

CSMCIR: CoT-Enhanced Symmetric Alignment with Memory Bank for Composed Image Retrieval

Zhipeng Qian^{1*}, Zihan Liang^{1*}, Yufei Ma^{1*}, Ben Chen^{1†}, Huangyu Dai¹,
Yiwei Ma², Jiayi Ji², Chenyi Lei¹, Han Li¹, Xiaoshuai Sun²

¹Kuaishou Technology

²Key Laboratory of Multimedia Trusted Perception and Efficient Computing,
Ministry of Education of China, Xiamen University, 361005, P.R. China
qianzhipeng@stu.xmu.edu.cn ✉ benchen4395@gmail.com

Abstract

Composed Image Retrieval (CIR) enables users to search for target images using both a reference image and manipulation text, offering substantial advantages over single-modality retrieval systems. However, existing CIR methods suffer from representation space fragmentation: queries and targets comprise heterogeneous modalities and are processed by distinct encoders, forcing models to bridge misaligned representation spaces only through post-hoc alignment, which fundamentally limits retrieval performance. As evidenced by t-SNE visualization in Fig. 2(a), this architectural asymmetry manifests as three distinct, well-separated clusters in the feature space, directly demonstrating how heterogeneous modalities and architectural asymmetry create fundamentally misaligned representation spaces from initialization. In this work, we propose CSMCIR, a unified representation framework that achieves efficient query-target alignment through three synergistic components. First, we introduce a Multi-level Chain-of-Thought (MCoT) prompting strategy that guides Multimodal Large Language Models to generate discriminative, semantically compatible captions for target images, establishing modal symmetry. Building upon this, we design a symmetric dual-tower architecture where both query and target sides utilize the identical shared-parameter Q-Former for cross-modal encoding, ensuring consistent feature representations and further reducing the alignment gap. Finally, this architectural symmetry enables an entropy-based, temporally dynamic Memory Bank strategy that provides high-quality negative samples while maintaining consistency with the evolving model state. Extensive experiments on four datasets demonstrate that our CSMCIR achieves state-of-the-art performance with superior training efficiency. Our code is available at <https://github.com/qzp2018/CSMCIR>.

*These authors contributed equally.

†The corresponding author.

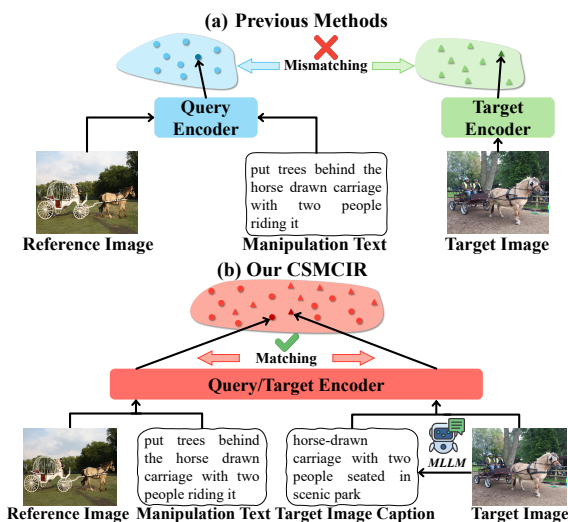


Figure 1: Workflows of existing CIR methods (a) and our proposed CSMCIR (b). Our approach achieves modal and structural symmetry for better alignment.

1 Introduction

Composed Image Retrieval (CIR) represents a significant advancement in multimodal search (Liang et al., 2025; Zhang et al., 2024b; Kim et al., 2025; Zheng et al., 2025). Unlike single-modality approaches, CIR integrates both reference images and manipulation text as inputs, enabling users to express search intents with enhanced precision. This multimodal paradigm has practical applications, such as e-commerce search, where users express nuanced preferences beyond a single modality.

Despite its promising prospects, existing CIR methods suffer from **representation space fragmentation** (Sun et al., 2024; Liu et al., 2023b; Yang et al., 2024b; Suo et al., 2024): queries (image and text) and targets (only image) comprise heterogeneous modalities and are processed by distinct encoders, forcing models to bridge misaligned representation spaces only through post-hoc alignment. Researchers have explored various fusion strategies for multimodal integration, including early fusion (Liu et al., 2023b; Levy et al., 2024),

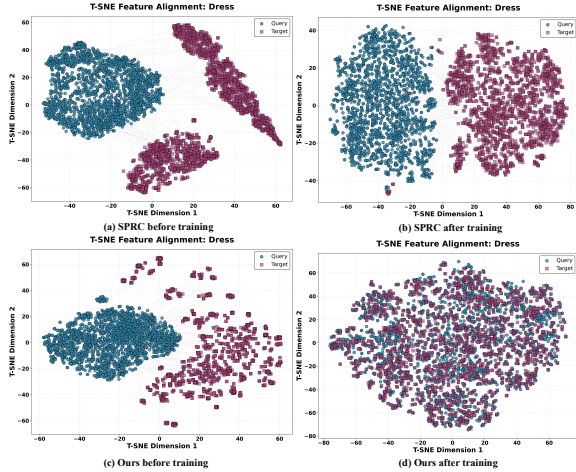


Figure 2: T-SNE visualization comparing SPRC and CSMCIR on Fashion-IQ dresses.

late fusion (Anwaar et al., 2021; Baldrati et al., 2023b; Wen et al., 2023b, 2024), and textual inversion (Baldrati et al., 2022; Saito et al., 2023; Baldrati et al., 2023a) approaches. However, as illustrated in Fig.1(a), these methods maintain asymmetric encoders for query and target sides, perpetuating the fragmentation that fundamentally limits contrastive learning effectiveness. Critically, as shown in Fig. 2(b), this fragmentation persists even after training, revealing the fundamental limitation of these asymmetric post-hoc alignment strategies.

Recent advances in Multimodal Large Language Models (MLLMs) enable transforming target images into textual representations through caption generation, addressing the modal inconsistency between query and target sides. However, there still exists a critical challenge: generated captions must be both discriminative to capture distinctive visual attributes and maintain comparable semantic detail with manipulation text. To address this, we propose Multi-level Chain-of-Thought (MCoT) prompting that guides MLLMs through structured reasoning to synthesize captions satisfying these requirements. Unlike approaches (Tian et al., 2025; Wen et al., 2024; Tang et al., 2024; Lin et al., 2025) that optimize manipulation prompts at query time and introduce inference latency, our method pre-generates target captions offline, making it practical for real-world deployment while establishing the foundation for modal symmetry.

With target captions generated via MCoT, both query and target sides exhibit consistent multimodal structure as image-text pairs, enabling a symmetric dual-tower architecture with shared-parameter Q-Former (shown in Fig. 1(b)). By processing both sides through the identical encoder,

this design directly addresses the encoder asymmetry in prior methods (Xu et al., 2024; Tian et al., 2025; Li et al., 2025a). The advantage of this symmetric design is evident in Fig. 2(c): even before training, our method demonstrates a smaller query-target distance, contrasting sharply with the severe fragmentation in the baseline approach. This validates that modal symmetry through consistent encoding establishes a solid foundation for progressive alignment, rather than relying on post-hoc bridging of misaligned spaces.

Finally, the achieved structural symmetry further enables Memory Bank(Wu et al., 2018) integration for enhanced contrastive learning. While prior work (Feng et al., 2024) deemed Memory Banks unsuitable for CIR due to architectural and modal asymmetry, our symmetric design naturally accommodates this mechanism. However, standard Memory Banks suffer from representation inconsistency in the CIR task. Due to limited training data and the need for rapid parameter updates, the stored representations in Memory Bank become misaligned with current batch representations as the model states evolve. To address this challenge, we propose an entropy-based Memory Bank strategy that incorporates temporal awareness and information-theoretic sample selection. The strategy dynamically updates stored representations using the current model state, ensuring diverse and informative negatives that enhance contrastive learning, improve query-target alignment, and boost training efficiency. As depicted in Fig. 2(d), after training, our method achieves tightly integrated query-target fusion, with embeddings becoming almost indistinguishable throughout the representation space, proving its effectiveness.

Taking the above designs into account, we introduce **CSMCIR**, a CoT-Enhanced Symmetric Alignment with Memory Bank for Composed Image Retrieval. Extensive experiments conducted across Fashion-IQ, CIRR, Shoes and LaSCO datasets demonstrate the effectiveness of our approach. In summary, our contributions include:

- We propose a unified symmetric framework that systematically addresses representation space fragmentation in CIR through MCoT-enhanced caption generation and parameter-shared dual-tower architecture.
- We introduce an entropy-based, temporally-aware Memory Bank that maintains representation consistency with evolving model states

for enhanced contrastive learning.

- Extensive experiments across four benchmarks demonstrate that our CSMCIR achieves state-of-the-art performance, with comprehensive ablation studies confirming the effectiveness of each component.

2 Related Work

2.1 Composed Image Retrieval

In recent years, there has been a surge of research and development at the intersection of multimodal learning (Zhang et al., 2024a, 2025; Yang et al., 2024a; Qian et al., 2024a,b) and search technologies (Chen et al., 2026). CIR enables the retrieval of target images matching both a reference image and manipulation text, requiring an effective understanding of complex semantic interactions between visual and textual modalities. Early approaches explored various strategies for multimodal fusion, including early-fusion (Liu et al., 2023b; Levy et al., 2024) and late-fusion (Anwaar et al., 2021; Baldrati et al., 2023b; Wen et al., 2023b, 2024; Chen et al., 2025; Li et al., 2025b). Another line of work introduced textual inversion modules (Gal et al., 2022; Baldrati et al., 2022; Saito et al., 2023; Baldrati et al., 2023a) to transform reference images into pseudo-word embeddings, which are subsequently concatenated with manipulation text for target retrieval. However, existing methods suffer from representation space fragmentation due to asymmetric query and target modalities. The query side combines a reference image with manipulation text, while the target side contains only an image. This fundamental asymmetry necessitates distinct encoders for each side (Jiang et al., 2024; Wen et al., 2023a; Xu et al., 2024; Levy et al., 2024; Liu et al., 2023b; Jang et al., 2024; Xing et al., 2025), forcing models to bridge misaligned representation spaces only through post-hoc alignment, which fundamentally limits contrastive learning effectiveness. To address this limitation, we propose a unified symmetric dual-tower architecture that establishes consistent representations, directly tackling the representation space fragmentation challenge.

2.2 Vision and Language Pre-training Models

Large Vision-Language Models (LVLMs) (Radford et al., 2021; Li et al., 2022; Lu et al., 2019; Li et al., 2023) such as CLIP (Radford et al., 2021) and BLIP (Li et al., 2022) have become foundational tools for CIR by enabling alignment between

reference images and manipulation text. Many works (Jiang et al., 2024; Xu et al., 2024; Wen et al., 2023b,a; Ventura et al., 2024; Lin et al., 2025) leverage LVLMs as encoders to enhance cross-modal matching. Recent research explores integrating Multimodal Large Language Models (MLLMs) (Li et al., 2023; Liu et al., 2023a; Bai et al., 2023) into CIR tasks. SPN (Feng et al., 2024) uses MLLMs to construct training triplets, while CIR-LVLM (Sun et al., 2024) employs them as user intent-aware encoders. Several approaches (Wen et al., 2024; Tang et al., 2024; Sun et al., 2023; Karthik et al., 2023) apply MLLMs to refine manipulation text at query time for improved matching. However, such query-time text refinement introduces significant inference latency, degrading user experience. In contrast, generating descriptive captions for target images, which can be performed offline, remains largely unexplored. We argue that MLLMs can produce detailed target image descriptions, enhancing alignment between query and target modalities. To this end, we leverage Chain-of-Thought (CoT) prompting (Wei et al., 2022; He et al., 2024; Zhang et al., 2023; Zheng et al., 2023) to generate enhanced target image captions, establishing unified representation spaces for effective alignment.

3 Methodology

3.1 Preliminary

CIR addresses the retrieval problem where a query Q combines a reference image I_r with a manipulation text T to search for the relevant target image I_t in a candidate set. Our approach tackles the fundamental limitation of **representation space fragmentation** in CIR: queries and targets comprise heterogeneous modalities processed by distinct encoders, creating misaligned representation spaces. By employing MLLMs to generate target image captions $T(I_t)$, we transform the paradigm from (I_r, T, I_t) to $(I_r, T, I_t, T(I_t))$ to establish modal symmetry. Notably, this transformation does not alter the essential task definition, as captions are derived solely from target images and generated automatically without any manual annotations. In practical scenario such as e-commerce platforms, these captions can be seamlessly substituted with existing product titles or descriptions.

3.2 Multi-level Chain-of-Thought Prompting

Despite the ease of generating captions for images using MLLMs, caption quality critically impacts

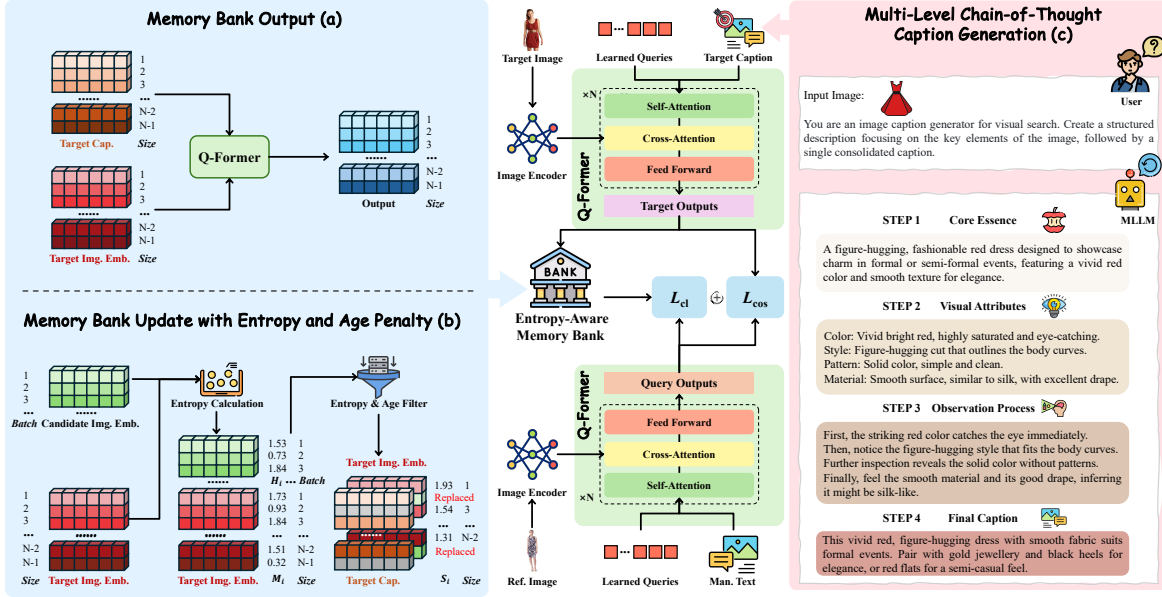


Figure 3: Overview of CSMCIR framework. **Right:** MCoT-based target caption generation. **Left:** (a) Entropy-Aware Memory Bank for negative sampling. (b) Memory Bank update via entropy-based and temporal scoring.

model performance. In CIR, users identify target images by describing distinctive differences between reference and target images. Consequently, captions must balance discriminative detail with conciseness: overly brief captions risk omitting critical visual attributes for identification, while exhaustive descriptions introduce redundancy and counterproductive noise that interferes with retrieval. Moreover, semantic and structural inconsistency between manipulation text and target captions may introduce misleading signals that undermine performance.

To address these challenges, we propose a Multi-level Chain-of-Thought (MCoT) prompting strategy that guides MLLMs to generate comprehensive yet discriminative captions with appropriate detail levels for target images. As depicted in Fig.3, our MCoT comprises four key steps as follows:

Step 1: Core Essence: The MLLM first generates concise descriptions capturing the essence of target images I_t , focusing on main objects and distinctive characteristics.

Step 2: Visual Attributes: Next, the MLLM identifies and describes key visual attributes of objects in the image, including color, material, shape, and spatial relationships.

Step 3: Observation Process: The MLLM then explicates its reasoning process for identifying the Core Essence, detailing which Visual Attributes were prioritized and why.

Step 4: Final Caption Formation: By synthesizing insights from the previous steps and a few

prompt examples, the MLLM generates a comprehensive final caption that incorporates both primary objects and discriminative details while avoiding redundant descriptions.

In summary, the complete caption generation process for target images can be formulated as:

$$T(I_t) = \Psi_{MCoT}(I_t) \quad , \quad (1)$$

where $T(I_t)$ denotes the generated caption for image I_t , and Ψ_{MCoT} represents the MCoT-based caption generation function.

Our MCoT enhances cross-modal alignment and mitigates representation space fragmentation between query and target sides, establishing a foundation for the symmetric dual-tower architecture and Memory Bank optimization in subsequent training stages. Caption construction requires only a few hours at most and, as a one-time preprocessing step, can be reused across all subsequent experiments.

3.3 Symmetrical Dual-tower Model Architecture

After obtaining target image captions, we establish structurally consistent multimodal pairs on both sides: (I_r, T) and $(I_t, T(I_t))$. This modal symmetry enables identical processing operations for both pairs. Specifically, inspired by (Xu et al., 2024; Jiang et al., 2024), we adopt the lightweight Querying Transformer (Q-Former) from BLIP-2 (Li et al., 2023) as our cross-modal encoder. As depicted in Fig.3, BLIP-2’s pretrained image encoder extracts visual features from both reference

and target images. These image features, along with their corresponding text (manipulation text for the reference image and generated caption for the target image), are fed into the identical Q-Former with fully shared parameters.

Additionally, a set of learnable query tokens q are introduced to facilitate cross-modal interaction between visual and textual representations on both sides. The encoding process can be formulated as:

$$Z_q = \text{Q-Former}(I_r, T, q), \quad (2)$$

$$Z_t = \text{Q-Former}(I_t, T(I_t), q), \quad (3)$$

where Z_q and Z_t denote the encoded cross-modal representations from the query and target encoders respectively.

3.4 Entropy-Aware Memory Bank Strategy

Traditional Memory Bank approaches (Wu et al., 2018) enable efficient large-scale contrastive learning but suffer from temporal inconsistency: frozen representations become misaligned as models evolve, degrading negative sample quality. This limitation is particularly severe in CIR, where representation space fragmentation led prior work (e.g., SPN (Feng et al., 2024)) to deem Memory Banks incompatible. While our symmetric architecture addresses structural misalignment, temporal inconsistency remains a critical challenge.

Static-Dynamic Decoupling. To address this, we propose a static-dynamic decoupling strategy: the memory bank stores static inputs (captions and frozen image embeddings before being put into Q-Former), while embeddings for negative sampling are dynamically recomputed using the current Q-Former at each step. This ensures all representations remain consistent with the current model state.

Entropy-Aware Sample Selection. Let $\mathcal{B} = \{\mathbf{z}_1, \dots, \mathbf{z}_B\}$ denote the [CLS] token embeddings of images from the current batch obtained via ViT (where B is the batch size), and $\mathcal{M} = \{\mathbf{m}_1, \dots, \mathbf{m}_M\}$ denote the image embeddings stored in the memory bank (where M is the memory bank size). To enhance sample diversity and select informative hard negatives, we measure uncertainty via information entropy.

Specifically, we first compute similarity-based probability distributions. For batch sample i relative to all memory samples:

$$p_{i,j}^{\mathcal{B} \rightarrow \mathcal{M}} = \frac{\exp(\mathbf{z}_i^T \mathbf{m}_j)}{\sum_{k=1}^M \exp(\mathbf{z}_i^T \mathbf{m}_k)}, \quad (4)$$

and for memory sample i relative to samples in memory bank:

$$p_{i,j}^{\mathcal{M} \rightarrow \mathcal{M}} = \frac{\exp(\mathbf{m}_i^T \mathbf{m}_j)}{\sum_{k=1}^M \exp(\mathbf{m}_i^T \mathbf{m}_k)}. \quad (5)$$

Then we calculate the information entropy to measure the uncertainty of each sample’s similarity distribution:

$$H_i^{\mathcal{B}} = - \sum_{j=1}^M p_{i,j}^{\mathcal{B} \rightarrow \mathcal{M}} \log p_{i,j}^{\mathcal{B} \rightarrow \mathcal{M}}, \quad (6)$$

$$H_i^{\mathcal{M}} = - \sum_{j=1}^M p_{i,j}^{\mathcal{M} \rightarrow \mathcal{M}} \log p_{i,j}^{\mathcal{M} \rightarrow \mathcal{M}}. \quad (7)$$

Higher entropy indicates that the sample exhibits greater dissimilarity to memory bank samples, making it suitable as an informative negative.

Temporal Decay and Replacement Strategy. To prevent outdated representations from persistently occupying the memory bank, we define a retention score for each memory sample that jointly considers diversity and temporal freshness:

$$\hat{H}_i^{\mathcal{M}} = \underbrace{\max\left(0, 1 - \frac{\Delta t_i}{N_{\max}}\right)}_{\text{freshness factor}} \cdot \underbrace{H_i^{\mathcal{M}}}_{\text{diversity factor}}, \quad (8)$$

where Δt_i denotes the number of training steps since sample i was last updated, and $N_{\max} = 10$ controls the maximum staleness threshold before complete decay. This retention score ensures that both low-diversity and stale samples are prioritized for replacement. Finally, batch samples are inserted into the memory bank by replacing memory samples with lower retention entropy, i.e., we replace sample i in the memory bank with sample j in the batch when $H_j^{\mathcal{B}} > \hat{H}_i^{\mathcal{M}}$.

Overall, this entropy-aware replacement strategy maintains a diverse and temporally consistent set of negatives, addressing both representation inconsistency and quality degradation inherent in traditional Memory Bank approaches for CIR.

3.5 Learning Objectives

Following previous works, contrastive loss is introduced to achieve alignment between the query and target sides of the CIR task. Specifically, following (Xu et al., 2024), we utilize the [CLS] token e_{cls} from the query side’s output Z_q as our query embedding u , which encapsulates global information of the query encoder output. For the target embedding v , we select the query token from the target

| Method | Dress | | Shirt | | Toptee | | Average | | |
|----------------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | R@10 | R@50 | R@10 | R@50 | R@10 | R@50 | R@10 | R@50 | Avg. |
| CoPE (Tang et al., 2025) | 39.85 | 66.98 | 45.03 | 66.81 | 48.61 | 72.01 | 44.50 | 68.60 | 56.55 |
| CaLa (Jiang et al., 2024) | 42.38 | 66.08 | 46.76 | 68.18 | 50.93 | 73.42 | 46.69 | 69.22 | 57.96 |
| CoVR-BLIP (Ventura et al., 2024) | 44.55 | 69.03 | 48.43 | 67.42 | 52.60 | 74.31 | 48.53 | 70.25 | 59.39 |
| CASE (Levy et al., 2024) | 47.44 | 69.36 | 48.48 | 70.23 | 50.18 | 72.24 | 48.79 | 70.68 | 59.74 |
| Re-ranking (Liu et al., 2023b) | 48.14 | 71.43 | 50.15 | 71.25 | 55.23 | 76.80 | 51.17 | 73.13 | 62.15 |
| FashionERN (Chen et al., 2024) | 43.93 | 68.77 | 52.70 | <u>75.07</u> | 56.09 | 78.38 | 50.91 | 74.07 | 62.49 |
| SPRC (Xu et al., 2024) | 49.18 | 72.43 | 55.64 | 73.89 | <u>59.35</u> | 78.58 | 54.72 | 74.97 | 64.85 |
| CCIN (Tian et al., 2025) | 49.38 | <u>72.58</u> | 55.93 | 74.14 | 57.93 | 77.56 | 54.41 | 74.76 | 64.59 |
| TME (Li et al., 2025a) | <u>49.73</u> | 71.69 | <u>56.43</u> | 74.44 | 59.31 | <u>78.94</u> | <u>55.15</u> | <u>75.02</u> | <u>65.09</u> |
| CSMCIR(Ours) | 52.45 | 74.81 | 57.70 | 75.76 | 61.14 | 80.98 | 57.07 | 77.27 | 67.17 |

Table 1: **Quantitative** comparison on the **Fashion-IQ** validation set. Overall 1st / 2nd in bold/underline.

| Method | Recall@K | | | | Recalls _{subset} @K | | | $(R@5 + R_{sub}@1)/2$ |
|----------------------------------|--------------|--------------|--------------|--------------|------------------------------|--------------|--------------|-----------------------|
| | K=1 | K=5 | K=10 | K=50 | K=1 | K=2 | K=3 | |
| CoPE (Tang et al., 2025) | 49.18 | 80.65 | 89.86 | 98.05 | 72.34 | 88.65 | 95.30 | 76.49 |
| CASE (Levy et al., 2024) | 48.00 | 79.11 | 87.25 | 97.57 | 75.88 | 90.58 | 96.00 | 77.50 |
| CaLa (Jiang et al., 2024) | 49.11 | 81.21 | 89.59 | 98.00 | 76.27 | 91.04 | 96.46 | 78.74 |
| CoVR-BLIP (Ventura et al., 2024) | 49.69 | 78.60 | 86.77 | 94.31 | 75.01 | 88.12 | 93.16 | 80.81 |
| Re-ranking (Liu et al., 2023b) | 50.55 | 81.75 | 89.78 | 97.18 | 80.04 | 91.90 | 96.58 | 80.90 |
| SPRC (Xu et al., 2024) | 51.96 | 82.12 | 89.74 | 97.69 | 80.65 | 92.31 | 96.60 | 81.39 |
| CCIN (Tian et al., 2025) | 53.41 | <u>84.05</u> | <u>91.17</u> | 98.00 | - | - | - | - |
| TME (Li et al., 2025a) | <u>53.42</u> | 82.99 | 90.24 | <u>98.15</u> | 81.04 | <u>92.58</u> | <u>96.94</u> | <u>82.01</u> |
| CSMCIR(Ours) | 53.88 | 85.06 | 91.69 | 98.58 | <u>80.10</u> | 92.24 | 96.92 | 82.58 |

Table 2: **Quantitative** comparison on the **CIRR** test set. Overall 1st / 2nd in bold/underline.

| Method | R@1 | R@10 | R@50 | Avg. |
|--------------------------------|--------------|--------------|--------------|--------------|
| FashionERN (Chen et al., 2024) | - | 55.59 | 81.71 | - |
| Prog. Lrn. (Zhao et al., 2022) | 22.88 | 58.83 | 84.16 | 55.29 |
| CAFF (Wan et al., 2024) | 25.21 | 60.17 | 80.79 | 55.39 |
| TG-CIR (Wen et al., 2023b) | 25.89 | 63.20 | 85.07 | 58.05 |
| CCIN (Tian et al., 2025) | 25.95 | 65.76 | 86.54 | 59.42 |
| CSMCIR(Ours) | 27.83 | 68.37 | 87.90 | 61.37 |

Table 3: Results on Shoes validation set. The best results are highlighted in bold.

| Method | R@1 | R@5 | R@10 | R@50 | R@500 | Avg. |
|---------------------------------|-------------|--------------|--------------|--------------|--------------|--------------|
| Random | 0.00 | 0.01 | 0.03 | 0.13 | 1.26 | 0.72 |
| LF-CLIP (Baldrati et al., 2022) | 4.01 | 10.23 | 14.68 | 32.08 | 72.69 | 26.74 |
| LF-BLIP (Baldrati et al., 2022) | 4.26 | 12.01 | 17.11 | 36.54 | 74.62 | 28.91 |
| CASE (Levy et al., 2024) | 7.08 | 18.50 | 26.16 | 50.25 | 85.46 | 37.49 |
| CSMCIR(Ours) | 7.83 | 23.89 | 32.90 | 58.44 | 90.31 | 42.68 |

Table 4: Results on LaSCO validation set. The best results are highlighted in bold.

side’s output Z_t with the highest similarity to the query embedding u .

3.5.1 Memory Bank Enhanced Contrastive Loss

For query-to-target alignment, since we implement the Memory Bank approach, the Memory Bank Enhanced contrastive loss is defined as follows:

$$\mathcal{L}_{cl} = -\frac{1}{|B|} \sum_{i \in B} \log \frac{\exp(\tau u_i^T v_i)}{\sum_{j \in B^*} \exp(\tau u_i^T v_j)}, \quad (9)$$

where B denotes the standard batch and B^* represents the batch expanded by our Entropy-Aware Memory Bank.

3.5.2 Adaptive Cosine Loss Alignment

Since on the target side, the above-mentioned contrastive loss only selects the query token with the highest similarity with the query embedding u as our target embedding v , this loss focuses solely on the most relevant token while potentially overlooking valuable information contained in other query tokens on the target side. To better leverage the multimodal information from the target side, we further introduce alignment between the query embedding u and all query tokens on the target side. Specifically, we introduce a learnable tensor α (initialized to 1) to adaptively weight the importance of different query tokens on the target side, then apply average pooling to aggregate the query tokens, and finally adopt a cosine loss to ensure comprehensive alignment between the two sides:

$$\mathcal{L}_{cos} = \frac{1}{B} \sum_{i=1}^B \left(1 - \frac{\left(\frac{1}{K} \sum_{k=1}^K \alpha_k \cdot v_k^i \right) \cdot u^i}{\left\| \frac{1}{K} \sum_{k=1}^K \alpha_k \cdot v_k^i \right\| \cdot \|u^i\|} \right), \quad (10)$$

where K represents the number of query tokens.

Finally, the overall loss \mathcal{L} is formulated as:

$$\mathcal{L} = \mathcal{L}_{cl} + \mathcal{L}_{cos}. \quad (11)$$

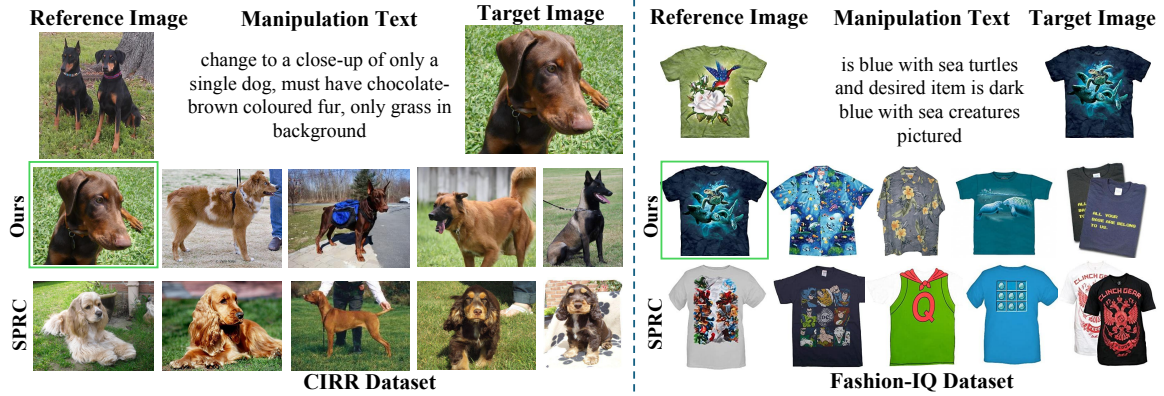


Figure 4: Qualitative comparison with SPRC. Green outlines indicate successfully retrieved targets.

4 Datasets and Evaluation Metrics

We evaluate our approach on three standard CIR benchmarks: Fashion-IQ, CIRR, Shoes and LaSCO. Fashion-IQ contains 77,684 fashion images forming 30,134 triplets across three categories (Dress, Toptee, and Shirt), with performance measured using Recall@10, Recall@50 and Recall_{mean} on the validation set. CIRR offers a more diverse benchmark with 36,554 triplets from 21,552 natural images, featuring everyday object interactions. CIRR evaluation uses Recall@1,5,10,50 for general performance and includes a specialized Recall_{subset}@1,2,3 metric for a challenging subset containing visually similar distractors, testing fine-grained discrimination capabilities. Shoes dataset is divided into 10K triplets for training and 4.6K for testing, performance measured using Recall@1, Recall@10, Recall@50 and Recall_{mean}. LaSCO is a dataset based on COCO images and VQA2.0 annotations, containing 389,305 queries on 121,479 natural images. Compared to CIRR, LaSCO offers x10 more queries, x2 more unique tokens, and x17 more corpus images across an open and broad domain of natural images with rich text. Performance on LaSCO is evaluated using Recall@1,5,10,50,500 and Recall_{mean} metrics. For LaSCO dataset, we use its own image captions as target image captions.

5 Experiments

5.1 Implementation Details

Our framework was implemented in PyTorch and ran on a single NVIDIA RTX A100 (40GB) GPU. Following (Xu et al., 2024), we adopt ViT-G as the image encoder, which remained frozen throughout training. We utilized the AdamW optimizer (Loshchilov and Hutter, 2017) with a weight decay of 0.05 and a learning rate of 1e-5 on a co-

sine decay schedule. For our model configuration, we set the number of query tokens to 32 and employed a batch size of 128. The model was trained for 15 epochs with Memory Bank capacities of 512 for CIRR, Fashion-IQ and Shoes datasets, and the Memory Bank capacity is 640 for the LaSCO dataset, age threshold N_{max} is 10. For Fashion-IQ, Shoes and CIRR datasets, we utilize Qwen2.5-VL-7B-Instruct (Bai et al., 2025) for target image caption generation. For the LaSCO dataset, we directly used the captions from its VQA2.0 annotations as target image descriptions.

5.2 Comparison Results

5.2.1 Quantitative Results

Tab. 1 shows CSMCIR achieves the highest recall across all Fashion-IQ metrics. Compared to TME(Li et al., 2025a) (which also uses QFormer), our method delivers substantial gains (65.09 vs. 67.17 in Avg. metric). As shown in Tab. 3, on the Shoes dataset, CSMCIR significantly outperforms CCIN (Tian et al., 2025) by 1.95 points. Tab. 2 and Tab. 4 show that CSMCIR also achieves state-of-the-art performance on open-domain CIRR and LaSCO datasets. Despite CIRR’s increased complexity, our method excels particularly in the $(R@5+R_{sub}@1)/2$ metric, substantially outperforming TME(Li et al., 2025a) (82.58 vs. 82.01), demonstrating strong generalizability. On LaSCO, CSMCIR achieves impressive gains of 6.74 and 8.19 in R@10 and R@50 metrics respectively. These consistent improvements across diverse datasets confirm CSMCIR’s effectiveness and robustness for composed image retrieval tasks.

Beyond performance metrics, training efficiency is crucial for practical utility. Our streamlined architecture with Memory Bank not only enhances performance but also significantly reduces training

| Method | Fashion-IQ | | | CIRR | | |
|-------------------|--------------|--------------|--------------|--------------|---------------------|--------------|
| | R@10 | R@50 | Avg. | R@5 | R _{sub} @1 | Avg. |
| baseline | 54.92 | 74.97 | 64.85 | 82.12 | 80.65 | 81.39 |
| +MCoT | 55.41 | 75.25 | 65.33 | 83.61 | 79.68 | 81.65 |
| +SA | 55.26 | 76.57 | 65.92 | 83.89 | 79.96 | 81.93 |
| +EAMB | 56.52 | 76.78 | 66.65 | 84.15 | 80.14 | 82.15 |
| +L _{cos} | 57.07 | 77.27 | 67.17 | 85.06 | 80.10 | 82.58 |

Table 5: The results obtained after ablating different modules on the Fashion-IQ and CIRR datasets.

| Method | Fashion-IQ | | | CIRR | | |
|----------|--------------|--------------|--------------|--------------|---------------------|--------------|
| | R@10 | R@50 | Avg. | R@5 | R _{sub} @1 | Avg. |
| baseline | 54.92 | 74.97 | 64.85 | 82.12 | 80.65 | 81.39 |
| w/o TD | 55.79 | 76.49 | 66.14 | 84.25 | 80.06 | 82.16 |
| w/o EA | 56.33 | 76.95 | 66.64 | 84.35 | 80.16 | 82.26 |
| Full | 57.07 | 77.27 | 67.17 | 85.06 | 80.10 | 82.58 |

Table 6: Ablation studies on EAMB Strategy.

costs. Using identical configurations (Q-Former with ViT-G backbone on a single A100 GPU), our method requires only **1.2** GPU hours on Fashion-IQ compared to SPRC’s **1.8** hours, and **2.1** GPU hours on CIRR versus SPRC’s **4.6** hours. Our CSM-CIR also achieves superior inference efficiency on A100 GPUs, with a latency of 0.03s per item (faster than SPRC’s 0.035s) and memory consumption of 6845MB (lower than SPRC’s 7140MB), making it more practical for real-world deployment.

5.2.2 Qualitative Results

Fig. 4 visualizes top-5 retrieval results on CIRR and Fashion-IQ datasets. Our CSMCIR demonstrates superior attention to image details compared to SPRC. On CIRR, we not only correctly predict the target at R@1, but also retrieve dogs with consistent breeds across all top-5 results, while SPRC shows significant breed variations. On Fashion-IQ, our top-1, top-2, and top-4 results closely match the manipulation text description, whereas SPRC fails to do so.

5.3 Ablation Study

To evaluate the effectiveness of our designed model architecture and modules, we conducted extensive experiments on the test set of CIRR and the validation set of Fashion-IQ.

5.3.1 Modules Ablation Study

As shown in Tab.5, since our model is built on SPRC(Xu et al., 2024), we set its performance as our baseline. Specifically, our Multi-level Chain-of-Thought(MCoT) successfully establishes cross-modal alignment by transforming the retrieval paradigm from asymmetric to symmetric, thereby achieving superior query-target alignment. Building upon this foundation, as shown in line3, our



Figure 5: Ablation study on Memory Bank strategies.

Symmetrical Structure (SA) further enhances alignment by employing identical shared-parameter encoders on both sides, delivering remarkable performance improvements. This validates the effectiveness of our unified architecture in addressing the modal alignment challenges inherent in CIR tasks. Furthermore, we validate our Entropy-Aware Memory Bank strategy (EAMB), which generates high-quality negative samples and expands batch size, achieving an improvement from 65.92 to 66.65 in the Fashion-IQ dataset. When combining Memory Bank Enhanced Contrastive Loss and Adaptive Cosine Loss, our framework further achieves performance gains, confirming their effectiveness.

5.3.2 Memory Bank Ablations

We conducted ablation experiments comparing our Memory Bank strategy with the naive FIFO approach from MoCo (He et al., 2020). Unlike traditional representation learning, CIR involves fewer samples, more complex modalities, and faster model updates. Directly storing Q-former-generated cross-modal features proves problematic: rapid Q-former updates across batches create significant feature distribution disparities, rendering them unsuitable as negatives for subsequent batches. As shown in Fig. 5, the FIFO strategy exhibits rapid performance degradation when Memory Bank size exceeds batch size. Our strategy addresses this by storing captions and image embeddings from previous batches, then constructing negatives through Q-Former in the current batch to maintain feature distribution consistency. Performance plateaus at Memory Bank size 512, likely sufficient to capture the feature distributions of the Fashion-IQ and CIRR datasets. We also conduct ablation studies on the EAMB strategy components, with results shown in Table 6. Removing either Temporal Decay (TD) or Entropy-Aware (EA) degrades performance on both Fashion-IQ and CIRR benchmarks, with TD showing greater impact. The full model achieves the best results, confirming the effectiveness of both components.

| Method | Fashion-IQ | | | CIRR | | | Shoes | | | | Lasco | | | |
|---------------------------------|--------------|--------------|--------------|--------------|---------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | R@10 | R@50 | Avg. | R@5 | R _{sub} @1 | Avg. | R@1 | R@10 | R@50 | Avg. | R@5 | R@10 | R@50 | Avg. |
| CoPE (ViT-L)(Tang et al., 2025) | 44.50 | 68.60 | 56.55 | 80.65 | 72.34 | 76.49 | - | - | - | - | - | - | - | - |
| SPRC (ViT-L)(Xu et al., 2024) | 51.04 | 73.38 | 62.21 | 80.65 | 79.59 | 80.12 | - | - | - | - | - | - | - | - |
| QURE (ViT-L)(Kwak et al., 2025) | 52.60 | 73.48 | 63.04 | 82.53 | 78.51 | 80.52 | - | - | - | - | - | - | - | - |
| Ours (ViT-L) | 54.25 | 74.66 | 64.45 | 83.04 | 78.82 | 80.93 | 25.50 | 64.74 | 87.11 | 59.11 | 22.45 | 31.35 | 56.93 | 36.91 |
| Ours (ViT-G) | 57.07 | 77.27 | 67.17 | 85.06 | 80.10 | 82.58 | 27.83 | 68.37 | 87.90 | 61.37 | 23.89 | 32.90 | 58.44 | 38.14 |

Table 7: ViT type ablations. Our method demonstrates superior performance across different ViT architectures, including ViT-L, proving its robustness and stability. Furthermore, our method achieves significant performance improvements when scaling from ViT-L to ViT-G architectures.

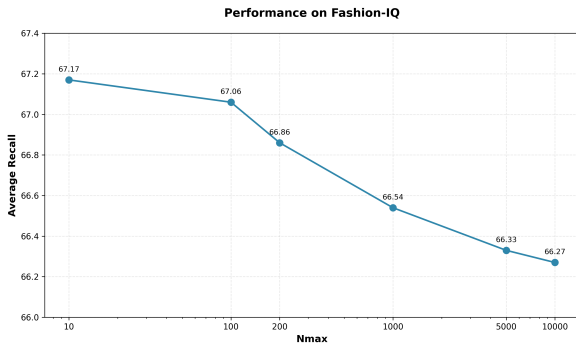


Figure 6: Ablation study on N_{max} parameter in our Memory Bank strategy on Fashion-IQ dataset.

5.3.3 ViT Ablation Study

The image encoder plays a vital role in image feature extraction, significantly impacting overall model performance. As shown in Tab.7, across all three datasets: Fashion-IQ, CIRR, Shoes and LaSCO, our model’s performance with ViT-G substantially outperforms the ViT-L version, confirming the importance of high-quality visual features in the CIR task. Nevertheless, even with the ViT-L version, our model still achieves excellent results, significantly outperforms QURE(Kwak et al., 2025) and CoPE(Tang et al., 2025) under equivalent conditions, demonstrating the effectiveness of our approach.

5.3.4 N_{max} parameter Ablation

As shown in Fig. 6, the performance drops as N_{max} increases, which validates the effectiveness of our time-step-based update strategy. This indicates that although entropy-based selection can identify high-quality negative samples to some extent, relying on these samples without timely updates significantly degrades the effectiveness of the memory bank strategy.

6 Conclusion

In this work, we propose CSMCIR, a unified symmetric framework that systematically addresses representation space fragmentation in CIR. Our approach achieves state-of-the-art performance through three synergistic innovations: Multi-level Chain-of-Thought prompting, parameter-shared dual-tower architecture, and Entropy-Aware Memory Bank. Extensive experiments across four benchmarks demonstrate CSMCIR’s superior performance, with comprehensive ablation studies confirming the effectiveness of each component. Our code is available at <https://github.com/qzp2018/CSMCIR>.

7 Limitations

While our CSMCIR framework achieves state-of-the-art performance on multiple benchmarks, several limitations warrant discussion:

7.1 Caption Generation Dependency:

Our approach leverages MLLMs to pre-generate captions for target images in an offline manner. While this effectively eliminates inference latency during runtime, it entails additional upfront time investment to generate the required captions for datasets where such target image captions are not readily available.

7.2 Dataset Scope:

Our evaluation is centered on the fashion and general object domains, with experiments conducted on benchmark datasets including Fashion-IQ, CIRR, Shoes, and LaSCO. However, performance in specialized domains—such as medical imaging and fine art—remains uninvestigated, primarily due to the scarcity of task-specific, high-quality datasets.

References

- Muhammad Umer Anwaar, Egor Labintsev, and Martin Kleinsteuber. 2021. Compositional learning of image-text query for image retrieval. In *Proceedings of the IEEE/CVF Winter conference on Applications of Computer Vision*, pages 1140–1149.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, and 1 others. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Alberto Baldrati, Lorenzo Agnolucci, Marco Bertini, and Alberto Del Bimbo. 2023a. Zero-shot composed image retrieval with textual inversion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15338–15347.
- Alberto Baldrati, Marco Bertini, Tiberio Uricchio, and Alberto Del Bimbo. 2022. 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.
- Alberto Baldrati, Marco Bertini, Tiberio Uricchio, and Alberto Del Bimbo. 2023b. Composed image retrieval using contrastive learning and task-oriented clip-based features. *ACM Transactions on Multimedia Computing, Communications and Applications*, 20(3):1–24.
- Ben Chen, Siyuan Wang, Yufei Ma, Zihan Liang, Xuxin Zhang, Yue Lv, Ying Yang, Huangyu Dai, Lingtao Mao, Tong Zhao, and 1 others. 2026. Onesearch-v2: The latent reasoning enhanced self-distillation generative search framework. *arXiv preprint arXiv:2603.24422*.
- Yanzhe Chen, Huasong Zhong, Xiangteng He, Yuxin Peng, Jiahuan Zhou, and Lele Cheng. 2024. Fashionern: enhance-and-refine network for composed fashion image retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 1228–1236.
- Zhiwei Chen, Yupeng Hu, Zixu Li, Zhiheng Fu, Xuemeng Song, and Liqiang Nie. 2025. Offset: Segmentation-based focus shift revision for composed image retrieval. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 6113–6122.
- Zhangchi Feng, Richong Zhang, and Zhijie Nie. 2024. Improving composed image retrieval via contrastive learning with scaling positives and negatives. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 1632–1641.
- Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. 2022. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738.
- Liqi He, Zuchao Li, Xiantao Cai, and Ping Wang. 2024. Multi-modal latent space learning for chain-of-thought reasoning in language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18180–18187.
- Young Kyun Jang, Donghyun Kim, Zihang Meng, Dat Huynh, and Ser-Nam Lim. 2024. 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.
- Xintong Jiang, Yaxiong Wang, Mengjian Li, Yujiao Wu, Bingwen Hu, and Xueming Qian. 2024. Cala: Complementary association learning for augmenting composed image retrieval. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2177–2187.
- Shyamgopal Karthik, Karsten Roth, Massimiliano Mancini, and Zeynep Akata. 2023. Vision-by-language for training-free compositional image retrieval. *arXiv preprint arXiv:2310.09291*.
- Sungyeon Kim, Xinliang Zhu, Xiaofan Lin, Muhammet Bastan, Douglas Gray, and Suha Kwak. 2025. Genius: A generative framework for universal multi-modal search. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 19659–19669.
- Jaehyun Kwak, Ramahdani Muhammad Izaaz Inhar, Se-Young Yun, and Sung-Ju Lee. 2025. Qure: Query-relevant retrieval through hard negative sampling in composed image retrieval. *arXiv preprint arXiv:2507.12416*.
- Matan Levy, Rami Ben-Ari, Nir Darshan, and Dani Lischinski. 2024. Data roaming and quality assessment for composed image retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 2991–2999.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding

- and generation. In *International conference on machine learning*, pages 12888–12900. PMLR.
- Shuxian Li, Changhao He, Xiting Liu, Joey Tianyi Zhou, Xi Peng, and Peng Hu. 2025a. Learning with noisy triplet correspondence for composed image retrieval. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 19628–19637.
- Zixu Li, Zhiwei Chen, Haokun Wen, Zhiheng Fu, Yupeng Hu, and Weili Guan. 2025b. Encoder: Entity mining and modification relation binding for composed image retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 5101–5109.
- Zihan Liang, Yufei Ma, Zhipeng Qian, Huangyu Dai, Zihan Wang, Ben Chen, Chenyi Lei, Yuqing Ding, and Han Li. 2025. **Uniecs: Unified multimodal e-commerce search framework with gated cross-modal fusion**. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, CIKM '25, page 1788–1797. ACM.
- Weihuang Lin, Yiwei Ma, Jiayi Ji, Xiaoshuai Sun, and Rongrong Ji. 2025. Cir-cot: Towards interpretable composed image retrieval via end-to-end chain-of-thought reasoning. *arXiv preprint arXiv:2510.08003*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023a. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916.
- Zheyuan Liu, Weixuan Sun, Damien Teney, and Stephen Gould. 2023b. Candidate set re-ranking for composed image retrieval with dual multi-modal encoder. *arXiv preprint arXiv:2305.16304*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32.
- Zhipeng Qian, Yiwei Ma, Jiayi Ji, and Xiaoshuai Sun. 2024a. X-refseg3d: Enhancing referring 3d instance segmentation via structured cross-modal graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pages 4551–4559.
- Zhipeng Qian, Yiwei Ma, Zhekai Lin, Jiayi Ji, Xiaowu Zheng, Xiaoshuai Sun, and Rongrong Ji. 2024b. Multi-branch collaborative learning network for 3d visual grounding. In *European Conference on Computer Vision*, pages 381–398. Springer.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, and 1 others. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR.
- Kuniaki Saito, Kihyuk Sohn, Xiang Zhang, Chun-Liang Li, Chen-Yu Lee, Kate Saenko, and Tomas Pfister. 2023. Pic2word: Mapping pictures to words for zero-shot composed image retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19305–19314.
- Shitong Sun, Fanghua Ye, and Shaogang Gong. 2023. Training-free zero-shot composed image retrieval with local concept reranking. *arXiv preprint arXiv:2312.08924*.
- Zelong Sun, Dong Jing, Guoxing Yang, Nanyi Fei, and Zhiwu Lu. 2024. Leveraging large vision-language model as user intent-aware encoder for composed image retrieval. *arXiv preprint arXiv:2412.11087*.
- Yucheng Suo, Fan Ma, Linchao Zhu, and Yi Yang. 2024. 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.
- Haomiao Tang, Jinpeng Wang, Yuang Peng, Guanghao Meng, Ruisheng Luo, Bin Chen, Long Chen, Yaowei Wang, and Shu-Tao Xia. 2025. Modeling uncertainty in composed image retrieval via probabilistic embeddings. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1210–1222.
- Yuanmin Tang, Xiaoting Qin, Jue Zhang, Jing Yu, Gaopeng Gou, Gang Xiong, Qingwei Ling, Saravan Rajmohan, Dongmei Zhang, and Qi Wu. 2024. Reason-before-retrieve: One-stage reflective chain-of-thoughts for training-free zero-shot composed image retrieval. *arXiv preprint arXiv:2412.11077*.
- Likai Tian, Jian Zhao, Zechao Hu, Zhengwei Yang, Hao Li, Lei Jin, Zheng Wang, and Xuelong Li. 2025. Ccin: Compositional conflict identification and neutralization for composed image retrieval. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 3974–3983.
- Lucas Ventura, Antoine Yang, Cordelia Schmid, and Gül Varol. 2024. Covr: Learning composed video retrieval from web video captions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 5270–5279.
- Yongquan Wan, Wenhai Wang, Guobing Zou, and Bofeng Zhang. 2024. Cross-modal feature alignment and fusion for composed image retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8384–8388.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Haokun Wen, Xuemeng Song, Xiaolin Chen, Yinwei Wei, Liqiang Nie, and Tat-Seng Chua. 2024. Simple

- but effective raw-data level multimodal fusion for composed image retrieval. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 229–239.
- Haokun Wen, Xuemeng Song, Jianhua Yin, Jianlong Wu, Weili Guan, and Liqiang Nie. 2023a. Self-training boosted multi-factor matching network for composed image retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5):3665–3678.
- Haokun Wen, Xian Zhang, Xuemeng Song, Yinwei Wei, and Liqiang Nie. 2023b. Target-guided composed image retrieval. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 915–923.
- Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742.
- Eric Xing, Pranavi Kolouju, Robert Pless, Abby Stylianou, and Nathan Jacobs. 2025. Context-cir: Learning from concepts in text for composed image retrieval. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 19638–19648.
- Xinxing Xu, Yong Liu, Salman Khan, Fahad Khan, Wangmeng Zuo, Rick Siow Mong Goh, Chun-Mei Feng, and 1 others. 2024. Sentence-level prompts benefit composed image retrieval. In *The Twelfth International Conference on Learning Representations*.
- Danni Yang, Ruohan Dong, Jiayi Ji, Yiwei Ma, Haowei Wang, Xiaoshuai Sun, and Rongrong Ji. 2024a. Exploring phrase-level grounding with text-to-image diffusion model. In *European Conference on Computer Vision*, pages 161–180. Springer.
- Xingyu Yang, Daqing Liu, Heng Zhang, Yong Luo, Chaoyue Wang, and Jing Zhang. 2024b. Decomposing semantic shifts for composed image retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 6576–6584.
- Jinlu Zhang, Jiji Tang, Rongsheng Zhang, Tangjie Lv, and Xiaoshuai Sun. 2025. Storyweaver: A unified world model for knowledge-enhanced story character customization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 9951–9959.
- Jinlu Zhang, Yiyi Zhou, Qiancheng Zheng, Xiaoxiong Du, Gen Luo, Jun Peng, Xiaoshuai Sun, and Rongrong Ji. 2024a. Fast text-to-3d-aware face generation and manipulation via direct cross-modal mapping and geometric regularization. *arXiv preprint arXiv:2403.06702*.
- Xin Zhang, Yanzhao Zhang, Wen Xie, Mingxin Li, Ziqi Dai, Dingkun Long, Pengjun Xie, Meishan Zhang, Wenjie Li, and Min Zhang. 2024b. Gme: Improving universal multimodal retrieval by multimodal llms. *arXiv preprint arXiv:2412.16855*.
- Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alex Smola. 2023. Multimodal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*.
- Yida Zhao, Yuqing Song, and Qin Jin. 2022. Progressive learning for image retrieval with hybrid-modality queries. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, pages 1012–1021.
- Ge Zheng, Bin Yang, Jiajin Tang, Hong-Yu Zhou, and Sibe Yang. 2023. Ddcot: Duty-distinct chain-of-thought prompting for multimodal reasoning in language models. *Advances in Neural Information Processing Systems*, 36:5168–5191.
- Zexin Zheng, Huangyu Dai, Lingtao Mao, Xinyu Sun, Zihan Liang, Ben Chen, Yuqing Ding, Chenyi Lei, Wenwu Ou, Han Li, and 1 others. 2025. Onevision: An end-to-end generative framework for multi-view e-commerce vision search. *arXiv preprint arXiv:2510.05759*.

A Entropy-Aware Memory Bank Strategy Flowchart

Our approach employs a **static-dynamic decoupling mechanism**: the memory bank stores static inputs (image captions $T(I_t)$ generated by MLLMs and frozen image embeddings from ViT), while during training, embeddings for negative sampling are dynamically recomputed using the current Q-Former. This ensures all representations remain consistent with the current model state, resolving the temporal inconsistency problem in traditional memory banks.

Algorithm 1 presents our entropy-aware memory bank update strategy in five key steps: First, it extracts batch embeddings \mathcal{B} via ViT (Step 1). Second, it computes similarity-based probability distributions between batch samples and memory samples, as well as within the memory bank itself (Step 2). Third, information entropy $H_i^{\mathcal{B}}$ and $H_i^{\mathcal{M}}$ are calculated to quantify the diversity of each sample—higher entropy indicates greater dissimilarity to existing memory samples, making them informative hard negatives (Step 3). Fourth, a retention score $\hat{H}_i^{\mathcal{M}}$ is computed by jointly considering both diversity (entropy factor) and temporal freshness (decay factor based on Δt_i), ensuring stale samples are prioritized for replacement (Step 4). Finally, high-entropy batch samples replace low-scoring memory samples when $H_j^{\mathcal{B}} > \hat{H}_i^{\mathcal{M}}$, maintaining a diverse and temporally consistent negative sample pool (Step 5). This comprehensive strategy effectively maintains memory bank quality while addressing representation inconsistency, enabling efficient large-scale contrastive learning for CIR tasks.

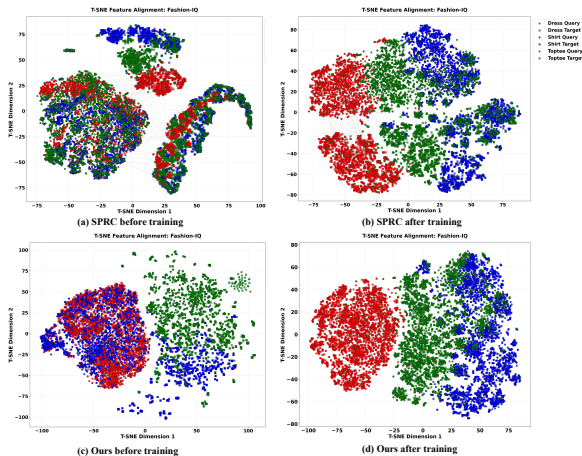


Figure 7: T-SNE visualization comparing SPRC and CSMCIR on the Fashion-IQ’s all three categories.

Algorithm 1 Entropy-Aware Memory Bank Update

```

1: Input: Batch images  $\{I_i\}_{i=1}^B$ , image captions  $T(I_t)$ ,
   memory bank  $\mathcal{M}$ 
2: Output: Updated memory bank  $\mathcal{M}$ 
3: // Step 1: Extract batch embeddings via ViT
4:  $\mathcal{B} = \{\mathbf{z}_1, \dots, \mathbf{z}_B\}$  where  $\mathbf{z}_i = \text{ViT}(I_i)$ 
5: // Step 2: Compute similarity-based probability distributions
6: for each batch sample  $i \in [1, B]$  do
7:    $p_{i,j}^{\mathcal{B} \rightarrow \mathcal{M}} = \exp(\mathbf{z}_i^T \mathbf{m}_j) / \sum_{k=1}^M \exp(\mathbf{z}_i^T \mathbf{m}_k), \forall j \in [1, M]$ 
8: end for
9: for each memory sample  $i \in [1, M]$  do
10:   $p_{i,j}^{\mathcal{M} \rightarrow \mathcal{M}} = \exp(\mathbf{m}_i^T \mathbf{m}_j) / \sum_{k=1}^M \exp(\mathbf{m}_i^T \mathbf{m}_k), \forall j \in [1, M]$ 
11: end for
12: // Step 3: Calculate information entropy for diversity
13: for each batch sample  $i \in [1, B]$  do
14:    $H_i^{\mathcal{B}} = -\sum_{j=1}^M p_{i,j}^{\mathcal{B} \rightarrow \mathcal{M}} \log p_{i,j}^{\mathcal{B} \rightarrow \mathcal{M}}$ 
15: end for
16: for each memory sample  $i \in [1, M]$  do
17:    $H_i^{\mathcal{M}} = -\sum_{j=1}^M p_{i,j}^{\mathcal{M} \rightarrow \mathcal{M}} \log p_{i,j}^{\mathcal{M} \rightarrow \mathcal{M}}$ 
18: end for
19: // Step 4: Compute retention score with temporal decay
20: for each memory sample  $i \in [1, M]$  do
21:    $\hat{H}_i^{\mathcal{M}} = \max(0, 1 - \Delta t_i / N_{\max}) \cdot H_i^{\mathcal{M}}$ 
22: end for
23: // Step 5: Replace low-scoring memory samples
24: Sort batch samples by  $H^{\mathcal{B}}$  (descending)
25: Sort memory samples by  $\hat{H}^{\mathcal{M}}$  (ascending)
26: for each high-entropy batch sample  $j$  do
27:   if  $H_j^{\mathcal{B}} > \hat{H}_i^{\mathcal{M}}$  for lowest-scoring memory sample  $i$  then
28:      $\mathbf{m}_i \leftarrow \mathbf{z}_j, T(I_t)_i \leftarrow T(I_t)_j, \Delta t_i \leftarrow 0$ 
29:   end if
30: end for
31: return Updated  $\mathcal{M}$ 

```

B Visualization

B.1 MCoT Cases Visualization

Fig. 8 shows examples of our MCoT-generated captions for target images across Fashion-IQ, CIRR, and Shoes datasets. Through MCoT, we successfully generate concise yet informative descriptions that accurately capture object characteristics while effectively mimicking the style of manipulation text. This approach ensures that generated captions maintain format consistency with query text and prevent the introduction of misleading or extraneous content.

B.2 T-SNE Visualization on Fashion-IQ dataset

To validate that our method effectively mitigates the representation space fragmentation problem, we visualize query-target feature distributions on the Fashion-IQ dataset using t-SNE. As shown in



Figure 8: Visualization of our MCoT-generated captions for target images.

Fig. 7 a-b, directly evidencing the space fragmentation caused by heterogeneous modalities and asymmetric encoders. In contrast, our method achieves significantly tighter clustering both before and after training (Fig. 7 c-d), demonstrating that introducing target captions establishes modal symmetry and bridges the modality gap. Moreover, query and target embeddings are positioned substantially closer under our symmetric architecture compared to SPRC, confirming superior alignment through consistent feature representations enabled by the shared-parameter Q-Former.