

Prototype-Based Semantic Consistency Alignment for Domain Adaptive Retrieval

Tianle Hu¹, Weijun Lv², Na Han³, Xiaozhao Fang^{2*}, Jie Wen⁴, Jiaying Li⁵, Guoxu Zhou²

¹School of Computer Science and Technology, Guangdong University of Technology, Guangzhou 510006, China

²School of Automation, Guangdong University of Technology, Guangzhou 510006, China

³School of Computer Science, Guangdong Polytechnic Normal University, Guangzhou 510665, China

⁴School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen 518055, China

⁵School of Artificial Intelligence, Guangdong University, Guangzhou 510006, China

{hutianlegdut, lvweijun0201}@163.com, {hannagdut, xzhfang168, jiewen_pr}@126.com, jiaying.li.cs@gmail.com, gx.zhou@gdut.edu.cn

Abstract

Domain adaptive retrieval aims to transfer knowledge from a labeled source domain to an unlabeled target domain, enabling effective retrieval while mitigating domain discrepancies. However, existing methods encounter several fundamental limitations: 1) neglecting class-level semantic alignment and excessively pursuing pair-wise sample alignment; 2) lacking either pseudo-label reliability consideration or geometric guidance for assessing label correctness; 3) directly quantizing original features affected by domain shift, undermining the quality of learned hash codes. In view of these limitations, we propose Prototype-Based Semantic Consistency Alignment (PSCA), a two-stage framework for effective domain adaptive retrieval. In the first stage, a set of orthogonal prototypes directly establishes class-level semantic connections, maximizing inter-class separability while gathering intra-class samples. During the prototype learning, geometric proximity provides a reliability indicator for semantic consistency alignment through adaptive weighting of pseudo-label confidences. The resulting membership matrix and prototypes facilitate feature reconstruction, ensuring quantization on reconstructed rather than original features, thereby improving subsequent hash coding quality and seamlessly connecting both stages. In the second stage, domain-specific quantization functions process the reconstructed features under mutual approximation constraints, generating unified binary hash codes across domains. Extensive experiments validate PSCA's superior performance across multiple datasets.

Introduction

Hashing receives extensive attention in the field of image retrieval due to its merits of compact storage and computational efficiency. The main purpose of hashing is to develop effective hash functions that preserve similarity relationships of original data in binary Hamming space. Several methods, such as Spectral hashing (SH) (Weiss, Torralba, and Fergus 2008), Density Sensitive Hashing (DSH) (Liu et al. 2016) and Scalable Graph Hashing (SGH) (Jiang and Li 2015) endeavor to preserve pair-wise similarity of original data within the Hamming space. Ordinal Constraint

Hashing (OCH) (Liu et al. 2018) introduces the ordinal relation in learning to hash. Iterative Quantization (ITQ) (Gong et al. 2012) focuses on maintaining the locality structure by improving the consistency between generated discrete codes and their corresponding continuous representations.

Nonetheless, these aforementioned methods assume that queries and retrieved images share consistent data distributions, limiting their applicability in complex real-world scenarios. For instance, online shopping platforms showcase product images shot under ideal conditions, whereas user-submitted query photos typically contain cluttered backgrounds. To bridge this non-negligible domain gap (Hu et al. 2025a), Domain Adaptation (DA) (Zhang et al. 2023c,b) is combined with hashing, giving rise to a promising research field, i.e., Domain Adaptive Retrieval (DAR).

DAR encompasses two retrieval scenarios, i.e., single-domain retrieval and cross-domain retrieval. The former supposes both queries and retrieved samples originate from the target domain. In the context of cross-domain retrieval, the source domain serves as the retrieved dataset while queries stem from the target domain. Recently, several DAR methods are proposed. Probability Weighted Compact Feature (PWCF) (Huang et al. 2020) utilizes a focal-triplet constraint to mitigate the domain gap in a domain-invariant subspace. Domain Adaptation Preconceived Hashing (DAPH) (Huang, Zhang, and Gao 2021) introduces Maximum Mean Discrepancy (MMD) (Gretton et al. 2012) to prompt the domain marginal distribution alignment. These geometry-oriented methods lack consideration of semantic relationships between features and labels, resulting in suboptimal performance when significant semantic misalignment exists. Consequently, subsequent methods shift their focus toward incorporating semantic guidance. Two-Step Strategy (TSS) (Chen et al. 2023) proposes a discriminative semantic fusion for hash learning. Semantic Guided Hashing Learning (SGHL) (Zhang et al. 2023a) and Dynamic Confidence Sampling and Label Semantic Guidance (DCS-LSG) (Zhang et al. 2024) further align the cross-domain conditional distributions by integrating category labels.

Despite their promising performance, we identify certain critical limitations of current DAR methods: 1) excessive focus on pair-wise sample alignment. Specifically, PWCF,

*Corresponding author

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

TSS, SGHL and DCS-LSG primarily minimize distribution discrepancies between semantically consistent sample pairs, which suffer from computational inefficiency and limited distributional coverage of data (Yuan et al. 2020). 2) inadequate handling of pseudo-label reliability. Pseudo-labeling serves to predict the latent semantic associations between classes and unlabeled data, consequently providing fully annotated data to facilitate knowledge transfer. However, existing methods typically adopt off-the-shelf strategies, neglecting correction mechanisms for erroneous predictions. This inevitably leads to biased domain alignment and degraded hash codes quality. Although the most recent method, DCS-LSG, considers pseudo-label noise, it relies solely on semantic consensus between dual projections, without incorporating geometric knowledge for reliability assessment. 3) directly mapping original features with imperfect domain alignment to Hamming space, resulting in high quantization errors and limited discriminative power of generated codes.

To systematically tackle the limitations mentioned above, we propose the Prototype-Based Semantic Consistency Alignment (PSCA) framework. The core innovation lies in a semantic consistency alignment that evaluates pseudo-label reliability by comparing geometric proximity with semantic predictions, adaptively weighting the pseudo-labels. To be precise, in the first stage, PSCA establishes orthogonal class prototypes within a domain-shared subspace, where the semantic consistency alignment performs as follows: when geometric and semantic indicators agree, semantic weights are adjusted based on decision margins, as larger margins reflect stronger prediction confidence; when they conflict, semantic contribution is reduced proportionally to the disagreement level. This process derives a soft membership matrix that guides prototype learning in turn, thereby capturing more reliable semantic connections that mitigate error propagation.

After stage one, the membership matrix and prototypes reconstruct enhanced features, ensuring superior coding quality by circumventing direct quantization strategies. In the second stage, domain-specific quantization functions quantize the reconstructed features under mutual approximation constraints, capturing domain-specific characteristics while facilitating unified hash learning. Figure 1 illustrates the PSCA framework, and the primary contributions are:

- An orthogonal prototype learning method is proposed that achieves effective class-level semantic alignment instead of pursuing pair-wise sample alignment.
- A semantic consistency alignment is designed to dynamically correct unreliable pseudo-labels by combining geometric proximity with semantic predictions, addressing the semantic error accumulation problem.
- A feature reconstruction strategy leverages discriminative prototypes and membership matrix to create enhanced feature representations, ensuring hash quantization on reliable rather than noisy information.
- Comprehensive experiments demonstrate that PSCA outperforms the existing state-of-the-art DAR methods.

Related Work

Hash Learning

Hashing seeks to encode high-dimensional data into low-dimensional binary codes with maintaining similarity relationships, offering the merits of low storage demand and high retrieval efficiency (Hu et al. 2025b). Hashing can be divided into two categories according to whether the semantic information is available, i.e., unsupervised hashing and supervised hashing. Supervised methods like DCMVH (Zhu et al. 2020) and MSDH (Luo et al. 2019) have achieved noteworthy performance. Nevertheless, the significant cost related to annotating labels brings a major obstacle to their scalability. By contrast, benefiting from no need for explicit semantics, unsupervised hashing methods serve as the more suitable retrieval manner for real-world applications. Some representative methods among this category include SGH (Jiang and Li 2015) and GraphBit (Wang et al. 2023b). However, unsupervised algorithms may pose a potential risk of limiting the discriminative power of hash codes due to the lack of explicit semantics. While admiring the positive performance of these methods, they are limited for DAR scenarios due to their assumption of retrieving in a single domain.

Domain Adaptive Retrieval

DAR assumes that there exists a well-labeled source domain and an unlabeled target domain, where these two domains are related yet different in distributions. It aims to mitigate the domain discrepancies and preserving semantic similarity in the learned hash codes, thereby achieving effective retrieval. PWCF (Huang et al. 2020) and DAPH (Huang, Zhang, and Gao 2021) geometrically constrain the distribution of samples to bridge domain gap. Since two domains share a common semantic space, TSS, SGHL and DCS-LSG further align pseudo-labels to target domain samples to encourage the transfer of domain-specific valuable insights. Note that aforementioned methods and the proposed PSCA are on the basis of machine learning principles.

With the rise of deep learning, several advanced deep learning-based DAR methods have emerged recently. PEACE (Wang et al. 2023a) considers the uncertainty of pseudo-labels and progressively boosts their reliability. CPH (Cui et al. 2024) innovatively constructs a domain-shared unit subspace, then aligns domains through prototype contrastive learning. COUPLE (Luo et al. 2025) employs graph flow diffusion to simulate the cross-domain knowledge transfer, dynamically identifying lower noise clusters.

Proposed Method

Problem Definition

Assume we have a source domain $\mathcal{D}_s = \{\mathbf{x}_{s_i}, \mathbf{y}_{s_i}\}_{i=1}^{n_s}$ comprising n_s samples, a target domain $\mathcal{D}_t = \{\mathbf{x}_{t_i}\}_{i=1}^{n_t}$ including n_t samples. Let $\mathbf{X}_s \in \mathbb{R}^{d \times n_s}$ and $\mathbf{X}_t \in \mathbb{R}^{d \times n_t}$ as the original sample matrices of two domains respectively, where d denotes the feature dimension. We define $\mathbf{X} = [\mathbf{X}_s, \mathbf{X}_t] \in \mathbb{R}^{d \times n}$ as the total sample matrix, where $n = n_s + n_t$. For each target sample \mathbf{x}_{t_i} , we obtain pseudo-label probabilities $\boldsymbol{\pi}_i = [\pi_{i1}, \pi_{i2}, \dots, \pi_{ic}] \in \mathbb{R}^{1 \times c}$ using

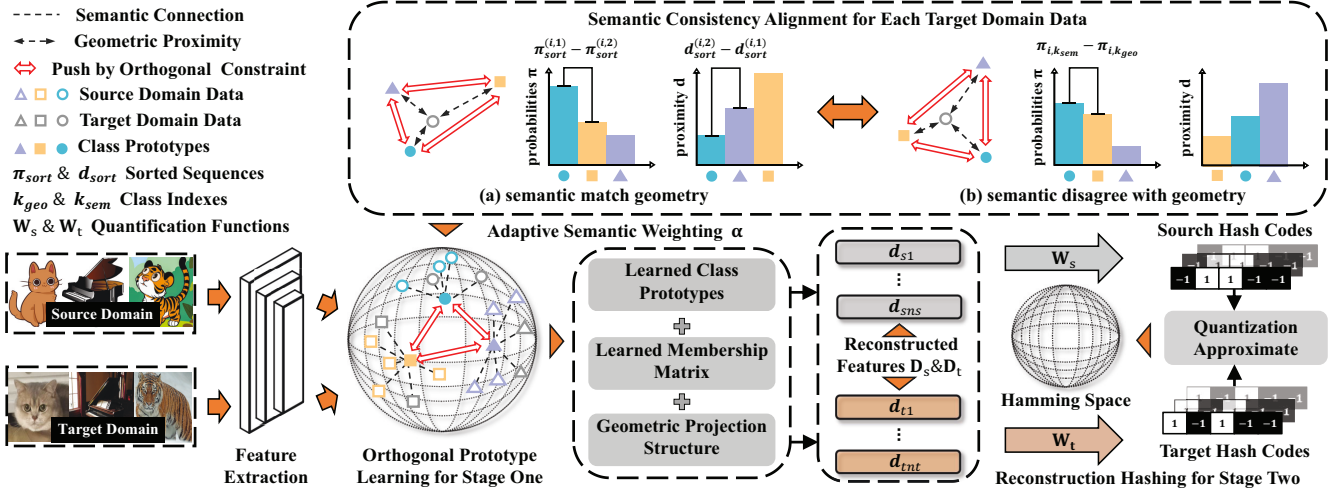


Figure 1: The framework of the proposed PSCA

Nearest Class Prototype (NCP) and Structured Prediction (SP) (Wang and Breckon 2020), where c is the number of categories. Let $\pi_{\text{sort}}^{(i,e)}$ be the e -th largest element in π_i , where $\pi_{\text{sort}}^{(i,1)} \geq \pi_{\text{sort}}^{(i,2)} \geq \dots \geq \pi_{\text{sort}}^{(i,c)}$. Then the pseudo-label is assigned according to $\pi_{\text{sort}}^{(i,1)}$. The detailed pseudo-labeling process is presented in the supplementary materials. Define the overall label set as $\mathbf{Y} = [\mathbf{Y}_s, \hat{\mathbf{Y}}_t]^\top \in \{0, 1\}^{n \times c}$, where $\hat{\mathbf{Y}}_t \in \mathbb{R}^{n_t \times c}$ are the pseudo-labels. The ultimate objective of DAR is to learn a set of similarity-preserving hash codes $\mathbf{B} = [\mathbf{b}_{s_1}, \dots, \mathbf{b}_{s_{n_s}}, \mathbf{b}_{t_1}, \dots, \mathbf{b}_{t_{n_t}}] \in \{-1, 1\}^{r \times n}$ for effective retrieval, where r represents the hash code length.

Prototype-Based Semantic Consistency Alignment

Due to domain shift, cross-domain samples with similar semantics may exhibit significantly distinct distributions. To bridge this gap, we first employ Maximum Mean Discrepancy (MMD) to align domain marginal distributions:

$$\min_{\mathbf{P}} \left\| \frac{1}{n_s} \sum_{i=1}^{n_s} \mathbf{P}^\top \mathbf{x}_i - \frac{1}{n_t} \sum_{q=n_s+1}^{n_s+n_t} \mathbf{P}^\top \mathbf{x}_q \right\|_2^2 \quad (1)$$

$$= \text{Tr}(\mathbf{P}^\top \mathbf{X} \mathbf{H} \mathbf{X}^\top \mathbf{P})$$

where $\mathbf{P} \in \mathbb{R}^{d \times q}$ is the projection matrix that maps original features into a common q -dimensional subspace ($q \ll d$), where the domain gap is expected to be bridged. \mathbf{H} is the MMD centering matrix defined as:

$$h_{ij} = \begin{cases} \frac{1}{n_s n_s}, & \mathbf{x}_i, \mathbf{x}_j \in \mathcal{D}_s \\ \frac{-1}{n_s n_t}, & \mathbf{x}_i \in \mathcal{D}_s \wedge \mathbf{x}_j \in \mathcal{D}_t \\ \frac{1}{n_t n_t}, & \mathbf{x}_i, \mathbf{x}_j \in \mathcal{D}_t \end{cases} \quad (2)$$

Apart from marginal discrepancies, semantic structure within data distributions requires consideration. Existing

methods achieve this by aligning conditional distributions or semantically consistent sample pairs, yet suffer from computational inefficiency and sensitivity to outliers.

To circumvent these issues, we introduce a prototype-based approach that learns discriminative class centers $\mathbf{O} = [\mathbf{o}_1, \dots, \mathbf{o}_c] \in \mathbb{R}^{q \times c}$, directly modeling class-level semantic connections to facilitate effective class alignment:

$$\min_{\mathbf{P}, \mathbf{O}} \sum_{i=1}^n \sum_{j=1}^c y_{ij} \|\mathbf{P}^\top \mathbf{x}_i - \mathbf{o}_j\|_2^2, \quad \text{s.t. } \mathbf{O}^\top \mathbf{O} = \mathbf{I}_c, \quad (3)$$

where \mathbf{I}_c represents a c -dimensional identity matrix. Ideally, \mathbf{O} can effectively reduce the intra-class divergence by gathering samples with identical semantics. Meanwhile, $\mathbf{O}^\top \mathbf{O} = \mathbf{I}_c$ ensures that different prototypes are maximally separated by forcing the inner product between distinct prototypes to be zero. Nevertheless, Eq. (3) exhibits two critical limitations: 1) incorrect $\hat{\mathbf{y}}_t$ shifts these domain-shared prototypes away from their latent ground-truth positions; 2) binary pseudo-labels fail to convey prediction confidence, hampering assessment of assignment reliability.

Addressing the above challenges, we propose a semantic consistency alignment that tackles them simultaneously. Since samples sharing the identical prototype have consistent labels, samples closer to prototypes are inherently more reliable for prediction than those at cluster boundaries. Therefore, $\mathbf{R} \in \mathbb{R}^{n_t \times c}$ is designed as a soft membership matrix, where each element r_{ij} provides a more reliable membership degree than \hat{y}_{ij} . This is achieved by incorporating the geometric proximity $d_{ij} = \|\mathbf{P}^\top \mathbf{x}_i - \mathbf{o}_j\|_2^2$. In particular, \mathbf{R} is optimized via the following objective:

$$\min_{\mathbf{R} \geq 0, \mathbf{R} \mathbf{1}_c = \mathbf{1}_{n_t}} \sum_{i=1}^{n_t} \sum_{j=1}^c \{r_{ij}^\sigma d_{ij} - \psi_{ij} \log(r_{ij})\} \quad (4)$$

where $\mathbf{1}$ denotes all-ones vector with dimension indicated by its subscript. In the first term, r_{ij} measures geometric proximity of projected target samples to prototypes.

The coefficient $\sigma > 1$ amplifies the penalty differences as $r_{ij}^\sigma < r_{ij}$, creating non-linear down-weighting. The second term serves as a semantic-aware term where ψ are elements of an adaptive weighting matrix $\Psi = \alpha \odot \hat{\mathbf{Y}}_t$, with $\alpha = [\alpha_1, \dots, \alpha_{n_t}]^\top \in \mathbb{R}^{n_t \times 1}$ broadcasted to match dimensions. This term activates only when $\hat{y}_{ij} = 1$, prompting larger r_{ij} values for the semantically predicted class. Here, α_i adaptively adjusts the semantic fusion strength by:

$$\alpha_i = \begin{cases} \frac{\pi_{\text{sort}}^{(i,1)} - \pi_{\text{sort}}^{(i,2)}}{d_{\text{sort}}^{(i,2)} - d_{\text{sort}}^{(i,1)} + \text{eps}}, & \text{if } k_{\text{geo}} = k_{\text{sem}} \\ \pi_{\text{sort}}^{(i,1)} (1 - |\pi_{i,k_{\text{geo}}} - \pi_{i,k_{\text{sem}}}|), & \text{otherwise} \end{cases} \quad (5)$$

where eps prevents the situation of division by zero. Let $\mathbf{d}_i = [d_{i1}, d_{i2}, \dots, d_{ic}] \in \mathbb{R}^{1 \times c}$ denote the distances from \mathbf{x}_{t_i} to all prototypes, we have the sorted sequence as $d_{\text{sort}}^{(i,1)} \leq \dots \leq d_{\text{sort}}^{(i,c)}$, where $d_{\text{sort}}^{(i,1)}$ indicates the smallest value in \mathbf{d}_i . $k_{\text{geo}} = \arg \min_j d_{ij}$ and $k_{\text{sem}} = \arg \max_j \pi_{ij}$ denote the indexes of geometrically closest and semantically preferred prototypes. When geometric and semantic indicators are consistent ($k_{\text{geo}} = k_{\text{sem}}$), larger semantic margins encourage greater reliance on semantic information by increasing α_i , while larger geometric margins favor geometric knowledge by decreasing α_i . In the conflicting case, α_i is reduced based on the disagreement intensity, i.e., $|\pi_{i,k_{\text{geo}}} - \pi_{i,k_{\text{sem}}}|$, thereby decreasing reliance on potentially erroneous assignments.

With the optimized \mathbf{R} from Eq. (4), we form the unified semantic matrix as $\hat{\mathbf{Y}} = [\mathbf{Y}_s, \mathbf{R}]^\top$. Integrating all key components above, stage one is formulated as:

$$\min_{\mathbf{P}, \mathbf{O}^\top \mathbf{O} = \mathbf{I}_c} \sum_{i=1}^n \sum_{j=1}^c \tilde{y}_{ij} \|\mathbf{P}^\top \mathbf{x}_i - \mathbf{o}_j\|_2^2 + \lambda_1 \text{Tr}(\mathbf{P}^\top \mathbf{X} \mathbf{H} \mathbf{X}^\top \mathbf{P}) + \lambda_2 \|\mathbf{P}\|_{2,1} \quad (6)$$

where λ_1 and λ_2 are the trade-off parameters. Here, $\|\mathbf{P}\|_{2,1}$ serves as a $\ell_{2,1}$ -norm with row-sparsity constraint.

Feature Reconstruction Hashing

Given that projection \mathbf{P} focuses on domain alignment, directly quantizing $\mathbf{P}^\top \mathbf{X}$ neglects the semantic discriminability established by prototype learning, inevitably reducing the hash coding quality. Thus, as an intermediate operation to avoid prematurely using features lacking semantic enhancement and seamlessly bridge both stages, we reconstruct semantically enhanced features by leveraging the reliable prototypes \mathbf{O} and learned memberships \mathbf{R} from stage one.

Since each row $\mathbf{r}_i = [r_{i1}, r_{i2}, \dots, r_{ic}]$ in \mathbf{R} indicates the confidences of \mathbf{x}_{t_i} belonging to each category, and the learned class prototypes \mathbf{O} exhibit stronger discriminability than the original data due to the orthogonal constraint and semantic consistency alignment. Consequently, we reconstruct target samples \mathbf{x}_{t_i} as $\tilde{\mathbf{x}}_{t_i} = \sum_m^c r_{im} \mathbf{o}_m^\top$. This confidence-weighted combination of prototypes encodes semantically enhanced representations. Meanwhile, the reliable labels are embedded into the source reconstruction, i.e., $\tilde{\mathbf{x}}_{s_i} = \sum_m^c y_{im} \mathbf{o}_m^\top$. This yields the novel representation $\tilde{\mathbf{X}} = [\tilde{\mathbf{x}}_{s_1}, \dots, \tilde{\mathbf{x}}_{s_{n_s}}, \tilde{\mathbf{x}}_{t_1}, \dots, \tilde{\mathbf{x}}_{t_{n_t}}]^\top \in \mathbb{R}^{n \times q}$ that better reveals the underlying ground-truth semantic structure.

However, while $\tilde{\mathbf{X}}$ excels in semantic discriminability, it may lose certain geometric structures of the original data, making sole reliance on $\tilde{\mathbf{X}}$ insufficient for comprehensive representation during hash coding. Hence, we fuse $\tilde{\mathbf{X}}$ and the projected features $\mathbf{P}^\top \mathbf{X}$. Let $\mathcal{C} = 2q$, the overall reconstructed features are expressed as $\mathbf{D} = [\mathbf{D}_s, \mathbf{D}_t] = [\mathbf{d}_{s1}, \dots, \mathbf{d}_{s_{n_s}}, \mathbf{d}_{t1}, \dots, \mathbf{d}_{t_{n_t}}] = [\tilde{\mathbf{X}}, \mathbf{X}^\top \mathbf{P}]^\top \in \mathbb{R}^{\mathcal{C} \times n}$. This concatenation both retains the reconstructed semantics and geometric projection structures. Building upon the reconstructed features \mathbf{D} , we proceed to develop discriminative hash codes in stage two. Towards further mitigating the adverse effects of pseudo-labeling and domain discrepancies, meanwhile for better transferring domain-specific information, we design distinct functions for both domains, i.e., $\mathbf{W}_s \in \mathbb{R}^{r \times \mathcal{C}}$ and $\mathbf{W}_t \in \mathbb{R}^{r \times \mathcal{C}}$. Given our objective is to learn a unified Hamming space instead of two separate domain-specific hash codes. Hence, the dual functions are required to be approximated during hash learning. Ultimately, the objective of stage two is given below:

$$\min_{\mathbf{W}_s, \mathbf{W}_t, \mathbf{B}_s, \mathbf{B}_t} \|\mathbf{W}_s \mathbf{D}_s - \mathbf{B}_s\|_F^2 + \|\mathbf{W}_t \mathbf{D}_t - \mathbf{B}_t\|_F^2 + \lambda_3 \|\mathbf{W}_s - \mathbf{W}_t\|_F^2 \quad (7)$$

s.t. $\mathbf{W}_s \mathbf{W}_s^\top = \mathbf{I}_r, \mathbf{W}_t \mathbf{W}_t^\top = \mathbf{I}_r, \mathbf{B}_s, \mathbf{B}_t \in \{-1, 1\}^{r \times *}$,

where the orthogonal constraints ensure independence hash bit encoding. λ_3 is a trade-off parameter and $* \in \{n_s, n_t\}$.

Optimization

Solution Process

As shown in Eqs. (6) and (7), there exist multiple variables that require to be solved, i.e., $\mathbf{P}, \mathbf{O}, \mathbf{R}, \mathbf{W}_s, \mathbf{W}_t, \mathbf{B}_s$ and \mathbf{B}_t . Due to space limitations, the detailed solution process of PSCA is presented in the supplementary materials. Meanwhile, an algorithm analysis subsection is also provided, including the computational complexity, the convergence analysis and running time comparison.

Out-of-Sample Extension

To facilitate generalization of unseen samples, a linear transformation matrix Φ is derived to model the regression mapping between features and learned binary hash codes:

$$\min_{\Phi} \|\Phi \mathbf{X} - \mathbf{B}\|_F^2 + \beta \|\Phi\|_F^2 \quad (8)$$

where $\|\Phi\|_F^2$ denotes a regularization term and β serves as a balancing parameter. Eq. (8) can be easily solved with:

$$\Phi = \mathbf{B} \mathbf{X}^\top (\mathbf{X} \mathbf{X}^\top + \beta \mathbf{I})^{-1} \quad (9)$$

Thereafter, the corresponding binary hash code of any query sample can be obtained by $\mathbf{b}_{\text{query}} = \text{sgn}(\Phi \mathbf{x}_{\text{query}})$.

Experiment

Dataset and Evaluation Metric

We carry out extensive comparative experiments on four public benchmark datasets, namely Office-31 (Saenko et al. 2010), Office-Home (Venkateswara et al. 2017), COIL20

Case	MNIST→USPS				COIL1→COIL2				A→D				A→W			
Code length	16	32	64	128	16	32	64	128	16	32	64	128	16	32	64	128
SH	15.56	13.67	13.54	12.95	40.18	44.64	42.84	38.36	14.08	13.62	12.02	10.91	12.04	11.97	9.83	9.90
DSH	20.60	22.21	24.28	26.50	37.92	44.85	46.38	46.05	11.48	13.86	16.66	19.88	9.58	12.14	15.09	18.05
SGH	14.24	16.69	19.70	21.95	51.04	52.31	51.77	50.53	19.92	21.19	24.86	27.50	16.95	20.13	22.47	25.44
OCH	13.73	17.22	20.18	23.34	46.50	50.67	55.25	56.54	14.29	20.43	24.86	27.50	14.85	20.24	22.49	25.86
ITQ+	22.84	21.20	19.15	18.52	46.68	50.49	50.80	50.63	15.42	16.74	17.99	16.59	14.94	16.19	15.00	15.21
LapITQ+	24.26	24.03	24.59	22.73	44.44	39.26	34.20	29.07	17.53	19.38	19.96	18.12	15.10	17.80	18.24	16.36
GTH-h	19.47	16.52	15.33	16.46	44.89	50.47	52.83	52.93	14.42	20.38	23.13	24.08	13.55	19.89	22.21	23.32
PWCF	43.90	50.94	52.51	53.17	65.29	64.98	67.34	67.00	24.53	29.57	32.46	34.55	21.98	32.38	34.14	35.21
DAPH*	30.22	35.14	38.18	39.36	75.77	77.58	78.49	81.96	29.14	29.69	28.17	27.83	21.22	25.36	26.85	28.74
SGHL	62.95	65.93	69.52	71.46	70.81	78.71	80.59	83.00	43.92	51.55	54.89	59.91	44.49	52.48	55.31	55.64
TSS	<u>64.19</u>	<u>69.11</u>	<u>72.59</u>	<u>73.88</u>	78.07	82.08	<u>85.27</u>	<u>87.55</u>	17.83	33.41	44.86	45.23	25.59	39.77	48.97	53.23
DCS-LSG	48.83	53.31	54.22	59.88	82.44	83.36	84.38	85.70	55.13	58.82	63.32	64.59	48.89	53.90	56.33	57.13
Ours	86.05	86.47	87.35	88.71	84.74	87.36	88.79	90.76	56.44	65.51	68.85	67.41	56.72	60.86	62.13	65.78

Table 1: Cross-domain retrieval performance (MAP%) on MNIST→USPS, COIL1→COIL2, A→D and A→W with varying code lengths. The bolded figures indicate the highest scores, and underlined figures indicate the second-highest scores.

(Nene et al. 1996), and MNIST-USPS (LeCun et al. 1998; Hull 1994). **Office-31** contains 4,110 images shot in office environments across three domains: Amazon (A), Webcam (W), and DSLR (D). For experiments, we establish two transferable retrieval cases: A→D and A→W. **Office-Home** consists of images across 65 categories found in office and domestic environments, categorized into four domains: the Artistic (Ar) with 2,427 samples, the Clipart (Cl) with 4,365 samples, the Product (Pr) with 4,439 samples, and Real-world (Rw) with 4,357 samples. Consistent with previous baselines, we select six domain permutation cases: R→P, R→C, A→R, P→R, C→R, and R→A. **COIL20** includes 1,440 images of 20 different objects. COIL1 and COIL2 are two subsets, each containing images captured from diverse angles. According to (Long et al. 2014), a 1,024-dimensional feature vector is extracted for per image, then we construct a retrieval case: COIL1→COIL2. **MNIST-USPS** are two handwritten datasets including digit images from 0 to 9. Following (Long et al. 2013), we resize MNIST to 16×16 pixels and create MNIST-USPS dataset by picking 2,000 MNIST and 1,800 USPS images. We create a retrieval case for experiments: MNIST→USPS.

Random 10% of target domain samples are selected as the testing set, while the remaining 90%, along with entire source samples, form the training set. For cross-domain retrieval, queries are expected to match the most similar images in source domain, whereas for single-domain retrieval, target domain is regarded as the retrieved database.

The mean Average Precision (MAP), Top-K Precision Curve, and Precision-Recall Curve are used to measure the hash coding quality. Note that higher values denote better performance for all evaluation metrics. Each trial is repeated 10 times, and we report the average MAP scores (%).

Baseline and Implementation Detail

We select several state-of-the-art methods as baselines for comparison with PSCA: SH (Weiss, Torralba, and Fergus 2008), DSH (Liu et al. 2016), SGH (Jiang and Li 2015), ITQ+ (Zhou et al. 2018), LapITQ+ (Zhou et al. 2018), GTH-h (Zhang et al. 2019), PWCF (Huang et al. 2020), DAPH* (Huang, Zhang, and Gao 2021), SGHL (Zhang et al. 2023a),

TSS (Chen et al. 2023), and DCS-LSG (Zhang et al. 2024). Where SH, DSH, and SGH belong to traditional hashing. ITQ+, LapITQ+, and GTH-h are transfer hashing. PWCF, DAPH*, SGHL, TSS, and DCS-LSG are hashing methods which aim at dealing with cross-domain scenarios. Note that DAPH* is the supervised variant of DAPH. ITQ+ and LapITQ+ achieve effective retrieval only when the source domain contains more samples than the target domain. Thus in some cases, their MAP scores are not reported. Additionally, we further compare with three advanced deep DAR baselines, i.e., PEACE (Wang et al. 2023a), CPH (Cui et al. 2024) and COUPLE (Luo et al. 2025).

The proposed PSCA includes three independent parameters, i.e., λ_1 to λ_3 , each of them controls the penalty weight of different objectives. We empirically explore the optimal combination of them by fixing one and adjusting others: within the range of [1, 10, 100] for λ_1 and λ_3 , range [0.1, 1, 10] for λ_2 . The specific parameters depend on the characteristics of datasets. β is fixed as 0.1 and σ is set to 2.

Experimental Analysis on Cross-Domain Retrieval

To evaluate PSCA’s cross-domain retrieval performance, we conduct comparison experiments with all baselines across varying code lengths. Tables 1 and 2 report the MAP results. By analyzing these tables, we derive the following observations: comprehensive experiments across four datasets demonstrate that PSCA consistently outperforms all baselines across different code lengths. The presented average MAP of PSCA exceeds that of the second-best baselines by 17.21%, 3.94%, 4.08% and 7.33% on cases MNIST→USPS, COIL1→COIL2, A→D and A→W respectively. On Office-Home, PSCA achieves a remarkable performance improvement by 8.82% on average across six retrieval cases. Based on the above analyses, PSCA exhibits consistent competitive advantages over other baselines. We conclude that PSCA performs well in handling cross-domain issues, whether on small-scale (MNIST-USPS, COIL20), medium-scale (Office-31), or large-scale datasets (Office-Home). To further evaluate the performance on cross-domain retrieval, we plot the Top-K Precision and Precision-Recall curves of PSCA and three most recent competi-

Case	P→R			C→R			R→A			R→P			R→C			A→R		
	16	64	128	16	64	128	16	64	128	16	64	128	16	64	128	16	64	128
SH	10.96	15.03	14.08	6.27	8.77	7.97	9.47	12.87	11.62	11.37	16.13	15.08	5.75	8.24	7.68	10.28	13.71	12.30
DSH	5.61	8.49	9.79	3.57	5.47	6.55	5.43	9.67	10.54	5.70	8.26	10.20	3.62	5.28	6.29	5.95	9.69	11.52
SGH	16.68	24.51	26.38	7.22	13.62	14.82	11.92	22.53	24.69	15.85	25.73	27.89	7.05	13.51	14.83	13.32	22.93	25.14
OCH	11.52	18.65	20.98	6.15	10.27	11.21	9.45	17.54	19.81	11.18	20.15	22.27	5.95	10.05	11.46	10.30	18.09	20.65
ITQ+	11.25	17.61	17.74	6.58	9.55	9.34	9.41	14.25	15.53	-	-	-	-	-	-	-	-	-
LapITQ+	11.99	16.89	16.02	7.27	10.37	10.87	8.89	13.56	13.75	-	-	-	-	-	-	-	-	-
GTH-h	10.68	19.80	22.44	6.70	11.41	12.69	9.57	17.54	19.87	12.01	22.21	23.94	5.97	11.63	13.08	11.00	19.94	21.24
PWCF	21.41	35.44	35.85	12.79	21.97	10.39	22.57	32.20	31.25	21.21	35.51	35.38	13.79	21.96	20.67	22.02	32.63	31.25
DAPH*	30.45	48.77	44.93	12.83	23.45	23.87	29.69	45.77	43.18	25.33	44.77	43.32	18.60	30.78	31.83	26.01	34.43	37.97
SGHL	27.92	49.38	53.72	18.11	29.96	34.34	22.27	38.83	42.81	28.73	49.89	53.78	16.08	28.45	30.37	20.01	35.49	39.13
TSS	9.52	48.24	61.41	6.64	39.59	50.90	10.74	38.28	49.11	10.61	50.65	62.58	5.96	24.54	<u>32.60</u>	12.57	47.72	57.11
DCS-LSG	47.76	69.72	70.21	37.39	57.36	61.59	33.42	48.81	50.06	54.53	68.00	70.27	19.28	30.45	31.48	47.52	67.81	67.73
Ours	55.84	76.06	78.15	47.59	68.61	68.82	44.84	64.02	68.77	58.83	73.04	75.06	28.34	39.65	42.87	54.17	72.78	74.81

Table 2: Cross-domain retrieval performance (MAP%) on Office-Home with varying code lengths.

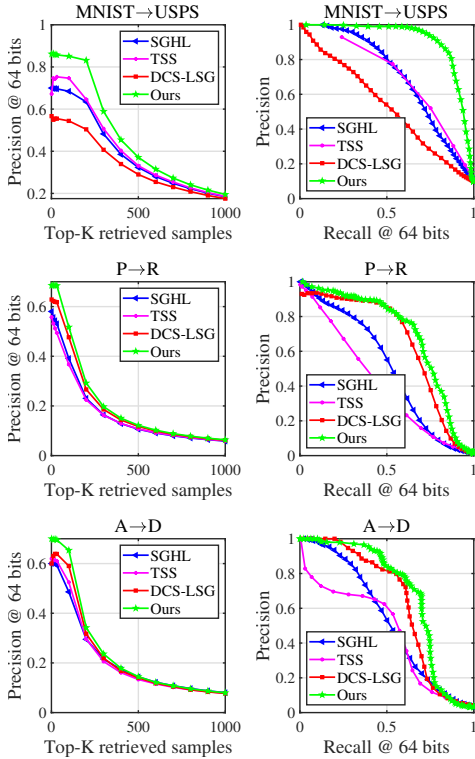


Figure 2: The Precision-Recall Curves and Top-K Precision Curves of SGHL, TSS, DCS-LSG and PSCA.

tive methods, i.e., SGHL, TSS and DCS-LSG. Specifically, experiments are conducted on three retrieval tasks across datasets of different scales. As shown in Figure 2, it can be observed from the Top-K Precision Curves that PSCA consistently maintains advantages as the number of retrieved samples increases. The Precision-Recall Curves demonstrate that PSCA outperforms the comparison baselines.

Experimental Analysis on Single-Domain Retrieval

To evaluate the effectiveness of PSCA in the realm of single-domain retrieval, four representative cases (MNIST→USPS,

COIL1→COIL2, A→D, P→R) are selected for experiments. We can conclude the following observations by analyzing Table 3: PSCA outperforms comparison baselines with the average MAP 6.81%, 5.39% and 12.55% higher than the second-best baselines on MNIST→USPS, COIL1→COIL2 and P→R respectively. Notably, PSCA shows the least performance improvement on A→D by 2.25%. This may be due to the fact that our domain-shared prototypes over-smooth reconstructed target domain features, hindering optimal single-domain retrieval.

Comparison with Deep Baseline

To further validate the superiority of PSCA, we compare it with three advanced deep learning-based DAR methods, i.e., PEACE (Wang et al. 2023a), CPH (Cui et al. 2024) and COUPLE (Luo et al. 2025). The comparison results are illustrated in Figure 3. As shown in Figure 3(a), PSCA markedly outperforms other competitors on MNIST-USPS, boosting the performance by 15.89% compared to the second-best method, COUPLE. Note that PSCA is based on traditional machine learning principles and adopts shallow features for this experiment. Figure 3(b) shows the comparison results on Office-Home. Here, the 4,096-d deep features utilized for PSCA are extracted from a pre-trained VGG-16 model (Huang et al. 2020). On average across six cases, PSCA surpasses the suboptimal baseline CPH by 1.98%. This improvement stems from PSCA’s capacity to address multiple limitations of other methods: CPH neglects error accumulation caused by erroneous pseudo-labels. PEACE and COUPLE perform domain alignment within Hamming space, where dimension is typically much smaller than original data, resulting in limited semantic information preservation.

Ablation Study

To clearly highlight the significance of each component in PSCA, several variants are designed. **PSCA-v1** denotes that the semantic-aware fusion is removed, i.e., $\alpha = \mathbf{0}$. This means that \mathbf{R} is dominated by geometric structure knowledge. **PSCA-v2** denotes that we omit the semantic consistency alignment component, solely using Eq. (3) to conduct semantic alignment. For validity, we reconstruct target samples as $\tilde{\mathbf{x}}_{t_i} = \sum_m^c \hat{y}_{im} \mathbf{o}_m^\top$. **PSCA-v3** means that prototype

Case	MNIST→USPS				COIL1→COIL2				A→D				P→R			
Code length	16	32	64	128	16	32	64	128	16	32	64	128	16	32	64	128
SH	46.30	47.82	49.12	47.81	52.91	57.07	57.23	52.61	30.54	35.66	42.50	45.64	13.15	18.71	22.57	20.66
DSH	41.42	45.30	47.85	50.76	43.44	52.85	58.06	58.84	22.45	33.38	40.09	46.31	6.10	11.44	16.61	14.45
SGH	15.60	30.78	35.55	41.78	54.30	59.25	59.97	58.49	38.67	45.59	53.57	<u>57.37</u>	18.97	26.18	32.61	34.97
OCH	24.23	32.90	36.34	44.36	54.24	61.08	65.56	65.98	33.30	41.65	50.78	53.74	13.45	21.14	25.34	28.02
ITQ+	50.22	49.66	44.38	43.21	58.74	60.53	61.86	60.94	35.03	42.62	43.12	39.12	15.60	20.60	24.96	24.05
LapITQ+	54.19	55.24	55.77	54.08	53.05	48.90	40.92	34.58	37.60	42.91	44.55	38.87	16.78	22.26	22.29	21.85
GTH-h	43.38	40.09	34.14	32.80	58.84	59.65	63.68	63.71	39.88	46.60	50.74	54.72	13.37	22.03	26.40	28.99
PWCF	57.57	64.00	65.45	65.63	69.36	70.81	72.41	70.43	24.01	31.25	38.65	43.35	23.84	33.83	38.26	37.91
DAPH*	63.95	70.53	72.38	73.00	71.93	74.87	75.35	75.54	42.18	46.86	47.73	50.26	20.43	31.02	33.30	31.50
SGHL	64.25	69.64	70.97	71.22	72.37	74.24	76.67	78.93	46.40	<u>53.66</u>	<u>58.05</u>	57.30	16.13	29.04	36.49	40.75
TSS	<u>69.55</u>	<u>74.15</u>	<u>77.26</u>	<u>78.13</u>	67.20	76.63	76.53	<u>82.84</u>	27.47	43.88	54.07	51.66	15.67	29.92	45.07	<u>53.97</u>
DCS-LSG	56.91	63.30	61.55	62.06	<u>73.77</u>	<u>77.07</u>	<u>77.93</u>	<u>79.77</u>	42.46	47.26	51.53	51.12	<u>33.45</u>	44.13	48.50	50.62
Ours	80.61	81.09	81.53	83.07	78.90	80.69	83.54	86.00	48.83	55.91	59.54	60.11	45.70	55.55	60.13	62.85

Table 3: Single-domain retrieval performance (MAP%) on MNIST→USPS, COIL1→COIL2, A→D and P→R with varying code lengths. The bolded figure indicates the highest score, and underlined figure indicates the second-highest score.

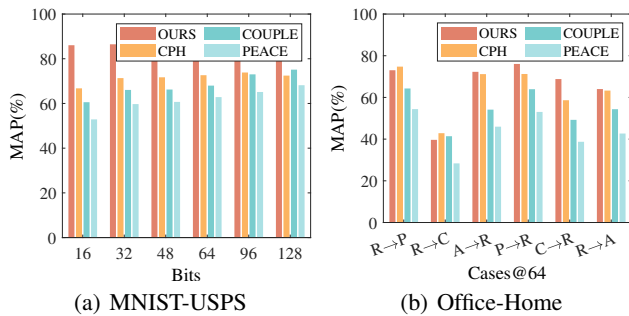


Figure 3: Comparison with deep methods on MNIST-USPS across all code lengths and Office-Home with 64 bits.

Case	MNIST→USPS					
Code length	16	32	48	64	96	128
PSCA-v1	58.74	59.79	60.09	61.68	62.41	61.38
PSCA-v2	77.11	79.09	81.89	82.77	82.95	83.24
PSCA-v3	38.25	40.26	42.08	43.19	43.40	44.56
PSCA-v4	76.35	78.76	79.31	79.87	80.10	80.92
PSCA	86.05	86.47	87.04	87.35	88.09	88.71

Table 4: Ablation study results on MNIST-USPS.

learning is completely removed, subsequently we fuse samples as $\mathbf{D}_* = [\mathbf{X}, \mathbf{X}^\top \mathbf{P}]^\top \in \mathbb{R}^{C_* \times n}$, where $C_* = d + q$. **PSCA-v4** indicates that feature reconstruction \mathbf{D} is omitted, directly quantizing $\mathbf{P}^\top \mathbf{X}_s$ and $\mathbf{P}^\top \mathbf{X}_t$ in stage two. The results of ablation experiments are reported in Table 4. It's evident that all components contribute to performance improvements of PSCA. The inferior performance of PSCA-v1 and PSCA-v3 demonstrates that prototypes indeed capture fundamental correct semantic patterns, validating the rationality of our geometric proximity for semantic correction.

Visualization Analysis

Figure 4 presents the t-SNE analyses of original data and hash codes generated by TSS, DCS-LSG and PSCA on MNIST→USPS. As shown in Figure 4(a), affected by domain gap, original data with the identical class generally distribute in two different areas. Compared to other baselines,

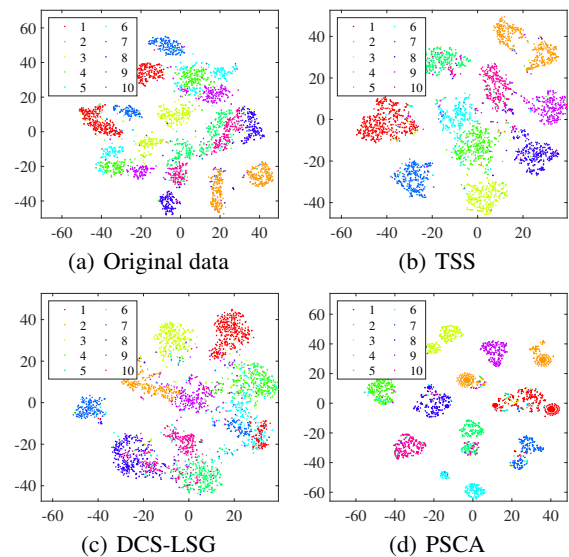


Figure 4: Distribution visualization on case MNIST→USPS. Different colors indicate different categories.

PSCA achieves tighter intra-class clustering and clearer inter-class boundaries, demonstrating the effectiveness of our PSCA in generating more discriminative hash codes.

Conclusion

While admiring the merits of previous methods, we identify their limitations and propose PSCA. Despite deviations in domain-shared class prototypes caused by incorrect pseudo-labels, they still capture the vast majority of accurate semantic patterns. Towards fully utilizing them, PSCA conducts orthogonal prototype learning and introduces the geometric proximity to weight potentially unreliable semantics via a novel semantic consistency alignment. By feature reconstruction, PSCA avoids directly quantizing original features affected by domain shift, achieving superior hash coding with remarkable retrieval performance. Comprehensive experiments validate that PSCA achieves SOTA performance.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant 62176065.

References

- Chen, Y.; Fang, X.; Liu, Y.; Zheng, W.; Kang, P.; Han, N.; and Xie, S. 2023. Two-step strategy for domain adaptation retrieval. *IEEE Transactions on Knowledge and Data Engineering*, 36(2): 897–912.
- 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.
- Gretton, A.; Borgwardt, K. M.; Rasch, M. J.; Schölkopf, B.; and Smola, A. 2012. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1): 723–773.
- Hu, T.; Chen, Y.; Cheng, C.; Xiao, J.; Sun, W.; and Fang, X. 2025a. Coarse-to-Fine Label Refinement for Domain Adaptive Retrieval. *Information Sciences*, 122532.
- Hu, T.; Chen, Y.; Lv, W.; Chen, Y.; and Fang, X. 2025b. Consistent coding guided domain adaptation retrieval. *Applied Intelligence*, 55(7): 706.
- 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. 1994. A database for handwritten text recognition research. *IEEE Transactions on pattern analysis and machine intelligence*, 16(5): 550–554.
- Jiang, Q.-Y.; and Li, W.-J. 2015. Scalable graph hashing with feature transformation. In *IJCAI*, volume 15, 2248–2254.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.
- Liu, H.; Ji, R.; Wang, J.; and Shen, C. 2018. Ordinal constraint binary coding for approximate nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 41(4): 941–955.
- Liu, H.; Wang, R.; Shan, S.; and Chen, X. 2016. Deep supervised hashing for fast image retrieval. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2064–2072.
- Long, M.; Wang, J.; Ding, G.; Sun, J.; and Yu, P. S. 2013. Transfer feature learning with joint distribution adaptation. In *Proceedings of the IEEE international conference on computer vision*, 2200–2207.
- Long, M.; Wang, J.; Ding, G.; Sun, J.; and Yu, P. S. 2014. Transfer joint matching for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1410–1417.
- Luo, J.; Zhao, Y.; Luo, X.; Xiao, Z.; Ju, W.; Shen, L.; Tao, D.; and Zhang, M. 2025. Cross-domain diffusion with progressive alignment for efficient adaptive retrieval. *IEEE Transactions on Image Processing*.
- Luo, X.; Zhang, P.-F.; Huang, Z.; Nie, L.; and Xu, X.-S. 2019. Discrete hashing with multiple supervision. *IEEE Transactions on Image Processing*, 28(6): 2962–2975.
- Nene, S. A.; Nayar, S. K.; Murase, H.; et al. 1996. Columbia object image library (coil-20).
- Saenko, K.; Kulis, B.; Fritz, M.; and Darrell, T. 2010. Adapting visual category models to new domains. In *Computer Vision—ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5–11, 2010, Proceedings, Part IV 11*, 213–226. Springer.
- 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, Q.; and Breckon, T. 2020. Unsupervised domain adaptation via structured prediction based selective pseudo-labeling. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 6243–6250.
- Wang, Z.; Xiao, H.; Duan, Y.; Zhou, J.; and Lu, J. 2023b. Learning Deep Binary Descriptors via Bitwise Interaction Mining. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2): 1919–1933.
- Weiss, Y.; Torralba, A.; and Fergus, R. 2008. Spectral hashing. *Advances in neural information processing systems*, 21.
- 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, L.; Liu, J.; Yang, Y.; Huang, F.; Nie, F.; and Zhang, D. 2019. Optimal projection guided transfer hashing for image retrieval. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(10): 3788–3802.
- Zhang, W.; Yang, X.; Teng, S.; and Wu, N. 2023a. Semantic-guided hashing learning for domain adaptive retrieval. *World Wide Web*, 26(3): 1093–1112.
- Zhang, W.; Zhou, K.; Teng, L.; Tang, F.; Wu, N.; Teng, S.; and Li, J. 2024. Dynamic Confidence Sampling and Label Semantic Guidance Learning for Domain Adaptive Retrieval. *IEEE Transactions on Multimedia*, 26: 2467–2479.
- Zhang, Y.; Tian, S.; Liao, M.; Zhang, Z.; Zou, W.; and Xu, C. 2023b. Fine-grained self-supervision for generalizable semantic segmentation. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(1): 371–383.

- Zhang, Y.; Tian, S.; Liao, M.; Zou, W.; and Xu, C. 2023c. A hybrid domain learning framework for unsupervised semantic segmentation. *Neurocomputing*, 516: 133–145.
- Zhou, J. T.; Zhao, H.; Peng, X.; Fang, M.; Qin, Z.; and Goh, R. S. M. 2018. Transfer hashing: From shallow to deep. *IEEE transactions on neural networks and learning systems*, 29(12): 6191–6201.
- 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.