the CIR model is trained with human-annotated training triplets. In contrast, the second group, *Unseen*, comprises scenarios in which the CIR model is trained without human-annotated triplets. Our implementations of Combiner and BLIP are denoted by grey box. Specifically, methods under (a) are trained with supervised triplets, while those notated with the symbol (†) are additionally augmented with VDG-generated visual deltas from supervised image pairs (*i.e.*, the same setting as Human + VDG in Table 1). The methods in (b) have been pretrained on external datasets for CIR tasks, resulting in enhanced performances. In our semi-supervised setup (c), the models are trained with augmented supervised triplets and pseudo triplets sub-sampled from the auxiliary gallery as described in Secs. 3.2 and 4.1. Methods in (d) are trained with large-scale datasets with the intention of producing zero-shot performances. In (e), we only use pseudo triplets generated from the auxiliary gallery for training.

The results clearly demonstrate that our VDG implementation significantly improves the performance of both Combiner and BLIP_{tdm} baselines. In (a), a detailed comparison shows that models augmented with VDG samples (with †) consistently outperform their counterparts (without †) in all

Table 4. Ablation results on BLIP_{tdm} baseline for CIRR validation set with NLVR2_{se} as auxiliary gallery. Best viewed with bold.

grouping	concat.	tdm loss	R@1	R@5	R@10	R@50
			52.45	81.18	88.62,	96.88
✓			55.66	83.45	91.21	97.33
✓	✓		56.28	84.39	91.41	97.92
✓	✓	✓	57.88	85.58	93.21	98.33




Figure 7. Analysis on the scale for (c) Supervised + Auxiliary Gallery with VDG of NLVR2_{se} and FashionIQ_{se}. The x-axis denotes the ratio of additional images to training images.

Table 5. Analysis on the scale for (e) Auxiliary Gallery with VDG. The dataset is scaled from one eighth (1/8) to the full set (1). Evaluated on CIRR Test set.

	R@1	1/8	1/4	1/2	1	R@10	1/8	1/4	1/2	1
COCO	42.15	42.89	43.10	43.49		79.81	81.01	81.16	81.59	
NLVR2	44.31	44.99	45.13	45.74		82.48	82.75	82.43	82.52	

recall metrics. Moreover, in (c), when we enhance our CIR baselines with additional pseudo triplets, there is a notable performance improvement. Specifically, when BLIP_{tdm} is combined with the auxiliary gallery, it surpasses the previous state-of-the-art results in most recall metrics, showcasing the advantages of VDG. Additionally, it’s important to highlight that our semi-supervised CIR method achieves these improvements with considerably fewer images than methods like CASE or CoVR, demonstrating its efficiency.

In the unseen scenario, our implemented Combiner baseline not only outperforms Pic2Word, which uses the same CLIP backbone, but our BLIP_{tdm} also significantly exceeds the former best-performing SEARLE or CoVR in most metrics. VDG’s key advantage is its ability to generate visual deltas that align with the targeted domain, such as using COCO for natural images and DeepFashion for fashion images, by solely utilizing images without any accompanying captions. As a result, VDG significantly reduces the number of training samples required, yet achieves enhanced retrieval performances compared to zero-shot approaches like Pic2Word, CASE, or CoVR. This demonstrates that by focusing on similar domain, image-only datasets, we can substantially improve the efficacy of CIR model training.

4.4. Further Analyses

Ablation study on Each Component. To showcase the effectiveness of our proposed learning strategies, we conducted an ablation study and presented the findings in Table 4. We consider three scenarios: (1) without grouping, where we randomly sample images from the auxiliary gallery to create pairs; (2) without concat., where we remove the concatenation step used in Eqn. 2 and use pseudo triplets only;

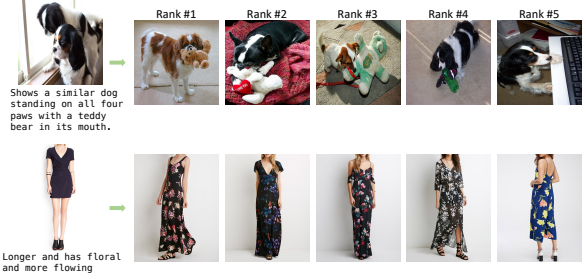


Figure 8. Retrieval results on auxiliary galleries, COCO, and DeepFashion. Actual user intents are used as text queries.

and (3) without tdm loss, where we exclude the tdm loss during training. The results indicate that each element sufficiently contributes to enhancing the performance.

Impact of Generated Data Scale. Figs. 7a and 7b demonstrate the effect of the size of VDG-generated data. In a semi-supervised setup, we find that the performance peaks when the number of pseudo-selected images approximates the size of the training set, and then saturates. This saturation might arise because additional pseudo triplets fail to represent the test sample distribution, particularly as we rigorously remove overlapping images with the test set when forming G' . For the unseen case as shown in Table 5, performance is enhanced with an increase in VDG-generated data. This suggests that even in the absence of supervised triplets for CIR model training, a larger set of pseudo triplets not only expands the model’s comprehension of visual compositions but also aligns more closely with the distribution of supervised triplets. This alignment facilitates improved targeted domain retrieval performances, striking a balance between generalization and performance enhancement.

Qualitative Retrieval Results. We facilitate retrieval by providing user intents along with query images, drawing from auxiliary galleries. The retrieval results, as depicted in Fig. 8, demonstrate accurate and relevant image retrieval. Additional results are detailed in the appendix.

5. Conclusion

In this study, we investigate a novel semi-supervised learning approach in the context of Composed Image Retrieval (CIR). Our findings reveal that integrating human-annotated data with pseudo triplets generated by the Visual Delta Generator (VDG) significantly enhances the generalization capacity of CIR models. The VDG approach not only streamlines the generation of visual deltas but also emerges as a cost-efficient and effective alternative to extensive human annotation. This work paves the way for future research in semi-supervised learning and showcases VDG as a promising direction for advancing CIR systems.

Acknowledgment: This work was partially supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2019-0-00079, Artificial Intelligence Graduate School Program(Korea University)).

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *NeurIPS*, 35: 23716–23736, 2022. [3](#)
- [2] Alberto Baldrati, Marco Bertini, Tiberio Uricchio, and Alberto Del Bimbo. Effective conditioned and composed image retrieval combining clip-based features. In *CVPR*, 2022. [1](#), [2](#), [5](#), [6](#), [7](#)
- [3] Alberto Baldrati, Lorenzo Agnolucci, Marco Bertini, and Alberto Del Bimbo. Zero-shot composed image retrieval with textual inversion. In *ICCV*, 2023. [1](#), [2](#), [3](#), [7](#)
- [4] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. *NeurIPS*, 32, 2019. [2](#)
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 33:1877–1901, 2020. [2](#)
- [6] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *ECCV*, pages 104–120. Springer, 2020. [2](#)
- [7] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, 2023. [3](#), [5](#)
- [8] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*, 2022. [2](#)
- [9] Niv Cohen, Rinon Gal, Eli A Meir, Gal Chechik, and Yuval Atzmon. “this is my unicorn, fluffy”: Personalizing frozen vision-language representations. In *ECCV*, pages 558–577. Springer, 2022. [2](#)
- [10] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning, 2023. [3](#), [4](#), [5](#)
- [11] Ginger Delmas, Rafael Sampaio de Rezende, Gabriela Csurka, and Diane Larlus. Artemis: Attention-based retrieval with text-explicit matching and implicit similarity. In *ICLR*, 2022. [1](#), [2](#), [7](#)
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [3](#)
- [13] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. Eva: Exploring the limits of masked visual representation learning at scale. In *CVPR*, pages 19358–19369, 2023. [5](#)
- [14] Sonam Goenka, Zhaoheng Zheng, Ayush Jaiswal, Rakesh Chada, Yue Wu, Varsha Hedau, and Pradeep Natarajan. Fashionvlp: Vision language transformer for fashion retrieval with feedback. In *CVPR*, pages 14105–14115, 2022. [7](#)
- [15] Yves Grandvalet and Yoshua Bengio. Semi-supervised learning by entropy minimization. *NeurIPS*, 17, 2004. [3](#)
- [16] Geonmo Gu, Sanghyuk Chun, Wonjae Kim, HeeJae Jun, Yooheon Kang, and Sangdoon Yun. Compodiff: Versatile composed image retrieval with latent diffusion. *arXiv preprint arXiv:2303.11916*, 2023. [2](#)
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [4](#)
- [18] Seunghoon Hong, Hyeonwoo Noh, and Bohyung Han. Decoupled deep neural network for semi-supervised semantic segmentation. *NeurIPS*, 2015. [2](#)
- [19] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *ICLR*, 2021. [2](#), [4](#), [6](#)
- [20] Young Kyun Jang and Nam Ik Cho. Generalized product quantization network for semi-supervised image retrieval. In *CVPR*, pages 3420–3429, 2020. [2](#)
- [21] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pages 4171–4186, 2019. [2](#), [3](#)
- [22] Jongseok Kim, Youngjae Yu, Hoesong Kim, and Gunhee Kim. Dual compositional learning in interactive image retrieval. In *AAAI*, pages 1771–1779, 2021. [7](#)
- [23] Jing Yu Koh, Ruslan Salakhutdinov, and Daniel Fried. Grounding language models to images for multimodal inputs and outputs. 2023. [3](#)
- [24] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *NeurIPS*, 25, 2012. [4](#)
- [25] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016. [2](#)
- [26] Semi-Supervised Learning. Semi-supervised learning. *CSZ2006.html*, 2006.
- [27] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, page 896. Atlanta, 2013. [2](#)
- [28] Matan Levy, Rami Ben-Ari, Nir Darshan, and Dani Lischinski. Data roaming and early fusion for composed image retrieval. *arXiv preprint arXiv:2303.09429*, 2023. [1](#), [7](#)
- [29] Junnan Li and et al. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022. [3](#), [4](#), [6](#)
- [30] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi.

- Align before fuse: Vision and language representation learning with momentum distillation. *NeurIPS*, 34:9694–9705, 2021. [5](#)
- [31] 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, 4](#)
- [32] 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 *ECCV*, pages 740–755. Springer, 2014. [6](#)
- [33] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. [3, 5, 1](#)
- [34] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *CVPR*, 2016. [6](#)
- [35] Zheyuan Liu, Cristian Rodriguez-Opazo, Damien Teney, and Stephen Gould. Image retrieval on real-life images with pre-trained vision-and-language models. In *ICCV*, 2021. [1, 2, 7](#)
- [36] Zheyuan Liu, Cristian Rodriguez-Opazo, Damien Teney, and Stephen Gould. Image retrieval on real-life images with pre-trained vision-and-language models. In *CVPR*, pages 2125–2134, 2021. [1, 4, 6](#)
- [37] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2018. [1](#)
- [38] Jiayi Ma, Xingyu Jiang, Aoxiang Fan, Junjun Jiang, and Junchi Yan. Image matching from handcrafted to deep features: A survey. *IJCV*, 129:23–79, 2021. [1, 2](#)
- [39] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *TPAMI*, 41(8):1979–1993, 2018. [2](#)
- [40] Filip Radenovic, Abhimanyu Dubey, Abhishek Kadian, Todor Mihaylov, Simon Vandenhende, Yash Patel, Yi Wen, Vignesh Ramanathan, and Dhruv Mahajan. Filtering, distillation, and hard negatives for vision-language pre-training. In *CVPR*, pages 6967–6977, 2023. [5](#)
- [41] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. [3, 4](#)
- [42] 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 *CVPR*, pages 19305–19314, 2023. [1, 2, 7](#)
- [43] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*, pages 2556–2565, 2018. [5](#)
- [44] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *NeurIPS*, 33:596–608, 2020. [2, 3](#)
- [45] Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. A corpus for reasoning about natural language grounded in photographs. In *ACL*, 2019. [6](#)
- [46] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *NeurIPS*, 30, 2017. [2](#)
- [47] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. [2](#)
- [48] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. [2, 3, 6](#)
- [49] Lucas Ventura, Antoine Yang, Cordelia Schmid, and Gül Varol. CoVR: Learning composed video retrieval from web video captions. *arXiv:2308.14746*, 2023. [1, 3, 7](#)
- [50] Nam Vo, Lu Jiang, Chen Sun, Kevin Murphy, Li-Jia Li, Li Fei-Fei, and James Hays. Composing text and image for image retrieval—an empirical odyssey. In *CVPR*, pages 6439–6448, 2019. [1, 2](#)
- [51] Haokun Wen, Xueming Song, Jianhua Yin, Jianlong Wu, Weili Guan, and Liqiang Nie. Self-training boosted multifaceted matching network for composed image retrieval. *arXiv preprint arXiv:2305.09979*, 2023. [2](#)
- [52] Hui Wu, Yupeng Gao, Xiaoxiao Guo, Ziad Al-Halah, Steven Rennie, Kristen Grauman, and Rogerio Feris. Fashion iq: A new dataset towards retrieving images by natural language feedback. In *CVPR*, pages 11307–11317, 2021. [1, 6](#)
- [53] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023. [3](#)