Robust Domain Adaptive Hashing via Structural Noise Modeling and Correction

Junsheng Wang[1], Tiantian Gong[1] *, Yeyun Wu[2], Xiaobing Sun[1]

[1] College of Information and Artificial Intelligence, Yangzhou University
[2] School of Artificial Intelligence and Big Data, Jiangxi Institute of Technology
{js_wang, gongtiantian, xbsun}@yzu.edu.cn, wuyeyun@jxut.edu.cn

Abstract

Deep hashing offers efficient storage and fast retrieval capabilities. As a result, it has been extensively applied to large-scale retrieval tasks. To alleviate the dependence on high-quality annotated data, recent research has focused on unsupervised domain adaptive hashing methods, which aim to transfer knowledge from a label-rich source domain to a label-scarce target domain. However, in open-world scenarios, source domain labels are often inevitably noisy, which tends to undermine the quality of learned hash codes and induce considerable performance deterioration. To this end, we introduce a novel *Robust Domain Adaptive Hashing* (RDAH) method to jointly mitigate the adverse effects of label noise and domain discrepancy. Specifically, we first model the loss distribution of training samples using a two-component Gaussian mixture model to estimate each sample's confidence, based on which the data is partitioned. Subsequently, we introduce a neighbor consistency-guided correction strategy, which leverages the semantic structure of high-confidence neighbors to perform weighted correction on noisy samples. Moreover, we design a dual-level cross-domain alignment mechanism that jointly mitigates domain shift from two complementary perspectives. Extensive experimental results validate the effectiveness and robustness of RDAH across multiple benchmark datasets.

Introduction

With the explosive growth of large-scale multimedia data, how to achieve efficient and accurate similarity retrieval from massive samples has become a core challenge. Hashing-based retrieval methods have been widely adopted for approximate nearest neighbor (ANN) search tasks across modalities such as images and texts due to their low storage overhead and high computational efficiency. For example, multimodal medical hashing retrieval (Zhang et al. 2025), large language models (Chen et al. 2024), recommendation systems (Wu et al. 2025), and retrieval-augmented generation (Luo et al. 2025a; Zhao et al. 2024). Hashing methods accelerate the retrieval process significantly by compressing original features into compact binary codes while preserving semantic similarity. In deep hashing learning, existing methods can be broadly divided into supervised and unsupervised approaches. Supervised hashing methods leverage label information to guide the learning of discriminative hash codes, typically achieving superior retrieval performance. In contrast, unsupervised methods lack semantic supervision, leading to poor discriminative hash codes, thereby limiting their effectiveness in open-world scenarios.

*Corresponding author.

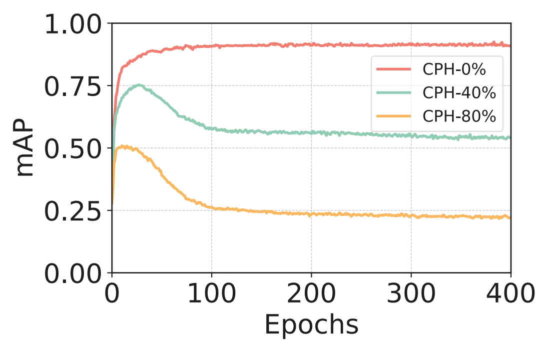

Figure 1: The retrieval performance of the CPH method (Cui et al. 2024) deteriorates substantially as the noise ratio increases, highlighting its vulnerability to label noise and the lack of robustness in noisy environments.

In real-world applications, retrieval systems frequently encounter cross-domain challenges, where query samples and database samples originate from different data distributions. Traditional hashing methods often fail to maintain hash code consistency in such settings, as they lack the ability to model distribution shifts across domains, leading to notable performance degradation. As a result, unsupervised cross-domain hashing retrieval (UCDHR) (Cui et al. 2024; Wang et al. 2023b,c) has attracted increasing attention. It aims to construct a shared semantic space across domains, where semantically similar samples from different distributions are projected into similar hash codes. Although existing cross-domain hashing methods alleviate domain shift to a certain extent, they typically rely on the strong assumption that training data is clean and accurately labeled, which rarely holds in the real world. Due to the limitations in data collection and annotation costs, along with the prevalence of weakly supervised methods, source domain data in real-world applications often contain substantial noisy labeled samples. The presence of these noisy labels inevitably leads to the overfitting problem during training, thus transferring noisy knowledge to the target domain. Specifically, the presence of noisy information misguides the model to learn incorrect hash codes, which hinders the learning of semantic consistency and ultimately degrades the performance of unsupervised hash retrieval. As illustrated in Figure 1, state-of-the-art cross-domain hashing methods suffer a significant drop in mAP retrieval performance when confronted with varying levels of label noise. Particularly under high noise rates, the model's performance fluctuates drastically, revealing poor robustness and inadequate cross-domain generalization. These observations highlight the urgent need to address label noise in the context of robust domain adaptive hash retrieval.

To this end, we propose a robust cross-domain hashing retrieval method that aims to mitigate the interference of noisy labels by leveraging data structure information and semantic alignment mechanisms, thereby achieving more stable and accurate retrieval performance across domains. More specifically, we propose a novel noise modeling and correction learning framework, named Robust Domain Adaptive Hashing (RDAH) as shown in Figure 2. First, we introduce an innovative structural noise modeling approach that fits the sample-wise loss distribution using a Gaussian mixture model (GMM), uncovering the underlying noise distribution characteristics. Subsequently, the confidence of each sample is adaptively modeled based on the posterior probability, and the source data is dynamically divided into clean and noisy sample sets. For noisy samples, we further design a neighbor consistency correction strategy that leverages the semantic information of their corresponding high-confidence clean neighbors for correction in the feature space, thereby enhancing both discriminability and semantic consistency. Finally, to further mitigate the adverse effects of noisy labels on hash code learning, we propose a dual-level alignment module to align cross-domain semantic structures from both the global category-level and domain-level perspectives, thereby effectively improving the robustness and transferability of the learned hash representations.

Our contributions can be summarized as follows:

- We reveal and summarize the negative impact of noisy labels in the source domain on unsupervised domain adaptive hash retrieval. To address this issue, we propose a robust domain adaptive hash retrieval method to learn domain-invariant hash representations. To the best of our knowledge, this work is the first to investigate label noise in domain adaptive hash retrieval.
- To mitigate the impact of noisy labels in the source domain, we propose a dynamic structural noise modeling and correction mechanism that reliably distinguishes and corrects noisy samples instead of discarding them arbitrarily, further reducing the risk of overfitting caused by limited training data in the source domain.
- To better alleviate the semantic misalignment problem caused by domain shift and noisy labels, we align cross-domain semantic structures at both the class-level and domain-level perspectives, effectively improving the robustness and generalization of hash representations.
- Extensive experimental results show that our method, RDAH, consistently outperforms state-of-the-art unsupervised domain adaptive hashing approaches under various levels of label noise, achieving significant gains in key metrics such as mAP.

Related Work

Learning to Hash

In recent years, hash learning has received increasing attention in the field of image retrieval because it can better balance computational efficiency and retrieval accuracy. Current hashing techniques can be generally divided into two categories: unsupervised and supervised hashing methods. Unsupervised hashing methods mainly consist of similarity reconstruction-based methods (Tu, Mao, and Wei 2020; Luo et al. 2020), pseudo-label-based methods (Zhang et al. 2020; Yuan et al. 2020), and self-supervised learning-based methods (Jang and Cho 2021; Qiu et al. 2021). These methods directly or indirectly learn the semantic information of the hash codes. Supervised hashing approaches can be grouped into pairwise methods (Li et al. 2017; Chen et al. 2019), ranking-based methods (Lai et al. 2015; Cakir et al. 2019), and pointwise methods (Hoe et al. 2021; Yuan et al. 2020). The primary difference among these methods lies in the granularity of the supervision signals they exploit. Pairwise methods or ranking-based methods construct a similarity structure that guides the hash code learning. However, pointwise methods directly take advantage of label information instead of similarity information. Although these methods have achieved encouraging results on standard benchmarks, in real applications, when the training data and query data come from different distributions, we urgently need to focus on domain adaptive hash retrieval methods.

Domain Adaptive Hashing

Domain adaptive hashing learning has received widespread attention in recent years. Its core goal is to effectively transfer rich source domain knowledge to the label-poor target domain, thereby improving image retrieval performance in cross-domain scenarios. Traditional methods (Gong et al. 2012; Huang, Zhang, and Gao 2021; Huang et al. 2020; Liu and Zhang 2019) use machine learning principles to align the feature distributions or structures of the source and target domains, and learn hash representations that are domain-invariant and discriminative. Thanks to the widespread application of deep learning in domain adaptation, a series of deep hashing-based domain adaptation methods have emerged in recent years. PEACE (Wang et al. 2023a) progressively minimizes the uncertainty of pseudo-labels through alternative optimization under the guidance of the domain discrepancy, and removes domain discrepancy in the Hamming space from two views. ROSE (Yang et al. 2024) proposes relational and prototypical structure learning to handle the universal domain adaptive problem. CPH (Cui et al. 2024) proposes learning a domain-shared unit hypersphere space through prototype contrastive learning, while preserving inter-sample similarity relationships within both domains. COUPLE (Luo et al. 2025b) proposes graph diffusion to achieve a stable target domain adaptation process.

Learning with Noisy Labels

As noisy labels severely degrade the generalization performance of deep neural networks, various methods have been proposed to improve robustness against noisy labels. Existing methods can be roughly divided into three categories: 1) Model-based methods (Li et al. 2022; Yao et al. 2020) introduce specialized structural modules or parameterized modeling mechanisms to simulate the transition matrices from clean labels to noisy ones, enabling the network to adapt to label corruption during training. 2) Loss-based methods (Zhang and Sabuncu 2018; Ren et al. 2018; Jiang et al. 2018) attempt to design robust loss functions or reweight samples, which employ automatic learning to assign weights to training samples. 3) Sample-based methods (Li et al. 2020; Wang et al. 2018; Jiang et al. 2020) attempt to separate clean samples from corrupted samples, and train the model on clean samples. Although these methods perform well on single-domain learning, they cannot be directly applied to domain adaptive hash learning under noisy labels.

Methodology

Problem Definition

To ensure clarity and facilitate understanding, let $\mathcal{D}_s = \{x_i^s, y_i^s\}_{i=1}^{N_s}$ represent the source domain data with noisy labels, where N_s is the sample number. $\mathcal{D}_t = \{x_i^t\}_{i=1}^{N_t}$ represents the target domain data, where x_i^t represents the i-th sample from the target domain, and N_t is the sample number. Cross-domain hashing retrieval employs knowledge transfer to learn similarity-preserving hash codes $B^t = [b_1^t, ..., b_{N_t}^t] \in \{-1, 1\}^{N_t \times l}$, where b_i^t denotes a compact l-bit binary code. In single-domain retrieval, both the queries and the database originate from D_t, whereas cross-domain retrieval requires matching the queries from D_t with samples from D_s. Noisy labels in the source domain not only destroy the similarity-preserving hash codes but also combine with domain discrepancy to exacerbate domain shift. Our goal is to learn a noise-invariant robust hashing method.

Warm-up Training on Noisy Source Data

Recent research (Arpit et al. 2017) has shown that after a short warm-up period, deep neural networks (DNNs) are more likely to produce accurate predictions for simple samples that are close to the ground-truth labels, while consistently yielding incorrect predictions for more noisy samples. To alleviate the overfitting problem of cross-entropy (CE), inspired by (Zhang and Sabuncu 2018), we adopt the robust generalized cross entropy (GCE) loss to mitigate the impact of noise in warm-up stage. Thus, we compute the discriminative loss for each sample as follows:

$$L_{gce}(f(x_i^s), y_i^s; q) = \frac{(1 - f_{y_i^s}(x_i^s)^q)}{q}, \tag{1}$$

where $f_{y_i^s}(x_i^s)$ represents the predicted probability of sample x_i^s on the true label y_i^s, and $q \in (0, 1]$. GCE loss degenerates to CE when q approaches 0, and reduces to the Mean Absolute Error (MAE) when $q = 1$. For simplicity, we denote the GCE loss of the i-th source sample $L_{gce}(f(x_i^s), y_i^s; q)$ as L_i.

Progressive Noisy Filtering and Correction

To successfully separate source samples into clean and noisy labeled data, and subsequently reassign pseudo labels to the identified noisy samples, we propose a progressive noisy filtering and correction module (PNFC) in Figure 2. The PNFC consists of two stages: GMM-based Noise Separation (GNS) and Neighbor Consistency Correction (NCC).

GMM-based Noise Separation. Considering the memory effect of DNNs, the loss of clean samples converges faster than that of samples with noisy labels (Arpit et al. 2017; Zhang et al. 2021). Therefore, we adopt a two-component Gaussian Mixture Model (GMM) to separate the noisy source samples into clean and noisy labeled data. Specifically, for each sample loss L_i in the source domain, the probability density of the two-component GMM is calculated as follows:

$$\psi(L_i) = \sum_{r=1}^{R} \beta_r \psi(L_i|r), \tag{2}$$

where $R = 2$ defines the number of Gaussian components, and β_r denotes the mixing coefficient in the convex combination. $\psi(L_i|r)$ is the probability density function. The posterior probability $\psi(r^{'}|L_i)$ indicates the possibility that the source sample x_i^s belongs to the set of clean samples, where $r^{'}$ has a smaller mean. $\psi(r^{'}|L_i)$ is formulated as:

$$\psi(r^{'}|L_i) = \frac{\psi(r^{'})\psi(L_i|r^{'})}{\psi(L_i)}. \tag{3}$$

Next, a threshold is applied to the posterior probability $\psi(r^{'}|L_i)$ to dynamically separate the source data into clean samples and noisy labeled samples, as follows:

$$\begin{cases} \mathcal{D}_s^c \supseteq (x_i^s, y_i^s), & \psi(r^{'}|L_i) \geq \gamma, \\ \mathcal{D}_s^n \supseteq (x_i^s, y_i^s), & \psi(r^{'}|L_i) < \gamma, \end{cases} \tag{4}$$

where $\mathcal{D}_s^c$ represents the clean data set, $\mathcal{D}_s^n$ denotes the noisy label data set. Subsequently, we reassign pseudo labels to the noisy sample set $\mathcal{D}_s^n$ to facilitate effective adaptive learning.

Neighbor Consistency Correction. To enhance semantic alignment between the source and target domains and prevent the model from overfitting to noisy labeled data, we propose a neighbor consistency correction (NCC) module to rectify noisy labels. Specifically, for the noise instance $x_i^s \in \mathcal{D}_s^n$, we use the k-nearest neighbor algorithm to search for the neighbors of k clean samples, formulated as:

$$\mathcal{N}(x_i^s, k) = \{(x_j^s, y_j^s),\ j = 1, ..., k,\ \forall x_j^s \in \mathcal{D}_s^c\}. \tag{5}$$

Here, $\mathcal{N}(x_i^s, k)$ denotes the clean sample neighbors of the noisy instance x_i^s, which helps ensure the accuracy of the reassigned pseudo labels. Furthermore, we perform the Jensen Shannon divergence (JSD) metric between the model predicted label distribution $\hat{y}_i^s$ of the noisy instance x_i^s and the

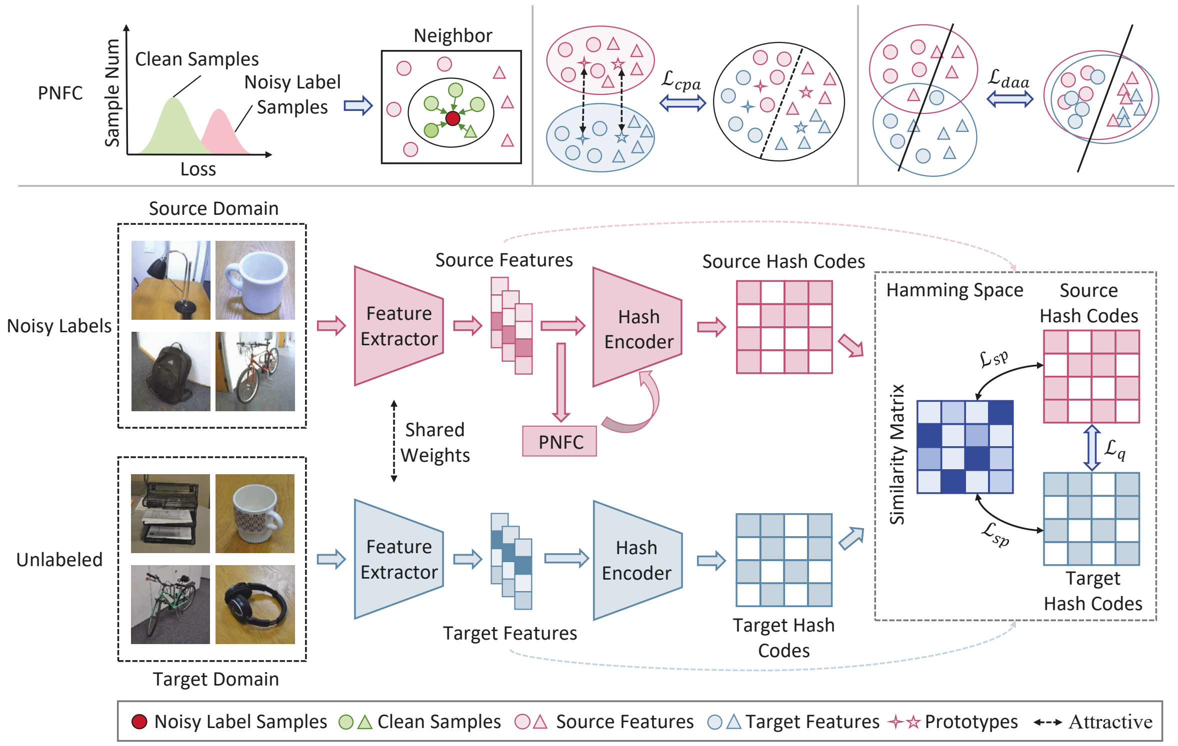

Figure 2: The overview of RDAH. Our goal is to learn noise-tolerant domain adaptive hash codes. Specifically, we model the noise distribution using a Gaussian mixture model and correct for noisy labels. In addition, we propose a dual-level cross-domain alignment strategy that collaboratively mitigates the challenges of domain shift and label noise.

true label distribution y_j^s of its k neighbors as follow:

$$\sum_{j=1}^{k} JSD(\hat{y}_i^s, y_j^s) = \sum_{j=1}^{k}(\frac{1}{2}KLD(y_j^s \parallel \frac{y_j^s + \hat{y}_i^s}{2}) + \frac{1}{2}KLD(\hat{y}_i^s \parallel \frac{y_j^s + \hat{y}_i^s}{2})), \quad (6)$$

where KLD(.) defines the Kullback-Leibler divergence function. The smaller $\mathcal{Q}(x_i^s) = \sum_{j=1}^{k} JSD(\hat{y}_i^s, y_j^s)$ is, the more consistent the distribution between $\hat{y}_i^s$ and y_j^s becomes. In addition, a threshold is set for $\mathcal{Q}(x_i^s) < \eta$ to ensure reliable pseudo labels. For the noisy instance x_i^s, we adopt the weighted label correction strategy to generate credible pseudo labels as follows:

$$\tilde{y}_i^s = \arg\max_c \sum_{j=1}^{k}(1 - JSD(\hat{y}_i^s, y_j^s)) \cdot y_j^s, \quad (7)$$

where $c = 1, 2, ..., C$ indicates that the c-th component of the label distribution vector has the highest value. Therefore, if both $\mathcal{Q}(x_i^s) < \eta$ and Eqn. (7) are satisfied, a pseudo label $\tilde{y}_i^s$ is assigned to the noisy instance x_i^s, denoted as $(x_i^s, \tilde{y}_i^s)$.

Dual-level Cross-domain Alignment

Following the progressive noisy filtering and correction, we employ dual-level cross-domain alignment to achieve domain alignment in Figure 2. To address the noisy sample issue in domain alignment, we first introduce a prototype contrastive learning objective. The goal is to pull positive pairs of prototypes from the same category across the source and target domains closer together, while pushing negative pairs from different categories further apart, thereby boosting the semantic compactness and discriminability of the learned representations in the shared feature space.

Cross-domain Prototype Alignment. Previous works adopted the Hamming center alignment (Wang et al. 2023a) to achieve semantic alignment, which compacts different samples to their corresponding semantic centers in the common feature space. However, due to noisy labels in the source, instance samples are misaligned with their corresponding centers, which makes it difficult to optimize the compactness of hash representations. To tackle these issues, we utilize the prototype contrastive loss to optimize hash code learning, formulated as:

$$L_{cpa} = \frac{1}{C}\sum_{c=1}^{C} -log \frac{exp(\frac{(p_c^s)^T p_c^t}{\tau})}{exp(\frac{(p_c^s)^T p_c^t}{\tau}) + \sum_{i \neq c} exp(\frac{(p_c^s)^T p_i^t}{\tau})}, \quad (8)$$

where τ represents the temperature hyperparameter. p_c^s and p_c^t are regarded as positive pairs, and prototypes of other categories are considered as negative samples. The source domain prototype p_c^s is calculated by averaging the samples with the same label. The target domain prototype p_c^t is obtained by averaging target domain samples with the same pseudo-label. The target domain pseudo-label assignment process is expressed as $\hat{y}_i^t = \arg\max_c cos(f(x_i^t), p_c^s)$. The prototype contrastive learning method is more suitable for cross-domain hash representation learning under noise, as it dilutes the negative impact of noise.

Domain Adversarial Alignment. To enforce consistency across domain representations, we use the domain adversarial alignment (DAA) loss as the global domain alignment loss. Specifically, a discriminator network d is trained in an adversarial manner, which attempts to distinguish whether features arise from the source domain or the target domain, while features attempt to confuse d. The adversarial loss for global feature alignment is defined as follows:

$$L_{daa}(\mathcal{D}_s, \mathcal{D}_t) = -\frac{1}{N_s}\sum_{i=1}^{N_s} log(d(f(x_i^s))) - \frac{1}{N_t}\sum_{i=1}^{N_t} log(1 - d(f(x_i^t))). \quad (9)$$

Similarity Preserving Loss. To preserve the intrinsic relationships among samples in the hash space, following (Cui et al. 2023, 2024), we utilize a Similarity Preserving Loss (SPL). Specifically, we first leverage the similarity matrix from the source domain labels to guide the hash learning process, which is formulated as follows:

$$L_{sp1} = \parallel \delta G^s - cos(H^s, H^s) \parallel_2^2, \quad (10)$$

where δ is defined as a scaling factor, and G^s denotes the similarity matrix of source domain labels. H^s signifies the source domain relaxed hash codes obtained through the hash encoder $H(\cdot)$. To transfer the learned knowledge from the source domain to the target domain, we enforce the features from both domains to be close to their corresponding hash codes, calculated as follows:

$$L_{sp2} = \parallel cos(Z^s, Z^t) - cos(H^s, H^t) \parallel_2^2, \quad (11)$$

where Z^s, Z^t represents the source and target domain features, H^s, H^t denotes the source and target domain hash codes. Therefore, the similarity preserving loss is formulated as follows:

$$L_{sp} = L_{sp1} + \lambda L_{sp2}, \quad (12)$$

where λ is a hyperparameter to balance the loss.

Quantization Loss. To ensure that the relaxed hash codes closely approximate the binary hash codes, we employ quantization loss for both the source and target domains. The quantization loss is calculated as follows:

$$L_q = \frac{1}{2} \parallel B^s - H^s \parallel_2^2 + \frac{1}{2} \parallel B^t - H^t \parallel_2^2 . \quad (13)$$

Here, B^s, B^t represent the binary hash codes.

Finally, the overall objective loss is formulated as follows:

$$L = \lambda 1 L_{cpa} + \lambda 2 L_q + \lambda 3 L_{sp} + \lambda 4 L_{daa}, \quad (14)$$

where $\lambda 1$, $\lambda 2$, $\lambda 3$, and $\lambda 4$ are trade-off parameters.

Experiments

Experimental Settings

Office-31 (Saenko et al. 2010): consists of three distinct domains: Amazon (A), Webcam (W), and DSLR (D). The dataset contains a total of 4,110 images from 31 different categories. Following previous research, we conduct six transfer image retrieval experiments: A→W, W→A, A→D, D→A, W→D, and D→W.

Office-Home (Venkateswara et al. 2017): consists of four domains: Product images (Pr), Real-World images (Rw), Artistic images (Ar), and Clip Art (Cl). It is a larger dataset containing 15,500 images across 65 categories. Following previous works, we conduct experiments on six cross-domain image retrieval tasks:Pr→Rw, Rw→Pr, Rw→Ar, Ar→Rw, Cl→Ar, Ar→Cl.

Digits: This dataset consists of two digit datasets: MNIST (LeCun et al. 2002) and USPS (Hull 2002). Either dataset can be selected as the source or target domain, resulting in two transfer tasks: MNIST(M) → USPS(U) and USPS(U) → MNIST(M).

Implementation Details. We select nine state-of-the-art domain adaptive hashing methods for comparison, including ITQ (Gong et al. 2012), GTH-g (Liu and Zhang 2019), GTH-h (Liu and Zhang 2019), DAPH (Huang, Zhang, and Gao 2021), PWCF (Huang et al. 2020), PEACE (Wang et al. 2023a), ROSE(Yang et al. 2024), CPH (Cui et al. 2024), and COUPLE (Luo et al. 2025b). Our method is evaluated using three common metrics: mean Average Precision (mAP), precision-recall curve, and Top-K precision curve.

Following prior studies (Huang et al. 2020; Cui et al. 2024), each dataset consists of 10% randomly selected target samples for testing, and the rest is used for training. This experiment employs a VGG-16 as the backbone for feature extraction. Specifically, our hash encoder employs a two-layer MLP model to project the visual features into the Hamming space. The entire network is optimized by the Adam update rule with the learning rate 0.0001 and batch size 256. The training epochs for all datasets are set to 500. For the hyperparameters in RDAH, we set $\lambda 1 = 0.1$, $\lambda 2 = 0.01$, $\lambda 3 = 0.1$, $\lambda 4 = 0.1$, $\lambda 1 = 1$, $\lambda 2 = 1$, $\lambda 3 = 100$, $\lambda 4 = 10$, and $\lambda 1 = 0.01$, $\lambda 2 = 0.01$, $\lambda 3 = 0.01$, $\lambda 4 = 0.01$ on Office-31, Office-Home and Digits, respectively. The η is set to 0.5 and γ is set to 0.5. The warmup epochs for Office-31, Office-Home, and Digits are set to $5, 5, 15$. q is set 0.75.

Comparison with the State-of-the-Art Methods

To quantitatively show the performance of our proposed RDAH, we emulate noisy transfer by randomly shuffling labels at specific noise ratios. We report results with noise ratios of 40%, 60%, and 80% for comprehensive comparison with current SOTA methods (e.g., CPH and COUPLE).

Table 1 details the results of Office-31 and Office-Home at different noise ratios. We find that the SOTA methods (i.e., CPH and COUPLE) significantly outperformed previous cross-domain hash retrieval methods. However, concerningly, as the noise ratio increases, the performance of these methods drops sharply due to their failure to model noisy data. A significant enhancement is achieved by our proposed RDAH on cross-domain retrieval tasks. Specifically, our proposed RDAH method achieves average performance gains of approximately 9.77%, 15.78% and 12.73% over CPH on the Office-Home and Office-31 datasets under

Noise	Methods	Office-31						Office-Home						Avg.
		A → W	W → A	A → D	D → A	W → D	D → W	Pr → Rw	Rw → Pr	Rw → Ar	Ar → Rw	Cl → Ar	Ar → Cl	
40%	ITQ	22.03	22.47	25.20	27.12	54.48	45.42	35.97	35.91	28.09	27.96	13.96	12.85	29.29
	GTH-g	9.71	9.37	10.70	9.90	20.49	22.21	12.64	13.59	11.04	11.70	4.92	5.40	11.81
	GTH-h	11.74	9.64	13.91	9.53	18.47	18.45	13.91	12.75	10.75	11.37	5.05	5.78	11.78
	PWCF	10.80	11.81	15.39	10.13	18.89	20.22	16.00	18.77	16.98	13.97	7.75	3.93	13.72
	DAPH	12.99	11.51	13.13	10.81	20.56	20.28	15.36	16.13	14.30	12.29	6.44	6.15	13.33
	PEACE	27.14	33.67	41.47	38.37	82.65	64.55	50.01	64.98	36.85	56.01	8.95	15.26	43.33
	ROSE	43.74	29.35	47.38	28.67	45.15	41.90	20.52	16.64	20.66	20.58	17.56	11.58	28.64
	CPH	50.08	50.97	57.51	49.33	88.98	77.17	55.04	53.28	38.06	43.02	29.76	19.07	51.02
	COUPLE	36.24	34.45	39.88	37.46	78.96	65.79	43.99	46.51	32.81	36.74	22.18	14.08	40.76
	Ours	**60.17**	**59.04**	**66.88**	**56.40**	**93.14**	**80.18**	**71.42**	**69.09**	**54.98**	**55.69**	**39.18**	**23.62**	**60.79**
60%	ITQ	23.06	21.59	24.87	24.08	58.53	48.07	34.53	36.19	28.10	28.33	13.87	11.79	29.42
	GTH-g	6.29	6.85	6.26	6.79	10.68	10.47	6.87	7.50	6.29	5.95	3.07	3.05	6.67
	GTH-h	7.10	6.25	8.03	6.23	10.57	10.41	6.99	7.76	7.68	6.29	3.30	3.49	7.01
	PWCF	9.10	6.65	9.34	6.32	12.20	8.53	9.14	11.26	11.34	7.04	4.79	2.35	8.17
	DAPH	7.54	7.43	7.68	8.55	11.01	12.61	8.09	8.76	8.49	8.53	3.57	3.25	7.96
	PEACE	11.62	19.23	10.63	20.49	70.10	45.07	35.20	57.10	28.57	45.49	5.89	10.75	30.01
	ROSE	38.29	24.03	38.17	20.14	36.83	34.20	18.92	16.44	18.26	20.93	16.21	10.42	24.40
	CPH	37.72	40.59	43.08	37.44	75.90	61.01	39.59	38.37	28.18	26.55	18.93	12.47	38.32
	COUPLE	28.90	26.53	31.87	30.65	69.47	58.16	30.30	32.20	25.85	23.23	13.55	9.96	31.72
	Ours	**55.34**	**53.38**	**59.62**	**48.55**	**86.98**	**70.17**	**66.48**	**61.64**	**51.02**	**43.23**	**32.70**	**20.06**	**54.10**
80%	ITQ	21.95	21.24	24.58	26.71	58.89	47.25	34.50	36.62	28.55	26.82	13.61	12.42	29.43
	GTH-g	4.60	4.40	4.68	4.36	5.68	6.42	3.13	3.24	2.97	2.69	2.17	2.16	3.88
	GTH-h	4.70	4.57	4.85	4.90	6.29	6.38	3.02	3.29	2.95	2.88	2.19	2.26	4.02
	PWCF	4.72	4.95	4.20	4.81	5.96	5.27	3.29	3.42	3.28	3.18	2.45	2.33	3.99
	DAPH	4.56	4.29	5.01	4.64	5.96	6.59	3.29	3.42	3.28	3.18	2.44	2.33	4.08
	PEACE	7.97	11.73	5.54	9.03	31.43	18.44	10.94	31.19	6.95	23.77	4.03	3.68	13.73
	ROSE	23.47	18.64	23.81	17.33	27.94	30.72	11.86	13.25	12.96	11.71	11.75	7.46	17.58
	CPH	28.39	29.54	32.35	31.77	62.73	51.64	28.94	32.14	23.40	21.48	12.64	9.22	30.35
	COUPLE	26.17	24.10	29.47	29.47	60.63	50.85	27.15	29.55	23.59	19.32	11.76	9.28	28.45
	Ours	**46.98**	**38.51**	**52.09**	**37.25**	**72.44**	**60.19**	**50.45**	**47.20**	**43.07**	**33.84**	**19.01**	**15.96**	**43.08**

Table 1: Cross-domain retrieval performance comparison with baselines on the Office-31 and Office-Home benchmarks with distinct noise ratios. The best result is marked in bold, and the second-best result is underlined.

Methods	MNIST → USPS				
	16	32	64	128	Avg.
CPH	58.08	59.74	62.20	63.66	60.92
COUPLE	27.92	29.57	27.56	30.13	28.79
Ours	**66.90**	**62.32**	**69.69**	**68.01**	**66.73**
Methods	USPS → MNIST				
	16	32	64	128	Avg.
CPH	35.62	39.86	44.75	49.15	42.35
COUPLE	27.39	25.79	26.30	30.22	27.42
Ours	**44.72**	**45.78**	**49.43**	**50.14**	**47.52**

Table 2: Comparison of single-domain retrieval mAP performance on the MNIST and USPS with 40% label noise.

varying noise levels, indicating that our RDAH method is more stable in resisting different noise levels.

To comprehensively demonstrate the effectiveness of RDAH in single-domain retrieval, we present the single-domain mAP results on the Digits dataset in Table 2, where the hash code length varies from 16 to 128 bits. It is clear that our RDAH outperforms other methods regardless of M→U and U→M. Moreover, we present the Top-K precision curves and precision-recall curves of COUPLE, CPH, and our RDAH method for single-domain retrieval tasks on the Digits dataset with 64-bit hash codes in Figure 3. The precision-recall curve evaluates the trade-off between precision and recall at all retrieval thresholds, while the Top-K precision reflects the model's ability to prioritize relevant samples. In contrast, our RDAH method achieves greater coverage in the precision-recall curves, fully demonstrating the superior performance of our RDAH method when encountering noise. Similarly, in the Top-K precision curve, as the number of retrieved images increases, our method is able to retrieve more correct images, which further underscores the effectiveness and superiority of our proposed approach.

Ablation Study

To verify the effectiveness of our proposed modules, we investigate six variants of RDAH on Office-31 benchmark with a 0.8 noisy setting, as shown in Table 3. From the quantitative results, we can draw the following conclusions: 1) When the input data lacks noise-robustness processing, the mAP performance drops significantly by nearly 7%. After using the structured noise partitioning and correction strategy, we observe performance improvements across all re-

Warm-up	GNS	NCC	L_{cpa}	L_{sp}	L_q	L_{daa}	A → W	W → A	A → D	D → A	W → D	D → W	Avg.
✓	✓	✓	✓	✓	✓	✓	**46.98**	**38.51**	**52.09**	**37.25**	**72.44**	**60.19**	**51.24**
✓	✓		✓	✓	✓	✓	40.01	36.00	45.24	37.02	66.81	56.99	47.01
✓			✓	✓	✓	✓	34.79	34.54	38.05	36.07	64.23	55.92	43.93
✓	✓	✓		✓	✓	✓	33.51	30.92	31.69	27.44	71.62	57.73	42.15
✓	✓	✓	✓		✓	✓	14.96	10.04	19.57	13.87	52.03	31.55	23.67
✓	✓	✓	✓	✓		✓	40.20	33.77	38.68	26.29	67.05	57.29	43.88
✓	✓	✓	✓	✓	✓		43.03	34.28	41.59	34.92	71.97	58.41	47.37

Table 3: Ablation study on Office-31 dataset with 80% label corruption rate. The best results are marked in bold font.

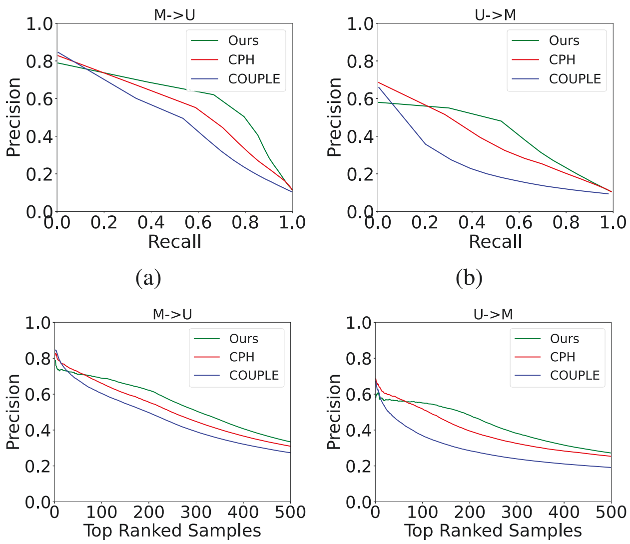

(a) (b) (c) (d)

Figure 3: Comparison of precision-recall and Top-K precision curves of different methods on MNIST and USPS.

(a) COUPLE (b) CPH (c) RDAH

Figure 4: The t-SNE visualization on the Office-31 dataset under 40% label noise for the W→D transfer task.

Figure 5: Comparison of Top-5 retrieved images between CPH and our method on Office-Home. Red boxes indicate incorrectly retrieved images.

sults. 2) The results demonstrate that cross-domain prototype alignment outperforms domain adversarial alignment. However, both have positive effects on cross-domain hashing learning, primarily attributed to the fact that prototype alignment can establish category alignment, while adversarial alignment can establish global alignment to eliminate domain shift. 3) The semantic similarity preservation module plays an important role in maintaining the semantic consistency between features and hash codes.

Feature Distribution Analysis

Figure 4 visualizes discriminative hash codes learned by COUPLE, CPH, and our RDAH during testing. It is convincing that the hash codes learned by RDAH are more discriminative, thus achieving more effective image retrieval. The superior performance of our method can be attributed to a key design aspect: the structural noise modeling and correction strategy can effectively identify and correct noisy samples, thereby learning more reliable semantic alignment information. Figure 5 visualizes the Top-5 retrieved images for the two methods. Our RDAH can retrieve more relevant images. For example, when the query is a laptop, our method retrieves more laptops than computers. This profoundly demonstrates that our method exhibits efficient and accurate hash retrieval capabilities in the presence of noise.

Conclusion

This paper investigates the important yet under-explored problem of domain adaptive hash retrieval under noisy conditions. To address this challenge, we propose a novel robust domain adaptive hash method named RDAH. Specifically, we first introduce structured noise modeling and neighborhood correction, which effectively alleviate the overfitting caused by noisy labels. Second, to tackle the dual challenge of domain shift, we enhance the semantic consistency of hash representations across domains by integrating prototype alignment with adversarial alignment. Finally, extensive experiments on three widely recognized domain adaptation benchmarks demonstrate the superior performance of the proposed RDAH method.

References

Arpit, D.; Jastrzebski, S.; Ballas, N.; Krueger, D.; Bengio, E.; Kanwal, M. S.; Maharaj, T.; Fischer, A.; Courville, A.; Bengio, Y.; et al. 2017. A closer look at memorization in deep networks. In *International conference on machine learning*, 233–242. PMLR.

Cakir, F.; He, K.; Bargal, S. A.; and Sclaroff, S. 2019. Hashing with mutual information. *IEEE transactions on pattern analysis and machine intelligence*, 41(10): 2424–2437.

Chen, Y.; Lai, Z.; Ding, Y.; Lin, K.; and Wong, W. K. 2019. Deep supervised hashing with anchor graph. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9796–9804.

Chen, Z.; Sadhukhan, R.; Ye, Z.; Zhou, Y.; Zhang, J.; Nolte, N.; Tian, Y.; Douze, M.; Bottou, L.; Jia, Z.; et al. 2024. Magicpig: Lsh sampling for efficient llm generation. *arXiv preprint arXiv:2410.16179*.

Cui, H.; Li, F.; Zhu, L.; Li, J.; and Zhang, Z. 2023. Online query expansion hashing for efficient image retrieval. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(3): 1941–1953.

Cui, H.; Zhao, L.; Li, F.; Zhu, L.; Han, X.; and Li, J. 2024. Effective comparative prototype hashing for unsupervised domain adaptation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 8329–8337.

Gong, Y.; Lazebnik, S.; Gordo, A.; and Perronnin, F. 2012. Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. *IEEE transactions on pattern analysis and machine intelligence*, 35(12): 2916–2929.

Hoe, J. T.; Ng, K. W.; Zhang, T.; Chan, C. S.; Song, Y.-Z.; and Xiang, T. 2021. One loss for all: Deep hashing with a single cosine similarity based learning objective. *Advances in Neural Information Processing Systems*, 34: 24286–24298.

Huang, F.; Zhang, L.; and Gao, X. 2021. Domain adaptation preconceived hashing for unconstrained visual retrieval. *IEEE Transactions on Neural Networks and Learning Systems*, 33(10): 5641–5655.

Huang, F.; Zhang, L.; Yang, Y.; and Zhou, X. 2020. Probability weighted compact feature for domain adaptive retrieval. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9582–9591.

Hull, J. J. 2002. A database for handwritten text recognition research. *IEEE Transactions on pattern analysis and machine intelligence*, 16(5): 550–554.

Jang, Y. K.; and Cho, N. I. 2021. Self-supervised product quantization for deep unsupervised image retrieval. In *Proceedings of the IEEE/CVF international conference on computer vision*, 12085–12094.

Jiang, L.; Huang, D.; Liu, M.; and Yang, W. 2020. Beyond synthetic noise: Deep learning on controlled noisy labels. In *International conference on machine learning*, 4804–4815. PMLR.

Jiang, L.; Zhou, Z.; Leung, T.; Li, L.-J.; and Fei-Fei, L. 2018. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *International conference on machine learning*, 2304–2313. PMLR.

Lai, H.; Pan, Y.; Liu, Y.; and Yan, S. 2015. Simultaneous feature learning and hash coding with deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3270–3278.

LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 2002. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.

Li, J.; Zhou, P.; Xiong, C.; and Hoi, S. C. 2020. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*.

Li, Q.; Sun, Z.; He, R.; and Tan, T. 2017. Deep supervised discrete hashing. *Advances in neural information processing systems*, 30.

Li, S.; Xia, X.; Zhang, H.; Zhan, Y.; Ge, S.; and Liu, T. 2022. Estimating noise transition matrix with label correlations for noisy multi-label learning. *Advances in Neural Information Processing Systems*, 35: 24184–24198.

Liu, J.; and Zhang, L. 2019. Optimal projection guided transfer hashing for image retrieval. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 8754–8761.

Luo, J.; Zhang, W.; Yuan, Y.; Zhao, Y.; Yang, J.; Gu, Y.; Wu, B.; Chen, B.; Qiao, Z.; Long, Q.; et al. 2025a. Large language model agent: A survey on methodology, applications and challenges. *arXiv preprint arXiv:2503.21460*.

Luo, J.; Zhao, Y.; Luo, X.; Xiao, Z.; Ju, W.; Shen, L.; Tao, D.; and Zhang, M. 2025b. Cross-domain diffusion with progressive alignment for efficient adaptive retrieval. *IEEE Transactions on Image Processing*.

Luo, X.; Wu, D.; Ma, Z.; Chen, C.; Deng, M.; Ma, J.; Jin, Z.; Huang, J.; and Hua, X.-S. 2020. Cimon: Towards high-quality hash codes. *arXiv preprint arXiv:2010.07804*.

Qiu, Z.; Su, Q.; Ou, Z.; Yu, J.; and Chen, C. 2021. Unsupervised hashing with contrastive information bottleneck. *arXiv preprint arXiv:2105.06138*.

Ren, M.; Zeng, W.; Yang, B.; and Urtasun, R. 2018. Learning to reweight examples for robust deep learning. In *International conference on machine learning*, 4334–4343. PMLR.

Saenko, K.; Kulis, B.; Fritz, M.; and Darrell, T. 2010. Adapting visual category models to new domains. In *European conference on computer vision*, 213–226. Springer.

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.

Venkateswara, H.; Eusebio, J.; Chakraborty, S.; and Panchanathan, S. 2017. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5018–5027.

Wang, H.; Sun, J.; Luo, X.; Xiang, W.; Zhang, S.; Chen, C.; and Hua, X.-S. 2023a. Toward effective domain adaptive retrieval. *IEEE Transactions on Image Processing*, 32: 1285–1299.

Wang, H.; Sun, J.; Wei, X.; Zhang, S.; Chen, C.; Hua, X.-S.; and Luo, X. 2023b. Dance: Learning a domain adaptive framework for deep hashing. In *Proceedings of the ACM Web Conference 2023*, 3319–3330.

Wang, H.; Wu, H.; Sun, J.; Zhang, S.; Chen, C.; Hua, X.-S.; and Luo, X. 2023c. Idea: An invariant perspective for efficient domain adaptive image retrieval. *Advances in Neural Information Processing Systems*, 36: 57256–57275.

Wang, Y.; Liu, W.; Ma, X.; Bailey, J.; Zha, H.; Song, L.; and Xia, S.-T. 2018. Iterative learning with open-set noisy labels. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8688–8696.

Wu, X.; Loveland, D.; Chen, R.; Liu, Y.; Chen, X.; Neves, L.; Jadbabaie, A.; Ju, M.; Shah, N.; and Zhao, T. 2025. GraphHash: Graph Clustering Enables Parameter Efficiency in Recommender Systems. In *Proceedings of the ACM on Web Conference 2025*, 357–369.

Yang, X.; Wang, H.; Sun, J.; Xiao, Y.; Wei, X.; Chen, C.; Hua, X.-S.; and Luo, X. 2024. ROSE: Relational and Prototypical Structure Learning for Universal Domain Adaptive Hashing. *IEEE Transactions on Information Forensics and Security*.

Yao, Y.; Liu, T.; Han, B.; Gong, M.; Deng, J.; Niu, G.; and Sugiyama, M. 2020. Dual t: Reducing estimation error for transition matrix in label-noise learning. *Advances in neural information processing systems*, 33: 7260–7271.

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, C.; Bengio, S.; Hardt, M.; Recht, B.; and Vinyals, O. 2021. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3): 107–115.

Zhang, W.; Wu, D.; Zhou, Y.; Li, B.; Wang, W.; and Meng, D. 2020. Deep unsupervised hybrid-similarity hadamard hashing. In *Proceedings of the 28th ACM international conference on multimedia*, 3274–3282.

Zhang, Y.; Hu, Y.; Cai, C.; Huang, Y.-A.; Huang, Z.-A.; and Tan, K. C. 2025. Anti-Confounding Hashing: Enhancing Radiological Image Retrieval via Debiased Weighting and Counterfactual Reasoning. *IEEE Transactions on Neural Networks and Learning Systems*.

Zhang, Z.; and Sabuncu, M. 2018. Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in neural information processing systems*, 31.

Zhao, P.; Zhang, H.; Yu, Q.; Wang, Z.; Geng, Y.; Fu, F.; Yang, L.; Zhang, W.; Jiang, J.; and Cui, B. 2024. Retrieval-augmented generation for ai-generated content: A survey. *arXiv preprint arXiv:2402.19473*.