

Distribution-Consistency-Guided Multi-modal Hashing

Jin-Yu Liu, Xian-Ling Mao*, Tian-Yi Che, Rong-Cheng Tu

School of Computer Science and Technology, Beijing Institute of Technology, Beijing, China
 {liujinyu1229, turongcheng}@gmail.com, {maoxl, ccty}@bit.edu.cn

Abstract

Multi-modal hashing methods have gained popularity due to their fast speed and low storage requirements. Among them, the supervised methods demonstrate better performance by utilizing labels as supervisory signals compared with unsupervised methods. Currently, for almost all supervised multi-modal hashing methods, there is a hidden assumption that training sets have no noisy labels. However, labels are often annotated incorrectly due to manual labeling in real-world scenarios, which will greatly harm the retrieval performance. To address this issue, we first discover a significant distribution consistency pattern through experiments, i.e., the 1-0 distribution of the presence or absence of each category in the label is consistent with the high-low distribution of similarity scores of the hash codes relative to category centers. Then, inspired by this pattern, we propose a novel **Distribution-Consistency-Guided Multi-modal Hashing (DCGMH)**, which aims to filter and reconstruct noisy labels to enhance retrieval performance. Specifically, the proposed method first randomly initializes several category centers, each representing the region's centroid of its respective category, which are used to compute the high-low distribution of similarity scores; Noisy and clean labels are then separately filtered out via the discovered distribution consistency pattern to mitigate the impact of noisy labels; Subsequently, a correction strategy, which is indirectly designed via the distribution consistency pattern, is applied to the filtered noisy labels, correcting high-confidence ones while treating low-confidence ones as unlabeled for unsupervised learning, thereby further enhancing the model's performance. Extensive experiments on three widely used datasets demonstrate the superiority of the proposed method compared to state-of-the-art baselines in multi-modal retrieval tasks.

Code — <https://github.com/LiuJinyu1229/DCGMH>

Introduction

With the rapid growth of multimedia data such as images, text, and videos, achieving effective retrieval from massive multi-modal data has become a significant challenge. To address this challenge, numerous information retrieval technologies have emerged (Jiang and Li 2017; Sung et al. 2018;

Zhang, Lai, and Feng 2018; Chen et al. 2020b), with multi-modal hashing methods (Liu, He, and Lang 2014; Chen et al. 2020a; Zhu et al. 2021; Lu et al. 2021; Wu et al. 2022) gaining widespread attention for their fast retrieval speed and low storage requirements. Unlike uni-modal hashing (Tu et al. 2018, 2021c; Guo et al. 2022; Tu, Mao, and Wei 2020; Tu et al. 2021b) and cross-modal hashing (Zhang, Peng, and Yuan 2018; Tu et al. 2022, 2023a, 2021a), multi-modal hashing maps data points from different modalities into a unified Hamming space for fusion, resulting in binary hash codes that facilitate efficient multi-multi retrieval. Compared to unsupervised methods (Song et al. 2013; Shen et al. 2015, 2018; Zheng et al. 2020b; Wu et al. 2021a), supervised multi-modal hashing methods (Yang, Shi, and Xu 2017; Xie et al. 2017; Yu et al. 2022; Zheng et al. 2024) generate more discriminative hash codes and achieve more accurate retrieval by utilizing labels as supervisory signals and there is a hidden assumption that training sets have no noisy labels for these supervised methods.

However, in real-world scenarios, labels may be incorrectly annotated due to manual labeling, such as an image that should be labeled as "tiger" being mistakenly labeled as "cat", which limits the applicability of existing supervised hashing methods in noisy label scenarios. While some works (Sun et al. 2022; Yang et al. 2022) in image hashing and cross-modal hashing have demonstrated that the presence of noisy labels in the training set can lead to model overfitting, resulting in indistinguishable hash codes and inaccurate retrieval, no work has yet focused on and resolved this issue in the field of multi-modal hashing.

To effectively tackle this issue, it is crucial to filter out noisy labels from the dataset. Existing single or cross-modal hashing methods have demonstrated that models initially learn effective hash mappings from clean labels but eventually overfit to noisy labels, resulting in degraded performance. This phenomenon suggests that even after a short training period, the generated hash codes are somewhat discriminative and tend to align with their corresponding category centers (Yuan et al. 2020; Tu et al. 2023b), as indicated by higher similarity scores. Thus, the 1-0 distribution of the category's presence or absence, i.e., the label vector, should be consistent with the high-low distribution of similarity scores between the hash codes and category centers. Specifically, if an instance belongs to a category, the corre-

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

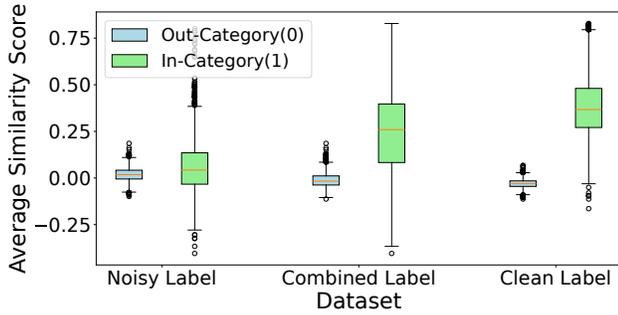


Figure 1: BoxPlot comparison of average similarity scores for in-category and out-category across clean and noisy label datasets, where "Out-Category(0)" represents the box plot distribution of the average similarity scores of hash codes to all categories it does not belong, while "In-Category(1)" represents the box plot distribution for the belonging categories, and the horizontal line within each box indicates the median of all average similarity scores.

sponding label bit should be '1', and the hash code should show a higher similarity score with the category center; otherwise, the bit should be '0', and the similarity score should be lower. Therefore, we hypothesize that this pattern of distribution consistency can effectively filter out noisy labels.

To validate this hypothesis, we conduct a BoxPlot statistical analysis comparing the average similarity scores between hash codes and their respective in-category and out-category centers. This analysis is performed on the MIR Flickr dataset which contains a training set of 5,000 instances with a noisy label ratio of 40%. After training the model for 10 epochs, we separate the dataset into noisy and clean subsets and analyze the distribution of similarity scores. The results, shown in Figure 1, reveal that in the clean label dataset, the similarity scores of hash codes to their respective category centers are significantly higher compared to non-belonging centers. In contrast, in the noisy label dataset, this difference is less pronounced due to the misalignment between assigned categories and actual hash code semantics. This observation confirms that there is a notable consistency between the 1-0 label distribution and the high-low similarity score distribution in clean labels, which is disrupted in noisy labels. Therefore, by exploiting these consistency differences, we can effectively filter out noisy labels, supporting our hypothesis.

Consequently, inspired by the distribution consistency pattern and the hypothesis, we propose a novel **Distribution-Consistency-Guided Multi-modal Hashing (DCGMH)** via the consistency between the 1-0 distribution of labels and the high-low distribution of similarity scores to filter and reuse noisy labels, thereby preventing the model from overfitting to noisy labels and improving its retrieval performance in real-world scenarios. Specifically, the proposed method first randomly initializes several category centers to compute the similarity scores of the hash codes relative to each category center, with each center representing the central region of its corresponding category; Then, based on the distribu-

tion consistency pattern, the noisy and clean labels are separately filtered out to mitigate the impact of noisy labels; Subsequently, for the noisy label set, we design a reconstruction strategy via the distribution consistency pattern to correct high-confidence noisy labels, while low-confidence noisy label instances are treated as unlabeled instances to facilitate unsupervised learning for extracting semantic information and further enhancing the performance of the multi-modal hashing model. In conclusion, the main contributions of the proposed method are as follows:

- We discover the consistency between the 1-0 distribution of the presence or absence of each category in the label and the high-low distribution of similarity scores of the hash codes relative to each category center and validate it through BoxPlot statistical analysis.
- We design a filter via the distribution consistency pattern to filter out noisy labels and improve the applicability of supervised multi-modal hashing methods in real-world scenarios. To the best of our knowledge, no similar work has been done.
- We design a corrector via the distribution consistency pattern to correct high-confidence noisy labels and utilize low-confidence ones for unsupervised learning.
- Extensive experiments on three benchmark datasets demonstrate the proposed method outperforms the state-of-the-art baselines in multi-modal retrieval tasks.

Related Work

Multi-modal Hashing

In multi-modal hashing, supervised methods utilize labels as supervisory information to help models learn richer semantic information, and can be categorized into shallow and deep approaches based on whether deep networks are utilized. Shallow supervised multi-modal hashing (Lu et al. 2019a,c; Zheng et al. 2019, 2022; An et al. 2022) often relies on linear mapping or matrix factorization to model the latent semantic associations between modalities. For instance, OMHDQ (Lu et al. 2019b) links hash code learning with both low-level data distribution and high-level semantic distribution based on paired semantic labels. SAPMH (Zheng et al. 2020a) employs paired semantic labels as parameter-free supervisory information to learn multidimensional latent representations. In contrast, deep supervised multi-modal hashing (Yan et al. 2020; Zhu et al. 2020; Shen et al. 2023; Lu et al. 2020; Tan et al. 2023) leverages deep networks to integrate feature extraction and hash code learning into a unified deep framework. For example, BSTH (Tan et al. 2022) introduces a bit-aware semantic transformation module to achieve fine-grained, concept-level alignment and fusion of multi-modal data, thereby generating high-quality hash codes. STBMH (Tu et al. 2024) addresses the issue of similarity transitivity broken in multi-label scenarios by designing an additional regularization term.

For almost all supervised multi-modal hashing methods, there is a hidden assumption that training sets have no noisy labels. However, in real-world scenarios, the labels are often annotated incorrectly due to manual labeling which greatly

harms the model’s performance. This issue necessitates effective strategies for noisy label learning to maintain model robustness and accuracy.

Noisy Label Learning

Noisy label learning has been extensively studied in tasks such as image classification. Existing approaches to handling noisy labels can be broadly categorized into two types: noise-robust modeling (Song et al. 2020, 2022; Shu et al. 2019) and label-noise cleaning (Zheng et al. 2020c; Wu et al. 2021b; Kim et al. 2021; Wei et al. 2022). Noise-robust modeling involves directly training robust models on noisy labels using noise-specific loss functions or regularization terms; for instance, NCR (Iscen et al. 2022) designs a regularization loss term based on the consistency between an instance and its neighboring nodes in the feature space. In contrast, label-noise cleaning focuses on filtering or correcting noisy labels directly, exemplified by FCF (Jiang et al. 2024), which introduces a fusion cleaning framework that combines correction and filtering to address different types of noisy labels.

In the context of hashing, there have been works addressing noisy labels in image hashing (Sun et al. 2022; Wang et al. 2023) and cross-modal hashing (Yang et al. 2022; Li et al. 2024) by using loss-based methods. For instance, DIOR (Wang et al. 2023) uses the concept of behavior similarity between original and augmented views to filter noisy labels, while CMMQ (Yang et al. 2022) designs a proxy-based contrastive loss to mitigate the impact of noisy labels. However, no existing work addresses the issue of noisy labels in multi-modal hashing retrieval. Based on the discovered distribution consistency pattern, we propose a novel approach called DCGMH to filter and reuse noisy labels, improving the robustness of the hashing model. The specific details of this approach will be discussed in the next section.

The Proposed Method

In the section, we first describe the problem definition and then provide detailed explanations of our proposed method’s architecture. Subsequently, we summarize the objective function and optimize the model’s training process. Finally, we demonstrate the out-of-sample extension.

Problem Definition

Similar to most existing multi-modal hashing methods, this work focuses on image-text datasets. Assuming that there is a dataset \mathcal{O} containing n instances, denoted as $\mathcal{O} = \{\mathbf{o}_i\}_{i=1}^n = \{\mathbf{x}_i, \mathbf{y}_i, \mathbf{l}_i\}_{i=1}^n$, where \mathbf{x}_i and \mathbf{y}_i represents the text and image modal data point, respectively. Moreover, $\mathbf{l}_i \in \{0, 1\}^m$ represents the label vector of the instance \mathbf{o}_i , where m is the total number of categories, and when an instance \mathbf{o}_i belongs to the category j , $l_{ij} = 1$; otherwise, $l_{ij} = 0$. Furthermore, the similarity matrix $\mathbf{S} \in \{-1, 1\}^{n \times n}$ is employed to represent the similarity between instances, such that when instances \mathbf{o}_i and \mathbf{o}_j share at least one category, $s_{ij} = 1$, indicating they are similar; otherwise, $s_{ij} = -1$, indicating they are dissimilar. Additionally, the instances in dataset \mathcal{O} are ultimately mapped to hash codes $\mathbf{B} \in \{-1, 1\}^{n \times k}$ in the Hamming space while preserving

the original semantics, where k represents the length of the hash codes.

Architecture

Multi-modal Hashing Network To generate high-quality hash codes, we use instance $\mathbf{o}_i = \{\mathbf{x}_i, \mathbf{y}_i, \mathbf{l}_i\}$ as the training data for the hashing network. Then, we employ a bag-of-words (BoW) model to extract the feature representation \mathbf{f}_i^x and a VGG model without the final classification layer to extract the feature representation \mathbf{f}_i^y for the text and image modal data, respectively. Subsequently, modality-specific multi-layer perceptions (MLP) are utilized to map the feature representation \mathbf{f}_i^* , $*$ $\in \{x, y\}$ of each modality to a unified space, denoted as:

$$\mathbf{u}_i^* = MLP^*(\mathbf{f}_i^*; \theta^*) \quad (1)$$

where \mathbf{u}_i^* is the resulting projected feature representation and θ^* denotes the learnable parameters of the MLP for the respective modality. Finally, the projected features from different modalities are fused by directly summing them, and this fused representation is then passed through a hashing function with a non-linear activation function ($\tanh(\cdot)$) to generate the fused hash code $\hat{\mathbf{b}}_i$, denoted as:

$$\hat{\mathbf{b}}_i = \mathcal{H}(\mathbf{u}_i^x + \mathbf{u}_i^y; \theta) \quad (2)$$

where \mathcal{H} refers to the hash function and θ is a set of learnable parameters. Finally, the final binary hash code \mathbf{b}_i for instance \mathbf{o}_i can be obtained as $\mathbf{b}_i = \text{sgn}(\hat{\mathbf{b}}_i)$, where $\text{sgn}(\cdot)$ is a function that maps positive values to 1 and negative values to -1. In summary, for instance \mathbf{o}_i , its binary hash code can be formulated as $\mathbf{b}_i = \text{sgn}(\mathcal{F}(\mathbf{x}_i, \mathbf{y}_i; \mathbf{P}))$, where \mathbf{P} is a set of learnable parameters.

Label Filter To filter noisy labels via our discovered distribution consistency pattern, we first randomly initialize several category centers $\mathbf{C} \in \mathcal{R}^{m \times k}$, i.e., $\mathbf{C} = \{\mathbf{c}_j\}_{j=1}^m$, where m is the number of categories and k is the length of hash code, ensuring that each category occupies a distinct and non-overlapping region. Then for each instance’s hash code $\hat{\mathbf{b}}_i$, we calculate its similarity scores \mathbf{d}_{ij} with each category center \mathbf{c}_j , denoted as:

$$\mathbf{d}_{ij} = \frac{\hat{\mathbf{b}}_i}{\|\hat{\mathbf{b}}_i\|} \left(\frac{\mathbf{c}_j}{\|\mathbf{c}_j\|} \right)^T \quad (3)$$

By aggregating the similarity scores, we obtain the similarity scores matrix $\mathbf{D} \in \mathcal{R}^{n \times m}$ which captures the similarity of hash code to each category center. In this matrix, when a hash code $\hat{\mathbf{b}}_i$ belongs to a particular category \mathbf{c}_j , the corresponding similarity score \mathbf{d}_{ij} tends to be high; otherwise, it is relatively low. Next, based on the consistency pattern between the 1-0 distribution of labels and the high-low distribution of similarity scores, we calculate the consistency level $\mathbf{T} = \{\mathbf{t}_i\}_{i=1}^n$ between the label distribution \mathbf{l}_i and similarity distribution \mathbf{d}_i . We then design a consistency-based criterion to filter the dataset \mathcal{O} into the clean label set \mathcal{O}_c and noisy label set \mathcal{O}_n , which can be formulated as:

$$\mathbf{t}_i = \frac{\sum_{j=1}^m \mathbf{l}_{ij} \mathbf{d}_{ij}}{\sum_{j=1}^m \mathbf{l}_{ij}} \quad (4)$$

$$\mathcal{O}_c = \{(x_i, y_i, l_i) | t_i > \epsilon(\tau)\} \quad (5)$$

$$\mathcal{O}_n = \{(x_i, y_i, l_i) | t_i \leq \epsilon(\tau)\} \quad (6)$$

where t_i is employed to measure the degree of consistency between the label distribution and similarity distribution, with higher values indicating greater consistency, τ represents the noise ratio and $\epsilon(\tau)$ denotes the filtering threshold to ensure τn instances are identified as noisy label instances.

Label Reconstructor To accurately learn the semantic knowledge and associations between labels and hash codes in the subsequent steps, we reconstruct the labels in \mathcal{O}_c and \mathcal{O}_n . On the one hand, since the labels in \mathcal{O}_c are clean, we directly use the initial labels as the reconstructed labels and adopt a standard measure in multi-modal hashing to establish semantic associations between hash codes and labels to learn implicit knowledge. On the other hand, for instances in \mathcal{O}_n , recognizing the significant advantage of labels as guiding information, we design a corrector via distribution consistency pattern to correct high-confidence noisy labels. Specifically, we treat the clean label set \mathcal{O}_c as a knowledge base, and for a given instance $\mathbf{o}_i = \{x_i, y_i, l_i\}$ in the noisy label set \mathcal{O}_n , we identify the two instances $\mathbf{o}_j = \{x_j, y_j, l_j\}$ and $\mathbf{o}_k = \{x_k, y_k, l_k\}$ from \mathcal{O}_c whose distribution of similarity scores d_j and d_k most closely match that of \mathbf{o}_i . Here if the labels l_j and l_k of \mathbf{o}_j and \mathbf{o}_k are consistent, we infer that \mathbf{o}_i has a high confidence of sharing the same label; otherwise, we consider \mathbf{o}_i as having low confidence and treat it as unlabeled data. This corrector can be represented as follows:

$$\mathbf{m}_i = d_i(\mathbf{D}_c)^T \quad (7)$$

$$\{\mathbf{o}_j, \mathbf{o}_k\} = \arg \min_{\mathbf{o}_j, \mathbf{o}_k \in \mathcal{O}_c} (\text{rank}(\mathbf{m}_i, j) + \text{rank}(\mathbf{m}_i, k)) \quad (8)$$

$$l_i = \begin{cases} l_j, & l_j = l_k; \\ \text{None}, & \text{otherwise.} \end{cases} \quad (9)$$

where \mathbf{D}_c is the similarity scores matrix of clean label set \mathcal{O}_c , \mathbf{m}_i represents the level of consistency between the similarity distribution of the noisy label instance \mathbf{o}_i and each instance in the clean label set \mathcal{O}_c , and $\text{rank}(\cdot, \cdot)$ indicates the ranking of consistency levels. As for low-confidence noisy labels, we discard the original labels and treat them as unlabeled. Consequently, the noisy label set \mathcal{O}_n is further divided into a corrected label set \mathcal{O}_r and an unlabeled set \mathcal{O}_u .

Objective Function and Optimization

Pointwise Learning for Clean Label Set To ensure that the generated fused hash codes accurately reflect label semantics, inspired by STBMH (Tu et al. 2024), we design the following pointwise loss on the clean label set \mathcal{O}_c to make the hash codes as close as possible to their corresponding categories while keeping them distant from non-relevant categories, denoted as:

$$\mathcal{L}_o = -\frac{1}{n_c} \sum_i \sum_{j \in \mathcal{Y}_i} \log \frac{\exp(\frac{1}{k} \hat{\mathbf{b}}_i^T \mathbf{c}_j)}{\exp(\frac{1}{k} \hat{\mathbf{b}}_i^T \mathbf{c}_j) + \sum_{h \in \mathcal{N}_i} \exp(\frac{1}{k} \hat{\mathbf{b}}_i^T \mathbf{c}_h)}, \quad (10)$$

where k is the length of hash code, n_c is the number of instances in \mathcal{O}_c , \mathcal{Y}_i contains the indices of all categories to which instance \mathbf{o}_i belongs and \mathcal{N}_i contains the indices of categories to which it does not belong. By minimizing \mathcal{L}_o , the value of $\exp(\frac{1}{k} \hat{\mathbf{b}}_i^T \mathbf{c}_j)$ becomes significantly higher than that of $\sum_{h \in \mathcal{N}_i} \exp(\frac{1}{k} \hat{\mathbf{b}}_i^T \mathbf{c}_h)$, indicating that the similarity between the hash code and its corresponding category center is high, while the similarity with non-relevant category centers is low, thereby achieving the desired objective.

Pairwise Learning for Corrector Label Set Meanwhile, for the corrected label set \mathcal{O}_r , since the labels within still have the potential for error correction, we adopt a pairwise loss to emphasize the relative similarity relationships among instances as a whole, rather than direct associations between instances and specific individual categories, expressed as follows:

$$\mathcal{L}_a = \sum_i^{n_r} \sum_j^{n_r} \|\cos(\hat{\mathbf{b}}_i, \hat{\mathbf{b}}_j) - s_{ij}\|_F^2 \quad (11)$$

$$\cos(\hat{\mathbf{b}}_i, \hat{\mathbf{b}}_j) = \frac{\hat{\mathbf{b}}_i^T \hat{\mathbf{b}}_j}{\|\hat{\mathbf{b}}_i\| \|\hat{\mathbf{b}}_j\|} = \frac{1}{k} \hat{\mathbf{b}}_i^T \hat{\mathbf{b}}_j \quad (12)$$

where k is the length of hash code, n_r is the number of instances in \mathcal{O}_r , $\cos(\cdot, \cdot)$ is employed to measure the cosine similarity between hash codes and s_{ij} represents the pairwise similarity defined by the labels. By minimizing \mathcal{L}_a , the cosine similarity between hash codes of similar instances defined by the labels will approach 1 and their Hamming distance $d_H(\hat{\mathbf{b}}_i, \hat{\mathbf{b}}_j)$ will decrease, where $d_H(\hat{\mathbf{b}}_i, \hat{\mathbf{b}}_j) = \frac{1}{2}(k - \hat{\mathbf{b}}_i^T \hat{\mathbf{b}}_j)$.

Unsupervised Learning for Unlabeled Set For the unlabeled set \mathcal{O}_u , given that text and image data points inherently contain rich semantic knowledge, we employ unsupervised contrastive learning to uncover hidden semantic relationships and further enhance model performance. Specifically, for an instance \mathbf{o}_i in \mathcal{O}_u , we first generate an augmented instance \mathbf{o}'_i through data augmentation. Then due to \mathbf{o}_i and \mathbf{o}'_i embody the same semantic meaning, their corresponding hash codes should be as consistent as possible. Therefore, we design the following contrastive loss to minimize the distance between hash code $\hat{\mathbf{b}}_i$ and $\hat{\mathbf{b}}'_i$:

$$\mathcal{L}_u = \frac{1}{n_u} \sum_{i=1}^{n_u} (1 - s_{ii}) + \frac{1}{n_u^2} \sum_{i=1}^{n_u} \sum_{j=1, j \neq i}^{n_u} \max(0, s_{ij} - \epsilon) \quad (13)$$

where $s_{ij} = \cos(\hat{\mathbf{b}}_i, \hat{\mathbf{b}}'_j)$ i.e. cosine similarity of hash codes, n_u is the number of instances in \mathcal{O}_u and ϵ is the threshold that restricts the similarity between different instance pairs. By minimizing \mathcal{L}_u , the similarity between the same instance pairs will approach 1, while the similarity between different instance pairs will be less than the threshold ϵ , thus ensuring that more similar instances are mapped to closer hash codes and achieving the goal of extracting the inherent semantics of the instance through unsupervised learning.

Method	MIR Flickr				NUS-WIDE				MS COCO			
	16	32	64	128	16	32	64	128	16	32	64	128
FOMH	0.697	0.726	0.733	0.742	0.586	0.604	0.657	0.662	0.425	0.440	0.468	0.485
OMHDQ	0.762	0.771	0.777	0.783	0.672	0.689	0.704	0.717	0.522	0.538	0.546	0.573
SDMH	0.772	0.782	0.794	0.798	0.666	0.690	0.698	0.713	0.396	0.398	0.399	0.399
SAPMH	0.751	0.783	0.800	0.806	0.639	0.670	0.677	0.692	0.438	0.442	0.458	0.463
DCMVH	0.715	0.725	0.742	0.754	0.590	0.612	0.667	0.685	0.399	0.414	0.477	0.499
BSTH	0.705	0.715	0.736	0.757	0.559	0.591	0.615	0.640	0.503	0.524	0.538	0.561
NCH	0.751	0.759	0.769	0.778	0.618	0.634	0.641	0.664	0.517	0.533	0.551	0.569
GCIMH	0.738	0.746	0.751	0.756	0.600	0.604	0.618	0.623	0.488	0.499	0.516	0.523
STBMH	0.748	0.767	0.779	0.795	0.622	0.656	0.681	0.691	0.509	0.540	0.564	0.589
DIOR	0.750	0.781	0.820	0.835	0.636	0.692	0.737	0.754	0.482	0.509	0.560	0.596
Ours	0.796	0.823	0.846	0.850	0.717	0.739	0.755	0.757	0.551	0.591	0.631	0.649

Table 1: MAPs at a noisy label ratio of 40% for different hash code lengths across the three benchmark datasets.

Center Learning and Quantization In addition, each category should occupy a distinct region without interference. To achieve this, the centers of the categories must be well-separated, which lead us to design the following center loss:

$$d_{ij} = \|\mathbf{c}_i - \mathbf{c}_j\|_F^2 \quad (14)$$

$$\mathcal{L}_c = -\frac{1}{|\{d_{ij} | i < j\}|} \sum_{i < j} d_{ij} - \min_{i < j} d_{ij} \quad (15)$$

where d_{ij} represents the distance between two category centers and $i < j$ indicates that only the upper triangular region is considered. By minimizing \mathcal{L}_c , sufficient separation between different categories is achieved which enhances the model’s ability to distinguish them. Furthermore, as the model ultimately relies on the $\text{sgn}(\cdot)$ function to convert the fused hash codes into binary hash codes, we introduce the following loss to control the quantization loss incurred during this process:

$$\mathcal{L}_q = \sum_i^n \|\hat{\mathbf{b}}_i - \mathbf{b}_i\|_F^2 \quad (16)$$

Overall Objective Function Finally, by combining the losses from the above components, the overall objective function of the multi-modal hashing network is defined as:

$$\mathcal{L} = \mathcal{L}_o + \alpha \mathcal{L}_a + \beta \mathcal{L}_u + \gamma \mathcal{L}_c + \eta \mathcal{L}_q \quad (17)$$

where α , β , γ and η are hyperparameters. By minimizing this loss during the training of the whole multi-modal hashing network, the model can generate high-quality and distinguishable hash codes even in the presence of noisy labels, which significantly enhances its performance.

Out-of-Sample Extension

After optimizing the multi-modal hashing network by minimizing \mathcal{L} , we can use the model to generate hash codes for instances outside the training set. Specifically, for a given unseen instance $\mathbf{o}_i = \{\mathbf{x}_i, \mathbf{y}_i, \mathbf{l}_i\}$, we obtain the binary hash code by feeding the instance into the optimized multi-modal hashing network, denoted as:

$$\mathbf{b}_i = \text{sgn}(\mathcal{F}(\mathbf{x}_i, \mathbf{y}_i; \mathbf{P})) \quad (18)$$

Experiments

Experimental Setting

Baselines To evaluate the superiority of our proposed method, we compare it with ten state-of-the-art multi-modal hashing methods, i.e., FOMH (Lu et al. 2019a), OMHDQ (Lu et al. 2019b), SDMH (Lu et al. 2019c), SAPMH (Zheng et al. 2020a), DCMVH (Zhu et al. 2020), BSTH (Tan et al. 2022), NCH (Tan et al. 2023), GCIMH (Shen et al. 2023), STBMH (Tu et al. 2024) and DIOR (Wang et al. 2023), where the first four adopt shallow frameworks and the latter six employ deep networks. Notably, DIOR is originally developed for noisy label handling in image hashing; here, we adapt it for multi-modal hashing while retaining its core structure.

Implementation Details Similar to DIOR, we first perform warm-up training without noisy label filtering and correction to allow the model to learn the basic hashing mapping capabilities, with the warm-up epochs set to 5, 5, and 30 for the MIR Flickr, NUS-WIDE, and MS COCO respectively. During the training of the hashing network on a single NVIDIA RTX 3090Ti GPU, the SGD optimizer with a batch size of 48 is adopted for parameter optimization, with an initial learning rate set to 0.005, 0.001, and 0.01 for the three datasets, respectively. Both the image and text modalities employ MLP architectures consisting of two linear projection layers ($d_m \rightarrow 4,096 \rightarrow 128k$), where d_m is the feature dimension of the modality data points and k is the length of the hash codes. Regarding hyper-parameters, α and β are set to 1, 1.5, 1.2 and 0.15, 0.05, 0.2 for the three datasets, respectively. γ and η are empirically set to 5 and 1 for all three datasets, respectively, which will be discussed later. Additionally, all experimental results are averaged over three runs with different random seeds.

Evaluation Protocol In our experiment, we evaluate the proposed model’s efficiency and effectiveness using two protocols: Hamming ranking and hash lookup. For the Hamming ranking protocol, which relies on Hamming distances, we use Mean Average Precision (MAP) to measure accuracy. Meanwhile, the hash lookup protocol’s accuracy is as-

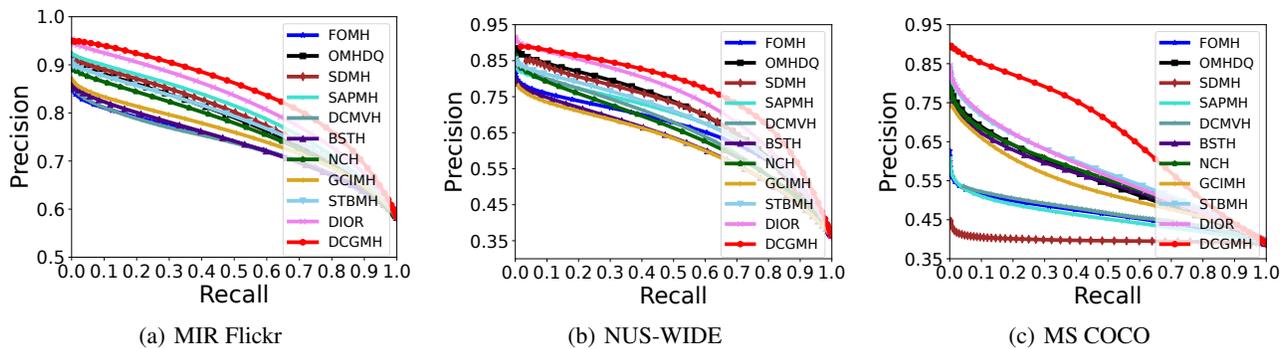


Figure 2: PR curves at a noisy label ratio of 40% on the three benchmark datasets with the code length of 64 bits.

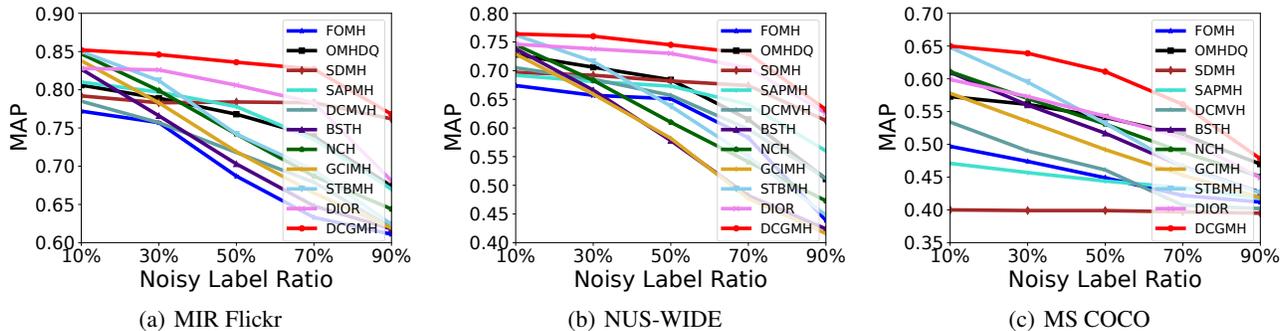


Figure 3: MAPs at different noisy label ratios on the training set of three benchmark datasets with the code length of 64 bits.

essed using the Precision-Recall (PR) curve within a specified Hamming radius. Furthermore, experiments are conducted across two sub-tasks: performance comparison at a fixed noisy label ratio, and performance variation with different noisy label ratios.

Experimental Results

Performance Comparison at a Fixed Noisy Label Ratio

To simulate the model’s retrieval performance in real-world noisy label scenarios, we conduct experiments on three datasets assuming 40% noisy labels in the training set, with the MAP and PR curve comparison with all baselines shown in Table 1 and Figure 2. The following conclusions can be drawn from the experimental results: (1) In the scenario with 40% noisy labels, our proposed method DCGMH achieves the best MAP performance compared to all baselines. Specifically, when compared to existing multi-modal hashing models that do not account for noisy labels, DCGMH provides average improvements of at least 4.9%, 6.7%, and 9.3% on the MIR Flickr, NUS-WIDE, and MS COCO, respectively. Additionally, DCGMH significantly outperforms the DIOR model designed to handle noisy labels. This demonstrates DCGMH’s effectiveness in generating high-quality hash codes in noisy label scenarios. (2) The PR curve of DCGMH outperforms all baselines, showing that the hash codes generated by the model maintain smaller Hamming distances for more similar instances, thereby maximizing the preservation of the instance’s intrinsic

semantic information and relationships. (3) As the length of the hash codes increases, both DCGMH and the baseline methods show improved performance. This indicates that longer hash codes enable a more detailed analysis of instance structure and features, allowing for better preservation of semantic information and consequently enhancing retrieval performance.

Performance Variation with Different Noisy Label Ratios

To evaluate how different noisy label ratios affect model performance, we examine the MAP metric under 64-bit hash codes as the noisy label ratio varied within the range of [10%, 30%, 50%, 70%, 90%]. The experimental results are depicted in Figure 3, from which we can obtain the following observations: (1) Our proposed method DCGMH consistently achieves the best retrieval performance across nearly all noisy label ratios, which demonstrates that DCGMH effectively mitigates the negative impact of noisy labels in the training set and showcases its exceptional robustness. (2) As the noisy label ratio increases, DCGMH shows the slowest performance decline compared to other models, indicating its lower sensitivity to noisy labels. Even with a high proportion of noisy labels, DCGMH can effectively filter and correct them to maintain superior performance. (3) Notably, when the noisy label ratio reaches 90%, the model’s performance drops sharply. This decline is likely due to the scarcity of clean label instances, which hinders the model’s ability to effectively learn basic hash

Method	MIR Flickr				NUS-WIDE				MS COCO			
	16	32	64	128	16	32	64	128	16	32	64	128
DCGMH-I	0.718	0.763	0.784	0.809	0.586	0.706	0.721	0.724	0.512	0.561	0.610	0.607
DCGMH-R	0.786	0.814	0.840	0.843	0.705	0.725	0.736	0.740	0.501	0.590	0.624	0.647
DCGMH-U	0.791	0.822	0.840	0.847	0.709	0.732	0.747	0.755	0.470	0.587	0.613	0.637
DCGMH-RU	0.781	0.811	0.836	0.839	0.699	0.722	0.723	0.731	0.477	0.587	0.611	0.635
DCGMH-D	0.770	0.805	0.828	0.827	0.678	0.693	0.701	0.710	0.469	0.583	0.622	0.638
Ours	0.796	0.823	0.846	0.850	0.717	0.739	0.755	0.757	0.551	0.591	0.631	0.649

Table 2: MAPs of the variants at a noisy label ratio of 40% on the three benchmark datasets.

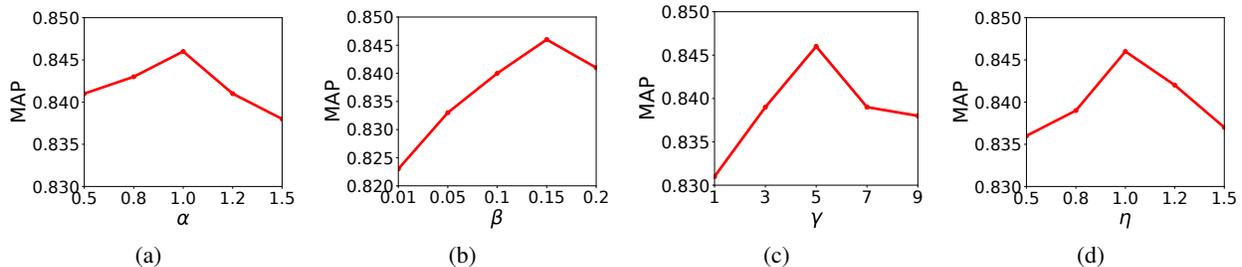


Figure 4: Sensitivity analysis of hyper-parameters on MIR Flickr with 64-bit code length and 40% noisy label ratio.

mappings during the warm-up phase, ultimately leading to overfitting and indistinguishable hash codes.

Ablation Study

To evaluate the effectiveness of each module design and filtering strategy, we define the following five variants: DCGMH-I is trained directly on the original dataset without applying any noise label filtering or reconstruction; DCGMH-R treats all noisy labels as unlabeled for unsupervised learning, without correcting high-confidence noisy labels; DCGMH-U directly discards low-confidence noisy labels; DCGMH-RU trains solely on clean labels and discards all filtered noisy labels; DCGMH-D filter noisy labels via the difference in loss between different augmented views like DIOR.

The experimental results are shown in Table 2, where the training set has 40% noisy labels. From these results, we can draw the following conclusions: (1) Compared to DCGMH-I trained directly with noisy labels, both DCGMH and the other variants show substantial performance improvements, highlighting the critical role of the noisy label filtering module in preventing model overfitting and enhancing retrieval performance. (2) Among DCGMH-R, DCGMH-U, DCGMH-RU, and DCGMH, the latter achieves the best retrieval performance, indicating that both the noisy label correction module and unsupervised semantic learning module contribute to enhancing the hash model’s ability to capture richer semantic information and generate more precise hash codes. (3) Compared to DCGMH-D, DCGMH delivers superior retrieval performance, demonstrating that our proposed filtering strategy more accurately and thoroughly filters out noisy labels to prevent model overfitting.

Sensitivity to Hyper-parameters

Taking the MIR Flickr dataset as an example, we explore the impact of different values of the hyperparameters α , β , γ , and η on model performance with the hash code length of 64 bits and the noisy label ratio of 40%. The experimental results are shown in Figure 4, where the β value varies between 0.01 and 0.2 based on empirical settings, while the α and η values range from 0.5 to 1.5, and the γ value ranges from 1 to 9. The results indicate that as the hyperparameter values increase, the model’s performance initially improves and then declines. Moreover, regardless of the values selected within the range, the performance consistently surpasses that of existing state-of-the-art multi-modal hashing models. To achieve optimal performance, we set the values of α , β , γ , and η to 1, 0.15, 5, and 1, respectively, in the experiments. Similar trends are observed on the NUS-WIDE and COCO, where the corresponding hyperparameters are set to 1.5, 0.05, 5, 1, and 1.2, 0.2, 5, 1, respectively.

Conclusion

In this paper, we propose a novel Distribution-Consistency-Guided Multi-modal Hashing (DCGMH), which can filter out noisy labels via our discovered consistency pattern between the 1-0 distribution of labels and the high-low distribution of similarity scores and simultaneously design a corrector to correct high-confidence noisy labels to generate high-quality and discriminative hash codes, thereby further enhancing the model’s retrieval performance in noisy label scenarios. Extensive experiments on three public datasets demonstrate that our proposed DCGMH outperforms the state-of-the-art baselines in multi-modal retrieval tasks.

Acknowledgments

The work is supported by National Natural Science Foundation of China (No. 62172039, 62402043, U21B2009 and 62276110) and MIIT Program(CEIEC-2022-ZM02-0247).

References

- An, J.; Luo, H.; Zhang, Z.; Zhu, L.; and Lu, G. 2022. Cognitive multi-modal consistent hashing with flexible semantic transformation. *Information Processing & Management*, 59(1): 102743.
- Chen, Y.; Zhang, H.; Tian, Z.; Wang, J.; Zhang, D.; and Li, X. 2020a. Enhanced discrete multi-modal hashing: More constraints yet less time to learn. *IEEE Transactions on Knowledge and Data Engineering*, 34(3): 1177–1190.
- Chen, Y.-C.; Li, L.; Yu, L.; El Kholy, A.; Ahmed, F.; Gan, Z.; Cheng, Y.; and Liu, J. 2020b. Uniter: Universal image-text representation learning. In *European conference on computer vision*, 104–120. Springer.
- Guo, J.-N.; Mao, X.-L.; Wei, W.; and Huang, H. 2022. Intra-category aware hierarchical supervised document hashing. *IEEE Transactions on Knowledge and Data Engineering*, 35(6): 6003–6013.
- 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, G.; Zhang, J.; Bai, X.; Wang, W.; and Meng, D. 2024. Which Is More Effective in Label Noise Cleaning, Correction or Filtering? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 12866–12873.
- Jiang, Q.-Y.; and Li, W.-J. 2017. Deep cross-modal hashing. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3232–3240.
- Kim, T.; Ko, J.; Choi, J.; Yun, S.-Y.; et al. 2021. Fine samples for learning with noisy labels. *Advances in Neural Information Processing Systems*, 34: 24137–24149.
- Li, L.; Shu, Z.; Yu, Z.; and Wu, X.-J. 2024. Robust online hashing with label semantic enhancement for cross-modal retrieval. *Pattern Recognition*, 145: 109972.
- Liu, X.; He, J.; and Lang, B. 2014. Multiple feature kernel hashing for large-scale visual search. *Pattern Recognition*, 47(2): 748–757.
- Lu, X.; Liu, L.; Nie, L.; Chang, X.; and Zhang, H. 2020. Semantic-driven interpretable deep multi-modal hashing for large-scale multimedia retrieval. *IEEE Transactions on Multimedia*, 23: 4541–4554.
- Lu, X.; Zhu, L.; Cheng, Z.; Li, J.; Nie, X.; and Zhang, H. 2019a. Flexible online multi-modal hashing for large-scale multimedia retrieval. In *Proceedings of the 27th ACM international conference on multimedia*, 1129–1137.
- Lu, X.; Zhu, L.; Cheng, Z.; Nie, L.; and Zhang, H. 2019b. Online multi-modal hashing with dynamic query-adaption. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*, 715–724.
- Lu, X.; Zhu, L.; Li, J.; Zhang, H.; and Shen, H. T. 2019c. Efficient supervised discrete multi-view hashing for large-scale multimedia search. *IEEE Transactions on Multimedia*, 22(8): 2048–2060.
- Lu, X.; Zhu, L.; Liu, L.; Nie, L.; and Zhang, H. 2021. Graph convolutional multi-modal hashing for flexible multimedia retrieval. In *Proceedings of the 29th ACM international conference on multimedia*, 1414–1422.
- Shen, X.; Chen, Y.; Pan, S.; Liu, W.; and Zheng, Y. 2023. Graph convolutional incomplete multi-modal hashing. In *Proceedings of the 31st ACM international conference on multimedia*, 7029–7037.
- Shen, X.; Shen, F.; Liu, L.; Yuan, Y.-H.; Liu, W.; and Sun, Q.-S. 2018. Multiview discrete hashing for scalable multimedia search. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 9(5): 1–21.
- Shen, X.; Shen, F.; Sun, Q.-S.; and Yuan, Y.-H. 2015. Multi-view latent hashing for efficient multimedia search. In *Proceedings of the 23rd ACM international conference on Multimedia*, 831–834.
- Shu, J.; Xie, Q.; Yi, L.; Zhao, Q.; Zhou, S.; Xu, Z.; and Meng, D. 2019. Meta-weight-net: Learning an explicit mapping for sample weighting. *Advances in neural information processing systems*, 32.
- Song, H.; Dai, R.; Raskutti, G.; and Barber, R. F. 2020. Convex and non-Convex approaches for statistical inference with class-conditional noisy labels. *Journal of Machine Learning Research*, 21(168): 1–58.
- Song, H.; Kim, M.; Park, D.; Shin, Y.; and Lee, J.-G. 2022. Learning from noisy labels with deep neural networks: A survey. *IEEE transactions on neural networks and learning systems*, 34(11): 8135–8153.
- Song, J.; Yang, Y.; Huang, Z.; Shen, H. T.; and Luo, J. 2013. Effective multiple feature hashing for large-scale near-duplicate video retrieval. *IEEE Transactions on Multimedia*, 15(8): 1997–2008.
- Sun, J.; Wang, H.; Luo, X.; Zhang, S.; Xiang, W.; Chen, C.; and Hua, X.-S. 2022. Heart: Towards effective hash codes under label noise. In *Proceedings of the 30th ACM International Conference on Multimedia*, 366–375.
- Sung, F.; Yang, Y.; Zhang, L.; Xiang, T.; Torr, P. H.; and Hospedales, T. M. 2018. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1199–1208.
- Tan, W.; Zhu, L.; Guan, W.; Li, J.; and Cheng, Z. 2022. Bit-aware semantic transformer hashing for multi-modal retrieval. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, 982–991.
- Tan, W.; Zhu, L.; Li, J.; Zhang, Z.; and Zhang, H. 2023. Partial multi-modal hashing via neighbor-aware completion learning. *IEEE Transactions on Multimedia*, 25: 8499–8510.
- Tu, R.-C.; Ji, L.; Luo, H.; Shi, B.; Huang, H.-Y.; Duan, N.; and Mao, X.-L. 2021a. Hashing based efficient inference for image-text matching. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 743–752.

- Tu, R.-C.; Jiang, J.; Lin, Q.; Cai, C.; Tian, S.; Wang, H.; and Liu, W. 2023a. Unsupervised cross-modal hashing with modality-interaction. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(9): 5296–5308.
- Tu, R.-C.; Mao, X.; and Wei, W. 2020. MLS3RDUH: Deep Unsupervised Hashing via Manifold based Local Semantic Similarity Structure Reconstructing. In *IJCAI*, 3466–3472.
- Tu, R.-C.; Mao, X.-L.; Feng, B.-S.; Bian, B.-B.; and Ying, Y.-s. 2018. Object detection based deep unsupervised hashing. *arXiv preprint arXiv:1811.09822*.
- Tu, R.-C.; Mao, X.-L.; Guo, J.-N.; Wei, W.; and Huang, H. 2021b. Partial-softmax loss based deep hashing. In *Proceedings of the Web Conference 2021*, 2869–2878.
- Tu, R.-C.; Mao, X.-L.; Ji, W.; Wei, W.; and Huang, H. 2023b. Data-aware proxy hashing for cross-modal retrieval. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 686–696.
- Tu, R.-C.; Mao, X.-L.; Kong, C.; Shao, Z.; Li, Z.-L.; Wei, W.; and Huang, H. 2021c. Weighted gaussian loss based hamming hashing. In *Proceedings of the 29th ACM International Conference on Multimedia*, 3409–3417.
- Tu, R.-C.; Mao, X.-L.; Liu, J.; Wei, W.; Huang, H.; et al. 2024. Similarity Transitivity Broken-Aware Multi-Modal Hashing. *IEEE Transactions on Knowledge and Data Engineering*.
- Tu, R.-C.; Mao, X.-L.; Tu, R.-X.; Bian, B.; Cai, C.; Wang, H.; Wei, W.; and Huang, H. 2022. Deep cross-modal proxy hashing. *IEEE Transactions on Knowledge and Data Engineering*, 35(7): 6798–6810.
- Wang, H.; Jiang, H.; Sun, J.; Zhang, S.; Chen, C.; Hua, X.-S.; and Luo, X. 2023. Dior: Learning to hash with label noise via dual partition and contrastive learning. *IEEE Transactions on Knowledge and Data Engineering*.
- Wei, Q.; Sun, H.; Lu, X.; and Yin, Y. 2022. Self-filtering: A noise-aware sample selection for label noise with confidence penalization. In *European Conference on Computer Vision*, 516–532. Springer.
- Wu, X.; Zhu, L.; Xie, L.; Zhang, Z.; and Zhang, H. 2021a. Multi-modal discrete tensor decomposition hashing for efficient multimedia retrieval. *Neurocomputing*, 465: 1–14.
- Wu, X.-M.; Luo, X.; Zhan, Y.-W.; Ding, C.-L.; Chen, Z.-D.; and Xu, X.-S. 2022. Online enhanced semantic hashing: Towards effective and efficient retrieval for streaming multi-modal data. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 4263–4271.
- Wu, Y.; Shu, J.; Xie, Q.; Zhao, Q.; and Meng, D. 2021b. Learning to purify noisy labels via meta soft label corrector. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 10388–10396.
- Xie, L.; Shen, J.; Han, J.; Zhu, L.; and Shao, L. 2017. Dynamic Multi-View Hashing for Online Image Retrieval. In *IJCAI*, volume 78, 122.
- Yan, C.; Gong, B.; Wei, Y.; and Gao, Y. 2020. Deep multi-view enhancement hashing for image retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(4): 1445–1451.
- 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.
- Yang, R.; Shi, Y.; and Xu, X.-S. 2017. Discrete multi-view hashing for effective image retrieval. In *Proceedings of the 2017 ACM on international conference on multimedia retrieval*, 175–183.
- Yu, J.; Huang, W.; Li, Z.; Shu, Z.; and Zhu, L. 2022. Hadamard matrix-guided multi-modal hashing for multi-modal retrieval. *Digital Signal Processing*, 130: 103743.
- Yuan, L.; Wang, T.; Zhang, X.; Tay, F. E.; Jie, Z.; Liu, W.; and Feng, J. 2020. Central similarity quantization for efficient image and video retrieval. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3083–3092.
- 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.
- Zhang, X.; Lai, H.; and Feng, J. 2018. Attention-aware deep adversarial hashing for cross-modal retrieval. In *Proceedings of the European conference on computer vision (ECCV)*, 591–606.
- Zheng, C.; Zhu, L.; Cheng, Z.; Li, J.; and Liu, A.-A. 2020a. Adaptive partial multi-view hashing for efficient social image retrieval. *IEEE Transactions on Multimedia*, 23: 4079–4092.
- Zheng, C.; Zhu, L.; Lu, X.; Li, J.; Cheng, Z.; and Zhang, H. 2019. Fast discrete collaborative multi-modal hashing for large-scale multimedia retrieval. *IEEE Transactions on Knowledge and Data Engineering*, 32(11): 2171–2184.
- Zheng, C.; Zhu, L.; Zhang, S.; and Zhang, H. 2020b. Efficient parameter-free adaptive multi-modal hashing. *IEEE Signal Processing Letters*, 27: 1270–1274.
- Zheng, C.; Zhu, L.; Zhang, Z.; Duan, W.; and Lu, W. 2024. LCEMH: Label Correlation Enhanced Multi-modal Hashing for efficient multi-modal retrieval. *Information Sciences*, 659: 120064.
- Zheng, C.; Zhu, L.; Zhang, Z.; Li, J.; and Yu, X. 2022. Efficient semi-supervised multimodal hashing with importance differentiation regression. *IEEE Transactions on Image Processing*, 31: 5881–5892.
- Zheng, S.; Wu, P.; Goswami, A.; Goswami, M.; Metaxas, D.; and Chen, C. 2020c. Error-bounded correction of noisy labels. In *International Conference on Machine Learning*, 11447–11457. PMLR.
- Zhu, L.; Lu, X.; Cheng, Z.; Li, J.; and Zhang, H. 2020. Deep collaborative multi-view hashing for large-scale image search. *IEEE Transactions on Image Processing*, 29: 4643–4655.
- Zhu, L.; Zheng, C.; Lu, X.; Cheng, Z.; Nie, L.; and Zhang, H. 2021. Efficient multi-modal hashing with online query adaptation for multimedia retrieval. *ACM Transactions on Information Systems (TOIS)*, 40(2): 1–36.