

PCSR: Pseudo-label Consistency-Guided Sample Refinement for Noisy Correspondence Learning

Zhuoyao Liu¹, Yang Liu¹, Wentao Feng¹, Shudong Huang^{1,2*},

¹College of Computer Science, Sichuan University, China

²Engineering Research Center of Machine Learning and Industry Intelligence, China
liuzhuoyao@stu.scu.edu.cn

Abstract

Cross-modal retrieval aims to align different modalities via semantic similarity. However, existing methods often assume that image-text pairs are perfectly aligned, overlooking *Noisy Correspondences* in real data. These misaligned pairs misguide similarity learning and degrade retrieval performance. Previous methods often rely on coarse-grained categorizations that simply divide data into clean and noisy samples, overlooking the intrinsic diversity within noisy instances. Moreover, they typically apply uniform training strategies regardless of sample characteristics, resulting in suboptimal sample utilization for model optimization. To address the above challenges, we introduce a novel framework, called *Pseudo-label Consistency-Guided Sample Refinement (PCSR)*, which enhances correspondence reliability by explicitly dividing samples based on pseudo-label consistency. Specifically, we first employ a confidence-based estimation to distinguish clean and noisy pairs, then refine the noisy pairs via pseudo-label consistency to uncover structurally distinct subsets. We further propose a *Pseudo-label Consistency Score (PCS)* to quantify prediction stability, enabling the separation of ambiguous and refinable samples within noisy pairs. Accordingly, we adopt *Adaptive Pair Optimization (APO)*, where ambiguous samples are optimized with robust loss functions and refinable ones are enhanced via text replacement during training. Extensive experiments on CC152K, MS-COCO and Flickr30K validate the effectiveness of our method in improving retrieval robustness under noisy supervision.

Introduction

Cross-modal retrieval (Anderson et al. 2018; Faghri et al. 2017), a fundamental task in multimodal learning, aims to bridge different modalities within large-scale multimedia data. Existing methods explore various strategies to bridge the semantic gap between modalities (Feng et al. 2023; Liu et al. 2020; Zhang et al. 2020; Liu et al. 2025b), but their effectiveness relies heavily on the assumption of accurately aligned image-text pairs. In practice, such precise alignment requires costly manual annotation, which is both challenging and error-prone. This results in the problem of

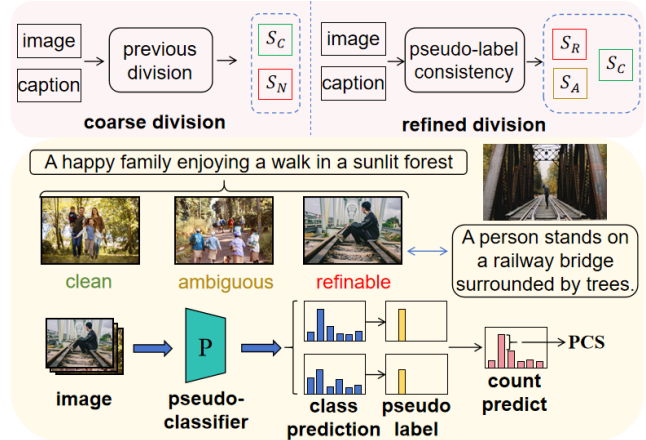


Figure 1: (Top) Difference between previous coarse division and our refined division method. (Bottom) The conceptual diagram of our proposed PCSR method. Specifically, we first obtain the pseudo-label of each sample using a pseudo-classifier, and then calculate PCS value by measuring the concentration of the pseudo-label prediction. Thereby distinguishing samples with consistent image predictions.

noisy correspondence, different from traditional noisy labels (Jiang et al. 2020; Bai et al. 2021; Iscen et al. 2022; Yan et al. 2023) that refer to incorrect category annotations. Noisy correspondence specifically refers to semantic misalignment between paired modalities, which consequently degrades the performance of existing cross-modal retrieval methods (Diao et al. 2021; Li et al. 2019; Lee et al. 2018).

In recent years, research addressing the problem of noisy correspondence remains limited. Earlier methods (Qin et al. 2022; Yang et al. 2023) typically divide the training data into noisy and clean subsets, and then employ strategies such as dynamic weight adjustment to mitigate the impact of incorrect supervision or introduce effective supervisory signals to enhance model optimization. However, these methods only consider a binary categorization and fail to account for the possibility that some pseudo-positive pairs within the noisy subset may still exhibit partial semantic alignment and thus can still provide informative supervision while training. To address this, (Ma et al. 2024) introduces a finer-grained

*Corresponding author.

classification that partitions samples into clean, vague, and noisy categories. However, such a division remains ambiguous, lacking quantifiable criteria such as prediction stability or semantic consistency to justify the classification. This approach merely refines existing confidence-based or loss-based frameworks without introducing fundamentally new classification principles. Moreover, in the subsequent training phase, coarse division approaches tend to apply uniform training strategies across all noisy pairs, which leads to two major drawbacks: stable noisy samples are underutilized, and ambiguous samples introduce incorrect supervision, thereby affecting the convergence of the model. The training strategies proposed after classification often only involve parameter adjustments compared to the original methods. Consequently, although these methods alleviate noisy supervision through finer categorization, the refinement exerts limited impact on the overall training strategy and offers restricted performance gains.

As illustrated in Figure 1, we propose a refined division strategy that categorizes samples into three subsets: clean, ambiguous, and refinable. We argue that there are intrinsic differences within the noisy subset: some are reliably identified as mismatches due to accurate model understanding and are thus termed as refinable pairs, while others are classified as noisy due to semantic ambiguity and are referred to as ambiguous pairs. For refinable pairs, which the model understands confidently, we perform text refinement to enhance their utility. For ambiguous pairs, where the model’s predictions are uncertain, we adopt a more robust training strategy. This approach enables us to fully exploit the distinct training value inherent in each category.

To address these challenges, we propose a novel method named *Pseudo-label Consistency-Guided Sample Refinement* (PCSR), which assesses correspondence reliability using confidence scores and pseudo-label consistency, thereby improving the robustness of cross-modal retrieval by enabling flexible sample division and targeted training strategies. Specifically, PCSR first estimates the confidence of each image-text pair based on the distribution of matching losses, using Gaussian Mixture Model (GMM) to distinguish clean pairs and noisy pairs. Considering some noisy pairs contain useful information that can be exploited through training, as illustrated in Figure 1, we further introduce a *Pseudo-label Consistency Score* (PCS) in our method to analyze the prediction stability of images within noisy pairs. By evaluating the consistency across multiple pseudo-label predictions, we separate noisy pairs into two different subsets: Ambiguous pairs that suffer from unstable pseudo-classifier predictions, implying limited training value. Conversely, refinable pairs exhibit stable predictions, suggesting that their images remain valuable for learning despite misaligned textual annotations. Building on this division, PCSR adopts *Adaptive Pair Optimization* (APO). Whereas clean pairs are trained using standard triplet loss and cross-entropy loss. For ambiguous pairs, to mitigate the impact of unstable supervision, we adopt a more robust Generalized Cross-Entropy (GCE) loss for training. For refinable pairs, benefiting from reliable image pseudo-prediction, we undergo a caption rematch strategy, where their original captions are

replaced by captions retrieved from clean pairs based on pseudo-label similarity. This allows the model to receive more reliable supervisory signals.

Our main contributions are summarized as follows:

- Inspired by the observation that noisy pairs exhibit intrinsic differences, we propose a novel *Pseudo-label Consistency-Guided Sample Refinement* (PCSR) framework, which explicitly models correspondence reliability through a refined sample division strategy.
- We propose an *Adaptive Pair Optimization* (APO) strategy, introducing distinct optimization processes tailored to clean, ambiguous, and refinable pairs, thereby fully exploiting the particular training value of each subset.
- We conduct extensive experiments on MS-COCO, Flickr30K and CC152K. Experimental results show that PCSR effectively obtains reliable supervisory signals from each identified category and significantly outperforms state-of-the-art methods under noisy conditions.

Related Work

Cross-Modal Matching Cross-modal matching (Wang et al. 2020; Yang et al. 2022) is a core task in information retrieval (Zhao et al. 2024) and multi-modal learning (Chuang et al. 2022), aiming to retrieve relevant content across different modalities using one modality as a query. Among these tasks, image-text retrieval is the most representative and is typically divided into two mainstream paradigms: coarse-grained and fine-grained alignment. Coarse-grained methods (Liu et al. 2025a; Chen et al. 2021; Li et al. 2019) map entire images and texts into a common embedding space using separate networks and compute their overall similarity primarily based on cosine distance. Fine-grained approaches (Lee et al. 2018; Zhang et al. 2022a,b; Pan, Wu, and Zhang 2023; Fu et al. 2024) emphasize cross-modal interaction and align local regions across modalities, subsequently combining local matching results to derive the final similarity score. Despite their effectiveness, these methods generally rely on large-scale, accurately aligned training pairs and largely ignore the critical issue of noisy correspondence in practical scenarios. Therefore, exploring how to effectively learn cross-modal retrieval in the noisy correspondence relationships is of great significance, but this challenging issue has rarely been touched upon in previous studies.

Noisy Correspondence Learning Noisy correspondence refers to the widespread phenomenon where semantically misaligned data pairs are erroneously regarded as matched pairs during training. NCR (Huang et al. 2021) represents the first pioneering work specifically dedicated to addressing this problem. NCR enhances the robustness of image-text matching models by using a Gaussian Mixture Model (GMM) to adaptively refine soft correspondence labels during training. Subsequently, BiCro (Yang et al. 2023) uses Beta Mixture Model (BMM) to refine binary noisy labels into soft labels, assuming similar images share similar texts and vice versa. MSCN (Han et al. 2023) uses clean data to train a meta-network for generating reliable soft labels to correct main network predictions. Other recent methods

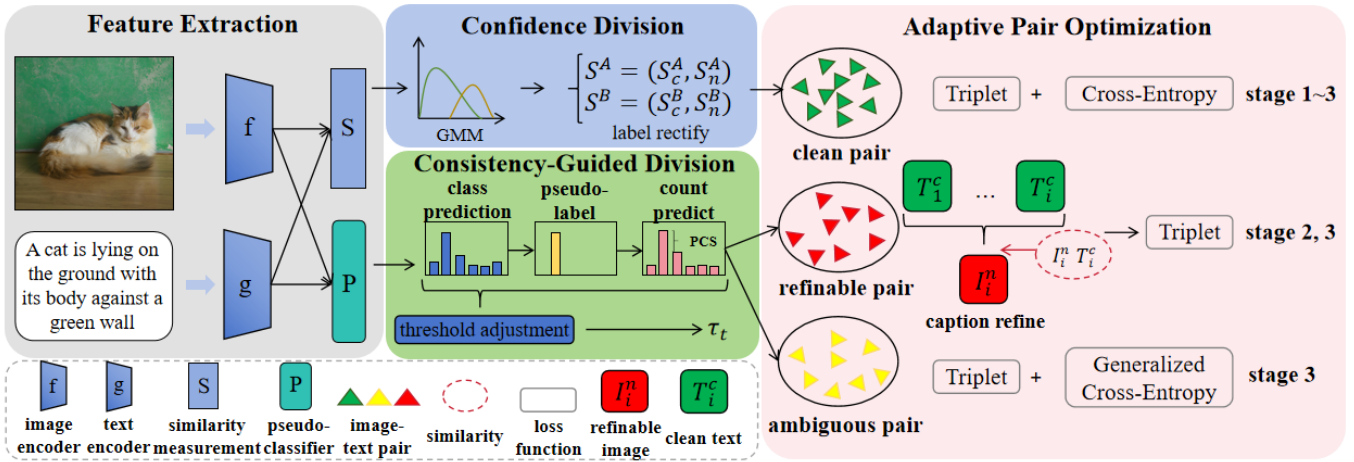


Figure 2: Overview of the proposed method. We compute the Pseudo-label Consistency Score (PCS) through repeated pseudo-label predictions, further dividing the noisy pairs into ambiguous pairs and refinable pairs. For clean pairs, we employ a combination of triplet loss and cross-entropy loss for optimization. For ambiguous pairs, we utilize triplet loss alongside generalized cross-entropy (GCE) loss to enhance robustness against label uncertainty. For refinable pairs, given their consistent pseudo-label predictions across iterations, we discard their original texts and replace them with captions selected from a subset of clean pairs that are most semantically similar to the noisy image’s pseudo-label, thereby enabling more reliable training.

(Qin et al. 2022; Li et al. 2024) enhance the loss function using dynamic weighting or bi-directional similarity. However, the existing methods mentioned above largely ignored structural differences among noisy pairs, training all noisy pairs using the same homogeneous method inherently limits the exploitation of their varied optimization potential. In contrast, we introduce a refined division strategy that effectively classifies noisy samples, allowing more precise distinction based on their inherent differences. Meanwhile, our method applies tailored training strategies to each category, enabling better utilization of noisy pairs based on their unique optimization potential.

Pseudo-Labeling Pseudo-labeling was initially introduced to deep learning by (Lee et al. 2013). It was first used in the field of computer vision and later widely adopted in semi-supervised learning and unsupervised learning. Semi-supervised methods (Chen, Dun, and Kyrillidis 2024; Chen and Wang 2024; Heidari, Zhang, and Guo 2024) often combined with consistency regularization, label propagation, and multi-model collaboration. These methods enhance pseudo-label quality and robustness by enforcing consistency, leveraging weak supervision, and exploiting model complementarity. Pseudo-labels are also used in unsupervised learning fields like self-supervised learning (Zhang et al. 2024) and knowledge distillation (Zhai et al. 2023; Wen, Lai, and Qian 2021), where models generate supervisory signals to guide representation learning or knowledge transfer. Further on, pseudo-labeling inherently tolerates label noise, as many of its techniques (Miyato et al. 2019; Wang et al. 2021) are inspired by or integrate noise-robust learning strategies to handle imperfect labels during training. Our method, building upon PC2 (Duan et al. 2024), applies pseudo-labeling to cross-modal retrieval and achieves significant performance improvements.

Method

Problem Overview

Given a query, cross-modal retrieval aims to identify semantically relevant instances from a dataset of a different modality. Without loss of generality, we take image-text retrieval as an example to discuss the problem of noisy correspondence in cross-modal retrieval. As illustrated in **Figure 2**, our proposed method comprises four key components. The feature extraction and confidence division parts are derived from the baseline strategy introduced in NCR (Huang et al. 2021). Specifically, given the image-text dataset, we use an image encoder $f(\cdot)$ and a text encoder $g(\cdot)$ to compute the feature embedding into a shared embedding space, then the similarity of image-text pairs is calculated within $S(f(I), g(T))$, where I and T represent images and texts, and S indicates the similarity function. Following, confidence division derives a probability distribution using a Gaussian Mixture Model (GMM) to divide the samples into clean samples ($y_i = 1$) and noisy samples ($y_i = 0$). To fully exploit the distinct training value of all samples, we propose a refined division method.

To better assess the reliability of image-text pairs under noisy correspondence, we adopt a *Consistency-Guided Division* (CGD) strategy. Inspired by semi-supervised learning, we propose a pseudo-classifier $C(\cdot)$ to predict pseudo categories of each image and text, based on its prediction consistency, we further divide noisy samples into refinable and ambiguous samples. Following this, we propose an *Adaptive Pair Optimization* (APO) strategy that performs targeted training based on the distinctive properties of the data.

Consistency-Guided Division

If the classifier consistently assigns the same category to an image across epochs, it suggests strong semantic confi-

dence, whereas fluctuating predictions indicate ambiguity. Based on this principle, we propose a refined division strategy named *Consistency-Guided Division* (CGD). By training a pseudo classifier within clean-only epochs, we define a Pseudo-label Consistency Score (PCS) to measure label stability and dynamically adjust a threshold τ to separate ambiguous and refinable samples. This module enables the model to selectively refine informative noisy pairs while avoiding misleading supervision from unstable ones. Thus, the model’s understanding of image-text pairs can be assessed by measuring the consistency of its prediction.

Pseudo-Classifier Training on clean data. Consistency-Guided Division relies heavily on category predictions from a pseudo-classifier, making it a critical component of the overall framework. Existing studies (Duan et al. 2024) applied pseudo-labeling in cross-modal retrieval, and proposed a method for training a pseudo-classifier. Using only the clean data, we train the pseudo-classifier by utilizing predicted category c for each image and the pseudo-label $l \in \{1, \dots, K\}$ of its corresponding caption. K represents the number of categories predicted by the Pseudo-Classifier.

Specifically, after the confidence division step, the dataset is divided into clean pairs and noisy pairs. Given a batch of clean pairs $(I_i^c, T_i^c)_{i=1}^B$, following the encoding process, the pseudo-classifier is subsequently employed to produce pseudo-label predictions for image $p_i^c = C(f(I_i^c))$ and text $q_i^c = C(g(T_i^c))$. We obtain the pseudo-label for the text by selecting the class with the highest predicted probability $\hat{q}_i^c = \arg \max(q_i^c)$, then optimize the pseudo-classifier using the standard cross-entropy loss with pseudo-label as the target:

$$\mathcal{L}_{\text{CE}} = \frac{1}{B} \sum_{i=1}^B \left(- \sum_{k=1}^K \hat{q}_{i,k}^c \log(p_{i,k}^c) \right) \quad (1)$$

To ensure diverse pseudo-label assignments, we regularize the classifier \mathbf{C} with an entropy loss that encourages uniform distribution over the label space:

$$\mathcal{L}_{\text{En}} = -\frac{1}{B} \sum_{i=1}^B p_i^c \cdot \log \left(\frac{1}{B} \sum_{j=1}^B p_j^c \right) \quad (2)$$

Pseudo-label Consistency Score. After obtaining the pseudo-predictions of images and texts, it is necessary to define a metric that quantifies the model’s level of understanding for each sample. We posit that the more consistent the model’s pseudo-label predictions for an image across multiple training iterations, the greater its semantic confidence in understanding the image. We first obtain the pseudo-label of the image $\hat{p}_i^c = \arg \max(p_i^c)$. In each training epoch, we count the occurrences of each label category as n_1, n_2, \dots, n_K , where n_K denotes the number of times class K is predicted. We define the *Pseudo-label Consistency Score* (PCS) as a metric to measure the prediction stability of the pseudo-classifier over multiple forward passes. Specifically, assume that $Q_i = [n_1, n_2, \dots, n_K]$ represents the cumulative pseudo-label prediction count of each class

for the image I_i over multiple training epochs. The Pseudo-label Consistency Score (PCS) for I_i is defined as:

$$\text{PCS}_i = n_{\max} - n_{\text{second}}, \quad (3)$$

where $n_{\max} = \max(Q_i)$ denotes the number of times the most frequently predicted class appears, and $n_{\text{second}} = \max(Q_i \setminus \{n_{\max}\})$ corresponds to the count of the second most frequent prediction. A larger PCS_i implies higher consistency in pseudo-label predictions, reflecting stronger semantic confidence from the model regarding I_i .

Dynamic Threshold Adjustment. To ensure a stable and effective separation of ambiguous and refinable samples, we design a dynamic threshold adjustment mechanism that adaptively updates the threshold τ during training. Among noisy samples with low confidence, we define *refinable* samples S_R as follows:

$$S_R = \{(I_i, T_i) \mid y_i = 0, \text{PCS}_i \geq \tau\}$$

and *ambiguous* samples S_A as follows:

$$S_A = \{(I_i, T_i) \mid y_i = 0, \text{PCS}_i < \tau\}$$

This mechanism is based on the target utilization rate λ_{target} , which gradually increases as training progresses:

$$\lambda_{\text{target}} = \lambda_{\min} + (\lambda_{\max} - \lambda_{\min}) \cdot \frac{t}{T} \quad (4)$$

where t is the current training iteration and T is the total number of training iterations. This encourages the model to utilize more samples progressively. Following, we compute the intermediate target threshold τ_{target} based on the discrepancy between current utilization λ and the expected target:

$$\tau_{\text{target}} = \tau - k \cdot (\lambda_{\text{target}} - \lambda) \quad (5)$$

Here, k is a sensitivity coefficient that controls the adjustment strength. If the actual utilization is lower than expected, τ will be reduced to include more samples, and vice versa. In this work, we update the threshold of current epoch τ_t using an Exponential Moving Average (EMA) to ensure smoothness:

$$\tau_t = (1 - \beta) \cdot \tau_{t-1} + \beta \cdot \tau_{\text{target}} \quad (6)$$

where β is the parameter controlling the update rate.

Adaptive Pair Optimization

To fully exploit the optimization potential of different types of pairs, we propose a training strategy named *Adaptive Pair Optimization* (APO). It progressively incorporates clean, refinable, and ambiguous pairs into the training process with tailored supervision objectives.

Stage 1: Clean-only Training. Training starts with only clean pairs \mathcal{C} to ensure stable initial convergence. The objective function for this phase incorporates standard triplet loss that encourages semantic alignment, along with cross-entropy loss for pseudo-label supervision and entropy regularization for promoting confident predictions. Specifically, the loss for clean data training is defined as:

$$\mathcal{L}_{\text{clean}} = \lambda_t \cdot \mathcal{L}_{\text{Triplet}}(I^c, T^c) + \lambda_{ce} \cdot \mathcal{L}_{\text{CE}} + \lambda_{en} \cdot \mathcal{L}_{\text{En}} \quad (7)$$

Noise	Methods	Flickr30K							MSCOCO 1K						
		Image→Text			Text→Image			Rsum	Image→Text			Text→Image			Rsum
		R@1	R@5	R@10	R@1	R@5	R@10		R@1	R@5	R@10	R@1	R@5	R@10	
20%	SCAN (ECCV'18)	59.1	83.4	90.4	36.6	67.0	77.5	414.0	66.2	91.0	96.4	45.0	80.2	89.3	468.1
	SGR (AAAI'21)	61.2	84.3	91.5	44.5	72.1	80.2	433.8	49.1	83.8	92.7	42.5	77.7	88.2	434.0
	NCR* (NIPS'21)	73.5	93.2	96.6	56.9	82.4	88.5	491.1	76.6	95.6	98.2	60.8	88.8	95.0	515.0
	DECL* (ACMMM'22)	77.5	93.8	97.0	56.1	81.8	88.5	494.7	77.5	95.9	98.4	61.7	89.3	95.4	518.2
	MSCN* (CVPR'23)	77.4	94.9	97.6	59.6	83.2	89.2	501.9	78.1	<u>97.2</u>	98.8	64.3	<u>90.4</u>	95.8	524.6
	BiCro* (CVPR'23)	78.1	94.4	97.5	60.4	84.4	89.9	504.7	78.8	96.1	98.6	63.7	90.3	<u>95.7</u>	523.2
	RCL* (TPAMI'23)	75.9	94.5	97.3	57.9	82.6	88.6	496.8	78.9	96.0	98.4	62.8	89.9	95.4	521.4
	CREAM* (TIP'24)	77.4	95.0	97.5	58.7	84.1	89.8	502.3	78.9	96.3	<u>98.7</u>	63.3	90.5	95.3	523.0
	L2RM* (CVPR'24)	77.9	<u>95.2</u>	97.8	59.8	83.6	89.5	503.8	<u>80.2</u>	96.3	98.5	64.2	90.1	95.4	<u>524.7</u>
	PCSR*	78.7	<u>95.1</u>	<u>97.9</u>	60.9	<u>83.7</u>	89.4	505.7	80.5	97.3	98.9	<u>63.8</u>	89.7	95.1	525.4
40%	SCAN (ECCV'18)	29.9	60.5	72.5	16.4	38.5	48.6	266.4	30.1	65.2	79.2	18.9	51.1	69.9	314.4
	SGR (AAAI'21)	47.2	76.4	83.2	34.5	60.3	70.5	372.1	43.9	78.3	89.3	37.0	72.8	85.1	406.4
	NCR* (NIPS'21)	75.3	92.1	95.2	56.2	80.6	87.4	486.8	76.5	95.0	98.2	60.7	88.5	95.0	513.9
	DECL* (ACMMM'22)	72.7	92.3	95.4	53.4	79.4	86.4	479.6	75.6	95.5	98.3	59.5	88.3	95.4	512.0
	MSCN* (CVPR'23)	74.4	94.4	96.9	57.2	81.7	87.6	492.2	74.8	94.9	98.0	60.3	88.5	94.3	510.9
	BiCro* (CVPR'23)	74.6	92.7	96.2	55.5	81.0	87.4	487.5	77.0	<u>95.9</u>	98.3	61.1	89.2	94.9	517.1
	RCL* (TPAMI'23)	72.7	92.7	96.1	54.8	80.0	87.1	483.4	77.0	<u>95.5</u>	<u>98.4</u>	61.2	90.2	95.1	515.3
	CREAM* (TIP'24)	<u>76.3</u>	<u>93.3</u>	<u>97.1</u>	<u>57.0</u>	82.6	<u>88.7</u>	<u>495.0</u>	76.5	95.0	98.2	61.7	<u>89.4</u>	95.6	516.8
	L2RM* (CVPR'24)	75.8	93.2	96.9	56.3	81.0	87.3	490.5	<u>77.5</u>	95.8	98.5	<u>62.0</u>	89.1	94.9	<u>517.7</u>
	PCSR*	76.6	94.5	97.3	57.4	<u>82.4</u>	88.9	497.1	77.7	96.0	98.1	63.1	<u>89.4</u>	<u>95.5</u>	519.8
60%	SCAN (ECCV'18)	16.9	39.3	53.9	2.8	7.4	11.4	131.7	27.8	59.8	74.8	16.8	47.8	66.4	293.4
	SGR (AAAI'21)	28.7	58.0	71.0	23.8	49.5	60.7	291.7	37.6	73.3	86.3	33.8	68.6	81.7	381.3
	NCR* (NIPS'21)	68.7	89.9	95.5	52.0	77.6	84.9	468.6	72.7	94.0	97.6	57.9	87.0	94.1	503.3
	DECL* (ACMMM'22)	65.2	88.4	94.0	46.8	74.0	82.2	450.6	73.0	94.2	97.1	57.0	86.8	94.0	502.5
	MSCN* (CVPR'23)	70.4	91.0	94.9	53.4	77.8	84.1	471.6	74.4	95.1	97.9	59.2	87.1	92.8	506.5
	BiCro* (CVPR'23)	67.6	90.0	94.4	51.2	77.6	84.7	466.3	73.9	94.4	97.8	58.3	87.2	93.5	505.5
	RCL* (TPAMI'23)	67.7	89.1	93.6	48.0	74.9	83.3	456.6	74.0	94.3	97.6	57.6	86.4	93.5	503.3
	CREAM* (TIP'24)	<u>70.6</u>	<u>91.2</u>	96.1	<u>53.3</u>	<u>79.2</u>	87.0	<u>477.4</u>	74.7	<u>94.7</u>	98.0	<u>59.7</u>	88.0	94.6	<u>509.9</u>
	L2RM* (CVPR'24)	70.0	90.8	95.4	51.3	76.4	83.7	467.6	<u>75.4</u>	<u>94.7</u>	<u>97.9</u>	59.2	87.4	93.6	508.4
	PCSR*	72.8	91.9	<u>95.8</u>	54.8	79.7	<u>86.5</u>	480.5	76.3	94.8	<u>97.9</u>	61.7	<u>87.5</u>	<u>94.0</u>	512.2
80%	SCAN (ECCV'18)	5.1	18.1	27.3	3.9	13.1	19.1	86.6	22.2	51.9	67.5	13.8	41.1	58.6	255.1
	SGR (AAAI'21)	13.7	35.1	47.6	12.1	30.9	41.9	181.3	26.7	60.7	75.6	25.3	58.2	72.6	319.1
	NCR* (NIPS'21)	1.4	7.1	11.7	1.5	5.4	9.3	36.4	21.6	52.6	67.6	15.1	38.1	49.8	244.8
	DECL* (ACMMM'22)	53.4	78.8	86.9	37.6	63.8	73.9	394.4	64.8	90.5	96.2	53.7	81.7	93.0	473.0
	MSCN* (CVPR'23)	1.0	4.4	9.1	0.4	1.4	2.5	18.8	66.8	91.6	96.2	52.7	83.0	90.9	481.2
	BiCro* (CVPR'23)	3.6	13.9	20.5	1.7	7.5	13.0	60.2	62.2	88.6	96.1	47.4	79.2	88.5	460.5
	RCL* (TPAMI'23)	51.7	75.8	84.4	34.5	61.2	70.7	378.3	67.4	90.8	96.0	50.6	81.0	90.1	475.9
	CREAM* (TIP'24)	<u>58.7</u>	<u>83.3</u>	<u>90.1</u>	<u>40.8</u>	<u>67.1</u>	<u>76.3</u>	<u>416.3</u>	68.6	<u>92.0</u>	<u>96.4</u>	52.4	84.8	<u>92.8</u>	<u>487.2</u>
	L2RM* (CVPR'24)	55.7	80.8	87.8	39.4	65.4	74.9	404.0	<u>69.0</u>	<u>92.0</u>	96.2	<u>52.7</u>	<u>82.4</u>	<u>92.8</u>	482.6
	PCSR*	63.1	87.1	93.2	44.6	70.9	78.6	437.5	71.6	92.9	96.7	55.3	<u>83.9</u>	93.3	493.7

Table 1: Image-Text Retrieval on Flickr30K and MS-COCO 1K datasets under different noise ratios. * indicates the noise robust method. The best and sub-optimal indicators are represented in **bold** and underline respectively.

where the triplet loss $\mathcal{L}_{Triplet}(I, T)$ is computed as:

$$\mathcal{L}_{Triplet}(I, T) = \sum_{(I, T) \in \mathcal{D}} \left[\hat{\alpha} - S(I, T) + S(I, \hat{T}) \right]^+ + \left[\hat{\alpha} - S(I, T) + S(\hat{I}, T) \right]^+ \quad (8)$$

$\hat{\alpha}$ is defined based on the work of (Duan et al. 2024):

$$\hat{\alpha} = \frac{m^{S_p(p_i^a, p_j^c)} - 1}{m - 1} \alpha, \quad (9)$$

where m and α are predefined parameters, S_p is a similarity function used to calculate the similarity of two distributions.

Stage 2: Clean + Refinable. In this stage, we incorporate refinable pairs \mathcal{R} into training. These samples show con-

sistent pseudo-labels, suggesting that although their paired captions may be misaligned, their images still provide useful training signals. To correct this, we apply a *caption refinement*: for each image $I_i^r \in \mathcal{R}$, we retrieve the most semantically similar caption T_j^c from clean texts T^c based on pseudo-label similarity.

$$j = \arg \max_b S_p(p_i^r, p_b^c) \quad (10)$$

Based on the rematched pairs (I_i^r, T_j^c) , the final training loss used for refinable data training is formulated as:

$$\mathcal{L}_{refinable} = \mathcal{L}_{Triplet}(I_i^r, T_j^c) \quad (11)$$

Stage 3: Clean + Refinable + Ambiguous. In the final training stage, we incorporate ambiguous pairs \mathcal{A} into the optimization process to further exploit potentially informative supervision. However, due to their inherent uncertainty,

we adopt a noise-robust loss function to stabilize training. Specifically, we utilize the *Generalized Cross-Entropy loss* (GCE), which combines the benefits of both mean absolute error and standard cross-entropy by controlling the influence of low-confidence predictions. The GCE loss is defined as:

$$\mathcal{L}_{\text{GCE}} = \frac{1}{B} \sum_{i=1}^B \left(\frac{1 - (p_i^a(\hat{q}_i^a))^\gamma}{\gamma} + \frac{1 - (q_i^a(\hat{p}_i^a))^\gamma}{\gamma} \right) \quad (12)$$

$\gamma \in (0, 1]$ is a hyperparameter controlling the robustness of the loss, set to 0.7 as recommended in (Zhang and Sabuncu 2018). The total loss used for ambiguous data training for this stage is formulated as:

$$\mathcal{L}_{\text{ambiguous}} = \lambda_t \cdot \mathcal{L}_{\text{Triplet}}(I^a, T^a) + \lambda_{gce} \cdot \mathcal{L}_{\text{GCE}} + \lambda_{en} \cdot \mathcal{L}_{\text{En}} \quad (13)$$

Experiments

Datasets

We evaluate our proposed method on three widely used multimodal datasets:

Flickr30K. This dataset comprises 31,000 images sourced from Flickr, each paired with five descriptive captions that provide detailed semantic annotations. The dataset encompasses a wide range of scenes, objects, and human activities, making it a standard benchmark for cross-modal retrieval. Following common practice, we utilize 29,000 images for training, 1,000 for validation, and 1,000 for testing.

MSCOCO. The MSCOCO dataset contains 123,287 images, each annotated with five captions describing various visual scenes and objects in daily life. It serves as a key benchmark for evaluating multimodal learning methods. In our experiments, we adopt the data split: 113,287 images for training, and 5,000 each for validation and testing.

Conceptual Captions. This large-scale dataset, curated from the web, consists of approximately 3.3 million image-text pairs. Each image is associated with an automatically mined caption, leading to a noise rate of 3%–20%. Following prior works (Huang et al. 2021; Duan et al. 2024), we adopt the CC152K subset, which includes 150,000 training images and 1,000 images each for validation and testing.

Evaluation Metrics

We adopt the standard Recall@K metrics (K = 1, 5, 10), denoted as R@1, R@5, and R@10, to evaluate image-text matching performance. These metrics measure the proportion of ground-truth instances successfully retrieved within the top-K ranked results, reflecting the effectiveness of the retrieval model in ranking relevant pairs at higher positions.

Implementation Details

For fair comparison, we use SGR as our cross-modal retrieval backbone, same as NCR(Huang et al. 2021). We set the number of pseudo-label categories to $K = 256$; the corresponding ablation results are presented in Table 3. For other training details, we set the batch size $B = 128$, parameters $k = 0.2$, $\lambda_{max} = 0.9$, $\lambda_{min} = 0.4$, $\beta = 0.7$. Following previous studies (Duan et al. 2024), we set the

Methods	Image-Text			Text-Image			Rsum
	R@1	R@5	R@10	R@1	R@5	R@10	
SCAN	30.5	55.3	65.3	26.9	53.0	64.7	295.7
SGR	11.3	29.7	39.6	13.1	30.1	41.3	165.4
NCR*	39.5	64.5	73.5	40.3	64.5	73.2	355.5
DECL*	39.0	66.1	75.5	40.7	66.3	76.7	364.3
MSCN*	40.1	65.7	76.0	40.6	67.1	76.4	366.7
BiCro*	40.8	67.2	76.4	42.1	67.6	76.4	370.2
RCL*	41.7	66.0	73.6	41.6	66.4	75.1	364.4
CREAM*	40.3	68.5	<u>77.1</u>	40.2	68.2	78.3	373.1
L2RM*	<u>43.0</u>	67.5	75.7	<u>42.8</u>	<u>68.0</u>	<u>77.2</u>	<u>374.2</u>
PCSR(Ours)	43.7	<u>67.7</u>	77.2	43.1	67.7	76.3	375.3

Table 2: Image-Text Retrieval on CC152K. * indicates the noise robust method. The best and sub-optimal indicators are represented in **bold** and underline respectively.

loss weights $\lambda_{en} = 10$, $\lambda_t = \lambda_{ce} = \lambda_{gce} = 1$ and parameter $m = 10$, $\alpha = 0.2$. Specifically, we warm up the model for 5 epochs on Flickr30K and MSCOCO datasets, 10 epochs on CC152K dataset. After warming up, we train our model for 50 epochs, Stage 1 corresponds to epochs 1–25, Stage 2 to epochs 20–40, and Stage 3 to epochs 41–50.

Main Results

We conduct experiments on three benchmark datasets to validate the effectiveness of our method in both synthetic and real-world noisy correspondence scenarios. To highlight the robustness of our framework, we compare it against two groups of baselines: (1) conventional cross-modal methods that do not account for noisy correspondence, including SCAN (Lee et al. 2018) and SGR (Diao et al. 2021); and (2) recent noise-aware methods, including NCR (Huang et al. 2021), DECL (Qin et al. 2022), MSCN (Han et al. 2023), BiCro (Yang et al. 2023), RCL (Hu et al. 2023), CREAM (Ma et al. 2024), and L2RM (Han et al. 2024).

Evaluation under Simulated Noise. To simulate noisy training conditions, we introduce synthetic noise by randomly mismatching a portion of image-caption pairs in Flickr30K and MS-COCO datasets, with noise ratios set to 20%, 40%, 60% and 80%. As shown in Table 1, our approach consistently achieves superior performance across all noise levels. Notably, when compared with the strongest baseline CREAM(Ma et al. 2024), our method delivers absolute improvements on Flickr30K, and MSCOCO 1K under the corresponding noise ratios. Notably, in the 80% noise setting, our PCSR method outperforms CREAM by +4.4 R@1 on image-to-text retrieval (63.1 vs. 58.7) and +3.8 R@1 on text-to-image retrieval (44.6 vs. 40.8), demonstrating superior robustness against noisy correspondences.

Evaluation under Real-World Noise. To assess performance in realistic scenarios, we test our method on the CC152K dataset, which contains naturally occurring image-text mismatches due to imperfect web-crawled annotations. This dataset presents a more challenging setting as the noise is unlabeled and often semantically subtle. As shown in Table 2, our method demonstrates strong retrieval capabilities

K	Image \rightarrow Text			Text \rightarrow Image			Rsum
	R@1	R@5	R@10	R@1	R@5	R@10	
32	69.0	90.0	95.0	53.1	78.3	85.9	471.3
64	71.3	90.9	95.3	52.6	78.9	85.5	474.4
128	71.4	91.5	95.9	52.8	79.2	85.6	476.4
256	72.8	91.9	95.8	53.8	79.7	86.5	480.5
512	70.0	91.3	95.4	53.2	79.1	86.1	475.1

Table 3: Effect of different K values on retrieval performance, conducted on Flickr30k under 60% noise.

Subsets			Image \rightarrow Text			Text \rightarrow Image			Rsum
S_C	S_A	S_R	R@1	R@5	R@10	R@1	R@5	R@10	
✓	✓	✓	72.8	91.9	95.8	53.8	79.7	86.5	480.5
✓		✓	71.8	92.0	95.7	53.9	79.3	86.1	478.8
✓	✓		70.9	91.5	95.9	52.8	78.7	85.6	475.4
✓			70.8	90.6	94.8	53.1	78.8	85.5	473.6

Table 4: Component ablation study on retrieval performance, conducted on Flickr30k under 60% noise.

ties under real-world noise. It surpasses the best-performing baseline L2RM by a margin of 1.1 points in overall retrieval metrics. This demonstrates our framework’s ability to effectively model and utilize noisy data, maintaining robust cross-modal alignment even in complex environments.

Ablation Studies

To assess the effectiveness of each PCSR component, we conduct ablation on Flickr30K with 60% noise ratio.

Effect of K value. Our PCSR framework relies on the category predictions generated by the pseudo-classifier, making the value of K particularly critical. As shown in Table 3, a moderately sized K allows our pseudo-classifier to model image semantics with appropriate granularity, which facilitates more accurate supervision from clean samples and enhances the quality of pseudo-caption selection.

The impact of each subset. To evaluate the contribution of different subsets in our training strategy, we conduct an ablation study on the inclusion of clean (S_C), ambiguous (S_A), and refinable (S_R) pairs. As shown in Table 4, training with all three subsets achieves the best performance, with an Rsum of 480.5. Removing either ambiguous or refinable pairs leads to a noticeable drop in retrieval accuracy, confirming the complementary benefits of both pseudo-label-guided strategies. Notably, excluding S_A slightly degrades performance compared to the full setting, while removing S_R causes a larger drop in R@1 for both image-to-text and text-to-image. This confirms the effectiveness of exploiting confident refinable pairs via semantic realignment, as well as the robust handling of ambiguous data via the L_{GCE} loss.

The impact of method components. As shown in Table 5, T_R denotes the use of the refinable subset’s training

Components		Image \rightarrow Text			Text \rightarrow Image			Rsum
CGD	APO	R@1	R@5	R@10	R@1	R@5	R@10	
✓	✓	72.8	91.9	95.8	53.8	79.7	86.5	480.5
✓	T_R	71.8	91.7	95.5	53.4	78.8	85.6	476.9
✓	T_A	70.8	90.6	94.8	53.1	78.8	85.5	473.6
		70.8	90.3	94.4	53.1	79.0	85.9	473.5

Table 5: Component ablation study on retrieval performance, conducted on Flickr30k under 60% noise.

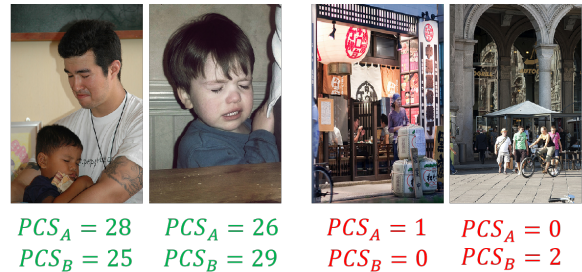


Figure 3: At epoch 40, we computed PCS values across all images using both model A and model B. These examples illustrate that semantically simpler images (left) lead to more stable pseudo-label predictions and thus higher PCS, complex semantics (right) correspond to lower PCS.

method (Eq. (11)) for both refinable and ambiguous subsets, while T_A applies the ambiguous subset’s training strategy (Eq. (13)) to train both refinable and ambiguous subsets. Our experiment proves that in the absence of accurate data partitioning, loss functions that were previously effective may no longer offer reliable supervisory signals.

Effect of PCS value. PCS measures the consistency of pseudo-classifier predictions, indicating an image’s semantic complexity. As shown in Figure 3, images with fewer distinctive features yield higher PCS, while semantically complex ones exhibit lower PCS due to prediction instability.

Conclusion

In this paper, we propose a novel Pseudo-label Consistency-Guided Sample Refinement (PCSR) framework for robust image-text retrieval under noisy correspondence. Our method explicitly distinguishes between clean, refinable, and ambiguous samples based on pseudo-label consistency. We design a three-stage optimization strategy: clean samples are directly trained, refinable samples are enhanced via text replacement, and ambiguous samples are optimized using noise-tolerant loss. This fine-grained division allows us to apply targeted training strategies according to sample characteristics, enabling better use of partially reliable supervision. Extensive experiments demonstrate that our method achieves substantial performance gains over existing state-of-the-art methods, especially under high-noise conditions.

Acknowledgments

This work was partially supported by the National Science Foundation of China under Grants 62376175, U2333211 and 22494712, the 111 Project under Grant B21044, and the Natural Science Foundation of Sichuan Province under Grants 2024NSFTD0035, 2021ZDZX0011, and 2025ZNS-FSC0480.

References

- Anderson, P.; He, X.; Buehler, C.; Teney, D.; Johnson, M.; Gould, S.; and Zhang, L. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 6077–6086.
- Bai, Y.; Yang, E.; Han, B.; Yang, Y.; Li, J.; Mao, Y.; Niu, G.; and Liu, T. 2021. Understanding and improving early stopping for learning with noisy labels. *Advances in Neural Information Processing Systems*, 34: 24392–24403.
- Chen, J.; Dun, C.; and Kyriillidis, A. 2024. Fast fixmatch: Faster semi-supervised learning with curriculum batch size. In *2024 IEEE International Symposium on Information Theory*, 1836–1841. IEEE.
- Chen, J.; Hu, H.; Wu, H.; Jiang, Y.; and Wang, C. 2021. Learning the best pooling strategy for visual semantic embedding. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 15789–15798.
- Chen, M.; and Wang, C. 2024. Multi-head co-training: An uncertainty-aware and robust semi-supervised learning framework. *Knowledge-Based Systems*, 302: 112325.
- Chuang, C.-Y.; Hjelm, R. D.; Wang, X.; Vineet, V.; Joshi, N.; Torralba, A.; Jegelka, S.; and Song, Y. 2022. Robust contrastive learning against noisy views. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 16670–16681.
- 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.
- Duan, Y.; Gu, Z.; Ying, Z.; Qi, L.; Meng, C.; and Shi, Y. 2024. Pc2: Pseudo-classification based pseudo-captioning for noisy correspondence learning in cross-modal retrieval. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 9397–9406.
- Faghri, F.; Fleet, D. J.; Kiros, J. R.; and Fidler, S. 2017. Vse++: Improving visual-semantic embeddings with hard negatives. In *British Machine Vision Conference*.
- Feng, Y.; Zhu, H.; Peng, D.; Peng, X.; and Hu, P. 2023. RONO: robust discriminative learning with noisy labels for 2D-3D cross-modal retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11610–11619.
- 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.
- Han, H.; Miao, K.; Zheng, Q.; and Luo, M. 2023. Noisy correspondence learning with meta similarity correction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7517–7526.
- Han, H.; Zheng, Q.; Dai, G.; Luo, M.; and Wang, J. 2024. Learning to rematch mismatched pairs for robust cross-modal retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26679–26688.
- Heidari, M.; Zhang, H.; and Guo, Y. 2024. Reinforcement learning guided semi-supervised learning. *Advances in Neural Information Processing Systems*, 37: 136990–137009.
- Hu, P.; Huang, Z.; Peng, D.; Wang, X.; and Peng, X. 2023. Cross-modal retrieval with partially mismatched pairs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8): 9595–9610.
- Huang, Z.; Niu, G.; Liu, X.; Ding, W.; Xiao, X.; Wu, H.; and Peng, X. 2021. Learning with noisy correspondence for cross-modal matching. *Advances in Neural Information Processing Systems*, 34: 29406–29419.
- Iscen, A.; Valmadre, J.; Arnab, A.; and Schmid, C. 2022. Learning with neighbor consistency for noisy labels. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 4672–4681.
- Jiang, L.; Huang, D.; Liu, M.; and Yang, W. 2020. Beyond synthetic noise: Deep learning on controlled noisy labels. In *International conference on Machine Learning*, 4804–4815. PMLR.
- Lee, D.-H.; et al. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *International Conference on Machine Learning*, volume 3, 896. Atlanta.
- 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*, 201–216.
- Li, K.; Zhang, Y.; Li, K.; Li, Y.; and Fu, Y. 2019. Visual semantic reasoning for image-text matching. In *Proceedings of the IEEE/CVF international conference on Computer Vision*, 4654–4662.
- Li, Y.; Huang, H.; Xu, J.; and Huang, S.-L. 2024. NAC: Mitigating Noisy Correspondence in Cross-Modal Matching Via Neighbor Auxiliary Corrector. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing*, 6815–6819. IEEE.
- Liu, C.; Mao, Z.; Zhang, T.; Xie, H.; Wang, B.; and Zhang, Y. 2020. Graph structured network for image-text matching. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 10921–10930.
- Liu, Y.; Feng, W.; Liu, Z.; Huang, S.; and Lv, J. 2025a. Aligning Information Capacity Between Vision and Language via Dense-to-Sparse Feature Distillation for Image-Text Matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 21679–21688.

- Liu, Y.; Liu, M.; Huang, S.; and Lv, J. 2025b. Asymmetric Visual Semantic Embedding Framework for Efficient Vision-Language Alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 5676–5684.
- Ma, X.; Yang, M.; Li, Y.; Hu, P.; Lv, J.; and Peng, X. 2024. Cross-modal retrieval with noisy correspondence via consistency refining and mining. *IEEE transactions on Image Processing*, 33: 2587–2598.
- Miyato, T.; Maeda, S.-I.; Koyama, M.; and Ishii, S. 2019. Virtual Adversarial Training: A Regularization Method for Supervised and Semi-Supervised Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8): 1979–1993.
- 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.
- 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.
- Wang, Q.; Han, B.; Liu, T.; Niu, G.; Yang, J.; and Gong, C. 2021. Tackling Instance-Dependent Label Noise via a Universal Probabilistic Model. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(11): 10183–10191.
- Wang, S.; Wang, R.; Yao, Z.; Shan, S.; and Chen, X. 2020. Cross-modal scene graph matching for relationship-aware image-text retrieval. In *Proceedings of the IEEE/CVF winter conference on Applications of Computer Vision*, 1508–1517.
- Wen, T.; Lai, S.; and Qian, X. 2021. Preparing lessons: Improve knowledge distillation with better supervision. *Neurocomputing*, 454: 25–33.
- Yan, J.; Luo, L.; Deng, C.; and Huang, H. 2023. Adaptive hierarchical similarity metric learning with noisy labels. *IEEE Transactions on Image Processing*, 32: 1245–1256.
- Yang, S.; Li, Q.; Li, W.; Li, X.; and Liu, A.-A. 2022. Dual-level representation enhancement on characteristic and context for image-text retrieval. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(11): 8037–8050.
- Yang, S.; Xu, Z.; Wang, K.; You, Y.; Yao, H.; Liu, T.; and Xu, M. 2023. Bicro: Noisy correspondence rectification for multi-modality data via bi-directional cross-modal similarity consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19883–19892.
- Zhai, X.; Mustafa, B.; Kolesnikov, A.; and Beyer, L. 2023. Sigmoid Loss for Language Image Pre-Training. In *2023 IEEE/CVF International Conference on Computer Vision*, 11941–11952.
- Zhang, H.; Mao, Z.; Zhang, K.; and Zhang, Y. 2022a. Show your faith: Cross-modal confidence-aware network for image-text matching. In *Proceedings of the AAAI conference on Artificial Intelligence*, volume 36, 3262–3270.
- Zhang, K.; Mao, Z.; Wang, Q.; and Zhang, Y. 2022b. Negative-aware attention framework for image-text matching. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 15661–15670.
- Zhang, Q.; Lei, Z.; Zhang, Z.; and Li, S. Z. 2020. Context-aware attention network for image-text retrieval. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 3536–3545.
- Zhang, S.; Li, Y.; Tian, J.; Man, Z.; Chung, C. Y.; and Shen, W. 2024. Improving Battery Life Prediction with Unlabeled Data: Confidence-Weighted Semi-Supervised Learning with Label Propagation. *IEEE Transactions on Transportation Electrification*.
- Zhang, Z.; and Sabuncu, M. 2018. Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in Neural Information Processing Systems*, 31.
- Zhao, X.; Li, D.; Zhong, Y.; Hu, B.; Chen, Y.; Hu, B.; and Zhang, M. 2024. SEER: Self-Aligned Evidence Extraction for Retrieval-Augmented Generation. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 3027–3041. Association for Computational Linguistics.