

# Cross Modal Fine-Grained Alignment via Granularity-Aware and Region-Uncertain Modeling

Jiale Liu<sup>1\*</sup>, Haoming Zhou<sup>1\*</sup>, Yishu Liu<sup>2</sup>, Bingzhi Chen<sup>3</sup>, Yuncheng Jiang<sup>1†</sup>

<sup>1</sup>South China Normal University, Guangzhou, China

<sup>2</sup>Harbin Institute of Technology, Shenzhen, China

<sup>3</sup>Beijing Institute of Technology, Zhuhai, China

liujiale, jiangyuncheng@m.scnu.edu.cn, zhouhaoming.smile@gmail.com, liuyishu@stu.hit.edu.cn, chenbingzhi@bit.edu.cn

## Abstract

Fine-grained image-text alignment is a pivotal challenge in multimodal learning, underpinning key applications such as visual question answering, image captioning, and vision-language navigation. Unlike global alignment, fine-grained alignment requires precise correspondence between localized visual regions and textual tokens, often hindered by noisy attention mechanisms and oversimplified modeling of cross-modal relationships. In this work, we identify two fundamental limitations of existing approaches: the lack of robust intra-modal mechanisms to assess the significance of visual and textual tokens, leading to poor generalization in complex scenes; and the absence of fine-grained uncertainty modeling, which fails to capture the one-to-many and many-to-one nature of region-word correspondences. To address these issues, we propose a unified approach that incorporates significance-aware and granularity-aware modeling and region-level uncertainty modeling. Our method leverages modality-specific biases to identify salient features without relying on brittle cross-modal attention, and represents region features as a mixture of Gaussian distributions to capture fine-grained uncertainty. Extensive experiments on Flickr30K and MSCOCO demonstrate that our approach achieves state-of-the-art performance across various backbone architectures, significantly enhancing the robustness and interpretability of fine-grained image-text alignment.

**Code** — <https://github.com/H3IloWorld/GRM>

## Introduction

Fine-grained image-text alignment is a challenging and essential task in multimodal learning which has significant applications in tasks like visual question answering (Antol et al. 2015; Yu et al. 2025), image captioning (Huang et al. 2019; Zeng et al. 2025), and vision-language navigation (Anderson et al. 2018; Chen et al. 2025). Unlike global image-text alignment, which considers the overall semantic similarity between an image and a text, fine-grained image-text alignment demands detailed alignment between localized visual representations and the corresponding textual tokens. Achieving this fine-grained alignment is challenging due to

\*These authors contributed equally.

†Corresponding author.

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

the need to reason over compositional details such as object attributes, spatial relations, and localized entities, which requires more precise modeling of the visual-textual correspondences.

Despite the remarkable advances in cross-modal alignment, existing methods (Wu et al. 2025; Zhao et al. 2024; Qiu, Ning, and He 2024) face two critical limitations that hinder their ability. First, while many approaches leverage cross-attention or fusion modules to model token-level interactions, the attention weights are typically driven by the retrieval objective. These attention maps are often noisy, lacking semantic grounding, and frequently focus on visually salient but semantically irrelevant regions. This results in a failure to explicitly identify which regions and tokens are truly critical for alignment, particularly in complex and ambiguous visual scenes. Without an effective mechanism to pinpoint the key areas for alignment, existing models struggle with misalignment, thereby impairing performance.

Second, uncertainty modeling has become increasingly recognized as crucial for capturing the inherent ambiguity in multimodal data. However, most recent methods treat uncertainty at the image-text pair level, assuming a one-to-one correspondence between image and text. In reality, the alignment is more nuanced: a single textual phrase may correspond to multiple regions, or a region may ambiguously match multiple text tokens. Yet, region-level uncertainty—capturing these one-to-many and many-to-one relationships—has been largely underexplored. This gap limits the ability of current methods to handle fine-grained uncertainty, which is critical for achieving robust and accurate cross-modal alignment.

Motivated by the limitations mentioned above, we raise the first fundamental question: *How can we effectively model the significance of embedding tokens within each modality?* Some existing approaches directly apply cross-modal attention to identify salient tokens (Fu et al. 2024). However, such methods are often inefficient and overly dependent on specific image-text pairs, which compromises generalizability. Therefore, we argue that significance modeling should be performed intra-modally, leveraging the inherent statistical biases within each modality rather than depending on cross-modal interactions. This modality-specific modeling can better adapt to diverse and complex downstream tasks. After obtaining modality-specific significance features, we

further raise a second fundamental question: *How can we model fine-grained uncertainty in the alignment process?* Inspired by the concept of probabilistic Distributional Representations, we adopt Gaussian distributions (Chang et al. 2020; Ji et al. 2023) to work out this problem. Specifically, we assume that each region feature can be represented as a multivariate Gaussian distribution, where the variance reflects the intrinsic uncertainty of the representation. To capture fine-grained uncertainty, we further model the image as a mixture of Gaussian distributions across multiple regions. Motivated by these two insights, we propose a unified framework **GRM** (Cross Modal Fine-Grained Alignment via **G**ranularity-Aware and **R**egion-Uncertain **M**odeling) that integrates granularity-aware significance modeling and region-level uncertainty modeling to enhance cross-modal alignment. Specifically, the main contributions of this work are threefold:

- We propose a Significance-aware and Granularity-aware Adapting approach that explicitly models intra-modal data biases to identify salient modality-specific features, thereby suppressing redundant information and enhancing alignment generalization.
- We introduce a prompt-driven mechanism for end-to-end region proposal, followed by the use of a mixture of Gaussians to model fine-grained region-level uncertainty, which improves the robustness of cross-modal alignment.
- We comprehensively evaluate our approach against existing fine-grained methods across diverse backbone architectures, consistently surpassing state-of-the-art performance on two benchmarks: Flickr30K and MS-COCO.

## Related Work

### Fine-grained Cross-modal Alignment

Fine-grained cross-modal alignment aims to establish precise correspondences between localized visual regions and textual elements, serving as a foundational task in multimodal understanding. Recent research (Faghri et al. 2017; Lee et al. 2018; Diao et al. 2021; Pan, Wu, and Zhang 2023; Fu et al. 2024; Liu et al. 2025; Pham et al. 2024; Fu et al. 2023) in this area typically follows a dual-stream architecture that separately encodes images and texts before aligning them through cross-modal interactions. Existing methods can be broadly categorized based on their visual encoders. One line of work relies on two-stage object detectors (Pham et al. 2024; Fu et al. 2023; Qin et al. 2022) (e.g., Faster R-CNN (Ren et al. 2016)) to extract region-level features as inputs to the alignment module. While such approaches offer explicit localization cues, they suffer from several drawbacks: the reliance on pre-trained detection models introduces error propagation and limits flexibility; in addition, adapting to downstream tasks often requires additional annotations and retraining, which increases the system’s complexity and domain sensitivity. In contrast, a more recent trend adopts end-to-end transformer-based vision encoders, particularly Vision Transformers (ViT) (Dosovitskiy et al. 2020; Liu et al. 2021), to directly model visual features with-

out explicit region proposals. This design benefits from unified optimization, reduced dependence on external annotations, and better scalability across datasets. Motivated by these advantages, we adopt a ViT-based architecture in our framework. Although CLIP-style (Lu et al. 2022; Wang et al. 2023; Jiang and Ye 2023; Zhang et al. 2024; Zheng et al. 2023) pretraining has proven effective for global image-text matching, its representation is typically optimized for coarse-level similarity and lacks the granularity required for fine-grained alignment. Therefore, we opt for independently pre-trained ViT and BERT (Devlin et al. 2019) models to retain modality-specific inductive biases and allow greater flexibility in downstream fine-grained alignment tasks.

### Uncertainty Quantification

Uncertainty modeling has recently gained attention in multimodal learning due to its potential to handle ambiguity, noise, and diverse interpretations inherent in image-text data. Prior works have explored uncertainty estimation through various paradigms, such as Bayesian neural networks (Kingma, Salimans, and Welling 2015; Molchanov, Ashukha, and Vetrov 2017; Kohl et al. 2018), probabilistic embeddings (Chen et al. 2022), and distributional representations (Gao et al. 2024). Broadly speaking, existing approaches can be categorized into two types: (1) Pair-level uncertainty modeling (Li et al. 2023), which estimates uncertainty at the global image-text pair level, often using confidence scores or variance-aware similarity metrics to account for ambiguity in the retrieval task. (2) Feature-level probabilistic modeling (Chen et al. 2022; Gao et al. 2024), which represents embeddings as probability distributions—typically Gaussians—to capture semantic variance and improve robustness in the embedding space. While these methods provide a useful first step toward modeling uncertainty, they predominantly focus on global or holistic representations, assuming a one-to-one alignment between an entire image and a sentence. This assumption fails to capture the fine-grained one-to-many and many-to-one relationships between visual regions and textual tokens, which are critical in tasks like phrase grounding and referring expression comprehension. As a result, global uncertainty modeling often overlooks localized ambiguities and leads to suboptimal alignment in complex scenes.

## Method

The overview of our GRM is illustrated in Figure 1. We first introduce the architecture of the modality encoders, followed by the detailed implementation of the Significance-aware and Granularity-aware Adapters. Finally, we present the methodological details of region prompting and fine-grained uncertainty modeling.

### Dual-encoder Feature Extraction

Our method adopts a dual-encoder architecture, wherein visual and textual features are independently extracted by modality-specific encoders without shared parameters. Given an input image instance  $I$ , we employ vision transformers (Dosovitskiy et al. 2020; Liu et al. 2021) as visual encoder  $f_v(\cdot)$  to obtain the initial visual representation

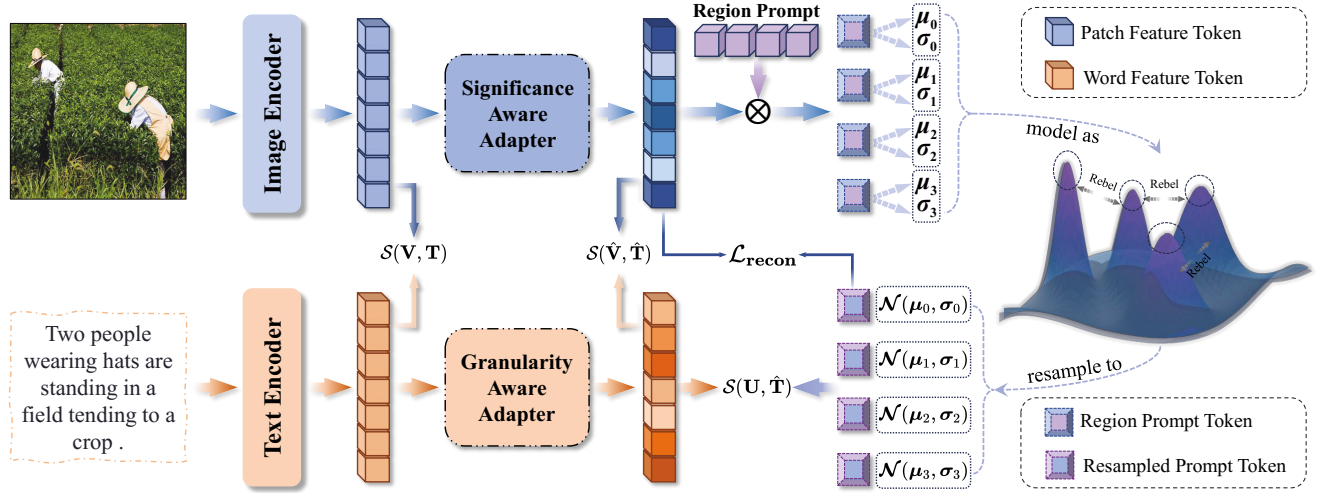


Figure 1: The Overview of the Proposed GRM. The visual encoder  $f_v(\cdot)$  and text encoder  $f_t(\cdot)$  independently encode input image and text instances to obtain their respective representations,  $\mathbf{V}$  and  $\mathbf{T}$ . These embeddings are then passed through two structurally identical but functionally distinct adapters: the Significance-aware Adapter and the Granularity-aware Adapter, which learn modality-specific distribution biases. Subsequently, the image embeddings undergo region-level prompt learning and uncertainty modeling to capture fine-grained semantic variations. Finally, a multi-level alignment strategy is applied to effectively align the cross-modal knowledge between images and texts.




Figure 2: The detailed architecture of the Significance-aware Adapter and the Granularity-aware Adapter.

$\mathbf{V} = f_v(I) \in \mathbb{R}^{L_v \times d}$ , where  $v_i$  denotes the  $i$ -th patch token generated by patch embedding of the visual encoder. Similarly, for a textual instance  $T$ , we utilize BERT (Devlin et al. 2019) as textual encoder  $f_t(\cdot)$  to obtain the corresponding textual representation  $\mathbf{T} = f_t(T) \in \mathbb{R}^{L_t \times d}$ , where  $t_i$  denotes the  $i$ -th word token generated by the tokenizer of the textual encoder.  $L_v$  and  $L_t$  denotes the length of image tokens and text tokens, and  $d$  indicates the latent dimension.

### Significance-aware and Granularity-aware Adapting

The architectures of the Significance-aware Adapter and Granularity-aware Adapter are illustrated in Figure 2. While both modules adopt an identical architectural design, they are instantiated independently and operate on different modalities, thereby fulfilling distinct functional roles. This modality-specific divergence, despite structural symmetry, reveals an intriguing dynamic that we empirically validate in the experimental section. Given the shared architecture be-

tween the two adapters, we present the detailed formulation using the visual input as a representative example in the following discussion. Given the visual representation  $\mathbf{V} \in \mathbb{R}^{L_v \times d}$  the feature adapting process is formulated as follows:

$$\begin{aligned} \mathbf{D}_v &= \mathbf{W}_2^{\mathbf{V}}(\psi(\mathbf{W}_1^{\mathbf{V}}\mathbf{V})) \in \mathbb{R}^{L_v \times 2} \\ \mathbf{A}_v &= f_g(\mathbf{D}_v, \tau)[:, 1] \in [0, 1]^{L_v} \\ \hat{\mathbf{V}} &= \mathbf{M} \odot \mathbf{A}_v \otimes \mathbf{1}_d \end{aligned} \quad (1)$$

Here,  $\odot$  and  $\otimes$  denotes element-wise multiplication and the outer product,  $\psi$  and  $f_g$  indicates the activation function (GELU in our implementation) and the Gumbel-Softmax function respectively. And  $\mathbf{1}_d$  is a  $d$ -dimensional all-one vector. The temperature parameter  $\tau$  is a scalar that controls the sharpness of the selection distribution.

### Region Prompting and Uncertainty Modeling

While prior studies (Gao et al. 2024; Li et al. 2023) have primarily modeled uncertainty at the holistic image-text level and achieved competitive performance, such formulations overlook the intrinsic granularity of cross-modal alignment. In contrast, we advocate for a fine-grained perspective, positing that matching uncertainty should be decomposed across localized image regions and individual textual tokens. This paradigm shift enables a more nuanced characterization of semantic ambiguity and forms the basis for our proposed uncertainty-aware fine-grained matching approach.

**Region Prompting.** Inspired by prompt learning (Jia et al. 2022; Pei et al. 2024; Wang et al. 2024), we first introduce a set of learnable prompts  $\mathbf{P} = \{p_0, \dots, p_{K-1}\} \in \mathbb{R}^{K \times d}$  to model the latent region representations within the image, where each prompt  $p_i$  serves as  $i$ -th semantic proxy




Figure 3: The computation process of bidirectional image-text alignment.

for a potential image region. We then transfer visual knowledge from image representations to these region prompts, thereby generating region-aware prompt embeddings with regional semantic information. To ensure scale invariance across queries, we apply L2 normalization for  $\mathbf{P}$  to get  $\hat{\mathbf{P}} = \{\hat{p}_0, \dots, \hat{p}_{K-1}\} \in \mathbb{R}^{K \times d}$ . Subsequently, we compute pairwise attention scores between each image patch token  $v_i$  and every region prompt  $p_k$  to capture their semantic affinity. Formally,

$$\mathbf{A}_r = \psi(\hat{\mathbf{V}} \cdot \hat{\mathbf{P}}^\top) \in \mathbb{R}^{L_v \times K},$$

$$\hat{\mathbf{A}}_r^{lk} = \frac{\mathbf{A}_r^{lk}}{\sum_{l'=1}^{L_v} \mathbf{A}_r^{l'k}}, \quad \hat{\mathbf{A}}_r \in \mathbb{R}^{L_v \times K}. \quad (2)$$

Here,  $\psi$  denotes the sigmoid function, as a single patch may be associated with multiple regions simultaneously. Finally, we leverage the normalized attention matrix  $\hat{\mathbf{A}}_r^{lk}$  to softly aggregate patch features into the region prompts, and compute the  $k$ -th mean prompt representation enriched with region-level semantics. The equation is expressed as:

$$\boldsymbol{\mu}_k = \sum_{l=1}^{L_v} \hat{\mathbf{A}}_r^{lk} \hat{\mathbf{V}}^l, \quad \boldsymbol{\mu} \in \mathbb{R}^{K \times d}. \quad (3)$$

**Uncertainty Modeling:** To enhance the expressiveness of region-level semantic modeling, we adopt a variational perspective (Kingma, Welling et al. 2013) by representing the semantics of each region as a Gaussian distribution. This formulation explicitly captures the inherent ambiguity and uncertainty in the semantic aggregation process. Through a reparameterization-based sampling mechanism, token-level semantics are injected into region-level representations, thereby improving the model’s adaptability and generalization capability to semantic structures.

Specifically, we employ a learnable prediction network  $\phi$  to estimate the log-variance  $\log \sigma^2$  for each region prompt, which serves as an indicator of the model’s confidence in the corresponding semantic representation. Formally,

$$\log \sigma_k^2 = \phi(\boldsymbol{\mu}_k), \quad \log \sigma^2 \in \mathbb{R}^{K \times d}. \quad (4)$$

We adopt the reparameterization trick to sample from each region’s distribution, thereby introducing diverse uncertainty and enhancing the expressiveness of region-level representations. Unlike conventional regional token embeddings that are treated as fixed vectors, we independently in-

ject region-level noise into each patch–region pair to construct patch-aware sampled representations. Formally,

$$\mathbf{z}_{lk} = \boldsymbol{\mu}_k + \boldsymbol{\epsilon}_{lk} \odot \exp\left(\frac{1}{2} \log \sigma_k^2\right), \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I}). \quad (5)$$

Here,  $\boldsymbol{\epsilon}$  denotes the noise sampled from a standard normal distribution  $\mathcal{N}(0, \mathbf{I})$ , where  $\mathbf{I}$  refers to the identity covariance matrix. Then we apply the normalized attention matrix to perform a weighted aggregation of the sampled features  $\mathbf{z}_{lk}$ , resulting in an uncertainty-aware semantic representation for each region  $\mathbf{u}_k$ :

$$\mathbf{u}_k = \sum_{l=1}^{L_v} \hat{\mathbf{A}}_{lk} \cdot \mathbf{z}_{lk}, \quad \mathbf{u} \in \mathbb{R}^d. \quad (6)$$

Finally, we obtain the uncertainty-aware region representation  $\mathbf{U} = \{\mathbf{u}_0, \dots, \mathbf{u}_{K-1}\} \in \mathbb{R}^{K \times d}$ .

### Multi-objective Optimization

#### Multi-level and Bidirectional Image-Text Alignment:

To comprehensively capture the semantic correspondences between images and texts, we introduce a multi-level, bidirectional token-wise alignment mechanism, which enables both fine-grained and semantically abstract interactions between the two modalities. We first define the token-level similarity between visual and textual features as  $\mathbf{S} = \mathbf{T} \cdot \mathbf{V}^\top \in \mathbb{R}^{L_v \times L_t}$ . To obtain a scalar instance-level similarity score  $\mathcal{S}$  between the image and the caption, we perform bidirectional aggregation over  $\mathbf{S}$ . The computation process is illustrated in Figure 3 and formally,

$$\mathcal{S} = \underbrace{\frac{1}{L_t} \sum_{l=1}^{L_t} \max_{L_v} \mathbf{S}_{l, L_v}}_{\text{Text2Image}} + \underbrace{\frac{1}{L_v} \sum_{l=1}^{L_v} \max_{L_t} \mathbf{S}_{l, L_t}}_{\text{Image2Text}} \in \mathbb{R}^{B \times B}, \quad (7)$$

where  $B$  denotes the batch size. For the original visual and textual feature pair  $(\mathbf{T}, \mathbf{V})$ , the significance-aware and granularity-aware feature pair  $(\hat{\mathbf{T}}, \hat{\mathbf{V}})$ , and the uncertainty-aware region-text pair  $(\hat{\mathbf{T}}, \mathbf{U})$ , we obtain three levels of similarity maps— $\mathcal{S}^{\text{ori}}$ ,  $\mathcal{S}^{\text{key}}$ , and  $\mathcal{S}^{\text{unc}}$ —computed via the equation 7. We then compute the contrastive loss with hard negative mining (Faghri et al. 2017) for each of the three similarity matrices via equation 8:

$$\mathcal{L} = \sum_{i=1}^B \sum_{\substack{j=1 \\ j \neq i}}^B [\alpha + \mathcal{S}_{ij} - \mathcal{S}_{ii}]_+ + [\alpha + \mathcal{S}_{ji} - \mathcal{S}_{ii}]_+, \quad (8)$$

resulting in three levels of contrastive objectives  $\mathcal{L}_{\text{con}}^{\text{ori}}$ ,  $\mathcal{L}_{\text{con}}^{\text{key}}$ , and  $\mathcal{L}_{\text{con}}^{\text{unc}}$ . Here,  $\alpha$  denotes the margin hyperparameter,  $\mathcal{S}_{ij}$  represents the similarity between the  $i$ -th image and the  $j$ -th text, and  $[x]_+ = \max(x, 0)$  denotes the hinge function. The overall multi-level loss is formulated as a weighted combination of the three components, governed by the hyperparameters  $a$ ,  $b$ , and  $c$ :

$$\mathcal{L}_{\text{con}} = a\mathcal{L}_{\text{con}}^{\text{ori}} + b\mathcal{L}_{\text{con}}^{\text{key}} + c\mathcal{L}_{\text{con}}^{\text{unc}}, \quad (9)$$

where  $a + b + c = 1$ .

Method	FG	Flickr30k-1k						MS-COCO 1k						MS-COCO 5k								
		Image to Text			Text to Image			Image to Text			Text to Image			Image to Text			Text to Image					
		R@1	R@5	R@10	R@1	R@5	R@10	rSum	R@1	R@5	R@10	R@1	R@5	R@10	rSum	R@1	R@5	R@10	R@1	R@5	R@10	rSum
<b>Faster R-CNN + BERT-base, 36 pre-computed regions</b>																						
VSE++	✗	81.7	95.4	97.6	61.4	85.9	91.5	513.5	79.7	96.4	98.9	64.8	91.4	96.3	527.5	58.3	85.3	92.3	42.4	72.7	83.2	434.2
MV-VSE	✗	82.1	95.8	97.9	63.1	86.7	92.3	517.5	80.4	96.6	99.0	64.9	91.2	96.0	528.1	59.1	86.3	92.5	42.5	72.8	83.1	436.3
CHAN	✓	80.6	96.1	97.8	63.9	87.5	92.6	518.5	81.4	96.9	98.9	66.5	92.1	96.7	532.5	59.8	87.2	93.3	44.9	74.5	84.2	443.9
HERM	✓	84.0	96.1	98.6	64.4	88.0	93.1	524.2	82.9	96.9	99.0	67.1	92.0	96.6	534.5	64.0	88.5	93.7	45.4	75.1	84.3	451.0
CORA	✓	83.4	95.9	98.6	64.1	88.1	93.1	523.3	82.8	97.3	99.0	67.3	92.4	96.9	535.6	64.3	87.5	93.6	45.4	74.7	84.6	450.1
<b>ViT-Base-224 + BERT-base, 14×14 patches</b>																						
SCAN	✓	69.5	90.9	95.6	56.4	83.1	90.0	485.6	76.0	95.4	98.1	64.5	90.8	95.8	520.6	53.9	81.8	90.0	42.9	72.3	82.5	423.5
VSE++	✗	71.8	92.8	96.5	59.4	84.7	90.9	496.1	75.0	94.6	98.0	62.7	89.4	94.9	514.6	52.4	80.3	88.8	40.6	70.4	81.1	413.4
SGR	✓	69.7	90.8	95.2	59.1	84.1	89.9	488.7	77.2	95.0	98.0	65.1	90.7	95.8	521.8	54.9	82.8	90.5	42.8	72.2	82.5	425.8
CHAN	✓	69.2	91.8	95.0	58.4	84.9	90.6	489.9	77.1	95.1	98.1	65.0	91.0	96.0	522.2	56.3	83.2	90.1	43.0	72.6	82.8	428.0
LAPS	✓	74.0	93.4	97.4	62.5	87.3	92.7	507.3	78.7	95.5	98.3	66.2	91.3	96.2	526.3	57.5	84.0	90.8	44.5	74.0	83.6	434.4
AVSE	✗	76.0	94.6	97.5	62.7	88.4	93.1	512.3	79.8	95.6	98.3	67.0	91.5	96.3	528.5	58.8	84.3	91.0	45.1	74.3	83.9	437.4
Ours	✓	<b>77.2</b>	<b>94.5</b>	<b>97.9</b>	<b>64.6</b>	<b>88.5</b>	<b>93.5</b>	<b>516.2</b>	<b>80.8</b>	<b>96.7</b>	<b>98.9</b>	<b>67.6</b>	<b>91.8</b>	<b>96.7</b>	<b>532.5</b>	<b>60.1</b>	<b>86.7</b>	<b>93.2</b>	44.4	<b>74.5</b>	<b>84.1</b>	<b>443.0</b>
<b>ViT-Base-384 + BERT-base, 24×24 patches</b>																						
SCAN	✓	75.4	94.4	96.9	63.6	88.6	93.5	512.5	77.0	95.7	98.4	64.6	91.1	96.2	523.0	54.9	82.8	90.4	42.4	72.4	82.8	425.8
VSE++	✗	77.1	95.7	97.5	65.8	90.2	94.3	520.5	76.1	95.5	98.5	65.1	91.6	96.3	523.1	53.3	81.8	90.0	42.6	72.6	82.9	423.1
SGR	✓	76.9	94.9	98.1	64.2	88.4	93.3	515.8	75.8	95.7	98.6	65.6	92.0	96.5	524.2	53.3	81.0	89.6	42.9	73.1	83.7	423.6
CHAN	✓	75.4	94.5	97.6	63.2	88.6	93.1	512.4	78.1	95.8	98.6	66.1	92.1	96.6	527.3	55.6	83.8	91.2	43.4	73.6	83.5	431.1
LAPS	✓	79.0	96.0	98.1	67.3	90.5	94.5	525.4	78.6	96.3	98.9	68.0	92.4	96.8	531.0	57.4	84.9	92.5	46.4	75.8	85.2	442.2
AVSE	✗	80.3	96.4	98.7	67.9	91.2	94.7	529.2	81.1	97.1	99.0	68.3	92.7	97.0	535.2	61.2	86.8	93.2	46.2	75.9	85.0	448.3
Ours	✓	<b>80.5</b>	<b>97.5</b>	<b>98.8</b>	<b>68.5</b>	<b>91.6</b>	<b>94.9</b>	<b>531.8</b>	<b>81.8</b>	<b>97.7</b>	<b>99.1</b>	<b>69.4</b>	<b>92.9</b>	<b>97.3</b>	<b>538.2</b>	<b>62.1</b>	<b>87.4</b>	<b>93.4</b>	<b>47.9</b>	<b>76.5</b>	<b>85.9</b>	<b>451.2</b>
<b>Swin-base-224 + BERT-base, 7×7 patches</b>																						
SCAN	✓	79.0	95.9	98.2	67.7	90.6	94.9	526.3	83.3	97.5	99.3	71.0	93.0	96.7	540.9	64.0	88.2	94.2	49.9	78.0	86.6	460.9
VSE++	✗	82.5	96.5	98.9	70.0	91.4	95.1	534.4	80.9	97.0	99.1	69.7	93.1	97.1	536.9	60.7	86.6	93.2	48.1	77.1	86.1	451.8
SGR	✓	80.4	97.0	98.7	66.9	90.2	94.5	527.6	81.2	97.1	99.1	69.9	93.2	97.2	537.7	61.0	86.7	93.2	48.6	77.2	86.3	453.1
CHAN	✓	81.4	97.0	98.6	68.5	90.6	94.5	530.6	81.6	97.2	99.3	70.6	93.7	97.6	539.8	64.1	87.9	93.5	49.1	77.3	86.1	458.0
LAPS	✓	82.4	97.4	99.5	70.0	91.7	95.4	536.3	84.0	97.6	99.3	72.1	93.7	97.3	544.1	64.5	89.2	94.4	51.6	78.9	87.2	465.8
AVSE	✗	83.9	97.4	99.4	70.0	92.4	95.6	538.7	84.9	<b>98.0</b>	99.3	72.1	<b>94.0</b>	97.4	545.7	66.2	89.8	94.7	51.7	79.2	87.3	468.9
Ours	✓	<b>85.6</b>	<b>98.7</b>	<b>99.5</b>	<b>73.0</b>	<b>92.9</b>	<b>96.2</b>	<b>546.0</b>	<b>85.2</b>	97.8	<b>99.6</b>	<b>73.2</b>	93.9	<b>97.8</b>	<b>547.5</b>	<b>66.3</b>	<b>90.1</b>	<b>94.9</b>	<b>52.1</b>	<b>79.8</b>	<b>87.6</b>	<b>470.8</b>
<b>Swin-base-384 + BERT-base, 12×12 patches</b>																						
SCAN	✓	81.9	96.9	98.9	70.0	92.7	95.8	536.1	82.9	97.7	99.4	71.3	93.5	97.3	542.1	63.0	88.5	94.3	50.1	78.9	87.4	462.2
VSE++	✗	83.3	97.5	99.2	71.1	93.2	96.2	540.6	81.6	96.8	99.1	69.1	92.7	96.7	536.1	61.1	87.3	93.3	47.8	76.9	85.9	452.4
SGR	✓	80.7	96.8	99.0	69.9	91.7	95.3	533.4	81.9	96.7	99.1	69.3	92.8	96.7	536.6	62.8	87.0	92.9	48.1	77.0	86.0	453.8
CHAN	✓	81.2	96.7	98.8	70.3	92.2	95.9	535.0	83.1	97.3	99.2	70.4	93.1	97.1	540.2	63.4	88.4	94.1	49.2	77.9	86.6	459.5
LAPS	✓	85.1	97.7	99.2	74.0	93.0	96.3	545.3	84.1	97.4	99.2	72.1	93.9	97.4	544.1	67.1	88.6	94.3	53.0	79.5	87.6	470.1
AVSE	✗	87.1	98.3	99.2	73.6	93.5	96.5	548.2	<b>85.1</b>	98.2	<b>99.5</b>	71.6	94.0	97.5	545.9	68.6	<b>90.2</b>	95.6	52.2	79.6	87.8	474.0
Ours	✓	<b>87.7</b>	<b>98.8</b>	<b>99.5</b>	<b>75.3</b>	<b>93.5</b>	<b>96.9</b>	<b>550.7</b>	85.0	<b>97.9</b>	99.4	<b>73.5</b>	<b>94.4</b>	<b>97.8</b>	<b>548.0</b>	<b>69.4</b>	90.0	<b>95.7</b>	<b>54.4</b>	<b>80.6</b>	<b>88.2</b>	<b>478.3</b>

Table 1: Comparison of image-text retrieval performance on the Flickr30K and MS-COCO test sets. We list detailed information about feature encoders, image resolution and the number of regions/patches obtained by the visual encoder (e.g., “ViT-Base-224” denotes the base version of the Vision Transformer with a 224×224 image resolution input, using 16×16 pixel patches, resulting in 14×14 visual patches per image). **FG** indicates whether fine-grained cross-modal alignment is employed. The best results are highlighted in **bold**.

**Semantic Consistency Constraint:** To enforce the global semantic consistency of the uncertainty-aware region embeddings, we first compute the average of all region-level representations  $\mathbf{U}$  to obtain a global semantic vector for the image. This vector serves as a holistic representation of the image’s semantic content. We then align this global vector with the global mean of the significance-aware patch features  $\hat{\mathbf{V}}$ , ensuring that the high-level semantic structures across both regions and patches are consistent. The reconstruction loss is defined as follows:

$$\mathcal{L}_{\text{recon}} = \left\| \frac{1}{K} \sum_{k=1}^K \mathbf{U}_k - \frac{1}{L_v} \sum_{l=1}^{L_v} \hat{\mathbf{V}}_l \right\|_2^2. \quad (10)$$

**Region Distribution and Diversity Regularization:** Since the standard normal distribution often serves as an ef-

fective model for smooth transformations and complex relationships in latent spaces (Kingma, Welling et al. 2013), especially in high-dimensional settings, we assume that the prior distribution of the latent variables for the region features follows a standard normal distribution  $\mathcal{N}(0, 1)$ . To encourage each region’s distribution to approximate this prior, we employ the Kullback-Leibler (KL) divergence as a regularization term. Formally,

$$\begin{aligned} \mathcal{L}_{\text{KL}} &= \sum_{k=1}^K \text{KL}(\mathcal{N}(\mu_k, \sigma_k^2) \| \mathcal{N}(0, 1)) \\ &= -\frac{1}{2} \sum_{k=1}^K (1 + \log \sigma^2 - \mu^2 - \sigma^2). \end{aligned} \quad (11)$$

Settings	Image to Text			Text to Image			rSum
	R@1	R@5	R@10	R@1	R@5	R@10	
w/o SA	73.6	93.2	96.5	60.7	86.6	92.2	502.8
w/o GA	74.9	93.8	97.0	61.5	87.3	92.5	507.0
w/o RP	73.5	93.1	96.8	60.7	86.9	92.3	503.3
w/o UM	74.4	93.4	97.3	62.1	88.0	93.1	508.3
Complete	77.2	94.5	97.9	64.6	88.5	93.5	516.2

Table 2: The ablation study on the contribution of each module in our proposed method.

To prevent attention collapse and promote the diversity of regions, we introduce an entropy regularization term:

$$\mathcal{L}_{\text{ent}} = -\frac{1}{K} \sum_{k=1}^K \mathbf{A}_r^k \log(\mathbf{A}_r^k), \quad (12)$$

where  $\mathbf{A}_r^k$  is from equation 2. The overall regularization objective is:

$$\mathcal{L}_{\text{reg}} = \mathcal{L}_{\text{KL}} + \mathcal{L}_{\text{ent}}. \quad (13)$$

## Experiment

### Comparison with State of the Arts

Following the standard evaluation protocols of the two benchmarks (Faghri et al. 2017; Zhang et al. 2022), we summarize the details of the feature encoders and cross-modal alignment strategies used in all compared methods. Specifically, we include six representative cross-modal alignment approaches (Faghri et al. 2017; Diao et al. 2021; Lee et al. 2018; Pan, Wu, and Zhang 2023; Fu et al. 2024; Liu et al. 2025), covering a range of alignment paradigms for comprehensive comparison.

As shown in Table 1, we report the quantitative results on Flickr30K and MS-COCO benchmarks. Our proposed method consistently outperforms all existing state-of-the-art approaches, demonstrating its strong capability in fine-grained cross-modal retrieval. Our method consistently and significantly outperforms all existing state-of-the-art approaches. When evaluated under identical visual and textual encoder configurations, our method demonstrates consistent performance gains in terms of rSum, with improvements ranging from 2.1%~5.6% on Flickr30K, 1.3%~4.0% on MS-COCO 1K, and 1.9%~5.6% on MS-COCO 5K. Notably, while the performance gain narrows slightly as the number of visual patches increases, our approach still maintains a significant margin across all settings. Compared to two-stage methods that rely on pre-trained object detectors as visual backbones, our approach offers both better performance and greater scalability. More importantly, it eliminates the need for additional detector training and enables end-to-end optimization for downstream tasks, making it more practical and efficient for real-world deployment. To further demonstrate the effectiveness of our proposed method, we provide heatmap visualizations for both patch-word alignment and image-text alignment, as shown in Figure 5 and 6.

Settings	Image to Text			Text to Image			rSum
	R@1	R@5	R@10	R@1	R@5	R@10	
w/o $\mathcal{L}_{\text{con}}^{\text{ori}}$	70.1	91.3	95.7	57.6	82.2	90.3	487.2
w/o $\mathcal{L}_{\text{con}}^{\text{key}}$	74.7	93.9	96.9	60.7	86.8	92.9	505.9
w/o $\mathcal{L}_{\text{con}}^{\text{unc}}$	73.8	93.1	96.1	59.9	86.1	92.6	501.6
w/o $\mathcal{L}_{\text{recon}}$	75.8	94.2	97.8	63.1	88.5	93.8	513.2
w/o $\mathcal{L}_{\text{reg}}$	76.2	94.7	97.5	63.5	88.5	93.8	514.2
Complete	77.2	94.5	97.9	64.6	88.5	93.5	516.2

Table 3: The ablation study on the contribution of each optimization objective in our proposed method.

### Ablation Studies

We conduct extensive ablation studies and robustness analyses to evaluate the effectiveness of our approach. All experiments are performed based on the Vision Transformer (ViT-Base) with  $224 \times 224$  input resolution and BERT-Base as the default text encoder.

**Module Gain.** To thoroughly assess the contribution of each component, we conduct a module-wise ablation study, as summarized in Table 2. The results demonstrate that both the Significance-aware Adapter (SA) and Region Prompts (RP) are essential components of our method. This performance gap can be attributed to the distribution mismatch between the pretraining corpus and downstream data. The observed improvements underscore the role of SA and RP in mitigating this distribution shift and enhancing model adaptability. Specifically, removing any module leads to a notable performance drop across multiple evaluation metrics, underscoring their individual importance, suggesting that the two components are not only effective in isolation but also highly complementary. This indicates that the design is not merely a sum of independent parts but rather a carefully coordinated structure where each module reinforces the other.

**Optimization Gain.** To assess the contribution of each loss component, we perform a series of ablation experiments, with the results presented in Table 3. It is evident that the three losses associated with our Multi-level and Bidirectional Image-Text Alignment— $\mathcal{L}_{\text{con}}^{\text{ori}}$ ,  $\mathcal{L}_{\text{con}}^{\text{key}}$  and  $\mathcal{L}_{\text{con}}^{\text{unc}}$ —play a pivotal role in our method. These objectives directly govern both the alignment of image-text pairs and the modeling of fine-grained uncertainty, thereby exerting a dominant influence on overall performance. In contrast, the loss  $\mathcal{L}_{\text{recon}}$  of Semantic Consistency Constraint and the loss  $\mathcal{L}_{\text{reg}}$  of Region Distribution and Diversity Regularization serve as auxiliary regularizers: they reinforce the knowledge learned by the region prompts but do not drive the core alignment process. This hierarchy of contributions is corroborated by the ablation results presented in Table 3, which demonstrate that removing any one of losses leads to the drop in performance, especially the three losses of Multi-level and Bidirectional Image-Text Alignment.

### Parameter Analysis

To investigate the impact of the number of region prompts  $P$  and the weighting proportions in multi-level alignment on

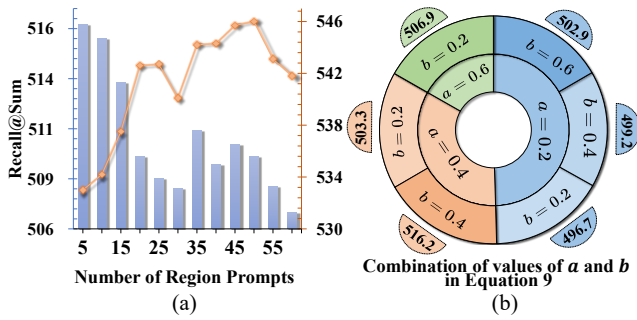


Figure 4: Comparison of hyperparameter experimental performance. The experiments are conducted on the Flickr30k dataset. **(a)** Parameter study on the number of region prompts  $P$ . The blue bar chart shows results with the visual backbone ViT-base-224, while the yellow line chart shows results with Swin-base-224. **(b)** Parameter study on different combinations of weights  $a$ ,  $b$  and  $c$  in Equation 9. Since  $a + b + c = 1$ , only  $a$  and  $b$  need to be specified. Six combinations are tested: when  $a = 0.2$  (blue blocks),  $b = \{0.2, 0.4, 0.6\}$ ; when  $a = 0.4$  (yellow blocks),  $b = \{0.2, 0.4\}$ ; and when  $a = 0.6$  (green blocks),  $b = 0.2$ .

model performance, we conducted corresponding hyperparameter experiments. The results are illustrated in Figure 4. **The number of region prompts.** When using ViT-Base-224 as the visual backbone, the optimal performance is achieved when the number of region prompts is set to 5. As the number of prompts increases beyond this point, the retrieval performance exhibits a noticeable decline. In contrast, when adopting Swin-Base-224 as the visual backbone, the best performance is obtained when the number of region prompts reaches 50. We speculate that this discrepancy stems from the architectural differences between the two backbones. Specifically, the Swin Transformer leverages local self-attention and a hierarchical structure, which tends to aggregate local visual information at an earlier stage. As a result, it requires a larger number of learnable region prompts to sufficiently capture and represent the finer-grained semantic prompts embedded in the image. In comparison, ViT’s global attention enables broader contextual modeling from the outset, thereby benefiting more from a smaller set of region prompts.

**Different value combinations of  $a$ ,  $b$ ,  $c$ .** The results in Figure 4(b) clearly demonstrate that the model performance is highly sensitive to the values of the weighting parameters  $a$ ,  $b$ , and  $c$  in the multi-level alignment. Specifically, the optimal performance is observed when  $a = b = 0.4$  and  $c = 0.2$ , while the poorest performance occurs when  $a = b = 0.2$  and  $c = 0.6$ . This trend suggests that the original similarity and the significance-aware similarity contribute equally and significantly to the effectiveness of image-text alignment. In contrast, excessive emphasis on uncertainty-aware similarity may introduce noise or overfitting, thereby degrading overall performance. These findings are in line with our ablation results, reinforcing the importance of balancing the contributions from different semantic alignment levels.




Figure 5: The visualization of fine-grained patch-word alignment with each linguistic word.

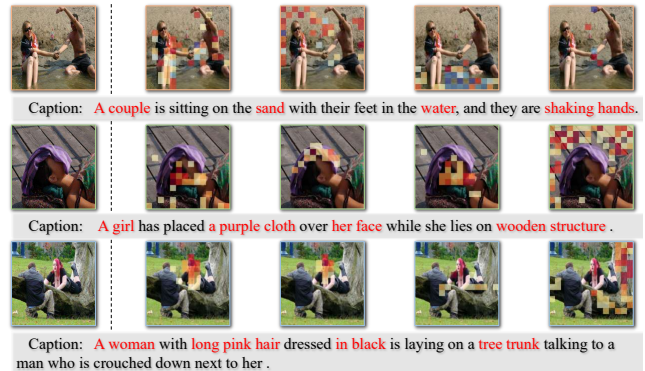


Figure 6: The visualization of fine-grained image-text alignment with some key words.

## Conclusion

In this paper, we address the critical challenges of fine-grained image-text alignment by identifying two key limitations in existing methods: the lack of effective intra-modal significance modeling and the absence of fine-grained uncertainty quantification. To this end, we propose GRM, a novel and effective approach that introduces new perspectives into cross-modal matching. By leveraging intra-modal semantic biases and modeling region features as a mixture of Gaussian distributions, GRM captures both salient content and region-level uncertainty in a principled manner. This approach enhances alignment robustness, semantic fidelity, and interpretability without over-reliance on global cross-modal attention. Extensive experiments across multiple benchmark datasets and diverse backbone architectures demonstrate the superior and consistent performance of our GRM, offering a strong foundation for future advancements in fine-grained cross-modal understanding and alignment.

## Acknowledgments

The works described in this paper are supported by The National Natural Science Foundation of China under Grant No. 62477015; Key Research and Development Program of Guangdong of China under Grant No. 2023B0303010004; The Innovation Team Project for Universities in Guangdong Province in China under Grant No. 2023KCXTD011.

## References

- Anderson, P.; Wu, Q.; Teney, D.; Bruce, J.; Johnson, M.; Sünderhauf, N.; Reid, I.; Gould, S.; and Van Den Hengel, A. 2018. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3674–3683.
- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Zitnick, C. L.; and Parikh, D. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, 2425–2433.
- Chang, J.; Lan, Z.; Cheng, C.; and Wei, Y. 2020. Data uncertainty learning in face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5710–5719.
- Chen, J.; Lin, B.; Liu, X.; Ma, L.; Liang, X.; and Wong, K.-Y. K. 2025. Affordances-oriented planning using foundation models for continuous vision-language navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 23568–23576.
- Chen, Y.; Zheng, Z.; Ji, W.; Qu, L.; and Chua, T.-S. 2022. Composed image retrieval with text feedback via multi-grained uncertainty regularization. *arXiv preprint arXiv:2211.07394*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 4171–4186.
- Diao, H.; Zhang, Y.; Ma, L.; and Lu, H. 2021. Similarity reasoning and filtration for image-text matching. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 1218–1226.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Faghri, F.; Fleet, D. J.; Kiros, J. R.; and Fidler, S. 2017. Vse++: Improving visual-semantic embeddings with hard negatives. *arXiv preprint arXiv:1707.05612*.
- Fu, Z.; Mao, Z.; Song, Y.; and Zhang, Y. 2023. Learning semantic relationship among instances for image-text matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15159–15168.
- Fu, Z.; Zhang, L.; Xia, H.; and Mao, Z. 2024. Linguistic-aware patch slimming framework for fine-grained cross-modal alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26307–26316.
- Gao, Z.; Jiang, X.; Xu, X.; Shen, F.; Li, Y.; and Shen, H. T. 2024. Embracing unimodal aleatoric uncertainty for robust multimodal fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26876–26885.
- Huang, L.; Wang, W.; Chen, J.; and Wei, X.-Y. 2019. Attention on attention for image captioning. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4634–4643.
- Ji, Y.; Wang, J.; Gong, Y.; Zhang, L.; Zhu, Y.; Wang, H.; Zhang, J.; Sakai, T.; and Yang, Y. 2023. Map: Multimodal uncertainty-aware vision-language pre-training model. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 23262–23271.
- Jia, M.; Tang, L.; Chen, B.-C.; Cardie, C.; Belongie, S.; Hariharan, B.; and Lim, S.-N. 2022. Visual prompt tuning. In *European conference on computer vision*, 709–727. Springer.
- Jiang, D.; and Ye, M. 2023. Cross-modal implicit relation reasoning and aligning for text-to-image person retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2787–2797.
- Kingma, D. P.; Salimans, T.; and Welling, M. 2015. Variational Dropout and the Local Reparameterization Trick. *Advances in Neural Information Processing Systems*, 28.
- Kingma, D. P.; Welling, M.; et al. 2013. Auto-encoding variational bayes.
- Kohl, S.; Romera-Paredes, B.; Meyer, C.; De Fauw, J.; Ledsam, J. R.; Maier-Hein, K.; Eslami, S.; Jimenez Rezende, D.; and Ronneberger, O. 2018. A probabilistic u-net for segmentation of ambiguous images. *Advances in neural information processing systems*, 31.
- Lee, K.-H.; Chen, X.; Hua, G.; Hu, H.; and He, X. 2018. Stacked cross attention for image-text matching. In *Proceedings of the European conference on computer vision (ECCV)*, 201–216.
- Li, H.; Song, J.; Gao, L.; Zhu, X.; and Shen, H. 2023. Prototype-based aleatoric uncertainty quantification for cross-modal retrieval. *Advances in Neural Information Processing Systems*, 36: 24564–24585.
- Liu, Y.; Liu, M.; Huang, S.; and Lv, J. 2025. Asymmetric Visual Semantic Embedding Framework for Efficient Vision-Language Alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 5676–5684.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.
- Lu, H.; Fei, N.; Huo, Y.; Gao, Y.; Lu, Z.; and Wen, J.-R. 2022. COTS: Collaborative Two-Stream Vision-Language Pre-Training Model for Cross-Modal Retrieval. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 15671–15680. IEEE.
- Molchanov, D.; Ashukha, A.; and Vetrov, D. 2017. Variational dropout sparsifies deep neural networks. In *International conference on machine learning*, 2498–2507. PMLR.
- Pan, Z.; Wu, F.; and Zhang, B. 2023. Fine-grained image-text matching by cross-modal hard aligning network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 19275–19284.

- Pei, W.; Xia, T.; Chen, F.; Li, J.; Tian, J.; and Lu, G. 2024. SA<sup>2</sup>VP: Spatially Aligned-and-Adapted Visual Prompt. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 4450–4458.
- Pham, K.; Huynh, C.; Lim, S.-N.; and Shrivastava, A. 2024. Composing object relations and attributes for image-text matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14354–14363.
- Qin, Y.; Peng, D.; Peng, X.; Wang, X.; and Hu, P. 2022. Deep evidential learning with noisy correspondence for cross-modal retrieval. In *Proceedings of the 30th ACM International Conference on Multimedia*, 4948–4956.
- Qiu, L.; Ning, S.; and He, X. 2024. Mining fine-grained image-text alignment for zero-shot captioning via text-only training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4605–4613.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2016. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6): 1137–1149.
- Wang, X.; Li, L.; Li, Z.; Wang, X.; Zhu, X.; Wang, C.; Huang, J.; and Xiao, Y. 2023. Agree: Aligning cross-modal entities for image-text retrieval upon vision-language pre-trained models. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, 456–464.
- Wang, Y.; Tang, X.; Ma, J.; Zhang, X.; Liu, F.; and Jiao, L. 2024. Cross-modal remote sensing image-text retrieval via context and uncertainty-aware prompt. *IEEE Transactions on Neural Networks and Learning Systems*.
- Wu, Y.; Zhou, R.; Li, H.; Zhou, Z.; Wei, F.; Li, M.; and He, G. 2025. Relation-aware Semantic Alignment Network for Text-to-Image Person Retrieval. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Yu, T.; Tong, Z.; Yu, J.; and Zhang, K. 2025. Fine-grained Adaptive Visual Prompt for Generative Medical Visual Question Answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 9662–9670.
- Zeng, D.; Shen, Y.; Lin, M.; Yi, Z.; and Ouyang, J. 2025. Zero-Shot Image Captioning with Multi-type Entity Representations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 22308–22316.
- Zhang, K.; Mao, Z.; Wang, Q.; and Zhang, Y. 2022. Negative-aware attention framework for image-text matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15661–15670.
- Zhang, Y.; Ji, Z.; Wang, D.; Pang, Y.; and Li, X. 2024. USER: Unified semantic enhancement with momentum contrast for image-text retrieval. *IEEE Transactions on Image Processing*, 33: 595–609.
- Zhao, Z.; Liu, B.; Lu, Y.; Chu, Q.; and Yu, N. 2024. Unifying multi-modal uncertainty modeling and semantic alignment for text-to-image person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 7534–7542.
- Zheng, X.; Wang, Z.; Li, S.; Xu, K.; Zhuang, T.; Liu, Q.; and Zeng, X. 2023. Make: Vision-language pre-training based product retrieval in taobao search. In *Companion Proceedings of the ACM Web Conference 2023*, 356–360.