

Heterogeneous Uncertainty-Guided Composed Image Retrieval with Fine-Grained Probabilistic Learning

Haomiao Tang¹, Jinpeng Wang^{2*}, Minyi Zhao³, Guanghao Meng¹, Ruisheng Luo¹, Long Chen⁴, Shu-Tao Xia¹

¹ Tsinghua Shenzhen International Graduate School, Tsinghua University

² Harbin Institute of Technology, Shenzhen

³ Fudan University

⁴ The Hong Kong University of Science and Technology
thm23@mails.tsinghua.edu.cn, wjp20@mails.tsinghua.edu.cn

Abstract

Composed Image Retrieval (CIR) enables image search by combining a reference image with modification text. Intrinsic noise in CIR triplets incurs intrinsic uncertainty and threatens model’s robustness. Probabilistic learning approaches have shown promise in addressing such issues; however, they fall short for CIR due to their instance-level holistic modeling and homogeneous treatments for queries and targets. This paper introduces a **Heterogeneous Uncertainty-Guided (HUG)** paradigm to overcome these limitations. HUG utilizes a *fine-grained* probabilistic learning framework, where queries and targets are represented by Gaussian embeddings capturing detailed concepts and uncertainties. We customize *heterogeneous* uncertainty estimations for multi-modal queries and uni-modal targets. Given a query, we capture uncertainties not only regarding uni-modal content quality but also multi-modal coordination, followed by a *provable* dynamic weighting mechanism to derive the comprehensive query uncertainty. We further design uncertainty-guided objectives, including query-target holistic contrast and fine-grained contrasts with comprehensive negative sampling strategies, which effectively enhance discriminative learning. Experiments on benchmarks demonstrate HUG’s effectiveness beyond state-of-the-art baselines, with faithful analysis justifying the technical contributions.

Code — <https://github.com/tanghme0w/AAAI26-HUG>

Introduction

Composed Image Retrieval (CIR) (Vo et al. 2019) is an emerging topic in multimedia retrieval that allows searching for images with multi-modal queries comprising reference images and modification texts. It allows users to articulate complex visual preferences that might be difficult to express through text or images alone, which facilitates personalization and is favorable in various e-commerce applications and social media (Wu et al. 2021). Despite the practical value, CIR is more challenging than classic uni-modal (Bowyer and Flynn 2000; Wan et al. 2014; Dubey 2020;

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

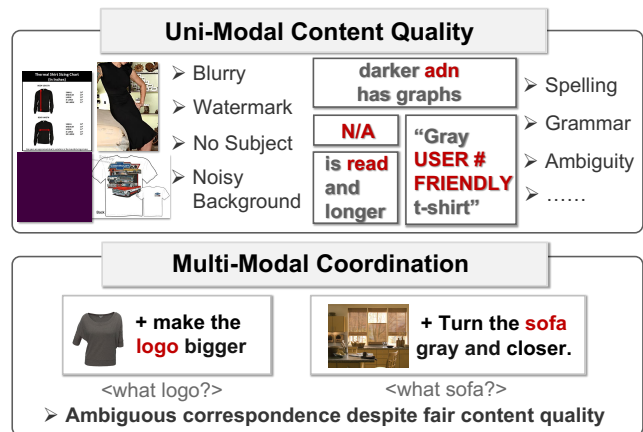


Figure 1: In Composed Image Retrieval, the uncertain multi-modal coordination between the reference image and modification text is also important in representation learning.

Lian et al. 2025) or cross-modal (Lee et al. 2018; Li et al. 2019; Wang et al. 2022b, 2024a; Meng et al. 2026) retrieval tasks on learning robust representations. This inherently results in uncertainty that threatens the robustness of search models.

The uncertainty in CIR is heterogeneous. We can characterize it by two typical forms, as exemplified in Figure 1:

- (i) **Content Quality.** Low-quality elements, such as blurry images or uninformative texts, are hard to avoid in CIR.
- (ii) **Multi-Modal Coordination within Queries.** In CIR, the multi-modal nature of queries raises a particular coordination issue. Even if an image and its accompanying text are considered high-quality individually, there may still be an ambiguous correspondence or mismatch.

Note that related works in multi-modal retrieval (Song and Soleymani 2019; Chun et al. 2021; Andrei, Chen, and Akata 2022; Chun 2024; Tang et al. 2025a) have provided some inspiration by probabilistic embedding learning (Abdar et al. 2020; Oh et al. 2019), which helps to identify and handle some of the above issues via uncertainty estimation. How-

ever, existing solutions still exhibit two major drawbacks when applied to CIR. Firstly, they typically operate at an *instance granularity*, failing to capture the complex and fine-grained user intents in CIR. Secondly, they apply *homogeneous strategies* to both query and target sides, since both are uni-modal. This may not be the best practice in CIR because the multi-modal coordination uncertainty issue at the query side requires further remedy.

In this paper, we propose a **Heterogeneous Uncertainty-Guided paradigm (HUG)** to comprehensively address these issues. HUG is carefully developed with a *fine-grained* probabilistic learning framework, representing each query and target image as a series of Gaussian embeddings. Each Gaussian aims to describe a fine-grained detail and capture a latent concept in the intricate matching space. The variance reflects the fine-grained uncertainty, allowing models to prioritize certain details while mitigating the adverse effects of fuzzy ones in the matching process. To better target CIR, we develop *heterogeneous* uncertainty estimation for the uni-modal target and the multi-modal query: while the target side only needs to model the content quality uncertainty, the query side further considers the multi-modal coordination uncertainty between the reference image and the modification text. In particular, to obtain overall query uncertainty, we combine the text- and image-specific content quality uncertainties as well as the multi-modal coordination uncertainty through a provable dynamic weighting mechanism. Guided by the established estimations, we introduce uncertainty-aware contrastive loss to learn discriminative holistic matching between queries and targets. Moreover, we further design uncertainty-guided fine-grained contrast for each Gaussian embedding, incorporating *component-, instance-, and modality-wise* negative sampling strategies to supplement robust learning signals.

We conduct extensive experiments on standard CIR benchmarks, showing HUG’s effectiveness against state-of-the-art baselines. Besides, we present a detailed model analysis, examining contributions of the key designs in HUG, including fine-grained representation, heterogeneous uncertainty estimation, and uncertainty-guided objectives. Moreover, the quantitative study of the learned representations via HUG reveals that each component of uncertainty can intuitively reflect image or text attributes, such as color, logo, or sleeve length, while the magnitude of uncertainty closely correlates with the ambiguity of these aspects. These intuitive findings highlight an intriguing interpretability in HUG.

To summarize, we make the following contributions:

- **Fine-grained probabilistic representation:** We represent each query and target image as a series of Gaussian embeddings to better capture attribute-level details, where variances reflect fine-grained uncertainties and help prioritize important details during the matching process.
- **Heterogeneous uncertainty estimation:** For the uni-modal target, we focus on content quality uncertainty; for the multi-modal query, we consider both content quality and multi-modal coordination uncertainty, which are integrated via a *provable* dynamic weighting mechanism.
- **Uncertainty-guided learning objectives:** Beyond holistic contrast, we introduce fine-grained contrastive loss, incor-

porating *component-, instance-, and modality-wise* negative sampling strategies to enhance learning efficacy.

- **Empirical results:** Benchmark results validate HUG’s superiority to state-of-the-art. Model analyses justify key designs. Quantitative study highlights the interpretability.

Related Works

Composed Image Retrieval (CIR)

CIR has two primary directions. Supervised CIR (Wang et al. 2022a; Zhang et al. 2022; Zhao, Song, and Jin 2022; Baldrati et al. 2022; Wen et al. 2023; Yang et al. 2023; Xu et al. 2023; bai et al. 2024) uses triplet training (reference image, modification text, target image) to fuse features and capture visual transformations. Zero-shot CIR (Baldrati et al. 2023; Saito et al. 2023; Tang et al. 2024; Lin et al. 2024; Suo et al. 2024; Wang et al. 2025; Li, Ma, and Yang 2025; Tang et al. 2025b,c) trains on independent image-text pairs, converting image features to pseudo-text but lacking triplet supervision, resulting in lower accuracy. We focus on supervised CIR.

Despite triplet supervision benefits, supervised CIR faces data quality issues (noise, ambiguity). Recent solutions include high-quality data refinement (Jang et al. 2024; Gu et al. 2024; Feng, Zhang, and Nie 2024; Ventura et al. 2024), semantic decomposition (Yang et al. 2024; Lin et al. 2024; Tian et al. 2025), LLM-based intent clarification (Baldrati et al. 2023; Karthik et al. 2024; Tian et al. 2025; Huynh et al. 2025), and regularization for ambiguous queries (Chen et al. 2024; Xu et al. 2024). Unlike methods eliminating uncertainty, our approach uses probabilistic embeddings to model uncertainties explicitly, incorporating them in training/inference for richer representations and robust training.

Uncertainty Learning

Uncertainty quantifies the likelihood that a model’s prediction may be incorrect. Two key uncertainty sources exist (Kiureghian and Ditlevsen 2009): (i) *Epistemic Uncertainty* (reduced by more data/improved architecture) and (ii) *Aleatoric Uncertainty* (from inherent data ambiguity, inevitable even with more data (Kendall and Gal 2017)). This work focuses on aleatoric uncertainty in CIR, aiming to quantify per-sample uncertainty under fixed data constraints.

To explore aleatoric uncertainty in computer vision, early image classification work (Shi and Jain 2019; Chang et al. 2020; Oh et al. 2019) used probabilistic distributions (instead of deterministic points) via lightweight uncertainty heads on pre-trained models, enhancing robustness and accuracy (Wang et al. 2024b; Fang et al. 2025). Subsequent research (Song and Soleymani 2019; Chun et al. 2021; Chun 2024) extended this to cross-modal retrieval. However, these methods use late fusion of independent unimodal predictions (Gao et al. 2024; Chen et al. 2024), neglecting modality interaction uncertainty and using coarse-grained instance-level estimation—limiting capture of complex dynamics critical to CIR (e.g., concept modification).

Closely related is (Xu et al. 2024), addressing CIR’s many-to-many correspondence and sparse annotations via identical uncertainty estimation for queries/targets and

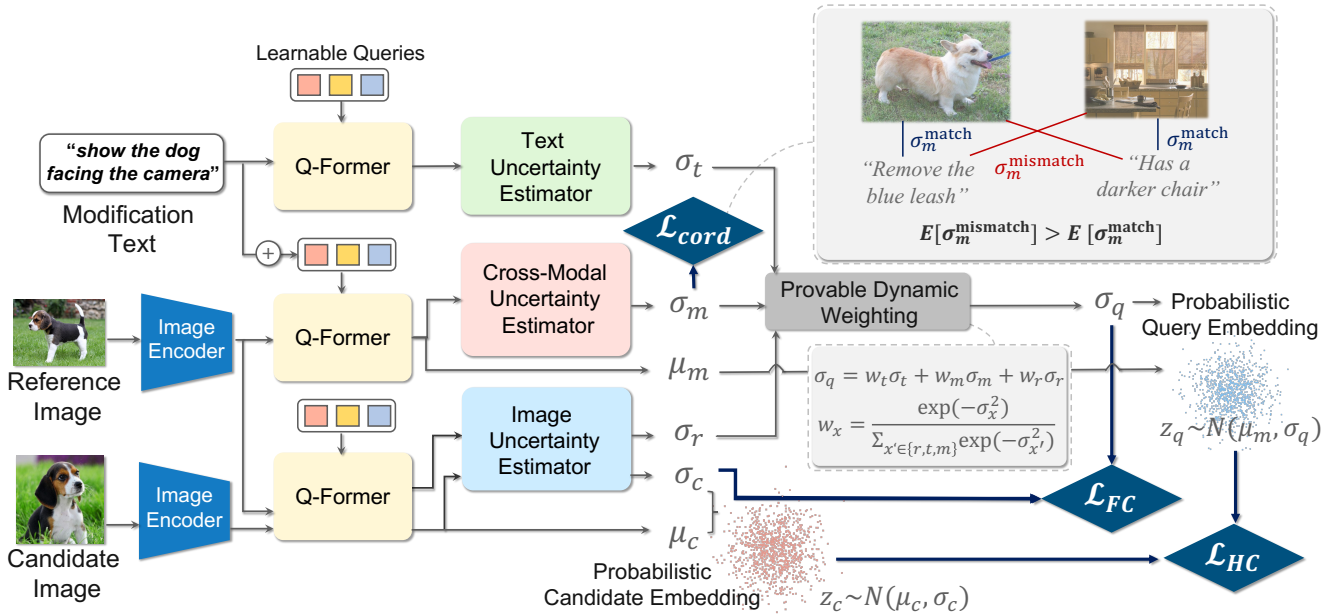


Figure 2: **H**eterogeneous **U**ncertainty-**G**uided (**HUG**) CIR. Modules with the same name share the same weights.

Monte Carlo sampling. Our approach differs in: (i) a heterogeneous uncertainty framework for queries-side capturing multi-modal coordination uncertainty; (ii) a closed-form uncertainty-aware distance metric computing expected query-target distance, improving efficiency and stability.

Our Solution

Problem Formulation and Method Overview

Composed Image Retrieval (CIR) operates on triplet data. Given a triplet (x_r, x_t, x_c) , where x_r denotes the reference image, x_t denotes the attached modification text, and x_c denotes the matched target image. The goal of CIR models is to learn a pair of encoders, f_q and f_c , producing multi-modal query representation $z_q = f_q(x_r, x_t)$ and image target representation $z_c = f_c(x_c)$, such that the query is closer to the target image than to any other candidate images:

$$d(z_q, z_c) < d(z_q, z_{c'}), \quad z_c \neq z_{c'}. \quad (1)$$

$d(\cdot, \cdot)$ denotes the distance metric.

Considering various forms of uncertainties caused by data noise in CIR, we propose a **H**eterogeneous **U**ncertainty-**G**uided paradigm (**HUG**) based on probabilistic learning. Specifically, we represent each query and target as a series of Gaussian embeddings. Take the query (x_r, x_t) as an example, its representation z_q is defined by $[z_q^1, z_q^2, \dots, z_q^K]$, where the k -th sub-representation is parameterized by a Gaussian, namely $z_q^k \sim \mathcal{N}(\mu_q^k, \Sigma_q^k)$, $\mu_q^k \in \mathbb{R}^D$, $\Sigma_q^k \in \mathbb{R}^{D \times D}$. For computation efficiency, we follow common practice (Song and Soleymani 2019; Chun et al. 2021; Chun 2024) that simplifies the covariance matrix Σ_q^k as a diagonal matrix by assuming dimensional mutual independence, and thus $z_q^k \sim \mathcal{N}(\mu_q^k, \sigma_q^{k,2} \mathbf{I})$. $\sigma_q^{k,2} \in \mathbb{R}^D$ is the variance vector reflecting the uncertainty and \mathbf{I} denotes the identity matrix.

As shown in Figure 2, we use BLIP-2's Q-Former (Li et al. 2023) to extract the fine-grained mean vectors via Q-Former's learnable query tokens, where each of the $K = 32$ tokens corresponds to a Gaussian. On the query side, the visual information is extracted with a pre-trained and fixed visual backbone and injected as the key and value in the cross-attention layers. We formulate this process by

$$\mu_q = h(x_{[LQ]}, x_t, x_r) \in \mathbb{R}^{32 \times D}, \quad (2)$$

where $h(\cdot, \cdot, \cdot)$ denotes the shared Q-Former. $x_{[LQ]} \in \mathbb{R}^{32 \times D}$ denotes learnable query tokens in Q-Former. On the target side, we leave the modification text blank and extract target mean vectors as

$$\mu_c = h(x_{[LQ]}, \emptyset, x_c) \in \mathbb{R}^{32 \times D}. \quad (3)$$

In the following sub-sections, we will first present the heterogeneous uncertainty estimation strategies on the query and target sides, after which we will introduce the uncertainty-guided learning framework and the objectives.

Heterogeneous Uncertainty Estimation

Unlike common multi-modal retrieval tasks, CIR features an asymmetric matching between multi-modal queries and unimodal targets. Thus, we estimate the uncertainties in a heterogeneous manner.

Target Uncertainty regarding Visual Content Quality

On the target side, variance parameters σ_c^2 indicate fine-grained content quality and visual informativeness from various aspects. Following (Chun 2024), we employ a 1-layer light-weight Transformer block upon Q-Former's output as the uncertainty (*i.e.*, variance) estimator,

$$\sigma_c^2 = g_V(\mu_c) \in \mathbb{R}^{32 \times D}, \quad (4)$$

where g_V denotes the visual uncertainty estimator.

Query-Side Uncertainties regarding Uni-modal Quality and Multi-modal Coordination On the query side, we consider more comprehensive uncertainty estimation. We regard the combination of the reference image and the modification text as the text-conditioned image representation:

$$z_q = f_q(x_r, x_t) = f_{x_t}(x_r), \quad (5)$$

where we characterize three uncertainties: **(i) Uncertainty in the reference image** x_r : content quality and visual informativeness of the reference image; **(ii) Uncertainty in the modification text** x_t : clarity and specificity of the textual modification; **(iii) Uncertainty in the text-conditioned function** $f_{x_t}(\cdot)$: coordination between modification intent and the reference image. We argue that uncertainty caused by the modifier function $f_{x_t}(\cdot)$ arises from intrinsic interactions between the reference image and the modification text, which is beyond the naive combination of uni-modal uncertainties. Accordingly, we extract the uncertainty factors as follows

$$\sigma_r^2 = g_V(h(x_{[\text{LQ}]}, \emptyset, x_r)) \in \mathbb{R}^{32 \times D}, \quad (6)$$

$$\sigma_t^2 = g_T(h(x_{[\text{LQ}]}, x_t, \emptyset)) \in \mathbb{R}^{32 \times D}, \quad (7)$$

$$\sigma_m^2 := \sigma_m^2(x_r, x_t) = g_M(\mu_q) \in \mathbb{R}^{32 \times D}. \quad (8)$$

g_V is the visual uncertainty estimator, sharing weights with its target-side counterpart. g_T and g_M are the textual and multi-modal uncertainty estimators, adopting the same model architecture but being independently parameterized.

To shape a more precise estimation of multi-modal coordination uncertainty, we further introduce regularization for the multi-modal uncertainty estimator. Intuitively, image-text pairs from the same triplet should exhibit lower coordination uncertainty than those from different triplets. Based on this insight, we design a ranking loss that discriminates the estimated uncertainty of image-text pairs within the same triplet from those across different triplets:

$$\mathcal{L}_{\text{Cord.}} = -\mathbb{E}_{(x_r, x_t, x_c) \neq (x'_r, x'_t, x'_c)} \log \mathcal{S}(\bar{\sigma}_m^2(x_r, x_t) - \bar{\sigma}_m^2(x_r, x'_t)), \quad (9)$$

where $\mathcal{S}(x) = \frac{1}{1+e^{-x}}$ is the sigmoid function. $\bar{\sigma}_m^2(x_c, x_t)$ denotes the mean value of multi-modal coordination uncertainty between x_c and x_t . The devised $\mathcal{L}_{\text{Cord.}}$ encourages the multi-modal uncertainty estimator to predict a higher uncertainty when the correspondence between the reference image and modification text is low or ambiguous.

Summarized Query Uncertainty via Dynamic Weighting

We combine the multi-modal coordination uncertainty σ_m^2 with the uni-modal uncertainties σ_r^2 (reference image) and σ_t^2 (text) in an element-wise manner to establish the overall query uncertainty. This combination is conducted on each of the fine-grained uncertainty components in parallel. Since the uncertainties in different aspects are expected to be decoupled and mutually independent, we combine them with a linear combination with dynamic weighting, formulated as

$$\sigma_q^k[i]^2 = \sum_{x \in \{r, t, m\}} w_x^k[i] \cdot \sigma_x^k[i]^2, \quad 1 \leq k \leq 32, 1 \leq i \leq D. \quad (10)$$

The fusion weights are input-adaptive:

$$w_x^k[i] = \frac{\exp(-\sigma_x^k[i]^2)}{\sum_{x' \in \{r, t, m\}} \exp(-\sigma_{x'}^k[i]^2)}, \quad (11)$$

and satisfy: $w_x^k[i] \geq 0$, $\sum_{x \in \{r, t, m\}} w_x^k[i] = 1$. Inspired by Zhang et al. (2023), we can prove that dynamic fusion using Equation (11) yields a tighter generalization error bound than using any static fusion weights.

Proposition 1 (Generalization Error Bounds). *Consider a loss function ℓ that is convex w.r.t. **scalar** variance values σ_x^2 , $x \in \{r, t, m\}$. Given a training set \mathcal{D} of size N , let $\hat{\mathbb{E}}[\ell(\sigma_x^2)] := \frac{1}{N} \sum_{n=1}^N \ell(\sigma_x^2(n))$ be the empirical estimate of the expected generalization loss across all data, then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, the following generalization error bound holds:*

$$\mathcal{E} \leq \sum_{x \in \{r, t, m\}} [\mathbb{E}(w_x) \cdot \hat{\mathbb{E}}[\ell(\sigma_x^2)] + \mathbb{E}(w_x) \cdot \mathfrak{R}_x(\ell(\sigma_x^2)) + \text{Cov}(w_x, \ell(\sigma_x^2))] + 3\sqrt{\frac{\ln(1/\delta)}{2N}}. \quad (12)$$

where $\mathbb{E}(w_x)$ is the expectation of fusion weights, $\mathfrak{R}_x(\ell(\sigma_x^2))$ is the Rademacher complexity, $\text{Cov}(w_x, \ell(\sigma_x^2))$ is the covariance between fusion weights and loss values.

Proof. See Appendix. \square

Corollary 1. *If all the following conditions hold: (i) ℓ is convex w.r.t. **scalar** σ_x^2 ; (ii) ℓ penalizes elements with large uncertainty values, i.e., $\rho(w_x^{\text{dynamic}}, \ell(\sigma_x^2)) < 0$, where ρ is the Pearson Correlation Coefficient; (iii) the expectation of dynamic weights is the same as the static weights for each modality, i.e., $\mathbb{E}(w_x^{\text{dynamic}}) = w_x^{\text{static}}$, then, dynamic weights fusion as eq. (11) will yield a strictly tighter generalization error bound than the static fusion, i.e., $\mathcal{E}_{\text{dynamic}} < \mathcal{E}_{\text{static}}$.*

Proof. See Appendix. \square

Remarks: Implications for HUG. We now analyze how the conditions stated in Corollary (1) are fulfilled by our proposed method. First, the sigmoid loss we adopted, which will be introduced as Equation (13), is convex and ensures the validity of condition (i). Condition (ii) is supported by the probabilistic learning scheme, which—according to the gradient-based analysis in (Chun et al. 2021)—leads to the down-weighting of items with higher predictive uncertainty during training. Lastly, the structure of Equation (11) guarantees the existence of a subset of dynamic weights w_x^{dynamic} such that the expectation satisfies $\mathbb{E}(w_x^{\text{dynamic}}) = w_x^{\text{static}}$, thereby meeting condition (iii). Taken together, these observations theoretically establish the superiority of dynamic weighting over static weighting.

Uncertainty-Guided Learning

Uncertainty-Guided Holistic Query-Target Contrast

We adopt a sigmoid contrastive loss (Zhai et al. 2023) to holistically align the query and target representations.

$$\begin{aligned} \mathcal{L}_{\text{HC}} = & -\mathbb{E}_{(x_r, x_t, x_c)} \log(\mathcal{S}(-a \cdot d(z_q, z_c) - b)) \\ & - B \cdot \mathbb{E}_{(x'_r, x'_t, x'_c) \neq (x_r, x_t, x_c)} \log(\mathcal{S}(a \cdot d(z_q, z'_c) + b)) \\ & - B \cdot \mathbb{E}_{(x'_r, x'_t, x'_c) \neq (x_r, x_t, x_c)} \log(\mathcal{S}(a \cdot d(z'_q, z_c) + b)), \end{aligned} \quad (13)$$

Method	Dress		Shirt		Top & Tee		Avg.		
	R@10	R@50	R@10	R@50	R@10	R@50	R@10	R@50	Avg.
CLIP4CIR (Baldrati et al. 2022)	33.81	59.40	39.99	60.45	41.41	65.37	38.40	61.74	50.07
ComqueryFormer (Li et al. 2024)	28.85	55.38	25.64	50.22	33.61	60.48	29.37	55.36	42.36
CRN (Yang et al. 2023)	32.67	59.30	30.27	56.97	37.74	65.94	33.56	60.74	47.15
FAME-ViL (Han et al. 2023)	42.19	67.38	47.64	68.79	50.69	73.07	46.84	69.75	58.29
MANME (Xu et al. 2023)	31.26	57.66	26.37	47.94	32.33	59.31	29.99	54.97	42.48
DWC (Huang et al. 2024)	32.67	57.96	35.53	60.11	40.13	66.09	36.11	61.39	48.75
CompoDiff [★] (Gu et al. 2024)	40.65	57.14	36.87	57.39	43.93	61.17	40.48	58.57	49.53
MGUR (Chen et al. 2024)	32.61	61.34	33.23	62.55	41.40	72.51	35.75	65.47	50.61
SSN (Yang et al. 2024)	34.36	60.78	38.13	61.83	44.26	69.05	38.92	63.89	51.40
BLIP4CIR+Bi (Liu et al. 2024)	42.09	67.33	41.76	64.28	46.61	70.32	43.49	67.31	55.40
SPIRIT (Chen, Zhou, and Peng 2024)	39.86	64.30	44.11	65.60	47.68	71.70	43.88	67.20	55.54
SADN (Wang et al. 2024c)	40.01	65.10	43.67	66.05	48.04	70.93	43.91	67.36	55.63
CaLa (Jiang et al. 2024)	42.38	66.08	46.76	67.28	50.93	74.11	46.69	69.16	57.92
CASE [★] (Levy et al. 2024)	48.48	70.23	47.44	69.36	50.18	72.24	48.70	70.61	59.66
CoVR [★] (Ventura et al. 2024)	-	-	-	-	-	-	49.40	70.98	60.19
VDG ^{★♠} (Jang et al. 2024)	47.89	69.81	51.36	71.08	53.29	74.65	50.85	71.85	61.35
QuRe (Kwak et al. 2025)	46.80	69.81	53.53	72.87	57.47	77.77	52.60	73.48	63.04
HUG (Ours)	<u>48.37</u>	71.56	<u>51.62</u>	74.41	58.26	78.22	52.75	74.73	63.74

Table 1: Comparison with existing methods on Fashion-IQ dataset. The best results are in bold font and second best results are underlined. Methods using extra data are marked with [★] and methods using an LLM with [♠].

where $\mathcal{S}(\cdot)$ is the Sigmoid function. B represents the proportion of negative samples to positive samples, typically set as the batch size. $d(\cdot, \cdot)$ denotes the uncertainty-aware holistic distance metric between queries and target images. a and b are two learnable parameters initialized by 1 and 0. Following related works (Shi and Jain 2019; Chang et al. 2020; Oh et al. 2019), we compute the uncertainty-aware distance as the expected Euclidean distance between two Gaussians. Consider two points $z_1 \sim \mathcal{N}(\mu_1, \sigma_1^2 \mathbf{I})$ and $z_2 \sim \mathcal{N}(\mu_2, \sigma_2^2 \mathbf{I})$, their expected Euclidean distance is:

$$\mathbb{E}_{z_1, z_2} [\|z_1 - z_2\|_2^2] = \|\mu_1 - \mu_2\|_2^2 + \|\sigma_1\|_2^2 + \|\sigma_2\|_2^2. \quad (14)$$

By applying the above distance metric to each fine-grained component of the query and target embeddings, we can derive the uncertainty-aware holistic distance metric as

$$d(z_q, z_c) = \|\mu_q - \mu_c\|_F^2 + \|\sigma_q\|_F^2 + \|\sigma_c\|_F^2. \quad (15)$$

Here, $\|\cdot\|_F$ denotes the Frobenius norm of the tensor.

Uncertainty-Guided Fine-Grained Contrast In order to align the fine-grained representations between the query and target, and promote the orthogonality and diversity of fine-grained uncertainty components, we introduce a contrastive strategy. Specifically, for the variance vector of the k -th fine-grained component, σ_M^k , the loss encourages differentiation:

$$\mathcal{L}_{FC} = -\sum_{M \in \{q, c\}} \sum_{k=1}^{32} \mathbb{E}_{\sigma_M^{k'} \neq \sigma_M^k} \left[\log \left(\mathcal{S} \left(a' \|\sigma_M^k - \sigma_M^{k'}\|_2^2 + b' \right) \right) \right], \quad (16)$$

where a' and b' are learnable. We employ three negative sampling strategies for $\sigma_M^{k'}$: (i) *Component-wise*: Negatives are other components of the same side and instance. (ii)




Figure 3: Model performance (average recall) on Fashion-IQ dataset under different settings of λ_{Cord} and λ_{FC} .

Instance-wise: Negatives are other components of the same side but different instances. (iii) *Modality-wise*: Negatives from other components of the other side and any instances.

Overall Learning Objectives Total learning objectives:

$$\mathcal{L}_{HUG} = \mathcal{L}_{HC} + \lambda_{FC} \mathcal{L}_{FC} + \lambda_{Cord} \mathcal{L}_{Cord}. \quad (17)$$

λ_{FC} and λ_{Cord} are loss balancing factors.

Experiments

Research Questions

We aim to answer the following research questions by conducting experiments on two standard CIR benchmarks:

RQ1: Compared to state-of-the-art approaches, how does HUG perform on CIR benchmarks?

RQ2: How does each design contribute in HUG?

RQ3: How to interpret the efficacy of our proposed HUG?

Method	Recall@K				Recall _{subset} @K			R@5+R _s @1 2
	K=1	K=5	K=10	K=50	K=1	K=2	K=3	
CIRPLANT (Liu et al. 2021)	19.55	52.55	68.39	92.38	39.20	63.03	79.49	45.88
CompoDiff [★] (Gu et al. 2024)	32.39	57.61	77.25	94.61	67.88	85.29	94.07	62.75
CASE [★] (Levy et al. 2024)	49.35	80.02	88.75	97.47	76.48	90.37	95.71	78.25
VDG ^{★♠} (Jang et al. 2024)	50.96	80.15	86.86	94.46	77.45	90.65	96.10	78.80
ComqueryFormer (Li et al. 2024)	25.76	61.76	75.90	95.13	51.86	76.26	89.25	56.81
CLIP4CIR (Baldrati et al. 2022)	38.53	69.98	81.86	95.93	68.19	85.64	94.17	69.09
MANME (Xu et al. 2023)	18.27	48.02	63.23	89.66	42.43	64.89	77.93	45.23
SPIRIT (Chen, Zhou, and Peng 2024)	40.32	75.10	84.16	96.88	73.74	89.60	<u>95.93</u>	74.42
SSN (Yang et al. 2024)	43.91	77.25	86.48	97.45	71.76	88.63	<u>95.38</u>	74.51
SADN (Wang et al. 2024c)	44.27	78.10	87.71	97.89	72.71	89.33	95.38	75.41
QuRe (Kwak et al. 2025)	52.22	<u>82.53</u>	<u>90.31</u>	98.17	<u>78.51</u>	<u>91.28</u>	96.48	<u>80.52</u>
HUG (Ours)	<u>51.09</u>	83.20	92.03	<u>97.89</u>	80.65	91.80	<u>95.93</u>	81.93

Table 2: Comparison with existing methods on CIRR dataset. The best results are in bold font and second best results are underlined. Methods using extra data are marked with [★] and methods using an LLM with [♠].

#	Experimental Variant	Dress		Shirt		Top & Tee		Average		Avg.	Avg. Time / Query (ms)
		R@10	R@50	R@10	R@50	R@10	R@50	R@10	R@50		
Baselines											
(0)	Point Matching	40.52	62.25	39.89	62.77	43.03	65.12	41.15	63.38	52.26	7.51
(1)	+ Probabilistic Embedding	42.74	64.40	44.74	65.71	47.52	67.55	45.00	65.89	55.44	10.08
Fine-Grained Uncertainty-Guided Learning											
(2)	+ Component-Wise Fine-Grained Contrast	44.28	64.86	48.89	67.70	51.61	70.45	48.26	67.67	57.97	20.69
(3)	+ Instance-Wise Fine-Grained Contrast	44.64	65.47	49.73	67.85	52.03	71.15	48.80	68.16	58.48	20.54
(4)	+ Modality-Wise Fine-Grained Contrast	45.13	66.83	49.27	68.97	53.85	71.92	49.42	69.24	59.33	20.73
Heterogeneous Uncertainty Estimation											
(5)	+ Cross-Modal Uncertainty	44.15	65.68	49.04	68.28	52.96	71.31	48.72	68.42	58.57	21.28
(6)	+ Multi-Modal Coordination Loss	47.82	70.28	51.27	73.96	57.69	77.62	52.26	73.95	63.11	21.19
(7)	+ Dynamic Weighting (Full Model)	48.37	71.56	51.62	74.41	58.26	78.22	52.75	74.73	63.74	21.35

Table 3: Ablation Study on the Fashion-IQ dataset. Variants in the table add components progressively from top to bottom. We conduct validation on a single A100 GPU and report the average retrieval time per query (inference + distance computation).

Interpretability of HUG (RQ3)

Experimental Setup

Datasets and Metrics We evaluate our model on two major benchmarks for composed image retrieval. Fashion-IQ (Wu et al. 2021) is a fashion dataset consisting of 18,000 training triplets and 6,016 validation triplets, with a total of 15,536 candidate images for validation. Model performance on this dataset is reported using the Recall@K metric for K=10 and K=50. CIRR (Liu et al. 2021) comprises 36,554 image triplets derived from 21,552 real-world photographs originally sourced from NLVR2. In addition to the conventional Recall@K metric, CIRR introduces a novel evaluation framework, Recall_{subset}@K, which assesses a model’s fine-grained ability to distinguish target images within small groups of six visually similar images.

Implementation Details We employ the pre-trained weights of BLIP-2 as the initial weights of the Q-Former. Training is conducted on a single A100-80G GPU with a batch size of 32 and an initial learning rate of 3×10^{-5} . We

implement an AdamW optimizer with parameters $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1.0 \times 10^{-7}$. Default hyper-parameter settings are $\lambda_{\text{Cord.}} = 0.1$ and $\lambda_{\text{FC}} = 0.5$ for eq. (17).

Comparison with State-of-The-Arts (RQ1)

We conduct comprehensive comparisons against state-of-the-art methods on both Fashion-IQ and CIRR datasets. As is shown in tables 1 and 2, HUG achieves significant improvement against existing SoTAs across both benchmarks, demonstrating the effectiveness of our proposed uncertainty-guided framework for composed image retrieval tasks. It is worth highlighting that HUG has outperformed methods that applies out-sourced data (e.g., videos, web images, AI generated images), as well as methods that utilize LLMs to refine or rewrite prompts. This indicates that under proper uncertainty-aware supervision, a CIR model can effectively identify noise within training data and achieve robust matching without the need for additional curated labels or LLM-based enhancements.

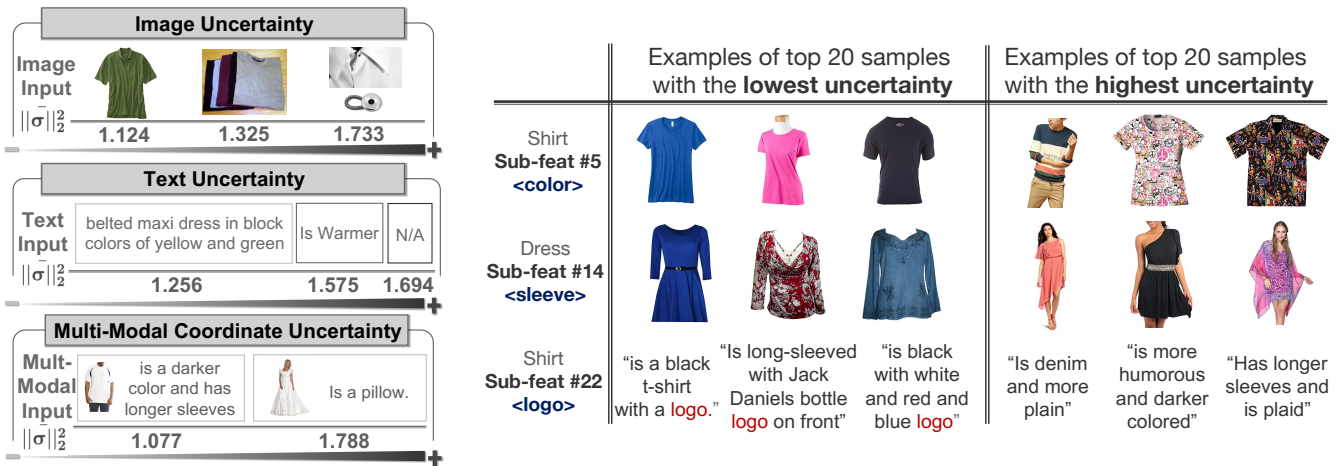


Figure 4: Qualitative analysis illustrating the meaning behind our learned uncertainty: (Left) Overall level of uncertainty reflects data quality. (Right) Different fine-grained uncertainty component corresponds to different sub-concepts.

Model Analyses (RQ2)

Ablation Study We compare configurations against a point matching baseline and a probabilistic embedding baseline. Baseline (0) aligns query image-text embeddings with target image embeddings using InfoNCE loss (He et al. 2020). Probabilistic baseline (1) uses equation (13) with generalized pooling (GPO) (Chen et al. 2021) for global uncertainty. Comparing (0) and (1) shows benefits of probabilistic uncertainty modeling.

Experiments (2,3,4) add fine-grained uncertainty-guided learning (equation (16)) using mean-pooled unimodal uncertainties. Consistent improvements over (1) confirm that finer uncertainty granularity enhances performance. The contrastive loss components are also verified.

Experiments (5,6,7) investigate heterogeneous uncertainty by integrating cross-modal with unimodal uncertainties. (4,5,6) show naive cross-modal inclusion degrades performance, while our multi-modal coordination loss is critical for improvement—highlighting the need to disentangle cross-modal and unimodal uncertainties. (6) vs (7) shows dynamic weighting outperforms static averaging in fusion.

Hyper-parameter Sensitivity We analyze the impact of two key hyper-parameters in our learning objective:

- *Coefficient of multi-modal coordination loss*, $\lambda_{\text{Cord.}}$. This coefficient balances the query-target ranking loss and the multi-modal coordination loss. As shown in Figures 3(a-c), setting $\lambda_{\text{Cord.}}$ to 0.1 allows it to act as an effective regularizer. However, increasing its weight too much leads to performance degradation.
- *Coefficient of fine-grained contrastive loss*, λ_{FC} . This coefficient controls the balance between the query-target ranking loss and the fine-grained uncertainty-contrastive loss. Figures 3(d-f) demonstrate that reducing λ_{FC} results in significant performance degradation, highlighting the importance of fine-grained contrastive loss in capturing meaningful fine-grained uncertainties.

Understanding overall uncertainty values. We analyze sample quality across uncertainty levels on Fashion-IQ. Overall uncertainty is defined as $\|\bar{\sigma}\|_2^2 = \frac{1}{K} \sum \|\sigma^k\|_2^2$ (average variance of fine-grained components). As shown in Figure 4 (left), higher overall uncertainty correlates with lower sample quality. Multi-modal coordinate uncertainty also reflects image-text correspondence ambiguity, confirming the model assesses both unimodal content quality and multi-modal interaction clarity.

Understanding fine-grained uncertainty values. We study sub-feature uncertainties via case studies: curating top/bottom 20 instances for each sub-feature (filtering high overall uncertainty samples). Qualitative analysis (Figure 4, right) reveals clear links between fine-grained uncertainties and real-world concepts. For example: Fashion-IQ Shirt’s 5th sub-feature connects to color; Similar phenomena also occur in Dress’s 14th sub-feature and Shirt’s 22nd sub-feature. This confirms the model effectively captures fine-grained concept uncertainty.

Conclusions

We propose a novel **Heterogeneous Uncertainty-Guided (HUG)** paradigm for Composed Image Retrieval (CIR). HUG represents both queries and targets as fine-grained Gaussian distributions, where the variances encode heterogeneous uncertainties. We apply a dynamic weighting mechanism that integrates uncertainty cues from content quality and cross-modal coordination, and formulate effective learning objectives for robust holistic and fine-grained matching. Extensive experiments demonstrate that HUG consistently outperforms prior approaches, offering resilience to noisy inputs. Our results highlight the critical role of uncertainty modeling in CIR, providing valuable insights for user-centric visual search systems and offering broader impact for related tasks like universal retrieval (Wei et al. 2023).

Acknowledgments

We sincerely thank the anonymous reviewers and chairs for their efforts and constructive suggestions, which have greatly helped us improve the manuscript. This work is supported in part by the National Natural Science Foundation of China under grants 624B2088 and 62571298. Long Chen was supported by the Hong Kong SAR RGC Early Career Scheme (26208924), and the National Natural Science Foundation of China Young Scholar Fund (62402408).

References

- Abdar, M.; Pourpanah, F.; Hussain, S.; Rezazadegan, D.; Liu, L.; Ghavamzadeh, M.; Fieguth, P. W.; Cao, X.; Khosravi, A.; Acharya, U. R.; Makarek, V.; and Nahavandi, S. 2020. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Inf. Fusion*.
- Andrei, N.; Chen, Y.; and Akata, Z. 2022. Probabilistic compositional embeddings for multimodal image retrieval. In *CVPR*.
- bai, Y.; Xu, X.; Liu, Y.; Khan, S.; Khan, F.; Zuo, W.; Goh, R. S. M.; and Feng, C.-M. 2024. Sentence-level prompts benefit composed image retrieval. In *ICLR*.
- Baldrati, A.; Agnolucci, L.; Bertini, M.; and Bimbo, A. D. 2023. Zero-shot composed image retrieval with textual inversion. In *ICCV*.
- Baldrati, A.; Bertini, M.; Uricchio, T.; and Bimbo, A. D. 2022. Conditioned and composed image retrieval combining and partially fine-tuning clip-based features. *CVPR Workshops*.
- Bowyer, K.; and Flynn, P. J. 2000. A 20th anniversary survey: Introduction to 'content-based image retrieval at the end of the early years'. *IEEE TPAMI*.
- Chang, J.; Lan, Z.; Cheng, C.; and Wei, Y. 2020. Data uncertainty learning in face recognition. In *CVPR*.
- Chen, J.; Hu, H.; Wu, H.; Jiang, Y.; and Wang, C. 2021. Learning the best pooling strategy for visual semantic embedding. In *CVPR*.
- Chen, Y.; Zheng, Z.; Ji, W.; Qu, L.; and Chua, T.-S. 2024. Composed image retrieval with text feedback via multi-grained uncertainty regularization. *ICLR*.
- Chen, Y.; Zhou, J.; and Peng, Y. 2024. Spirit: Style-guided patch interaction for fashion image retrieval with text feedback. *ACM TOMM*.
- Chun, S. 2024. Improved probabilistic image-text representations. In *ICLR*.
- Chun, S.; Oh, S. J.; Rezende, R. S. d.; Kalantidis, Y.; and Larlus, D. 2021. Probabilistic embeddings for cross-modal retrieval. In *CVPR*.
- Dubey, S. R. 2020. A decade survey of content based image retrieval using deep learning. *IEEE TCSVT*.
- Fang, H.; Zhou, C.; Kong, J.; Gao, K.; Chen, B.; Liang, T.; Ma, G.; and Xia, S.-T. 2025. Grounding Language with Vision: A Conditional Mutual Information Calibrated Decoding Strategy for Reducing Hallucinations in LVLMS. *NeurIPS*.
- Feng, Z.; Zhang, R.; and Nie, Z. 2024. Improving composed image retrieval via contrastive learning with scaling positives and negatives. In *MM*.
- Gao, Z.; Jiang, X.; Xu, X.; Shen, F.; Li, Y.; and Shen, H. T. 2024. Embracing unimodal aleatoric uncertainty for robust multimodal fusion. In *CVPR*.
- Gu, G.; Chun, S.; Kim, W.; Jun, H.; Kang, Y.; and Yun, S. 2024. Compodiff: Versatile composed image retrieval with latent diffusion. *TMLR*.
- Han, X.; Zhu, X.; Yu, L.; Zhang, L.; Song, Y.-Z.; and Xiang, T. 2023. Fame-vil: Multi-tasking vision-language model for heterogeneous fashion tasks. In *CVPR*.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *CVPR*.
- Huang, F.; Zhang, L.; Fu, X.; and Song, S. 2024. Dynamic weighted combiner for mixed-modal image retrieval. In *AAAI*.
- Huynh, C.; Yang, J.; Tawari, A.; Shah, M.; Tran, S. D.; Hamid, R.; Chilimbi, T.; and Shrivastava, A. 2025. Collm: A large language model for composed image retrieval. In *CVPR*.
- Jang, Y. K.; Kim, D.; Meng, Z.; Huynh, D.; and Lim, S.-N. 2024. Visual delta generator with large multi-modal models for semi-supervised composed image retrieval. In *CVPR*.
- Jiang, X.; Wang, Y.; Li, M.; Wu, Y.; Hu, B.; and Qian, X. 2024. Cala: Complementary association learning for augmenting composed image retrieval. In *SIGIR*.
- Karthik, S.; Roth, K.; Mancini, M.; and Akata, Z. 2024. Vision-by-language for training-free compositional image retrieval. In *ICLR*.
- Kendall, A.; and Gal, Y. 2017. What uncertainties do we need in bayesian deep learning for computer vision? In *NIPS*.
- Kiureghian, A. D.; and Ditlevsen, O. 2009. Aleatory or epistemic? does it matter? *Structural Safety 2009*.
- Kwak, J.; Inhar, R. M. I.; Yun, S.-Y.; and Lee, S.-J. 2025. Qure: Query-relevant retrieval through hard negative sampling in composed image retrieval. In *ICML*.
- Lee, K.-H.; Chen, X.; Hua, G.; Hu, H.; and He, X. 2018. Stacked cross attention for image-text matching. *ArXiv*.
- Levy, M.; Ben-Ari, R.; Darshan, N.; and Lischinski, D. 2024. Data roaming and quality assessment for composed image retrieval. In *AAAI*.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. C. H. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*.
- Li, K.; Zhang, Y.; Li, K.; Li, Y.; and Fu, Y. R. 2019. Visual semantic reasoning for image-text matching. In *ICCV*.
- Li, S.; Xu, X.; Jiang, X.; Shen, F.; Liu, X.; and Shen, H. T. 2024. Multi-grained attention network with mutual exclusion for composed query-based image retrieval. *IEEE TCSVT*.
- Li, Y.; Ma, F.; and Yang, Y. 2025. Imagine and seek: Improving composed image retrieval with an imagined proxy. In *CVPR*.

- Lian, N.; Li, J.; Wang, J.; Luo, R.; Wang, Y.; Xia, S.-T.; and Chen, B. 2025. AutoSSVH: Exploring Automated Frame Sampling for Efficient Self-Supervised Video Hashing. In *CVPR*.
- Lin, H.; Wen, H.; Song, X.; Liu, M.; Hu, Y.; and Nie, L. 2024. Fine-grained textual inversion network for zero-shot composed image retrieval. In *SIGIR*.
- Liu, Z.; Rodriguez-Opazo, C.; Teney, D.; and Gould, S. 2021. Image retrieval on real-life images with pre-trained vision-and-language models. In *ICCV*.
- Liu, Z.; Sun, W.; Hong, Y.; Teney, D.; and Gould, S. 2024. Bi-directional training for composed image retrieval via text prompt learning. In *WACV*.
- Meng, G.; Wang, J.; Wang, Q.-W.; Ren, X.; and Zhao, D. 2026. Imagine with Layout and Sketch: Enhancing Vision-Language Retrieval with Dual-Stream Multi-Modal Query Refinement. In *AAAI*.
- Oh, S. J.; Murphy, K. P.; Pan, J.; Roth, J.; Schroff, F.; and Gallagher, A. C. 2019. Modeling uncertainty with hedged instance embeddings. In *ICLR*.
- Saito, K.; Sohn, K.; Zhang, X.; Li, C.-L.; Lee, C.-Y.; Saenko, K.; and Pfister, T. 2023. Pic2word: Mapping pictures to words for zero-shot composed image retrieval. In *CVPR*.
- Shi, Y.; and Jain, A. K. 2019. Probabilistic face embeddings. In *ICCV*.
- Song, Y.; and Soleymani, M. 2019. Polysemous visual-semantic embedding for cross-modal retrieval. In *CVPR*.
- Suo, Y.; Ma, F.; Zhu, L.; and Yang, Y. 2024. Knowledge-enhanced dual-stream zero-shot composed image retrieval. In *CVPR*.
- Tang, H.; Wang, J.; Peng, Y.; Meng, G.; Luo, R.; Chen, B.; Chen, L.; Wang, Y.; and Xia, S.-T. 2025a. Modeling uncertainty in composed image retrieval via probabilistic embeddings. In *ACL*.
- Tang, Y.; Yu, J.; Gai, K.; Zhuang, J.; Xiong, G.; Gou, G.; and Wu, Q. 2025b. Missing target-relevant information prediction with world model for accurate zero-shot composed image retrieval. In *CVPR*.
- Tang, Y.; Yu, J.; Gai, K.; Zhuang, J.; Xiong, G.; Hu, Y.; and Wu, Q. 2024. Context-i2w: Mapping images to context-dependent words for accurate zero-shot composed image retrieval. In *AAAI*.
- Tang, Y.; Zhang, J.; Qin, X.; Yu, J.; Gou, G.; Xiong, G.; Lin, Q.; Rajmohan, S.; Zhang, D.; and Wu, Q. 2025c. Reason-before-retrieve: One-stage reflective chain-of-thoughts for training-free zero-shot composed image retrieval. In *CVPR*.
- Tian, L.; Zhao, J.; Hu, Z.; Yang, Z.; Li, H.; Jin, L.; Wang, Z.; and Li, X. 2025. Ccin: Compositional conflict identification and neutralization for composed image retrieval. In *CVPR*.
- Ventura, L.; Yang, A.; Schmid, C.; and Varol, G. 2024. Covr: Learning composed video retrieval from web video captions. In *AAAI*.
- Vo, N.; Jiang, L.; Sun, C.; Murphy, K.; Li, L.-J.; Fei-Fei, L.; and Hays, J. 2019. Composing text and image for image retrieval - an empirical odyssey. In *CVPR*.
- Wan, J.; Wang, D.; Hoi, S. C. H.; Wu, P.; Zhu, J.; Zhang, Y.; and Li, J. 2014. Deep learning for content-based image retrieval: A comprehensive study. *MM*.
- Wang, C.; Nezhadarya, E.; Sadhu, T.; and Zhang, S. 2022a. Exploring compositional image retrieval with hybrid compositional learning and heuristic negative mining. In *EMNLP*.
- Wang, J.; Chen, B.; Liao, D.; Zeng, Z.; Li, G.; Xia, S.-T.; and Xu, J. 2022b. Hybrid contrastive quantization for efficient cross-view video retrieval. In *WWW*.
- Wang, J.; Zeng, Z.; Chen, B.; Wang, Y.; Liao, D.; Li, G.; Wang, Y.; and Xia, S.-T. 2024a. Hugs bring double benefits: Unsupervised cross-modal hashing with multi-granularity aligned transformers. *IJCV*.
- Wang, L.; Ao, W.; Boddeti, V. N.; and Lim, S.-N. 2025. Generative zero-shot composed image retrieval. In *CVPR*.
- Wang, L.; Qin, Y.; Sun, Y.; Peng, D.; Peng, X.; and Hu, P. 2024b. Robust contrastive cross-modal hashing with noisy labels. In *MM*.
- Wang, Y.; Huang, W.; Li, L.; and Yuan, C. 2024c. Semantic distillation from neighborhood for composed image retrieval. In *MM*.
- Wei, C.; Chen, Y.; Chen, H.; Hu, H.; Zhang, G.; Fu, J.; Ritter, A.; and Chen, W. 2023. Uniir: Training and benchmarking universal multimodal information retrievers. *arXiv*.
- Wen, H.; Zhang, X.; Song, X.; Wei, Y.; and Nie, L. 2023. Target-guided composed image retrieval. In *MM*.
- Wu, H.; Gao, Y.; Guo, X.; Al-Halah, Z.; Rennie, S.; Grauman, K.; and Feris, R. 2021. Fashion iq: A new dataset towards retrieving images by natural language feedback. In *CVPR*.
- Xu, Y.; Bin, Y.; Wei, J.; Yang, Y.; Wang, G.; and Shen, H. T. 2023. Multi-modal transformer with global-local alignment for composed query image retrieval. *IEEE TMM*.
- Xu, Y.; Wei, J.; Bin, Y.; Yang, Y.; Ma, Z.; and Shen, H. T. 2024. Set of diverse queries with uncertainty regularization for composed image retrieval. *IEEE TCSVT*.
- Yang, Q.; Ye, M.; Cai, Z.; Su, K.; and Du, B. 2023. Composed image retrieval via cross relation network with hierarchical aggregation transformer. *IEEE TIP*.
- Yang, X.; Liu, D.; Zhang, H.; Luo, Y.; Wang, C.; and Zhang, J. 2024. Decomposing semantic shifts for composed image retrieval. In *AAAI*.
- Zhai, X.; Mustafa, B.; Kolesnikov, A.; and Beyer, L. 2023. Sigmoid loss for language image pre-training. In *ICCV*.
- Zhang, F.; Yan, M.; Zhang, J.; and Xu, C. 2022. Comprehensive relationship reasoning for composed query based image retrieval. In *MM*.
- Zhang, Q.; Wu, H.; Zhang, C.; Hu, Q.; Fu, H.; Zhou, J. T.; and Peng, X. 2023. Provable dynamic fusion for low-quality multimodal data. In *ICML*.
- Zhao, Y.; Song, Y.; and Jin, Q. 2022. Progressive learning for image retrieval with hybrid-modality queries. In *SIGIR*.