

Neighbor-aware Instance Refining with Noisy Labels for Cross-Modal Retrieval

Yizhi Liu^{1,2*}, Ruitao Pu^{1,2*}, Shilin Xu¹, Yingke Chen³, Quan-Hui Liu¹, Yuan Sun^{1,2,4†}

¹College of Computer Science, Sichuan University, Chengdu, China, 610065

²State Key Laboratory of AI Safety, Beijing, China, 100086

³Department of Computer and Information Sciences, Northumbria University, Newcastle upon Tyne, UK, NE1 8ST

⁴National Key Laboratory of Fundamental Algorithms and Models for Engineering Numerical Simulation, Sichuan University, Chengdu, China, 610207

liyuzhi@stu.scu.edu.cn, ruitaopu@gmail.com, xushilin990@gmail.com, yke.chen@gmail.com, quanhuiliu@scu.edu.cn, sunyuan_work@163.com.

Abstract

In recent years, Cross-Modal Retrieval (CMR) has made significant progress in the field of multi-modal analysis. However, since it is time-consuming and labor-intensive to collect large-scale and well-annotated data, the annotation of multi-modal data inevitably contains some noise. This will degrade the retrieval performance of the model. To tackle the problem, numerous robust CMR methods have been developed, including robust learning paradigms, label calibration strategies, and instance selection mechanisms. Unfortunately, they often fail to simultaneously satisfy model performance ceilings, calibration reliability, and data utilization rate. To overcome the limitations, we propose a novel robust cross-modal learning framework, namely Neighbor-aware Instance Refining with Noisy Labels (NIRNL). Specifically, we first propose Cross-modal Margin Preserving (CMP) to adjust the relative distance between positive and negative pairs, thereby enhancing the discrimination between sample pairs. Then, we propose Neighbor-aware Instance Refining (NIR) to identify pure subset, hard subset, and noisy subset through cross-modal neighborhood consensus. Afterward, we construct different tailored optimization strategies for this fine-grained partitioning, thereby maximizing the utilization of all available data while mitigating error propagation. Extensive experiments on three benchmark datasets demonstrate that NIRNL achieves state-of-the-art performance, exhibiting remarkable robustness, especially under high noise rates.

Code — <https://github.com/perquisite/NIRNL>

Introduction

With the rapid development of multimodal data on the Internet, cross-modal retrieval (CMR) has become a research hot topic in the field of multimodal learning. Recently, a large number of CMR methods have been proposed, which aim to retrieve semantically related samples across heterogeneous modalities, such as image-text (Wehrmann, Kolling, and Barros 2020; Ge et al. 2023) or video-image scenarios (Gorti et al. 2022; Fang et al. 2023). Although these methods obtain the desired performance, most of them (Zhen et al.

2019; Pu et al. 2025b) rely heavily on clean-annotated data to learn multi-modal representations in a shared semantic space. In practice, collecting perfectly labeled data is both expensive and time-consuming. Due to annotation mistakes and ambiguous semantics, this inevitably results in noisy labels. The presence of noisy labels can severely harm the learning model, thereby weakening retrieval performance. To this end, some weakly supervised or semi-supervised CMR methods (Mandal, Rao, and Biswas 2019) have been proposed, which attempt to alleviate the influence of perfect labeling. However, they still have an implicit assumption, that is, the available labels are entirely correct. Therefore, how to robustly learn from noisy data has become a key challenge.

Recently, various CMR methods (Feng et al. 2025a, 2023; Pu, Peng, and Hua 2024; Feng et al. 2025b; Pu et al. 2025c; Wang et al. 2024) have been proposed to robustly learn representations from multi-modal data with noisy labels. They could be roughly divided into three strategies, i.e., robust learning (Hu et al. 2021; Feng et al. 2023), label calibration (Pu, Peng, and Hua 2024; Okamura, Harakawa, and Iwahashi 2023), and instance selection (Pu et al. 2025c; Wang et al. 2024). To be specific, robust learning aims to design a robust loss, thereby directly tolerating the influence of noisy labels. To directly eliminate the influence of noisy labels, label calibration refines the labels from the source, thus improving the upper limit of model performance. To filter out instances with noisy labels, instance selection first identifies the wrong labels and then trains the model with clean data. Although they have demonstrated promising outcomes, these strategies generally suffer from some limitations. For example, robust learning relies on prior assumptions about noise distribution and can only tolerate noise, but cannot eliminate its limitations on the upper limit of model performance. Label calibration could introduce new noise or amplify the errors when class boundaries are ambiguous or the noise distribution heavily overlaps with the true distribution. Instance selection is sensitive to the pre-set threshold, which can easily lead to the filtering of clean instances or the omission of noisy instances. Meanwhile, it could also result in a significant waste of training data. In summary, dynamically coordinating model performance ceilings, calibration reli-

*Equal contribution.

†Corresponding author.

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

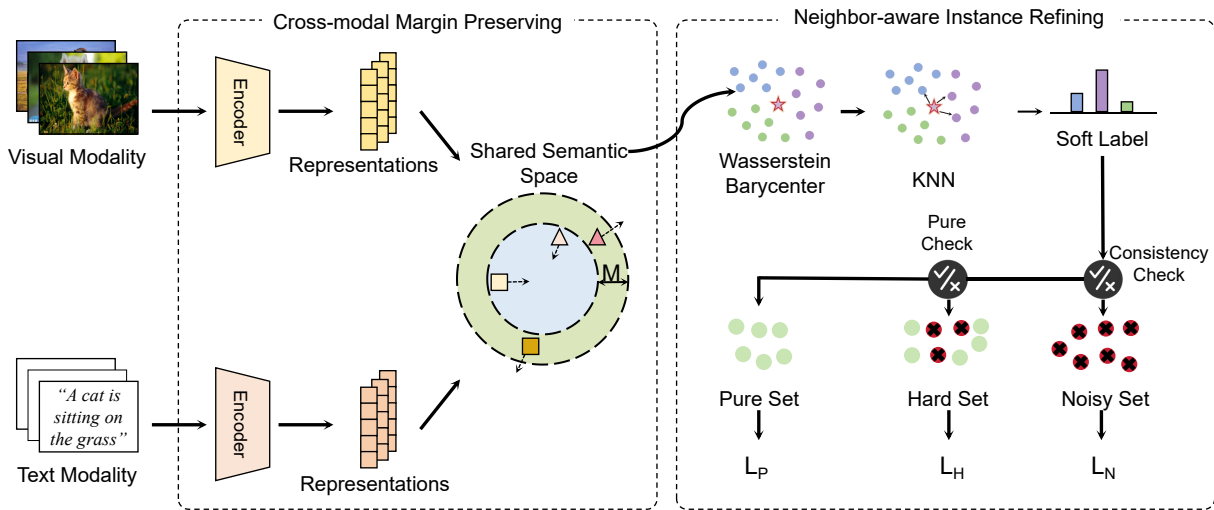


Figure 1: The framework of our proposed NIRNL. Our framework comprises two core modules operating in parallel: **the Cross-modal Margin Preserving (CMP)** module and **the Neighbor-aware Instance Refining (NIR)** module. The CMP module refines the global structure of the embedding space, promoting proximity between positive pairs (indicated by light yellow and light pink) while enforcing separation of negative pairs (indicated by dark yellow and dark pink). For clarity, only image samples are visualized in the NIR module. The NIR module initially computes the Wasserstein Barycenter of samples and generates soft labels through KNN. It subsequently partitions the dataset into pure, hard, and noisy subsets by evaluating the consistency between soft labels and ground-truth labels. Finally, we design three different loss functions for each subset to dig up as much semantic information as possible.

bility, and data utilization efficiency under complex noise scenarios remains a key yet challenging problem.

To overcome the aforementioned challenges, we propose a novel Neighbor-aware Instance Refining with Noisy Labels (NIRNL) framework to mitigate the negative effects of noisy labels. As shown in Fig.1, NIRNL is composed of two core modules, i.e., Cross-modal Margin Preserving (CMP) and Neighbor-aware Instance Refining (NIR). Firstly, CMP imposes constraints on the relative distances between positive and negative pairs to enhance the discriminability of representations in the shared semantic space. Then, NIR retrieves nearest neighbors and evaluates their consistency with ground-truth labels to generate soft labels, which could effectively partition these instances into three distinct subsets, i.e., pure, hard, and noisy subsets. For the pure subset, where labels are highly reliable, NIR optimizes instances directly to fully exploit their supervisory signals. For the hard subset, where label reliability is uncertain, NIR applies a weighted optimization strategy to balance the influence of potentially corrupted annotations while retaining useful information. For the noisy subset, where labels are deemed unreliable, NIR performs label calibration to recover informative content and mitigate the adverse effects of noise. The main contributions are summarized as follows:

- To achieve high tolerance to noisy labels, we propose a robust cross-modal learning framework (NIRNL), which unifies robust learning, label calibration, and instance selection. To the best of our knowledge, this is the first work to balance model performance ceiling, calibration reliability, and data utilization rate in multi-modal learning with noisy labels.

- We propose neighbor-aware instance refining that dynamically partitions training instances into pure, hard, and noisy subsets by perceiving the global neighborhood distribution. Further, we design a customized optimization strategy to explore all available information as much as possible, thereby enhancing the robustness.
- Extensive experiments comprehensively verify that our proposed NIRNL has remarkably superior performance over the current state-of-the-art methods.

Related Work

Cross-modal Retrieval

With the exponential growth of internet data, cross-modal retrieval (CMR) has emerged as a key technology for information retrieval across diverse data types (Su et al. 2025; Li et al. 2025b, 2024, 2025a). The fundamental challenge in CMR lies in bridging the heterogeneity gap between different modalities (Luo et al. 2025; Sun et al. 2023; Yin et al. 2025). Some methods focus on unsupervised methods. Early approaches employed single methods, such as UDCMH (Wu et al. 2018), which extract features and generate hash codes as pseudo-labels for learning. However, these methods heavily relied on the quality of the initial similarity matrix. DJSRH (Su, Zhong, and Zhang 2019) utilizes matrix fusion to incorporate neighborhood information from all modalities, but this introduced redundancy. Further advancing the field, researchers have adopted hybrid models, such as UGACH (Zhang, Peng, and Yuan 2018), UCCH (Hu et al. 2022), and UCGKANH (Lin et al. 2025). However, the performance of these methods is inherently constrained by their

reliance on pseudo-labels generated from the data itself.

Consequently, some supervised methods leverage explicit label information to learn a shared semantic space. These methods leverage uncertainty modeling. PCME (Chun et al. 2021) regards each sample as a probability distribution in the embedding space but heavily relies on paired data. Building upon this, DECL (Qin et al. 2022) employs evidential learning to quantify uncertainty caused by potentially noisy correspondences, yet it lacks uncertainty modeling for individual outcomes. Moreover, some methods have explored adversarial learning (Wang et al. 2017), graph neural networks DAGN (Qian et al. 2022; Liang et al. 2023), consistency learning DRCL (Pu et al. 2025a). Although the supervised methods discussed above have achieved impressive performance on various benchmarks, their success is largely predicated on a critical assumption: the availability of large-scale, meticulously annotated training datasets.

Cross-modal Retrieval with Noisy Labels

In practice, however, collecting and annotating such high-quality data is both prohibitively expensive and highly time-consuming. Consequently, researchers often turn to collecting web data, which offers scale, but is noisy, misguides training, corrupts semantic alignment, and hurts retrieval (Han et al. 2025). Noise-robust learning is now critical for real-world cross-modal search. These methods generally fall into three main categories.

Robust Learning methods design loss functions or training strategies that tolerate noise. For example, RONO (Feng et al. 2023) adopts discriminative center learning to pull clean samples closer and push noisy ones away. However, these methods suppress rather than correct noise. Instance Selection methods filter out noisy samples and train on clean subsets. RSHNL (Pu et al. 2025a) uses a self-paced learning strategy, while NRCH (Wang et al. 2024) applies the small-loss criterion. Yet, these methods often discard a large portion of data, including hard-but-clean samples that are essential for learning robust decision boundaries. Label Calibration methods take a direct approach by correcting noisy labels. UOT-RCL formulates label correction as a partial Optimal Transport problem to align noisy labels with corrected ones, while TCL (Li et al. 2021) leverages a small trusted dataset to guide co-learning. Although effective, these methods risk introducing new errors or amplifying existing ones.

However, these methods struggle to strike a balance between model performance ceilings, calibration reliability, and data utilization efficiency. To address this, our proposed NIRNL method uses consensus signals from heterogeneous sources to identify, refine, and fully exploit the data.

The Proposed Method

Preliminaries

Without loss of generality, taking the visual-text retrieval as an example, we define some denotations to present the learning with noisy labels (LNL) problem in cross-modal retrieval (CMR). Considering a dataset with noisy labels $D = \{\mathcal{V}_i, \mathcal{T}_i, \mathcal{Y}_i\}_{i=1}^N$ with N instances, where $(\mathcal{V}_i, \mathcal{T}_i)$ is the visual-text pair and $\mathcal{Y}_i = \{\mathcal{Y}_i^1, \mathcal{Y}_i^2, \dots, \mathcal{Y}_i^C\} \in \{0, 1\}^{1 \times C}$

denotes the corresponding noisy label over C categories. Further, if the i -th instance belongs to the c -th category, $\mathcal{Y}_i^c = 1$, otherwise $\mathcal{Y}_i^c = 0$. The fundamental goal of CMR is to learn a series of modality-specific sub-networks responsible for projecting multimodal data into the common semantic space where cross-modal similarity can be directly measured. Let the dimensionality of the common semantic space be d , sub-networks for the visual and text modalities are denoted as $\mathcal{F}^{\mathcal{V}}(\cdot, \Phi^{\mathcal{V}}) \in \mathbb{R}^{1 \times d}$ and $\mathcal{F}^{\mathcal{T}}(\cdot, \Phi^{\mathcal{T}}) \in \mathbb{R}^{1 \times d}$, respectively, where $\Phi^{\mathcal{V}}$ and $\Phi^{\mathcal{T}}$ are the learnable parameters. For brevity, we use $f_i^{\mathcal{V}}$ and $f_i^{\mathcal{T}}$ to their outputs, i.e., $f_i^{\mathcal{V}} = \mathcal{F}^{\mathcal{V}}(\mathcal{V}_i)$ and $f_i^{\mathcal{T}} = \mathcal{F}^{\mathcal{T}}(\mathcal{T}_i)$, respectively.

Cross-modal Margin Preserving

To strengthen the distinguishability among sample pairs, we propose Cross-modal Margin Preserving (CMP) to constrain the relative distance between positive and negative pairs, thereby yielding more compact representations for intra-class samples and differentiated representations for inter-class ones. Specifically, CMP enforces that the similarity of positive pairs exceeds that of negative pairs by at least a predefined margin. The CMP loss could be formulated as follows:

$$\begin{aligned} \mathcal{L}_{CMP} = & \frac{1}{N} \sum_{i=1}^N \sum_{\substack{j=1 \\ j \neq i}}^N |\Gamma(f_i^{\mathcal{V}}, f_j^{\mathcal{T}}) - \Gamma(f_i^{\mathcal{V}}, f_i^{\mathcal{T}}) + \mathcal{M}|_+ \\ & + \frac{1}{N} \sum_{i=1}^N \sum_{\substack{j=1 \\ j \neq i}}^N |\Gamma(f_i^{\mathcal{T}}, f_j^{\mathcal{V}}) - \Gamma(f_i^{\mathcal{T}}, f_i^{\mathcal{V}}) + \mathcal{M}|_+, \end{aligned} \quad (1)$$

where \mathcal{M} is the margin and $|\cdot|_+$ denotes the hinge function, which returns the input if it is positive and zero otherwise, i.e., $|x|_+ = \max(0, x)$.

Neighbor-aware Instance Refining

To mitigate the negative impact of noisy labels, previous methods (Yang et al. 2022; Pu et al. 2025c) typically select confident instances in small batches based on the small-loss criterion to guide model training. However, this strategy overlooks the global neighborhood structure among samples and discards potentially informative instances, making the performance of the model susceptible to the training data distribution and prone to performance bottlenecks. For this issue, we propose Neighbor-aware Instance Refining (NIR) to integrate neighborhood consistency for generating soft labels, thereby achieving a refined partitioning of samples and formulating corresponding learning strategies to maximize the use of all available information. First, NIR generates soft labels for the two modalities by using,

$$\begin{aligned} \hat{p}(c | \mathcal{V}_i) = & \frac{1}{K} \sum_{k=1, \mathcal{V}_k \in \mathcal{N}_i^{\mathcal{V}}}^K \mathbb{I}[y_k^c = 1], \quad c \in [1, 2, \dots, C], \\ \hat{p}(c | \mathcal{T}_i) = & \frac{1}{K} \sum_{k=1, \mathcal{T}_k \in \mathcal{N}_i^{\mathcal{T}}}^K \mathbb{I}[y_k^c = 1], \quad c \in [1, 2, \dots, C], \end{aligned} \quad (2)$$

where $\mathcal{N}_i^{\mathcal{V}}$ and $\mathcal{N}_i^{\mathcal{T}}$ denote the sets of the K nearest neighbors of \mathcal{V}_i and \mathcal{T}_i in the visual and textual modalities, respectively. Then, we evaluate the consistency between the

learned soft labels and ground-truth labels to divide instances into three subsets, i.e., pure subset, hard subset, and noisy subset. Among these, if an instance achieves consistency in two modalities, and it can be incorporated into the pure subset, which could be written as:

$$\mathcal{D}_P = \{(\mathcal{V}_i, \mathcal{T}_i, \mathcal{Y}_i) : \mathcal{Y}_i^{\text{argmax}_c \hat{p}(c|\mathcal{V}_i)} = 1 \ \& \ \mathcal{Y}_i^{\text{argmax}_c \hat{p}(c|\mathcal{T}_i)} = 1, i \in [1, 2, \dots, N]\}. \quad (3)$$

Further, if an instance only achieves consistency in one modality, and it can be incorporated into the hard subset, which could be formulated as:

$$\mathcal{D}_H = \{(\mathcal{V}_i, \mathcal{T}_i, \mathcal{Y}_i) : \mathcal{Y}_i^{\text{argmax}_c \hat{p}(c|\mathcal{V}_i)} = 1 \mid \mathcal{Y}_i^{\text{argmax}_c \hat{p}(c|\mathcal{T}_i)} = 1, i \in [1, 2, \dots, N]\} - \mathcal{D}_P. \quad (4)$$

At last, if an instance can not achieve consistency in any modalities, and it can be incorporated into the noisy subset, which could be given by:

$$\mathcal{D}_N = \{(\mathcal{V}_i, \mathcal{T}_i, \mathcal{Y}_i) : \mathcal{Y}_i^{\text{argmax}_c \hat{p}(c|\mathcal{V}_i)} = 0 \ \& \ \mathcal{Y}_i^{\text{argmax}_c \hat{p}(c|\mathcal{T}_i)} = 0, i \in [1, 2, \dots, N]\}. \quad (5)$$

For the three subsets, we deploy corresponding optimization strategies to learn semantic-consistency representations under the interference of noisy labels. To achieve this, we first extract semantic barycenters (Agueh and Carlier 2011) for each class based on the feature distribution, and this process could be expressed as:

$$\begin{aligned} \bar{u}_c &= \arg \min_u \sum_{i=1}^{N_c} \sum_{* \in \{\mathcal{V}, \mathcal{T}\}} \omega_i^* \cdot \mathcal{W}_{2, \lambda}^2(u, f_i^*), \\ \text{s.t.} \quad & \sum_{i=1}^{N_c} \sum_{* \in \{\mathcal{V}, \mathcal{T}\}} \omega_i^* = 1, \forall \omega_i^* \geq 0, c \in [0, 1, \dots, C], \end{aligned} \quad (6)$$

where N_c is the number of samples belonging to class c , ω_j^* denotes the weight for each sample (uniform weight in this paper), $\mathcal{W}_2^2(\cdot, \cdot)$ denotes the 2-Wasserstein distance, and λ is the regularization coefficient. In our paper, we use the expectation maximization (EM) algorithm (Dempster, Laird, and Rubin 1977; Feng, Ren, and Xie 2023) to converge the barycenters.

After mining barycenters, for the i -th sample in two modalities, we can obtain its probability that it belongs to its corresponding barycenter as follows:

$$s(*_i) = \sum_{c=1}^C \mathcal{Y}_i^c \frac{\Gamma(f_i^*, \bar{u}_c)}{\sum_{j=1}^C \Gamma(f_i^*, \bar{u}_j)}, * \in \{\mathcal{V}, \mathcal{T}\}, \quad (7)$$

where $\gamma(\cdot, \cdot)$ refers to the cosine similarity operator. Then, we adopt different optimization techniques for \mathcal{D}_P , \mathcal{D}_H , and \mathcal{D}_N . For the pure subset \mathcal{D}_P , we assume with high confidence that the labels are reliable and expect each sample to align closely with the corresponding semantic barycenter. To this end, we adopt the cross-entropy (CE) loss, which is known for its strong generalization ability (Ghosh, Kumar, and Sastry 2017) to optimize the model, i.e.,

$$\mathcal{L}_P = -\frac{1}{|\mathcal{D}_P|} \sum_{i=1}^{|\mathcal{D}_P|} \sum_{* \in \{\mathcal{V}, \mathcal{T}\}} \log s(*_i), \quad (8)$$

where $|\cdot|$ means the size of a set. For the hard subset \mathcal{D}_H , we believe that a majority of the samples are still clean, but due to label uncertainty, we cannot be fully confident. To address this, we adopt a loss-weighting strategy that allows the model to place greater emphasis on instances that are more likely to be clean, thereby improving robustness. The weight for each instance can be obtained as follows:

$$\ell_i = 1 - (1 - s(\mathcal{V}_i)) \cdot (1 - s(\mathcal{T}_i)). \quad (9)$$

Then, we guide the model optimization by adopting the weighted CE loss as follows:

$$\mathcal{L}_H = -\frac{1}{|\mathcal{D}_H|} \sum_{i=1}^{|\mathcal{D}_H|} \ell_i \sum_{* \in \{\mathcal{V}, \mathcal{T}\}} \log s(*_i). \quad (10)$$

For the noisy subset \mathcal{D}_N , we assume with high confidence that all samples are mislabeled and apply a label correction strategy to exploit useful information from them. For the i -th instance, we first fuse the soft labels from two modalities as follows:

$$\hat{p}_i^c = 1 - (1 - \hat{p}(c|\mathcal{V}_i)) \cdot (1 - \hat{p}(c|\mathcal{T}_i)), c \in [1, 2, \dots, C] \quad (11)$$

Then, we use the maximum value of \hat{p}_i^c as the corrected label, i.e., $\hat{y}_i = \text{argmax}_c \hat{p}_i^c$. Further, we construct the corresponding one-hot encoded label $\hat{\mathcal{Y}}_i \in \{0, 1\}^C$, where $\mathcal{Y}_i^c = 1$ if and only if $c = \hat{y}_i$. At last, recognizing that label correction may introduce biased predictions in complex data distributions, we mitigate the risk of overfitting caused by accumulated historical errors by the robust Mean Absolute Error (MAE) loss (Ghosh, Kumar, and Sastry 2017) as follows:

$$\mathcal{L}_N = -\frac{1}{|\mathcal{D}_N|} \sum_{i=1}^{|\mathcal{D}_N|} \sum_{* \in \{\mathcal{V}, \mathcal{T}\}} \left(1 - \frac{\Gamma(f_i^*, \bar{u}_{\hat{y}_i})}{\sum_{j=1}^C \Gamma(f_i^*, \bar{u}_j)} \right). \quad (12)$$

Optimization

Combining the aforementioned loss functions, the objective loss function of NIRNL is shown as follows:

$$\mathcal{L} = (\mathcal{L}_P + \mathcal{L}_H + \mathcal{L}_N) + \alpha \mathcal{L}_{CMP}, \quad (13)$$

where α is the balance coefficient to control the contribution of the corresponding loss term. Notably, the first three terms are applied only to samples within their corresponding subsets (i.e., \mathcal{D}_P , \mathcal{D}_H , and \mathcal{D}_N), while the last term is applied to all samples.

Experiments

Dataset

To verify the effectiveness of our NIRNL, we conduct extensive experiments on three available benchmark datasets, including Wikipedia (Rasiwasia et al. 2010), XMedia (Peng et al. 2015), and INRIA-Websearch (Krapac et al. 2010). A detailed description of all datasets is provided as follows:

- **Wikipedia** is made of 2,866 image-text pairs, which are labeled with one of 10 semantic categories. In our experiment, the image modality is represented by a 4,096-dimensional vector extracted by the pre-trained VGG-19

Method	Wikipedia								Mean
	0.2		0.4		0.6		0.8		
	I2T	T2I	I2T	T2I	I2T	T2I	I2T	T2I	
DGCPN	31.5	26.9	31.5	26.9	31.5	26.9	31.5	26.9	29.2
CIRH	23.6	22.5	23.6	22.5	23.6	22.5	23.6	22.5	23.1
UCCH	38.6	37.4	38.6	37.4	38.6	37.4	38.6	37.4	38.0
ALGCN	41.9	40.1	30.9	28.8	18.4	17.2	13.4	14.1	25.6
RONO	50.5	47.1	48.8	45.8	45.3	41.8	41.6	38.2	44.9
GNN4CMR	47.6	44.1	42.0	39.8	32.3	30.6	20.9	20.4	34.7
DHRL	46.2	42.8	45.0	42.2	35.5	33.3	24.0	24.5	36.7
NRCH	43.6	42.3	39.5	37.7	33.6	32.0	31.6	31.2	36.4
DRCL	47.5	41.6	36.5	32.3	23.4	20.4	16.1	14.7	29.1
RSHNL	49.1	45.4	44.3	41.6	38.3	36.4	27.8	26.8	38.7
NIRNL	51.6	46.6	51.7	46.5	49.2	46.1	41.7	39.4	46.6

Table 1: The MAP scores with different noise rates.

(Simonyan 2014) model, while the text modality is denoted as a 300-dimensional vector obtained by the pre-trained Doc2vec (Lau and Baldwin 2016) model. Following previous works (Pu et al. 2025a), we randomly partition it into three subsets: 2,173, 231, and 462 pairs for training, validation, and testing, respectively.

- **XMedia** is a benchmark dataset with 5 modalities belonging to 20 semantic labels. In this work, we just use 5,000 image-text pairs to conduct experiments. Specifically, the image modality is encoded as 4,096-dimensional vectors using a pre-trained VGG-19 model, while the text modality is represented by a 3,000-dimensional Bag-of-Words (BOW) embedding. Following previous works (Pu et al. 2025a), we randomly split it into three subsets: 4,000, 500, and 500 pairs for training, validation, and testing, respectively.
- **INRIA-Websearch** is a widely-used dataset that comprises over 70,000 image-text pairs from 353 semantic categories. In this work, we use a subset including 14,698 pairs within semantic categories to perform experiments. Specifically, we employ the 4096-dimensional features extracted by the pre-trained AlexNet (Krizhevsky, Sutskever, and Hinton 2017) as image input, while utilizing the 1000-dimensional features mined by the LDA (Hu et al. 2021) as text inputs. Like previous works (Pu et al. 2025a), we randomly divide it into three subsets: 9,000, 1,332, and 4,366 pairs for training, validation, and testing, respectively.

Baselines

To demonstrate the effectiveness and superiority of the proposed NIRNL, we compare it with 10 state-of-the-art (SOTA) cross-modal retrieval (CMR) methods, including three unsupervised methods (i.e., DGCPN (Yu et al. 2021), CIRH (Zhu et al. 2022), and UCCH (Hu et al. 2022)), four robust methods against noisy labels (i.e., RONO (Feng et al. 2023), DHRL (Shu et al. 2024), NRCH (Wang et al. 2024), and RSHNL (Pu et al. 2025c)), and three supervised methods (i.e., ALGCN (Qian et al. 2021), GNN4CMR (Qian et al. 2022), and DRCL (Pu et al. 2025a)). For the sake

of rigor, all comparisons are conducted using results reproduced under the same dataset settings.

Experimental settings and Evaluation Metric

To evaluate the performance of the proposed NIRNL and SOTA methods, we perform two common CMR tasks, i.e., using image modality samples as queries to retrieve semantically similar samples from the text modality (I2T), and using text modality samples as queries to retrieve semantically similar samples from the image modality (T2I). Following previous work (Pu et al. 2025c), we assess the robustness of each method under varying levels of symmetric label noise, with noise rates set to 0.2, 0.4, 0.6, and 0.8, respectively. To intuitively quantify retrieval performance, we adopt the Mean Average Precision (MAP) score (Qian et al. 2021; Pu et al. 2025a) as the evaluation metric. For fairness, we keep the original backbone networks frozen during training and report the MAP scores on the test set at the epoch where MAP reaches its peak on the validation set.

Comparisons with State-of-the-Art Methods

To verify the superiority of NIRNL, we compare it with 10 SOTA methods in terms of MAP scores. Tab.1 and Tab.2 report the MAP scores on the three datasets under varying noise rates. In the tables, the second-best scores are underlined, and the best scores are **highlighted in bold**. Then, we provide the precision-recall curves under the 0.6 noise rate in Fig.2 to further assess the retrieval performance of NIRNL. According to the detailed analysis of these results, we can draw the following conclusion:

- Although unsupervised methods (i.e., DGCPN, CIRH, and UCCH) are inherently unaffected by label noise due to the absence of supervision signals, their performance is limited by a clear bottleneck, as evidenced by their inferior results compared to some supervised methods (such as RONO and NIRNL).
- Some supervised methods (i.e., ALGCN, GNN4CMR, and DRCL) heavily rely on accurate annotations. As the noise rate increases, they suffer significant performance degradation or even fail due to their inability to handle noisy labels. In contrast, some robust supervised methods (i.e., RONO, NRCH, RSHNL, and NIRNL) demonstrate greater robustness by incorporating mechanisms specifically designed to mitigate the impact of label noise.
- From the PR-curves, it shows that the curve of NIRNL lies above that of other methods, demonstrating NIRNL can achieve higher retrieval precision at the same recall rates. In general, by refining training instances to capture the global neighborhood structure and exploit all potential information, our proposed NIRNL demonstrates superior and more robust performance.

Ablation Study

To investigate the contribution of specific components in our method, we conduct ablation experiments on the three datasets at a 0.6 noisy rate. Specifically, we construct four variants and the corresponding results are reported in Tab.3,

Method	Ref.	XMedia					INRIA-Websearch												
		0.2		0.4		0.6		0.8		Mean	0.2		0.4		0.6		0.8		Mean
		I2T	T2I	I2T	T2I	I2T	T2I	I2T	T2I		I2T	T2I	I2T	T2I	I2T	T2I	I2T	T2I	
DGCPN	AAAI'21	48.4	33.6	48.4	33.6	48.4	33.6	48.4	33.6	41.0	4.6	4.4	4.6	4.4	4.6	4.4	4.6	4.4	4.5
CIRH	TKDE'22	83.9	82.0	83.9	82.0	83.9	82.0	83.9	82.0	83.0	30.9	31.2	30.9	31.2	30.9	31.2	30.9	31.2	31.1
UCCH	TPAMI'23	78.6	79.5	78.6	79.5	78.6	79.5	78.6	79.5	79.1	7.4	6.2	7.4	6.2	7.4	6.2	7.4	6.2	6.8
ALGCN	TMM'21	85.6	84.6	74.7	72.3	44.3	40.2	14.1	12.1	53.5	29.6	28.1	17.4	16.0	6.7	6.2	2.4	2.1	13.6
RONO	CVPR'23	90.7	91.7	90.1	90.6	88.8	89.4	87.6	87.4	89.5	44.4	44.4	41.9	42.1	38.2	37.4	34.5	32.6	39.4
GNN4CMR	TPAMI'23	85.4	85.1	80.6	80.5	69.5	68.6	43.9	43.2	69.6	47.3	47.1	42.0	40.9	35.8	34.4	29.9	28.3	38.2
DHRL	TBD'24	90.6	91.1	86.6	86.5	73.5	73.3	43.8	45.5	73.9	37.1	37.3	24.3	24.1	14.0	13.7	6.2	7.1	20.5
NRCH	ACM MM'24	88.4	89.8	89.7	90.6	85.5	86.1	84.8	85.9	87.6	42.3	43.6	42.1	43.4	41.1	42.6	40.6	41.9	42.2
DRCL	TMM'25	83.2	81.5	64.0	64.2	35.8	37.1	12.2	13.4	48.9	47.2	46.5	34.7	32.4	17.4	16.2	5.5	5.1	25.6
RSHNL	AAAI'25	91.2	91.2	90.2	90.0	87.8	87.1	86.2	84.9	88.6	52.7	53.4	52.0	52.3	49.8	50.4	42.7	43.0	49.5
NIRNL	ours	92.5	92.0	92.1	91.6	91.6	92.0	91.3	91.2	91.8	52.9	53.3	52.2	52.8	51.6	52.5	49.9	50.8	52.0

Table 2: The MAP scores with different noise rates on the XMedia and INRIA-Websearch datasets.

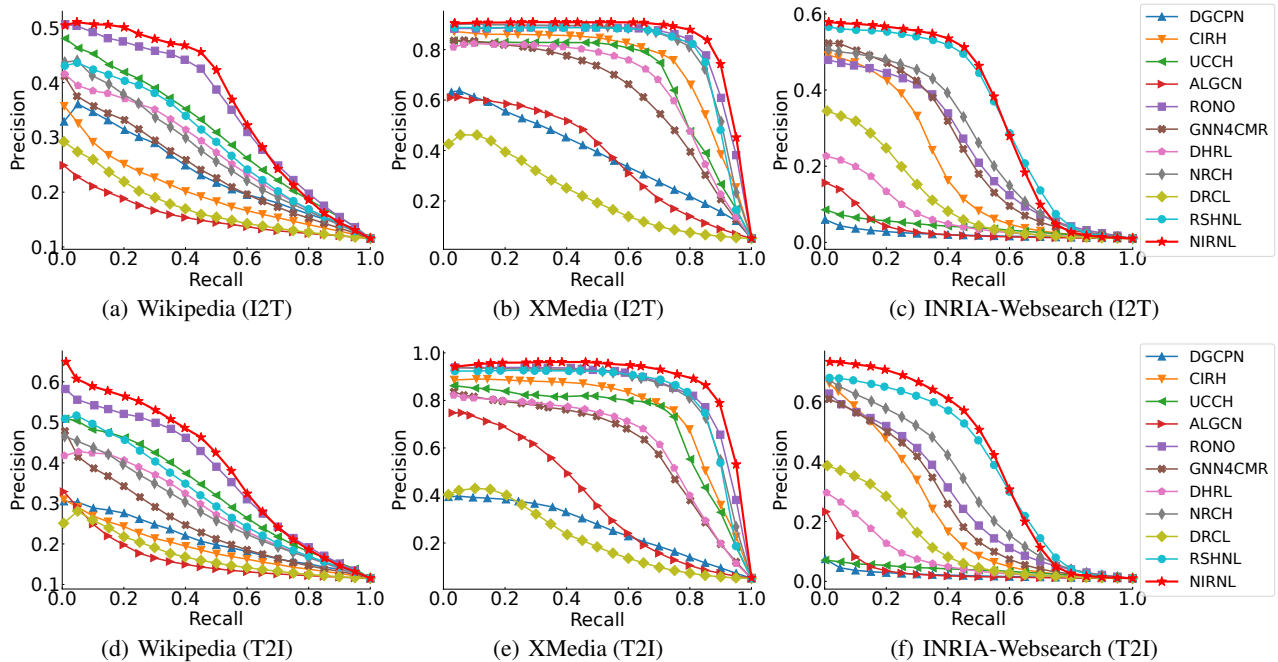


Figure 2: Precision-recall curves under the 0.6 noise rate.

where **NIRNL-1** denotes the removal of CMP, **NIRNL-2** represents the discarding of the noisy subset, **NIRNL-3** means treating the hard subset the same as the pure subset, without applying instance loss-weighting, and **NIRNL-4** refers to the exclusion of the barycenters alignment mechanism for all samples and only using CMP to optimize the model. From the results, it could be concluded that:

- The removal of CMP weakens the discriminability among inter-class samples, thereby leading to a drop in retrieval performance.
- The noisy subset contains potentially useful information, which can be leveraged through label correction to improve overall performance. In contrast, directly discarding these samples may lead to suboptimal results due to the loss of informative signals.

- When the loss-weighting mechanism designed for the hard subset is deleted, the model treats all samples as equally reliable, which amplifies the effect of noisy labels and results in performance degradation.
- Excluding the barycenter alignment mechanism is equivalent to discarding label information as a supervisory signal, leading to a noticeable decline in performance.

Overall, the ablation studies highlight the complementary roles of each component within the NIRNL framework.

Robustness Analysis

To intuitively demonstrate the robustness of the proposed NIRNL, we compare it with RSHNL in terms of MAP scores versus epochs under the 0.6 noise rate. As illustrated




Figure 3: The MAP scores versus epochs under the 0.6 noise rate.




Figure 4: The number of instances versus epochs under the 0.6 noise rate.

Dataset Task	Wikipedia			XMedia			Websearch		
	I2T	T2I	mean	I2T	T2I	mean	I2T	T2I	mean
NIRNL-1	25.9	22.9	24.4	40.0	41.6	40.8	7.9	8.7	8.3
NIRNL-2	47.1	42.4	44.8	88.8	88.4	88.6	46.3	47.1	46.7
NIRNL-3	48.4	45.8	47.1	90.0	90.6	90.3	50.5	51.6	51.1
NIRNL-4	41.2	39.8	40.5	90.1	91.4	90.8	49.9	50.9	50.4
NIRNL	49.2	46.1	47.7	91.6	92.0	91.8	51.6	52.5	52.1

Table 3: Ablation study under the 0.6 noise rate, where ‘Websearch’ denotes the INRIA-Websearch dataset.

in Fig.3, our NIRNL consistently outperforms RSHNL across all three datasets. Specifically, RSHNL exhibits lower performance on the XMedia and INRIA-Websearch datasets due to its failure to leverage noisy samples. On the Wikipedia dataset, RSHNL even shows signs of overfitting to noisy labels, as it fails to capture the global distributional structure of neighboring instances. In contrast, NIRNL employs the Neighbor-aware Instance Refining (NIR) module to adaptively refine all instances by capturing global distributional information, thereby maximizing the use of available sample information and achieving superior performance.

Additionally, to further evaluate the robustness of NIRNL, we track the number of three types of instances across different training epochs under a 0.6 noise rate, including the total number of clean instances in the training set and the number of truly clean instances within this selected pure subset. As shown in Fig.4, NIRNL can incorporate the majority

of clean samples into the training process over time. This may be attributed to the Neighbor-aware Instance Refining (NIR) mechanism, which enables the model to perceive the global distribution of instances and refine instances to maximize the use of all available information. These findings strongly support the effectiveness and robustness of the proposed NIRNL.

Conclusion

In this paper, to robustly learn with noisy labels, we propose a novel unified cross-modal learning framework, termed Neighbor-aware Instance Refining with Noisy Label (NIRNL). Our NIRNL inherits the advantages of the three existing paradigms to balance the upper limit of model performance, calibration reliability, and data utilization rate. Specifically, we first construct Cross-modal Margin Preserving (CMP) to structurally regularize the embedding space to boost representation discriminability. Then, we propose Neighbor-aware Instance Refining (NIR) to fully explore the semantic information of each instance. To be specific, NIR employs neighborhood consensus to tri-split data, thereby obtaining pure, hard, and noisy instances, respectively. Moreover, NIR designs three different loss functions to process these instances respectively, thereby achieving more robust learning with noisy labels. We conduct various experiments on three benchmarks with different noise rates. Experimental results show that our NIRNL obtains the best retrieval performance and robustness, especially under high noise rates.

Acknowledgments

This work is supported by the Central Government's Guide to Local Science and Technology Development Fund under Grant 2025ZYDF101.

References

- Agueh, M.; and Carlier, G. 2011. Barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2): 904–924.
- Chun, S.; Oh, S. J.; De Rezende, R. S.; Kalantidis, Y.; and Larlus, D. 2021. Probabilistic embeddings for cross-modal retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8415–8424.
- Dempster, A. P.; Laird, N. M.; and Rubin, D. B. 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the royal statistical society: series B (methodological)*, 39(1): 1–22.
- Fang, B.; Wu, W.; Liu, C.; Zhou, Y.; Song, Y.; Wang, W.; Shu, X.; Ji, X.; and Wang, J. 2023. Uatvr: Uncertainty-adaptive text-video retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13723–13733.
- Feng, C.; Ren, Y.; and Xie, X. 2023. OT-Filter: An optimal transport filter for learning with noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16164–16174.
- Feng, Y.; Li, Y.; Sun, Y.; Qin, Y.; Peng, D.; and Hu, P. 2025a. Interactive Cross-modal Learning for Text-3D Scene Retrieval. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- 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.
- Feng, Y.; Zhu, H.; Peng, D.; Peng, X.; Song, X.; and Hu, P. 2025b. Robust Cross-modal Alignment Learning for Cross-Scene Spatial Reasoning and Grounding. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Ge, X.; Chen, F.; Xu, S.; Tao, F.; and Jose, J. M. 2023. Cross-modal semantic enhanced interaction for image-sentence retrieval. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1022–1031.
- Ghosh, A.; Kumar, H.; and Sastry, P. S. 2017. Robust loss functions under label noise for deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.
- Gorti, S. K.; Vouitsis, N.; Ma, J.; Golestan, K.; Volkovs, M.; Garg, A.; and Yu, G. 2022. X-pool: Cross-modal language-video attention for text-video retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5006–5015.
- Han, H.; Luo, M.; Liu, H.; Nan, F.; and Liu, J. 2025. A Unified Optimal Transport Framework for Cross-Modal Retrieval With Noisy Labels. *IEEE Transactions on Neural Networks and Learning Systems*.
- Hu, P.; Peng, X.; Zhu, H.; Zhen, L.; and Lin, J. 2021. Learning cross-modal retrieval with noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5403–5413.
- Hu, P.; Zhu, H.; Lin, J.; Peng, D.; Zhao, Y.-P.; and Peng, X. 2022. Unsupervised contrastive cross-modal hashing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3): 3877–3889.
- Krapac, J.; Allan, M.; Verbeek, J.; and Juried, F. 2010. Improving web image search results using query-relative classifiers. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 1094–1101. IEEE.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2017. ImageNet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6): 84–90.
- Lau, J. H.; and Baldwin, T. 2016. An empirical evaluation of doc2vec with practical insights into document embedding generation. *arXiv preprint arXiv:1607.05368*.
- Li, S.; Liu, T.; Tan, J.; Zeng, D.; and Ge, S. 2021. Trustable co-label learning from multiple noisy annotators. *IEEE Transactions on Multimedia*, 25: 1045–1057.
- Li, Y.; Feng, Y.; Sun, Y.; Peng, D.; Peng, X.; and Hu, P. 2025a. Learning Source-Free Domain Adaptation for Visible-Infrared Person Re-Identification. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Li, Y.; Qin, Y.; Sun, Y.; Peng, D.; Peng, X.; and Hu, P. 2024. RoMo: Robust unsupervised multimodal learning with noisy pseudo labels. *IEEE Transactions on Image Processing*.
- Li, Y.; Sun, Y.; Qin, Y.; Peng, D.; Peng, X.; and Hu, P. 2025b. Robust Duality Learning for Unsupervised Visible-Infrared Person Re-Identification. *IEEE Transactions on Information Forensics and Security*.
- Liang, K.; Liu, Y.; Zhou, S.; Tu, W.; Wen, Y.; Yang, X.; Dong, X.; and Liu, X. 2023. Knowledge graph contrastive learning based on relation-symmetrical structure. *IEEE Transactions on Knowledge and Data Engineering*, 36(1): 226–238.
- Lin, H.; Shen, S.; Zhang, Y.; and Xia, R. 2025. Unsupervised Contrastive Graph Kolmogorov–Arnold Networks Enhanced Cross-Modal Retrieval Hashing. *Mathematics*, 13(11): 1880.
- Luo, X.; Chen, P.; Liu, C.; Jin, X.; Wen, J.; Liu, Y.; and Wang, J. 2025. Enhancing Multimodal Protein Function Prediction Through Dual-Branch Dynamic Selection with Reconstructive Pre-Training. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, 7598–7606.
- Mandal, D.; Rao, P.; and Biswas, S. 2019. Semi-supervised cross-modal retrieval with label prediction. *IEEE Transactions on Multimedia*, 22(9): 2345–2353.

- Okamura, D.; Harakawa, R.; and Iwahashi, M. 2023. LC-NME: Label correction using network prediction based on memorization effects for cross-modal retrieval with noisy labels. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(1): 590–602.
- Peng, Y.; Zhai, X.; Zhao, Y.; and Huang, X. 2015. Semi-supervised cross-media feature learning with unified patch graph regularization. *IEEE Transactions on Circuits and Systems for Video Technology*, 26(3): 583–596.
- Pu, R.; Peng, D.; and Hua, F. 2024. Deep noisy multi-label learning for robust cross-modal retrieval. In *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*, 304–317. Springer.
- Pu, R.; Qin, Y.; Peng, D.; Song, X.; and Zheng, H. 2025a. Deep Reversible Consistency Learning for Cross-modal Retrieval. *IEEE Transactions on Multimedia*.
- Pu, R.; Qin, Y.; Song, X.; Peng, D.; Ren, Z.; and Sun, Y. 2025b. SHE: Streaming-media Hashing Retrieval. In *Forty-second International Conference on Machine Learning*, volume 267, 49946–49958.
- Pu, R.; Sun, Y.; Qin, Y.; Ren, Z.; Song, X.; Zheng, H.; and Peng, D. 2025c. Robust Self-Paced Hashing for Cross-Modal Retrieval with Noisy Labels. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(19): 19969–19977.
- Qian, S.; Xue, D.; Fang, Q.; and Xu, C. 2021. Adaptive label-aware graph convolutional networks for cross-modal retrieval. *IEEE Transactions on Multimedia*, 24: 3520–3532.
- Qian, S.; Xue, D.; Fang, Q.; and Xu, C. 2022. Integrating multi-label contrastive learning with dual adversarial graph neural networks for cross-modal retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4): 4794–4811.
- 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.
- Rasiwasia, N.; Costa Pereira, J.; Coviello, E.; Doyle, G.; Lanckriet, G. R.; Levy, R.; and Vasconcelos, N. 2010. A new approach to cross-modal multimedia retrieval. In *Proceedings of the 18th ACM International Conference on Multimedia*, 251–260.
- Shu, Z.; Bai, Y.; Yong, K.; and Yu, Z. 2024. Deep Cross-Modal Hashing With Ranking Learning for Noisy Labels. *IEEE Transactions on Big Data*.
- Simonyan, K. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Su, C.; Zheng, H.; Peng, D.; and Wang, X. 2025. DiCA: Disambiguated Contrastive Alignment for Cross-Modal Retrieval with Partial Labels. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 20610–20618.
- Su, S.; Zhong, Z.; and Zhang, C. 2019. Deep joint-semantics reconstructing hashing for large-scale unsupervised cross-modal retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3027–3035.
- Sun, Y.; Ren, Z.; Hu, P.; Peng, D.; and Wang, X. 2023. Hierarchical consensus hashing for cross-modal retrieval. *IEEE Transactions on Multimedia*, 26: 824–836.
- Wang, B.; Yang, Y.; Xu, X.; Hanjalic, A.; and Shen, H. T. 2017. Adversarial cross-modal retrieval. In *Proceedings of the 25th ACM International Conference on Multimedia*, 154–162.
- Wang, L.; Qin, Y.; Sun, Y.; Peng, D.; Peng, X.; and Hu, P. 2024. Robust Contrastive Cross-modal Hashing with Noisy Labels. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 5752–5760.
- Wehrmann, J.; Kolling, C.; and Barros, R. C. 2020. Adaptive cross-modal embeddings for image-text alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 12313–12320.
- Wu, G.; Lin, Z.; Han, J.; Liu, L.; Ding, G.; Zhang, B.; and Shen, J. 2018. Unsupervised Deep Hashing via Binary Latent Factor Models for Large-scale Cross-modal Retrieval. In *IJCAI*, volume 1, 5.
- Yang, E.; Yao, D.; Liu, T.; and Deng, C. 2022. Mutual quantization for cross-modal search with noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7551–7560.
- Yin, Z.; Feng, Y.; Yan, M.; Song, X.; Peng, D.; and Wang, X. 2025. RoDA: Robust Domain Alignment for Cross-Domain Retrieval Against Label Noise. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 9535–9543.
- Yu, J.; Zhou, H.; Zhan, Y.; and Tao, D. 2021. Deep graph-neighbor coherence preserving network for unsupervised cross-modal hashing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 4626–4634.
- Zhang, J.; Peng, Y.; and Yuan, M. 2018. Unsupervised generative adversarial cross-modal hashing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Zhen, L.; Hu, P.; Wang, X.; and Peng, D. 2019. Deep supervised cross-modal retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10394–10403.
- Zhu, L.; Wu, X.; Li, J.; Zhang, Z.; Guan, W.; and Shen, H. T. 2022. Work together: Correlation-identity reconstruction hashing for unsupervised cross-modal retrieval. *IEEE Transactions on Knowledge and Data Engineering*, 35(9): 8838–8851.