

Unsupervised Cross-Domain Image Retrieval via Prototypical Optimal Transport

Bin Li¹, Ye Shi¹, Qian Yu², Jingya Wang^{1*}

¹ShanghaiTech University

²Beihang University

{libin3, shiye, wangjingya}@shanghaitech.edu.cn; qianyu@buaa.edu.cn

Abstract

Unsupervised cross-domain image retrieval (UCIR) aims to retrieve images sharing the same category across diverse domains without relying on labeled data. Prior approaches have typically decomposed the UCIR problem into two distinct tasks: intra-domain representation learning and cross-domain feature alignment. However, these segregated strategies overlook the potential synergies between these tasks. This paper introduces ProtoOT, a novel Optimal Transport formulation explicitly tailored for UCIR, which integrates intra-domain feature representation learning and cross-domain alignment into a unified framework. ProtoOT leverages the strengths of the K-means clustering method to effectively manage distribution imbalances inherent in UCIR. By utilizing K-means for generating initial prototypes and approximating class marginal distributions, we modify the constraints in Optimal Transport accordingly, significantly enhancing its performance in UCIR scenarios. Furthermore, we incorporate contrastive learning into the ProtoOT framework to further improve representation learning. This encourages local semantic consistency among features with similar semantics, while also explicitly enforcing separation between features and unmatched prototypes, thereby enhancing global discriminativeness. ProtoOT surpasses existing state-of-the-art methods by a notable margin across benchmark datasets. Notably, on DomainNet, ProtoOT achieves an average P@200 enhancement of 24.44%, and on Office-Home, it demonstrates a P@15 improvement of 12.12%. Code is available at <https://github.com/HCVLAB/ProtoOT>.

Introduction

Cross-domain image retrieval (CIR) aims at utilizing imagery data from one domain as queries to retrieve relevant samples from distinct domains. While significant progress has been made in supervised CIR studies (Huang et al. 2015; Ji et al. 2017), these advancements are constrained by the requirement for annotated labels. Unfortunately, this reliance on manual labels hampers the practical scalability of supervised CIR methods, due to the considerable costs associated with label acquisition in real-world scenarios. As a result, the usability of existing supervised CIR techniques becomes limited.

*Corresponding author.

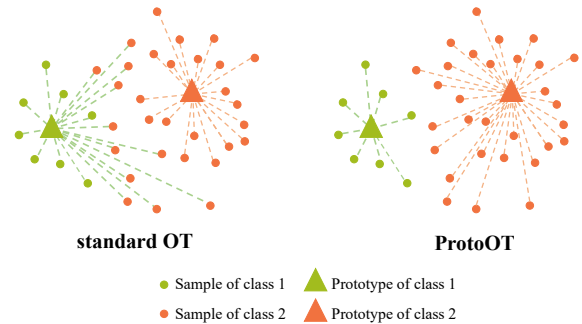


Figure 1: Comparison between the standard OT and our proposed ProtoOT to deal with distribution imbalance. Different colors represent different categories, and the dashed lines indicate that the samples are matched to related prototypes.

To address this challenge, recent studies (Hu and Lee 2022; Wang et al. 2023) have explored unsupervised cross-domain image retrieval (UCIR) without depending on labeled information. However, the performance of these UCIR methods tends to be suboptimal due to the intricate nature of domain-specific clustering representation learning and the complexities involved in accommodating noise-tolerant cross-domain associations. Moreover, these studies tend to treat UCIR as two distinct tasks: intra-domain representation learning and cross-domain feature alignment. Regrettably, this segregated approach overlooks the underlying correlations and potential synergies between these tasks, resulting in subpar performance.

In this paper, we investigate the UCIR problem by seamlessly integrating the intra-domain feature representation learning and cross-domain alignment. In particular, the former provides representative features for the alignment procedure in the latter, while the latter further enhances the representation in the former by integrating cross-domain knowledge. Notably, the core of both intra-domain representation learning and cross-domain alignment lies in minimizing the distance between the sample distribution and prototype distribution in order to obtain features that are both intra-class coherent and inter-class distinctive. Recent years have seen great success in applying Optimal Trans-

port (OT) (Villani et al. 2009) in both representation learning (Caron et al. 2020; Asano, Rupprecht, and Vedaldi 2019) and cross-domain alignment (Flamary et al. 2016; Courty et al. 2017). This naturally motivates us to a unified OT method for both intra-domain representation learning and cross-domain alignment for the UCIR problem. However, the UCIR task presents challenges stemming from significant imbalances in both intra-domain and cross-domain distributions. The standard OT methods lead to balanced assignments that are not suitable for UCIR, as shown in Figure 1. Existing unbalanced OT methods, such as (Séjourné, Peyré, and Vialard 2023; Zhan et al. 2021) also fall short in addressing UCIR’s intricacies. A critical issue arises in existing unbalanced OT methods where some samples may remain unassigned to prototypes, consequently hindering the training process of UCIR.

To address this issue, we develop a novel OT formulation named Prototypical Optimal Transport (ProtoOT), tailored explicitly for UCIR. ProtoOT leverages the strengths of the K-means clustering method to effectively handle distribution imbalances in UCIR. By utilizing K-means for generating initial prototypes and approximating class marginal distributions, we modify the constraints in OT accordingly, significantly enhancing its performance in UCIR scenarios. To further enhance the representation learning, we employ contrastive learning (Oord, Li, and Vinyals 2018) within the ProtoOT framework. This encourages local semantic consistency among features with similar semantics, while also explicitly enforcing separation between features and unmatched prototypes, thereby enhancing global discriminativeness. Our main contributions are summarized as follows:

- We address the UCIR problem by synergistically tackling intra-domain feature representation learning and cross-domain assignment, presenting the first unified Optimal Transport framework that simultaneously addresses both tasks.
- We introduce ProtoOT, a novel OT formulation that effectively balances the strengths of K-means and OT to handle substantial distribution imbalances in the context of UCIR.
- ProtoOT outperforms existing state-of-the-art techniques by a substantial margin across benchmark datasets. Notably, on DomainNet, ProtoOT achieves an average enhancement of 24.44% in terms of P@200, and on Office-Home, an improvement of 12.12% in terms of P@15.

Related Work

Cross-Domain Image Retrieval

Cross-domain image retrieval (CIR) stands as a well-explored task in computer vision. It involves retrieving images from one visual domain based on a query image from a different domain. CIR boasts extensive practical applications; in the fashion domain, for example (Gajic and Baldrich 2018; Bao et al. 2022), it finds utility in matching user-provided images with a product database, thus presenting users with target products or accessories. Most existing works (Sain et al. 2021; Ji et al. 2017; Lee et al. 2018) heavily

lean on well-annotated data within the supervised learning framework, yet sacrifice scalability for real-world CIR implementations. Recently, there have been efforts (Hu and Lee 2022; Wang et al. 2023) to move beyond relying on labels and tackle a tougher challenge known as unsupervised cross-domain image retrieval (UCIR). To bridge the gap between different domains, the preceding works (Hu and Lee 2022; Wang et al. 2023) split the UCIR task into intra-domain representation learning and cross-domain alignment, designing different algorithms for each part. However, they overlooked the essential correlations between these two parts, making it difficult for them to mutually enhance each other.

Unsupervised Domain Adaptation

Unsupervised domain adaptation (UDA) aims to address the challenge of transferring knowledge from a labeled source domain to a target domain without any annotations. To mitigate the domain shift, existing methods focus on reducing domain discrepancy through techniques like Maximum Mean Discrepancy (Gretton et al. 2012; Long et al. 2017), domain adversarial training (Zhang et al. 2018; Dai et al. 2020) and prototypes alignment (Lin et al. 2022; Li, Lü, and Li 2022; Yue et al. 2021; Chang et al. 2022). However, while these methods have demonstrated substantial efficacy in UDA, they cannot be used directly for UCIR tasks because there are no labels in the source domain. In this paper, we delve into a more challenging scenario in which both the source and target domains are unlabeled. The problem of exploring weak self-representation learning for both domains while concurrently transferring knowledge from the other remains unresolved.

Optimal Transport

Optimal transport (OT), initially proposed by Kantorovich (Villani et al. 2009), offers an efficient solution for transferring mass from one distribution to another. In the field of computer vision, OT has been utilized in diverse tasks like point cloud registration (Shen et al. 2021), shape matching (Saleh et al. 2022) and learning with noisy labels (Chang, Shi, and Wang 2023). These tasks involve seeking an optimal mapping between data distributions to minimize a cost function. Notably, (Asano, Rupprecht, and Vedaldi 2019; Caron et al. 2020) have achieved significant advancements in unsupervised representation learning through OT-driven matching of samples to prototypes. Furthermore, in the context of domain adaptation, OT-based approaches (Courty, Flamary, and Tuia 2014; Courty et al. 2017; Redko et al. 2019; Chang et al. 2022) have exhibited promise in mitigating the gap between source and target domains by aligning their respective probability distributions. However, jointly exploring self-supervised feature representation learning and domain adaptation within a unified OT framework remains unexplored.

Methodology

In the UCIR task, we deal with two distinct domains, denoted as $\mathcal{D}_A = \{(\mathbf{x}_i^A)\}_{i=1}^{N_A}$ and $\mathcal{D}_B = \{(\mathbf{x}_j^B)\}_{j=1}^{N_B}$, both of which lack labeled data. The goal is to retrieve images

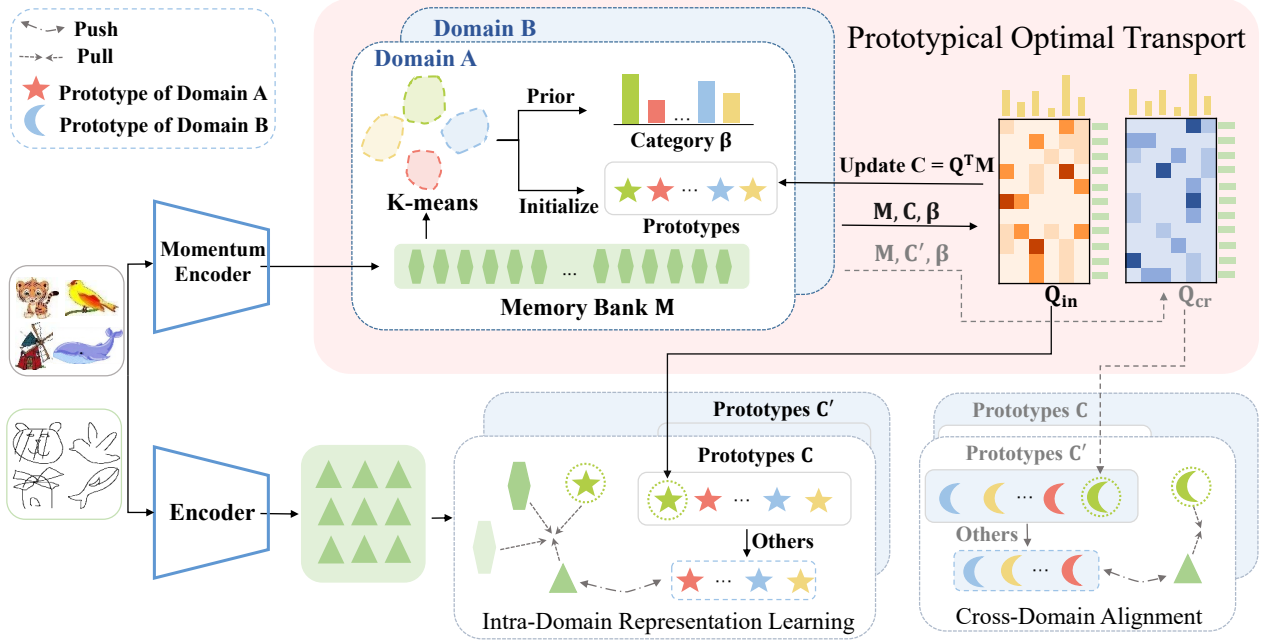


Figure 2: An overview of the proposed method for UCIR. We utilize K-means to generate initial prototypes and approximate class marginal distributions, which have modified the constraints of Optimal Transport. Both the intra-domain representation learning and cross-domain alignment are based on ProtoOT. Furthermore, the two are committed to enhancing the local consistency and global discriminability of features by employing the same form of contrastive loss.

of the category k in domain B based on a query image of category k in domain A . To address this challenging task, in this paper, we propose a unified OT framework, i.e. ProtoOT from two perspectives, i.e., intra-domain clustering and cross-domain alignment, as shown in Figure 2.

ProtoOT

Optimal Transport focuses on seeking the most efficient way to transform the mass of one distribution into another under the given marginal constraints. In unsupervised situations, where no information is available constraints in OT, SwAV (Caron et al. 2020) adopts the standard OT to obtain the representations of mini-clusters. Given a similarity matrix $S \in \mathbb{R}^{r \times c}$, the optimal coupling matrix Q^* of the standard OT can be obtained by solving the following optimization problem:

$$Q^*(S) = \operatorname{argmax}_{Q \in U} \operatorname{Tr}(Q^\top S) + \varepsilon H(Q), \quad (1)$$

$$U = \left\{ Q \in \mathbb{R}_+^{r \times c} \mid Q \mathbf{1}_c = \frac{1}{r} \mathbf{1}_r, Q^\top \mathbf{1}_r = \frac{1}{c} \mathbf{1}_c \right\}. \quad (2)$$

where $\varepsilon > 0$, the entropic regularization term $H(Q) = -\sum_{ij} Q_{ij} \log Q_{ij}$.

Since the semantic information of the categories in UCIR is essential, enforcing uniform matching blindly is clearly unreasonable. To this end, we propose ProtoOT tailored for

the UCIR task, which incorporates the results of K-means to the class marginal distribution for OT, instead of simply assuming a uniform distribution.

$$Q^*(M, C, \beta) = \operatorname{argmax}_{Q \in U(\beta)} \operatorname{Tr}(Q^\top M C^\top) + \varepsilon H(Q), \quad (3)$$

$$U(\beta) = \left\{ Q \in \mathbb{R}_+^{r \times c} \mid Q \mathbf{1}_c = \frac{1}{r} \mathbf{1}_r, Q^\top \mathbf{1}_r = \beta \right\}. \quad (4)$$

$$\beta = [\beta_1, \beta_2, \dots, \beta_c]^\top \quad (5)$$

where M is the set of the features, C is the prototypes, $\mathbf{1}_i$ is the i -dimensional vector of ones and $\beta_j \in \mathbb{R}$ is the probability of the j -th pseudo label generated by K-Means.

Intra-Domain Representation Learning via ProtoOT

Here we employ ProtoOT for intra-domain Representation Learning. First, we maintain a memory bank M by a momentum encoder f_m (He et al. 2020) for each domain. For each batch, we use the feature encoder f to obtain features q , whose augmented counterpart k encoded by f_m can be found in the memory bank. Follow (He et al. 2020; Chen et al. 2020), we update the parameters using gradient back-propagation only within f , and then utilize the parameters of f to momentum update f_m . We perform K-means on the

memory bank M and use the obtained centroids from K-means as the initialized prototypes of intra-domain clustering.

The UCIR task is a category-level cross-domain retrieval task, achieving a good intra-domain feature representation is important. The effectiveness of SwAV (Caron et al. 2020) in learning representations of mini-clusters highlights the superiority of OT-based methods in clustering, which motivates us to utilize OT to assist in the learning of category semantic features. As previously mentioned, ProtoOT integrates the distribution information from K-means, allowing it to acquire a confident transport plan Q_{in}^* even in the absence of any data labels. Therefore, we use ProtoOT to acquire pseudo labels $y_i \in \mathbf{Y}$ and corresponding prototypes C_{y_i} for each sample, where \mathbf{Y} is the set of all pseudo labels.

$$y_i = \operatorname{argmax}_{y \in \mathbf{Y}} Q_{in}^*(M, C, \beta). \quad (6)$$

To attain semantic clustering of features, aggregating features with local similarity is reasonable. However, merely focusing on the local similarity among features cannot effectively eliminate interference from features that are semantically similar yet not from the same category. This drives us to explore global discrimination as well. Specifically, to achieve local consistency in features, we bring together the feature q_i with its augmentation k_i in the embedding space. In addition to this, getting closer to the nearest neighbor n_i of the feature in the memory bank and converging with the matched prototype C_{y_i} is intuitive. We denote a set $\mathcal{P}_i^{\text{in}}$ to collect intra-domain positive triplets of sample i , i.e., $\mathcal{P}_i^{\text{in}} := \{k_i, n_i, C_{y_i}\}$. To explore global discrimination, it is crucial to ensure a large margin between prototypes of different categories. For this purpose, we consider prototypes that are not matched with features by ProtoOT as negative samples. Similarly, we define the set $\mathcal{N}_i^{\text{in}} = \{C_y \in \mathcal{C} | y \in \mathbf{Y}, y \neq y_i\}$.

Unlike K-means, which uses per-sample to obtain hard pseudo labels and form prototypes, here we utilize the transport plan $Q_{in}^*(M, C, \beta)$ of ProtoOT to get soft pseudo labels. This enables the prototypes C to be more globally representative, i.e.,

$$C = Q_{in}^*{}^T M. \quad (7)$$

Then, we integrate local consistency and global discrimination into a contrastive learning loss:

$$\mathcal{L}_{in} = \sum_{i \in \mathcal{D}} -\log \frac{\exp(q_i^T p_i^{\text{in}}/\tau)}{\exp(q_i^T p_i^{\text{in}}/\tau) + \sum_{n \in \mathcal{N}_i^{\text{in}}} \exp(q_i^T n/\tau)}, \quad (8)$$

where $p_i^{\text{in}} \in \mathcal{P}_i^{\text{in}}$ represents a positive sample of sample i , τ is the temperature value.

Cross-Domain Alignment via ProtoOT

Through intra-domain representation learning, prototypes with strong category semantics can be formed separately in domain \mathcal{D}_A and domain \mathcal{D}_B . In order to eliminate the domain gap, it is crucial to facilitate the mutual transfer of knowledge learned in the intra-domain clustering between the two domains. The achievements of OT (Courty, Flamary, and Tuia 2014; Courty et al. 2017) in domain adaptation lead

us to believe that OT-based methods are capable of addressing this fully unsupervised adaptation between the two unlabeled domains. However, it requires simultaneous consideration of the potential data imbalance in both domains, which makes the OT with enforced uniform matching even less suitable for the cases where imbalance is more severe. The natural tolerance of ProtoOT to data imbalance equips it with the ability to address this challenging problem.

Specifically, we utilize ProtoOT to match the features from domain \mathcal{D}_A to the prototypes C' of domain \mathcal{D}_B , where Q_{cr}^* is the cross-domain transport plan between the memory bank M of domain \mathcal{D}_A and prototypes C' .

$$y'_i = \operatorname{argmax}_{y' \in \mathbf{Y}'} Q_{cr}^*(M, C', \beta_{cr}), \quad (9)$$

where the symbol $'$ denotes a different domain, β_{cr} is the categories distribution of M obtained by K-means. We consider the matched prototype $C_{y'_i}$ as the positive sample and other prototypes are treated as negative samples $n' \in \mathcal{N}_i^{\text{cr}}$.

$$\mathcal{N}_i^{\text{cr}} = \{C'_{y'} \in \mathcal{C}' | y' \in \mathbf{Y}', y' \neq y'_i\}. \quad (10)$$

To mitigate domain discrepancies, it is important to aggregate features that exhibit semantic similarity across domains while disregarding the impact of other features that might share semantic resemblance but belong to different categories. In light of this, to achieve cross-domain semantic alignment, we bring the matched prototypes from one domain closer to learn semantic consistency, while pushing away the unmatched prototypes from the other domain to attain discrimination across domains. Due to the large domain gap, relying on the nearest neighbor feature across domains based on similarity becomes unreasonable. The loss function \mathcal{L}_{cr} of cross-domain alignment is as follows:

$$\mathcal{L}_{cr} = \sum_{i \in \mathcal{D}} -\log \frac{\exp(q_i^T C'_{y'_i}/\tau)}{\exp(q_i^T C'_{y'_i}/\tau) + \sum_{n' \in \mathcal{N}_i^{\text{cr}}} \exp(q_i^T n'/\tau)}. \quad (11)$$

A Unified ProtoOT Framework for UCIR

In UCIR, both the intra-domain representation learning and cross-domain alignment are based on ProtoOT. Furthermore, the two are dedicated to enhancing the local consistency and global discriminability of features through the same form of contrastive loss. Therefore, the two parts of the UCIR task can be unified within the framework based on ProtoOT. In summary, our total training objective can be written as:

$$\mathcal{L}_{total} = \mathcal{L}_{in} + \lambda \mathcal{L}_{cr} \quad (12)$$

Experiments

Datasets

We evaluate our proposed method on two datasets: Office-Home and DomainNet. The Office-Home (Venkateswara et al. 2017) dataset comprises 4 domains (Art, Clipart, Product, Real) encompassing 65 categories. For our experiments, we employ all available images. The DomainNet (Peng et al. 2019) dataset consists of 6 domains (Clipart, Infograph,

Method	Clipart→Sketch			Sketch→Clipart			Infograph→Real			Real→Infograph		
	P@50	P@100	P@200	P@50	P@100	P@200	P@50	P@100	P@200	P@50	P@100	P@200
ID(Wu et al. 2018)	49.46	46.09	40.44	54.38	47.12	37.73	28.27	27.44	26.33	39.98	31.77	24.84
ProtoNCE(Li et al. 2021)	46.85	42.67	36.35	54.52	45.04	35.06	28.41	28.53	28.50	57.01	41.84	30.33
CDS(Kim et al. 2021)	45.84	42.37	37.16	59.13	48.83	37.40	28.51	27.92	27.48	56.69	39.76	26.38
PCS(Yue et al. 2021)	51.01	46.87	40.19	59.70	50.67	39.38	30.56	30.27	29.68	55.42	42.13	30.76
DD(Hu and Lee 2022)	56.31	52.74	47.38	63.07	57.26	48.17	35.52	35.24	34.35	57.74	46.69	35.47
ProtoOT	82.70	82.26	80.32	86.69	85.57	81.59	48.74	49.28	49.60	75.36	68.49	56.20
Method	Infograph→Sketch			Sketch→Infograph			Painting→Clipart			Clipart→Painting		
	P@50	P@100	P@200	P@50	P@100	P@200	P@50	P@100	P@200	P@50	P@100	P@200
ID(Wu et al. 2018)	30.35	29.04	26.55	42.20	34.94	27.52	64.67	54.41	40.07	42.37	39.61	35.56
ProtoNCE(Li et al. 2021)	28.24	26.79	24.23	39.83	31.99	24.77	55.44	43.74	32.59	39.13	35.87	32.07
CDS(Kim et al. 2021)	30.55	29.51	27.00	46.27	36.11	27.33	63.15	47.30	32.93	37.75	35.18	32.76
PCS(Yue et al. 2021)	30.27	28.36	25.35	42.58	34.09	25.91	63.47	53.21	41.68	48.83	46.21	42.10
DD(Hu and Lee 2022)	31.29	29.33	26.54	43.66	36.14	28.12	66.42	56.84	46.72	52.58	50.10	46.11
ProtoOT	45.72	45.51	43.77	62.92	56.41	47.08	91.04	90.09	86.67	82.70	82.66	81.38
Method	Painting→Quickdraw			Quickdraw→Painting			Quickdraw→Real			Real→Quickdraw		
	P@50	P@100	P@200	P@50	P@100	P@200	P@50	P@100	P@200	P@50	P@100	P@200
ID(Wu et al. 2018)	20.34	19.59	18.79	21.12	19.81	18.48	28.27	27.46	26.32	23.45	22.79	22.01
ProtoNCE(Li et al. 2021)	21.63	21.24	20.56	23.95	22.84	21.56	26.38	25.70	24.45	25.10	24.81	23.78
CDS(Kim et al. 2021)	18.75	18.89	17.88	21.37	21.44	19.46	19.28	19.14	18.67	15.36	15.57	15.82
PCS(Yue et al. 2021)	25.12	24.65	23.80	24.03	23.24	22.13	34.82	33.92	31.73	28.98	28.85	28.16
DD(Hu and Lee 2022)	39.72	38.59	37.63	33.45	33.81	34.29	42.79	42.75	42.70	41.90	42.10	41.59
ProtoOT	63.06	59.92	56.73	56.51	55.66	52.04	61.55	61.60	61.63	67.24	67.02	65.41

Table 1: Unsupervised Cross-domain Retrieval Accuracy (%) on DomainNet

Painting, Quickdraw, Real, and Sketch). Based on the criteria established by (Hu and Lee 2022), we utilize 7 categories with over 200 images in each domain for our experiments.

Implementation Details

We employ the ResNet-50(He et al. 2016) architecture as the encoder f_θ , initializing its parameters using MoCov2(Chen et al. 2020) trained on the unlabeled ImageNet dataset(Deng et al. 2009). The extracted features undergo l_2 -normalization. Our optimization employs the Adam optimizer with a learning rate of 2.5×10^{-4} over 200 epochs, with a batch size of 64. The K-means clustering is performed on the memory bank at the end of each epoch. For the Sinkhorn Algorithm(Cuturi 2013), the entropic regularization coefficient ϵ is set to 0.05 and following (Caron et al. 2020) the iterations is 3. The number of prototypes corresponds to the number of classes in the training set: 65 for Office-Home and 7 for DomainNet. Implementation of our framework is in PyTorch(Paszke et al. 2019).

Evaluation Metrics

Follow (Hu and Lee 2022), our evaluation metric involves measuring precision across the top 1/5/15 retrieved images for the Office-Home dataset, considering a minimum of 15 images per category. For the subset of the selected 7 categories within DomainNet, we assess precision using the top 50/100/200 retrieved images. Our primary focus lies in category-level cross-domain retrieval. In this context, we consider retrieved images that correspond to the same semantic classes as the query as correct matches.

Baselines

In our experimental validation, we assess the effectiveness of ProtoOT by comparing it against the following baseline methods: **ID** (Wu et al. 2018): This approach learns instance-level information by discerning distinctions between instances. **ProtoNCE** (Li et al. 2021): This method introduces the integration of prototypes to effectively capture semantic information and enhance cluster capabilities. **CDS** (Kim et al. 2021): Taking both intra- and cross-domain discriminative aspects into consideration among instances, this method aims to address domain adaptation. **PCS** (Yue et al. 2021): Designed for few-shot unsupervised domain adaptation, this cross-domain self-supervised learning method employs K-means to assign prototypes and utilizes similarity vectors for cross alignment. **DD** (Hu and Lee 2022): This method proposes a distance-of-distance loss for quantifying domain discrepancy. These baselines serve as benchmarks for validating the performance of our ProtoOT.

Comparison with State-of-the-Art Methods

ProtoOT’s impressive performance is evident from the summarized results in Table 3. On the Office-Home dataset, ProtoOT outperforms the leading baseline(Hu and Lee 2022) with average improvements of 7.94%, 10.28%, and 12.12% in P@1, P@5, and P@15, respectively. The improvements are even more pronounced on DomainNet, showcasing substantial enhancements of 21.59%, 23.57%, and 24.44% in P@50, P@100, and P@200. Table 1 and 2 show the detailed retrieval performance for Office-Home and DomainNet, respectively. It is worth noting that the Painting → Clipart retrieval on DomainNet shows the most significant improvement, with an accuracy boost of 39.95% in terms

Method	Art→Real			Real→Art			Art→Product			Product→Art		
	P@1	P@5	P@15	P@1	P@5	P@15	P@1	P@5	P@15	P@1	P@5	P@15
ID(Wu et al. 2018)	35.89	33.13	29.60	39.89	34.42	27.65	25.88	24.91	22.49	32.17	25.94	20.23
ProtoNCE(Li et al. 2021)	40.50	36.39	34.00	44.53	39.26	32.99	29.54	27.89	25.75	35.73	30.61	24.55
CDS(Kim et al. 2021)	45.08	41.15	38.73	44.71	40.75	35.53	32.76	31.47	28.90	35.75	32.48	26.82
PCS(Yue et al. 2021)	41.70	38.51	36.22	44.96	39.88	33.99	33.29	31.50	29.53	39.24	34.77	28.77
DD(Hu and Lee 2022)	45.12	42.33	40.06	47.95	43.68	38.38	35.39	34.67	32.61	42.51	37.94	31.41
ProtoOT	49.94	49.24	49.55	54.90	51.81	48.17	44.25	45.95	45.71	53.43	52.07	46.89
Method	Clipart→Real			Real→Clipart			Product→Real			Real→Product		
	P@1	P@5	P@15	P@1	P@5	P@15	P@1	P@5	P@15	P@1	P@5	P@15
ID(Wu et al. 2018)	29.48	26.48	23.25	35.51	32.17	27.96	50.73	45.03	39.05	45.12	41.46	38.01
ProtoNCE(Li et al. 2021)	25.25	22.66	20.83	41.15	37.66	31.95	53.84	48.25	42.21	47.74	44.85	41.21
CDS(Kim et al. 2021)	32.51	30.30	27.80	38.88	36.48	33.16	54.00	50.07	45.60	49.39	47.27	43.98
PCS(Yue et al. 2021)	29.07	26.06	24.00	40.60	38.11	34.06	56.45	50.78	45.37	49.90	47.11	43.73
DD(Hu and Lee 2022)	33.31	30.57	28.14	44.66	41.47	37.41	57.42	52.69	47.90	51.71	48.48	44.95
ProtoOT	40.92	40.68	39.96	51.11	52.12	50.44	69.34	67.27	64.74	57.72	59.32	59.72
Method	Product→Clipart			Clipart→Product			Art→Clipart			Clipart→Art		
	P@1	P@5	P@15	P@1	P@5	P@15	P@1	P@5	P@15	P@1	P@5	P@15
ID(Wu et al. 2018)	31.52	28.55	24.15	24.01	22.42	20.60	26.78	24.79	21.64	21.17	17.86	14.71
ProtoNCE(Li et al. 2021)	36.13	33.99	28.24	21.17	20.63	20.47	28.97	26.15	22.98	21.33	17.40	14.46
CDS(Kim et al. 2021)	37.69	34.99	30.42	27.24	26.46	24.86	25.59	23.77	22.41	22.41	20.34	17.34
PCS(Yue et al. 2021)	39.51	37.51	32.81	26.39	25.86	24.92	31.23	28.74	26.11	24.51	21.27	17.54
DD(Hu and Lee 2022)	42.26	37.42	33.74	27.79	27.26	25.97	32.67	30.79	28.70	27.26	23.94	20.53
ProtoOT	51.92	50.66	49.58	39.01	37.95	37.16	40.54	38.27	36.01	30.24	29.31	27.34

Table 2: Unsupervised Cross-domain Retrieval Accuracy (%) on Office-Home

	Average	ID	ProtoNCE	CDS	PCS	DD	ProtoOT	Improvement
DomainNet	P@50	37.07	37.21	36.89	41.23	47.09	68.68	+21.59
	P@100	33.34	32.59	31.84	36.87	43.47	67.04	+23.57
	P@200	28.72	27.85	26.69	31.74	39.09	63.53	+24.44
Office-Home	P@1	33.18	35.49	37.17	38.07	40.67	48.61	+7.94
	P@5	29.76	32.15	34.63	35.01	37.60	47.88	+10.28
	P@15	25.78	28.30	31.30	31.42	34.15	46.27	+12.12

Table 3: Average Accuracy (%) for Unsupervised Cross-domain Retrieval on DomainNet and Office-Home

of P@200. Furthermore, ProtoOT maintains over 50% retrieval precision consistently, even in challenging tasks that involve Quickdraw, with an impressive average enhancement of 22.62%, 21.74%, and 19.90% in P@50, P@100, and P@200, respectively. This is attributed to the strong matching capabilities of ProtoOT.

Ablation Study

To thoroughly assess the effectiveness of our proposed ProtoOT, we meticulously compare its impact on both intra-domain feature representation learning and cross-domain alignment, as outlined in Table 4. We systematically replace each component with the standard OT to demonstrate the superiority of our model’s design. The experimental results clearly demonstrate the substantial contribution of cross-domain alignment, highlighting the significance of bridging the domain gap between distinct domains for the UCIR task. Furthermore, it’s evident that ProtoOT maintains its ad-

vantageous characteristics. In summary, the unified ProtoOT framework displays remarkable performance in both intra-domain clustering and cross-domain alignment, underscoring its exceptional effectiveness.

To further validate the superiority of our ProtoOT, we conducted a comparison with several common variants of OT. The Unbalanced OT(Séjourné, Peyré, and Vialard 2023) uses soft penalties to relax the conservation of marginal constraints. The Partial OT(Figalli 2010) focuses on transporting only a fraction of mass at the lowest cost. Due to their intrinsic characteristics that are not well suited for the UCIR task, as depicted in Table 5, they yield suboptimal performance.

To investigate the impact of the coefficient λ in Eq.(12), we conduct hyper-parameter analysis experiments on DomainNet as shown in Table 6. The favorable performance can be obtained within the range of 0.001 to 0.1. Based on the observation, we set $\lambda = 0.01$ in all experiments.

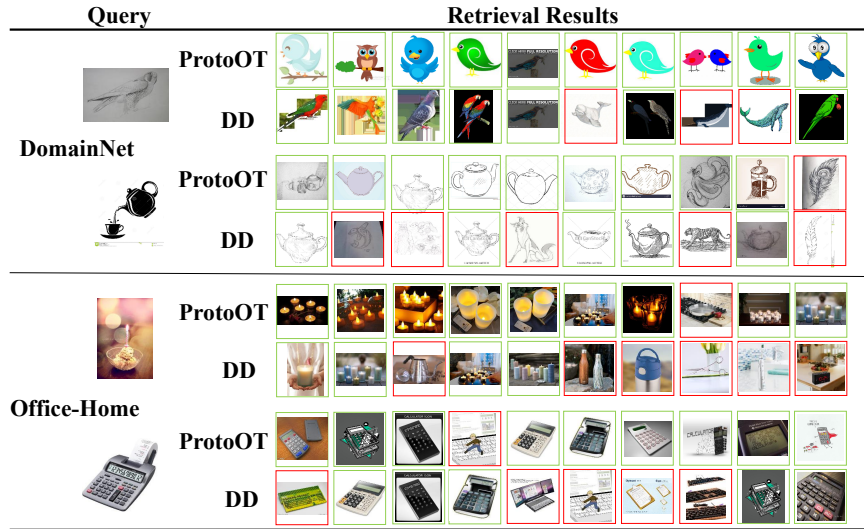


Figure 3: Top 10 retrieval results in DomainNet and Office-Home. The green and red boxes denote correct and incorrect retrievals, respectively.

Intra	Cross	P@50	P@100	P@200
SOT	X	54.92	51.57	46.13
SOT	SOT	62.23	59.91	56.90
SOT	ProtoOT	67.56	65.80	61.97
ProtoOT	X	55.22	52.57	47.68
ProtoOT	SOT	66.86	65.29	61.84
ProtoOT	ProtoOT	68.68	67.04	63.53

Table 4: Evaluation of the effectiveness of ProtoOT on both intra-domain feature representation learning and cross-domain alignment (Average on 12 tasks). SOT is an abbreviation for the standard Optimal Transport.

Intra	Cross	P@50	P@100	P@200
UOT	UOT	63.92	62.37	59.29
POT	POT	62.29	60.68	57.44
ProtoOT	ProtoOT	68.68	67.04	63.53

Table 5: The retrieval performance comparison for several OT variants(Average on 12 tasks on DomainNet). UOT and POT represent Unbalanced OT and Partial OT, respectively.

λ	P@50	P@100	P@200
0.001	62.79	61.07	57.25
0.005	66.77	65.21	61.64
0.01	68.68	67.04	63.53
0.05	65.58	64.06	60.69
0.1	65.02	63.66	60.33

Table 6: The effect of coefficient λ in equation 12 (Average on 12 tasks on DomainNet)

Visualization of the Retrieval Results

We offer a quantitative visualization of our proposed ProtoOT alongside the most promising state-of-the-art method, DD(Hu and Lee 2022), by showcasing cross-domain retrieval outcomes. The results presented in Figure 3, ordered from top to bottom, correspond to the retrieval results of the following four tasks: Sketch-Clipart, Clipart-Sketch, Art-Product, and Product-Art cross-domain retrieval. Upon observing the outcomes, it becomes evident that out ProtoOT is effective in eliminating the interference caused by samples that share visual similarity but belong to different categories.

Conclusion

In this paper, we have addressed the challenge of unsupervised cross-domain image retrieval (UCIR) by introducing ProtoOT, a novel Optimal Transport formulation explicitly designed for this task. We have shown that by unifying the intra-domain representation learning and the cross-domain feature alignment within the ProtoOT framework, we can capitalize on their synergistic potential and significantly enhance UCIR performance. Our experimental results have demonstrated the superiority of ProtoOT over existing state-of-the-art techniques across benchmark datasets.

Acknowledgments

This work was supported by Shanghai Local College Capacity Building Program (23010503100), Shanghai Sailing Program (21YF1429400, 22YF1428800), NSFC (No.62303319), Shanghai Frontiers Science Center of Human-centered Artificial Intelligence (ShangHAI), MoE Key Laboratory of Intelligent Perception and Human-Machine Collaboration (ShanghaiTech University), Shanghai Clinical Research and Trial Center and Shanghai Engineering Research Center of Intelligent Vision and Imaging.

References

- Asano, Y.; Rupprecht, C.; and Vedaldi, A. 2019. Self-labelling via simultaneous clustering and representation learning. In *International Conference on Learning Representations*.
- Bao, C.; Zhang, X.; Chen, J.; and Miao, Y. 2022. MMFL-net: multi-scale and multi-granularity feature learning for cross-domain fashion retrieval. *Multimedia Tools and Applications*, 1–33.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in neural information processing systems*, 33: 9912–9924.
- Chang, W.; Shi, Y.; Tuan, H.; and Wang, J. 2022. Unified optimal transport framework for universal domain adaptation. *Advances in Neural Information Processing Systems*, 35: 29512–29524.
- Chang, W.; Shi, Y.; and Wang, J. 2023. CSOT: Curriculum and Structure-Aware Optimal Transport for Learning with Noisy Labels. *arXiv preprint arXiv:2312.06221*.
- Chen, X.; Fan, H.; Girshick, R.; and He, K. 2020. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*.
- Courty, N.; Flamary, R.; Habrard, A.; and Rakotomamonjy, A. 2017. Joint distribution optimal transportation for domain adaptation. *Advances in neural information processing systems*, 30.
- Courty, N.; Flamary, R.; and Tuia, D. 2014. Domain adaptation with regularized optimal transport. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2014, Nancy, France, September 15–19, 2014. Proceedings, Part I 14*, 274–289. Springer.
- Cuturi, M. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26.
- Dai, Y.; Liu, J.; Ren, X.; and Xu, Z. 2020. Adversarial training based multi-source unsupervised domain adaptation for sentiment analysis. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 7618–7625.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Figalli, A. 2010. The optimal partial transport problem. *Archive for rational mechanics and analysis*, 195(2): 533–560.
- Flamary, R.; Courty, N.; Tuia, D.; and Rakotomamonjy, A. 2016. Optimal transport for domain adaptation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 1: 1–40.
- Gajic, B.; and Baldrich, R. 2018. Cross-domain fashion image retrieval. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 1869–1871.
- 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.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9729–9738.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hu, C.; and Lee, G. H. 2022. Feature representation learning for unsupervised cross-domain image retrieval. In *European Conference on Computer Vision*, 529–544. Springer.
- Huang, J.; Feris, R. S.; Chen, Q.; and Yan, S. 2015. Cross-domain image retrieval with a dual attribute-aware ranking network. In *Proceedings of the IEEE international conference on computer vision*, 1062–1070.
- Ji, X.; Wang, W.; Zhang, M.; and Yang, Y. 2017. Cross-domain image retrieval with attention modeling. In *Proceedings of the 25th ACM international conference on Multimedia*, 1654–1662.
- Kim, D.; Saito, K.; Oh, T.-H.; Plummer, B. A.; Sclaroff, S.; and Saenko, K. 2021. Cds: Cross-domain self-supervised pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9123–9132.
- Lee, T.; Lin, Y.-L.; Chiang, H.; Chiu, M.-W.; Hsu, W.; and Huang, P. 2018. Cross-domain image-based 3d shape retrieval by view sequence learning. In *2018 international conference on 3D vision (3DV)*, 258–266. IEEE.
- Li, J.; Lü, S.; and Li, Z. 2022. Unsupervised domain adaptation via softmax-based prototype construction and adaptation. *Information Sciences*, 609: 257–275.
- Li, J.; Zhou, P.; Xiong, C.; and Hoi, S. 2021. Prototypical Contrastive Learning of Unsupervised Representations. In *International Conference on Learning Representations*.
- Lin, H.; Zhang, Y.; Qiu, Z.; Niu, S.; Gan, C.; Liu, Y.; and Tan, M. 2022. Prototype-guided continual adaptation for class-incremental unsupervised domain adaptation. In *European Conference on Computer Vision*, 351–368. Springer.
- Long, M.; Zhu, H.; Wang, J.; and Jordan, M. I. 2017. Deep transfer learning with joint adaptation networks. In *International conference on machine learning*, 2208–2217. PMLR.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Peng, X.; Bai, Q.; Xia, X.; Huang, Z.; Saenko, K.; and Wang, B. 2019. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1406–1415.
- Redko, I.; Courty, N.; Flamary, R.; and Tuia, D. 2019. Optimal transport for multi-source domain adaptation under target shift. In *The 22nd International Conference on artificial intelligence and statistics*, 849–858. PMLR.

- Sain, A.; Bhunia, A. K.; Yang, Y.; Xiang, T.; and Song, Y.-Z. 2021. Stylemeup: Towards style-agnostic sketch-based image retrieval. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8504–8513.
- Saleh, M.; Wu, S.-C.; Cosmo, L.; Navab, N.; Busam, B.; and Tombari, F. 2022. Bending graphs: Hierarchical shape matching using gated optimal transport. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11757–11767.
- Séjourné, T.; Peyré, G.; and Vialard, F.-X. 2023. Unbalanced Optimal Transport, from theory to numerics. *Handbook of Numerical Analysis*, 24: 407–471.
- Shen, Z.; Feydy, J.; Liu, P.; Curiale, A. H.; San Jose Estepar, R.; San Jose Estepar, R.; and Niethammer, M. 2021. Accurate point cloud registration with robust optimal transport. *Advances in Neural Information Processing Systems*, 34: 5373–5389.
- 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.
- Villani, C.; et al. 2009. *Optimal transport: old and new*, volume 338. Springer.
- Wang, X.; Peng, D.; Yan, M.; and Hu, P. 2023. Correspondence-Free Domain Alignment for Unsupervised Cross-Domain Image Retrieval. In *The Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023*.
- Wu, Z.; Xiong, Y.; Yu, S. X.; and Lin, D. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3733–3742.
- Yue, X.; Zheng, Z.; Zhang, S.; Gao, Y.; Darrell, T.; Keutzer, K.; and Vincentelli, A. S. 2021. Prototypical cross-domain self-supervised learning for few-shot unsupervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13834–13844.
- Zhan, F.; Yu, Y.; Cui, K.; Zhang, G.; Lu, S.; Pan, J.; Zhang, C.; Ma, F.; Xie, X.; and Miao, C. 2021. Unbalanced feature transport for exemplar-based image translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15028–15038.
- Zhang, W.; Ouyang, W.; Li, W.; and Xu, D. 2018. Collaborative and adversarial network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3801–3809.