



Unsupervised dual-teacher knowledge distillation for pseudo-label refinement in domain adaptive person re-identification

Sidharth Samanta¹ · Debasish Jena¹ · Suwendu Rup²

Received: 11 May 2024 / Revised: 17 July 2024 / Accepted: 17 August 2024

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

Abstract

Unsupervised Domain Adaptation (UDA) in person re-identification (reID) addresses the challenge of adapting models trained on labeled source domains to unlabeled target domains, which is crucial for real-world applications. A significant problem in clustering-based UDA methods is the noise in pseudo-labels generated due to inter-domain disparities, which can degrade the performance of reID models. To address this issue, we propose the Unsupervised Dual-Teacher Knowledge Distillation (UDKD), an efficient learning scheme designed to enhance robustness against noisy pseudo-labels in UDA for person reID. The proposed UDKD method combines the outputs of two source-trained classifiers (teachers) to train a third classifier (student) using a modified soft-triplet loss-based metric learning approach. Additionally, a weighted averaging technique is employed to rectify the noise in the predicted labels generated from the teacher networks. Experimental results demonstrate that the proposed UDKD significantly improves performance in terms of mean Average Precision (mAP) and Cumulative Match Characteristic curve (Rank 1, 5, and 10). Specifically, UDKD achieves an mAP of **84.57** and **73.32**, and Rank 1 scores of **94.34** and **88.26** for Duke to Market and Market to Duke scenarios, respectively. These results surpass the state-of-the-art performance, underscoring the efficacy of UDKD in advancing UDA techniques for person reID and highlighting its potential to enhance performance and robustness in real-world applications.

Keywords Person re-identification · Unsupervised domain adaptation · Knowledge distillation · Soft triplet loss · mAP-weighted averaging

✉ Sidharth Samanta
c120005@iiit-bh.ac.in

Debasish Jena
Debasish@iiit-bh.ac.in

Suwendu Rup
srup.it@nitrr.ac.in

¹ Department of CSE, IIIT Bhubaneswar, Gothapatana, Bhubaneswar 751003, Odisha, India

² Department of IT, NIT Raipur, Amanaka, Raipur 492010, Chhattisgarh, India

1 Introduction

Visual surveillance plays a pivotal role in proactively managing risks and bolstering security through extensive person tracking across large areas. With the proliferation of CCTV cameras, traditional manual surveillance methods are increasingly becoming impractical, necessitating the adoption of automated video surveillance systems powered by computer vision. Facial features [1, 2] are widely considered as biometric identifiers due to their distinctiveness. Recent advancements in face recognition leveraging Robust Principal Component Analysis (RPCA) [3] and optimization techniques such as Grey Wolf Optimization [4, 5] have significantly propelled the field forward. Despite these advancements, challenges persist in CCTV environments where facial visibility can be compromised or completely obscured, prompting the use of person re-identification (reID) techniques to ensure reliable identification under such circumstances.

Person reID has emerged as a prominent area of research in computer vision in recent years, addressing the challenge of matching individuals across different visual contexts. This task is complex due to variations in poses, lighting conditions, and occlusions, resulting in significant differences in how individuals appear across images and videos. Applications of person reID include surveillance systems, human identification in videos, pedestrian tracking, and crowd analysis, highlighting its practical importance in various domains.

Although supervised reID methods [6] have achieved commendable performance, but they rely on labelled data from the same domain (dataset) for model training. However, when these supervised models are deployed in different domains, they frequently experience a significant performance drop [7].

In this context, Unsupervised Domain Adaptation (UDA) [9] plays a crucial role in adapting supervised reID models to new, unlabeled domains by mitigating the issue of domain shift. UDA is concerned with the transfer of knowledge from a labeled source domain to an unlabeled target domain. This knowledge transfer aims to enhance the model's performance in the target domain by leveraging insights gained from the source domain. Figure 1 provides a graphical representation of the UDA process for person reID. The schematic representation illustrates the adaptation process, which involves a feature extractor and classifier trained in a supervised manner. This adaptation leverages domain-invariant reID knowledge to generate pseudo-labels from target domain images.

Among various UDA methods, clustering-based UDA [10–12] has gained popularity over time, due to the ability to group similar looking individuals together. This approach involves by utilizing a pre-trained model to assign pseudo labels to target domain images, which are then employed for clustering. The advantage of clustering-based UDA is its ability to reduce the impact of noisy pseudo labels, which is considered as a common challenge in conventional clustering-based UDA methods [10].

One of the fundamental challenge in UDA models is handling label noise in the form of inaccuracies in pseudo labels generated for target domain data. Several recent studies have suggested different strategies to address this issue. Some approaches focus on refining pseudo labels to enhance the accuracy [13, 14], while others emphasize learning from soft pseudo labels [13, 15]. Furthermore, some efforts have been made towards developing noise-resilient methods [16, 17] to counter the adverse effects of inaccuracies in pseudo labels on reID model performance.

The motivation behind this study is to address the persistent issue of noisy pseudo labels in UDA for person reID. Noisy labels can significantly degrade the performance of reID models by introducing inaccuracies that lead to incorrect model learning, reducing the model's

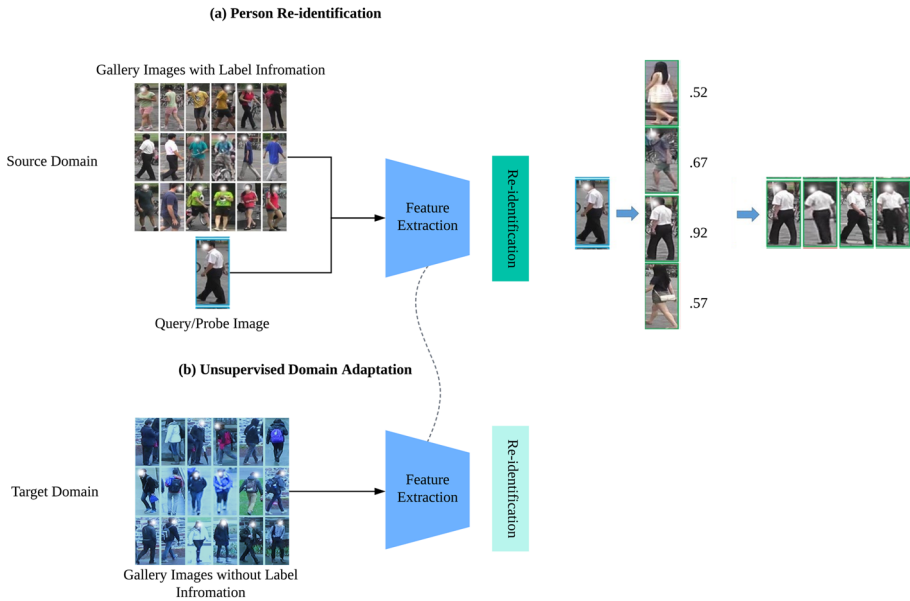


Fig. 1 A schematic representation of the Unsupervised Domain Adaptation process in person reID. The adaptation involves a feature extractor and classifier trained in a supervised manner, and it leverages domain-invariant reID knowledge to generate pseudo-labels from target domain images. The images of these persons are used from the Market1501 [8] dataset, and for the privacy concerns, the faces have been masked

ability to generalize across different domains. These inaccuracies stem from the erroneous assignment of labels during the pseudo-labeling process, which can cause the model to learn incorrect associations and patterns. To overcome this problem, we propose an Unsupervised Dual-Teacher Knowledge Distillation (UDKD) learning scheme that combines collaborative noise-rectification and learning from soft pseudo labels. This approach aims to enhance the robustness and accuracy of reID models in unsupervised settings by mitigating the adverse effects of noisy labels.

The workflow of the proposed method consists of three phases, namely teacher training (3.1), clustering (3.2), and knowledge distillation (3.3). For better clarity, the contribution of this paper can be summarized as follows:

1. An Unsupervised Dual-Teacher Knowledge Distillation (UDKD) method is proposed to enhance the robustness of person reID models.
2. An Enhancement has been suggested for the baseline model with a modified backbone architecture and feature learning technique.
3. To further improve robustness in person reID tasks, soft pseudo-labels from teachers are utilized, contributing to a more effective knowledge distillation process.
4. A modified soft triplet loss for UDKD is incorporated to optimize sample discrepancies and positive distances.
5. A comprehensive experimental evaluation has been conducted employing various benchmark datasets, and the proposed scheme shows superior performance than that of the state-of-the-art approaches.

The subsequent sections of this paper are structured as follows. Section 2 provides a comprehensive overview of recent advancements in cluster-based domain adaptive person reID.

In Section 3, the proposed method is extensively elucidated, highlighting its key components and underlying principles. The experimental setup and specific details are outlined in Section 4. To rigorously evaluate the effectiveness of the proposed scheme and its various components, we conduct a thorough testing and evaluation process. The results are presented and discussed in Section 5, along with a comparative analysis between our method and state-of-the-art approaches in UDA for person reID. Finally, the results are discussed in Section 6 and concluded in Section 7, summarizing the key findings and discussing potential avenues for future research.

2 Related works

In the field of person reID, Unsupervised Domain Adaptation (UDA) techniques are crucial due to the absence of human identity information in real-world deployment scenarios. UDA leverages reID knowledge from labelled source domains to compensate for the lack of labelled data in target or deployment domains. Recent research has underscored the competitive performance of clustering-based UDA methods when compared to supervised approaches. However, the existence of inter-domain disparities introduces noise into pseudo labels generated by pre-trained networks, impacting overall model performance. To address this challenge, we draw inspiration from existing works in the domain of person reID, clustering-based domain adaptation, and knowledge distillation.

In clustering-based Domain Adaptation approaches, [18] introduced a novel method for matching individuals across camera views by using spatiotemporal sequences. Furthermore, [19] improved pseudo-label accuracy through hierarchical clustering with hard-batch triplet loss. These developments have laid the foundation for research in this domain.

To tackle noisy pseudo labels, several strategies have been proposed. [13] introduced the Mutual Mean-Teaching (MMT), which employs a mutual teaching strategy for pseudo-label refinement. Similarly, [11] presented a two-branch architecture optimized for classification and metric learning to adapt to target domains, emphasizing domain adaptation and metric learning techniques.

A range of label refinement approaches, including Group-aware Label Transfer (GLT) by [20], probabilistic uncertainty-guided progressive label refinery (P²LR) by [21], and Noise Resistible Network (NRNet) by [22], have been suggested to tackle noisy pseudo labels. These approaches collectively address challenges related to intra-camera similarity, self-discrepancy, and feature distribution noise in different domains.

Knowledge distillation in UDA has been exemplified by methods such as Moving Semantic Transfer Network and Progressive Feature Alignment Network, both of which focus on enhancing semantic understanding and feature alignment, reducing domain divergence, and improving representation learning capabilities. Moreover, the Adversarial Double Mask based Pruning (ADMP) method by [23] employs advanced adversarial learning techniques to eliminate unnecessary features and enhance feature alignment in domain adaptation scenarios. These approaches underline the importance of leveraging knowledge from diverse sources to narrow domain gaps.

A major challenge in person re-identification (reID) is the need for extensive data annotation. The Color Prompting (CoP) method [24] addresses this issue by generating pseudo-supervision messages to facilitate data-free, continual, unsupervised domain adaptive reID. This approach leverages colour prompts to support continual learning and adaptation, effectively managing the domain shift over time. Similarly, the Adaptive Memorization with

Group labels (AdaMG) framework [25] emphasizes creating comprehensive sample descriptions and managing noisy pseudo labels in unsupervised reID tasks. By integrating group labels, AdaMG enhances model adaptability and robustness, resulting in improved performance under varying domain conditions.

In the context of multi-view, multi-person 3D pose estimation, the problem of domain shift is addressed using unsupervised domain adaptation with a dropout discriminator [26]. This method improves the accuracy of 3D pose estimation by aligning features across different views and individuals, significantly enhancing pose estimation robustness and generalization capabilities. Additionally, a global-local transformer-based framework for unsupervised reID [27] employs a multi-branch structure to learn robust features from pedestrian images. This approach highlights the effectiveness of transformer architectures in capturing both global and local features essential for accurate reID across diverse datasets and environments. Collectively, these studies demonstrate significant advancements in unsupervised domain adaptation techniques, encompassing innovative labelling strategies and transformative model architectures to enhance the reliability and applicability of reID systems.

Despite commendable efforts in addressing challenges related to clustering-based domain adaptation, pseudo-label generation, and knowledge distillation for person reID in existing literature, several persistent limitations remain. The presence of noisy pseudo-labels continues to affect model performance due to inter-domain disparities. Current strategies for pseudo-label refinement, although advanced, still suffer from issues like intra-camera similarity, self-discrepancy, and feature distribution noise. Additionally, extensive data annotation remains a significant challenge, particularly for unsupervised methods. Domain shifts in multi-view, multi-person 3D pose estimation and capturing both global and local features in transformer-based frameworks also present ongoing difficulties. These gaps underscore the necessity for more robust and effective solutions, such as our proposed UDKD scheme, to advance the field of unsupervised domain adaptation for person reID.

In response to this challenge, our UDKD scheme takes a novel approach by integrating the outputs of two teachers. This integration aims to minimize the impact of noisy pseudo-labels generated by one teacher, leveraging the complementary guidance provided by the other. An integral aspect of our UDKD method is the adoption of soft pseudo-labels, where the predicted probability distribution across all classes is considered for training the student network. This not only contributes to the robustness of the learning process but also addresses the issue of noisy pseudo-labels. To effectively handle these soft labels, we introduce a modified soft triplet loss into the scheme. This nuanced strategy collectively addresses the challenge of noisy pseudo-labels, thereby enhancing the efficacy of our proposed method in advancing unsupervised domain adaptation (UDA) for person reID. The subsequent section, 'Proposed Methodology,' delves into the intricate details of our approach, elucidating its components and demonstrating their impact on refining pseudo-labels in the domain adaptive setting.

3 Proposed methodology

Unsupervised domain adaptation is the process where data from one domain $D_s = (x_i^s, y_i^s)_{i=1}^m$ is used to train a model that can be applied to another domain $D_t = x_i^t|_{i=1}^n$, where x^s and x^t denote person images in the source and target domains, respectively, whereas m and n represent the number of images in the source and target domains. The goal is to learn a model that can effectively generalize to the new domain by predicting labels $y^{\sim t}$ for every target image x^t , even though no labelled data is available for that domain. This can be accomplished by either transferring knowledge from the source domain to the target domain or learning a model invariant to changes in the domain.

Clustering-based UDA involves pre-training the deep neural network $F(\cdot|\theta)$ with the target domain in order to encode the features $\{F(x_i^s|\theta)\}_{i=1}^m$. Followed by re-tuning the network parameters θ in order to transfer the encoded features $\{F(x_i^t|\theta)\}_{i=1}^n$ to the target domain.

The present work introduces an Unsupervised Dual-Teacher Knowledge Distillation (UDKD) learning scheme, aiming to ameliorate the challenges posed by the generation of pseudo labels during the clustering process. These pseudo-labels are produced to categorize data points based on their similarity, thus serving as a form of weak supervision. However, due to the inherent noise and imprecision associated with the clustering procedure, these labels may exhibit inaccuracies.

To mitigate the influence of such noise, the approach at hand capitalizes on the deployment of two independent classifiers. An averaging method is employed to minimize the potential impact of outliers within the pseudo labels, leading to the refinement of these labels and rendering them more dependable and robust. Subsequently, these refined soft pseudo labels assume a pivotal role in training a smaller network in a supervised manner, thereby augmenting the overall robustness of UDKD.

The methodology builds upon the combined architecture of the stronger baseline of Luo et al. [28] and the Group-aware Label Transfer network of Zheng et al. [29] as the initial method. This baseline forms the foundation for subsequent modifications with the introduction in two critical phases: Teacher Training (3.1) and Clustering (3.2). In Phase 1, we enhance the baseline by integrating two teacher models, denoted as C_1 and C_2 . These teachers undergo fine-tuning using the source domain, adding discriminative power to the original baseline. In Phase 2, a clustering mechanism is introduced to adapt the model to the target domain, refining features through k-means clustering. Architectural details of these modifications, along with a rationale and comparative analysis against the original stronger baseline, are provided in Appendix A. This approach aims to capitalize on the baseline's strengths while tailoring it to the specific demands of unsupervised domain adaptation for person reID.

For a comprehensive exposition of the proposed UDKD, inclusive of its procedural intricacies pertaining to training and knowledge distillation, refer to the end-to-end workflow diagram depicted in Fig. 2, which provides an illustrative visualization. In phase 1 (represented in blue and P1), image samples x^s of the source domain D_s are used to train the parameters θ_1 and θ_2 of classifiers C_1^s and C_2^s by using the source domain labels y^s and loss function \mathcal{L}^s . In phase 2 (represented in green and P2), the parameters of the same classifiers are re-tuned by using k-means clustering and loss function \mathcal{L}^c for every image sample x^t of the target domain D_t . In the last phase (represented in red and P3), the parameters θ_3 of the third classifier C^t are tuned by using the combined output $Y^{\sim t}$ of C_1^s and C_2^s using loss function \mathcal{L}_{st} for the target domain D_t . Additionally, the flowchart presented in Fig. 3 offers a step-by-step depiction of the sequential processes involved in the UDKD methodology. This flowchart provides a visual representation of the sequential steps involved in the UDKD process, including Dual-Teacher Training, Clustering, and Knowledge Distillation phases. Each phase is depicted with its corresponding input, processes, and output, highlighting the comprehensive approach of UDKD for effective unsupervised domain adaptation in person re-identification. Algorithm 1 succinctly encapsulates the essence of the training and knowledge distillation procedures for reference.

3.1 Phase 1: Dual-Teacher Training

UDKD begins by training two different CNNs separately with the source data D_s to model two feature transformation functions, $F(\cdot|\theta_1)$ and $F(\cdot|\theta_2)$. Here, each input sample x_i^s is trans-

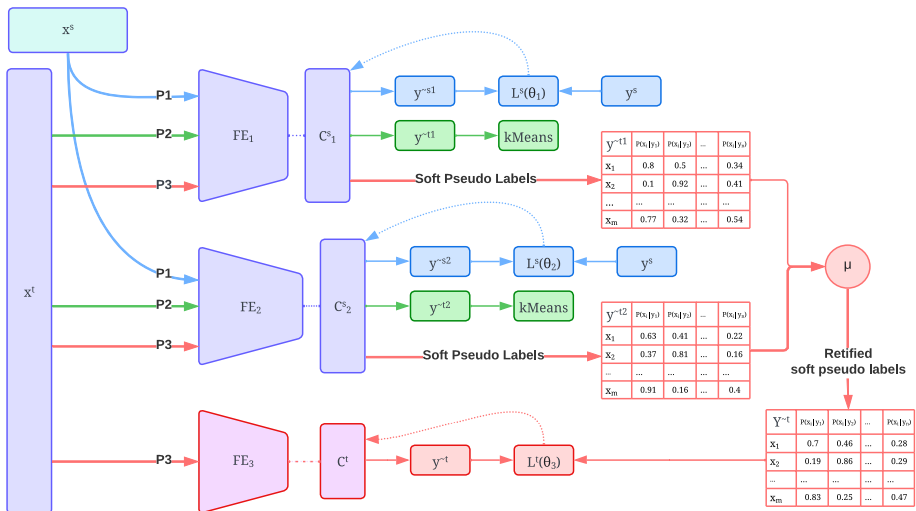


Fig. 2 End-to-end workflow diagram of the proposed UDKD learning scheme. In phase 1 (represented in blue and P1), image samples x^s of the source domain D_s are used to train the parameters θ_1 and θ_2 of classifiers C_1^s and C_2^s by using the source domain labels y^s and loss function \mathcal{L}^s . In phase 2 (represented in green and P2), the parameters of the same classifiers are re-tuned by using k-means clustering and loss function \mathcal{L}^c for every image sample x^t of the target domain D_t . In the last phase (represented in red and P3), the parameters θ_3 of the third classifier C^t are tuned by using the combined output $Y^{\sim t}$ of C_1^s and C_2^s using loss function \mathcal{L}_{st} for the target domain D_t

formed into two feature representations $y_i^{\sim s1}$ and $y_i^{\sim s2}$. Using these feature representations, the reID classifiers C_1^s and C_2^s produce two m -dimensional probability vectors corresponding to the predicted identities of the source domains. Here, m is the number of classes or identities present in the source domain. Using a combination of a triplet loss function \mathcal{L}_{trip}^s and an identity loss \mathcal{L}_{id} function, the CNNs are optimized to distinguish features belonging to distinct identities.

$$\mathcal{L}^s(\theta) = \mathcal{L}_{id}^s(\theta) + (\lambda^s \times \mathcal{L}_{id}^s(\theta)) \tag{1}$$

$$\mathcal{L}_{id}^s(\theta) = -\frac{1}{n} \sum_{i=1}^n y_i^s \log(y_i^{\sim s}) \tag{2}$$

$$\mathcal{L}_{trip}^s(\theta) = \frac{1}{n} \sum_{i=1}^n \max(0, \|y_i^{\sim s} - y_{i,p}^s\| + \chi - \|y_i^{\sim s} - y_{i,n}^s\|) \tag{3}$$

Where $y_{i,p}^s$ and $y_{i,n}^s$ represent positive and negative identities in each minibatch, \mathcal{L}_{ce} is the cross-entropy loss, $\|-\|$ represents the euclidean distance, χ is the margin and λ^s is the weighting parameter for source domain losses.

In summary of phase 1 of UDKD, the classifiers C_1^s and C_2^s are trained with the images x_i^s and ground truth y_i^s of source domain D_s to produce the predictions \hat{y}_i^s for each x_i^s . Further loss functions ($\mathcal{L}_{id}^s, \mathcal{L}_{id}^s$) are used to fine-tune the weights θ_1 and θ_2 to enhance the performance.

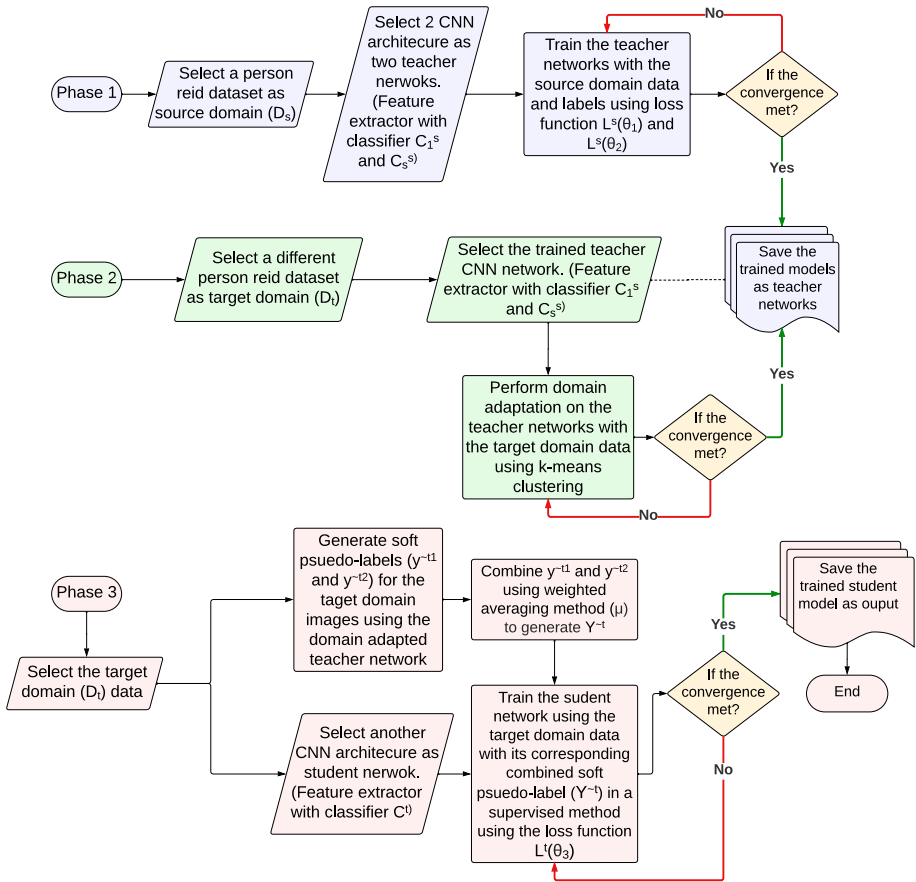


Fig. 3 Flowchart of the proposed Unsupervised Dual-Teacher Knowledge Distillation (UDKD) methodology. This flowchart provides a visual representation of the sequential steps involved in the UDKD process, including Dual-Teacher Training, Clustering, and Knowledge Distillation phases. Each phase is depicted with its corresponding input, processes, and output, highlighting the comprehensive approach of UDKD for effective unsupervised domain adaptation in person re-identification

3.2 Phase 2: Clustering

After training of two teacher networks, the classifiers C_1^s and C_2^s are used to generate soft pseudo-labels for k number of image samples x_i^t present in the target domain D_t . Each classifier generates n -dimensional probability vector for each x_i^t , in which each element represents the probability of the image x_i^t belongs to the pseudo id $y_i^{\sim t1}$ and $y_i^{\sim t2}$. (i.e. $P(x_i^t \in y_i^{\sim t1})$ and $P(x_i^t \in y_i^{\sim t2})$). Further, the k-means clustering algorithm is used to generate the probability vector (M_1, M_2, \dots, M_n) . Here, the parameters θ_1 and θ_2 of C_1^s and C_2^s are re-tuned using the \mathcal{L}_c loss function.

$$\mathcal{L}_c = \sum_{i=1}^k \sum_{j=1}^n \|x_i - M_j\|^2 \tag{4}$$

Algorithm 1 Unsupervised Dual-Teacher Knowledge Distillation (UDKD) training algorithm.

Require: Source and Target domain datasets D_s and D_t .
Require: Three CNN classifiers C_1^s , C_2^s , and C^t whose weights θ_1 , θ_2 , and θ_3 are pretrained on ImageNet.
Input : Values for S_s , S_t , λ^s , n , learning rate, $T'num_epochs$, $T''num_epochs$, $Snum_epochs$, χ for triplet loss and soft triplet loss.
Require: Assign the value of n as the value of k for k-means clustering.

- 1 **for** $p \leftarrow 1$ to $T'num_epochs$ **do**
- 2 **for** $i \leftarrow 1$ to m **do**
- 3 **Generate** the feature vectors $y_i^{\sim s1}$ and $y_i^{\sim s2}$, by using the classifiers C_1^s and C_2^s for image sample x_i^s of the source domain D_s ;
- 4 **Update** the parameters θ_1 and θ_2 of C_1^s and C_2^s by using the loss function (1);
- 5 **for** $q \leftarrow 1$ to $T''num_epochs$ **do**
- 6 **for** $i \leftarrow 1$ to n **do**
- 7 **Assign** each item of $y_i^{\sim s1}$ and $y_i^{\sim s2}$ to cluster M_i which has the closest mean;
- 8 **Calculate** new mean for each cluster;
- 9 **Update** θ_1 and θ_2 by using the loss function \mathcal{L}_c ;
- 10 **for** $r \leftarrow 1$ to $Snum_epochs$ **do**
- 11 **for** $i \leftarrow 1$ to n **do**
- 12 **Generate** soft pseudo label $y_i^{\sim t1}$ and $y_i^{\sim t2}$ for image sample x_i^t of D_t by using classifier C_1^s and C_2^s ;
- 13 **Combine** the output vectors $y_i^{\sim t1}$ and $y_i^{\sim t2}$ to generate a composite soft pseudo label $Y_i^{\sim t}$ by using (6) for $w = mAP$;
- 14 **Generate** the feature vector $y_i^{\sim t}$ for image sample x_i^t by using the student network classifier C^t ;
- 15 **Update** the parameters θ_3 of C^t by using the loss function (7);

3.3 Phase 3: knowledge distillation

In the final phase, the target domain image samples x^t are fed to the classifiers C_1^s and C_2^s to generate two probability vectors ($y_i^{\sim t1}$ and $y_i^{\sim t2}$). Subsequently, these two vectors are combined using an averaging function μ (as defined in (5)). The rationale behind this process lies in harnessing the strengths of multiple models to create a more robust ensemble. By leveraging the diverse perspectives and learned representations of each classifier, the ensemble aims to enhance overall performance, utilizing the complementary capabilities of the individual models. This ensemble-based approach contributes to improved accuracy and generalization across the target domain, ultimately achieving superior results in the context of unsupervised domain adaptation for person reID. The following reasons would justify the combination of the results of teacher models in ensemble learning, including:

- **Improved Prediction Accuracy** Combining the predictions of multiple models helps mitigate variance and enhance the overall accuracy of the ensemble. This proves beneficial, especially when individual models exhibit high bias or susceptibility to overfitting.
- **Increased Robustness** Utilizing multiple models contributes to the ensemble’s robustness against noise and outliers in the data. The presence of diverse models ensures that the impact of any single model is diluted, making the ensemble more resilient to erratic data patterns.

- **Leveraging Complementary Models** Different models may excel at predicting distinct facets of the data. Integrating predictions from these complementary models allows the ensemble to exploit their individual strengths, leading to an overall improvement in performance by capturing a broader spectrum of data characteristics.
- **Enhanced Generalization Ability** Training multiple models on different subsets of the data and amalgamating their predictions enhances the ensemble’s generalization ability. This approach enables the ensemble to grasp a wider range of patterns and relationships present in the data, ultimately contributing to superior performance across diverse scenarios.

$$Y_i^{\sim t} = \bigcup_{j=1}^n \mu(P(x_i^t \in y_j^{\sim t1}), P(x_i^t \in y_j^{\sim t2})) \tag{5}$$

In subsection 5.2, the simple averaging method μ_{sa} , as well as the weighted averaging method μ_w , are tested in different scenarios. We have considered the mAP and $R1$ scores of each teacher network C_1^s and C_2^s as weights in the weighted averaging method ($w = m$ and $w = r$). In conclusion, we utilize the mAP -weighted averaging method $\mu_{w=m}$ as a final approach because it performs better than the simple and $R1$ -weighted averaging methods in each scenario.

Simple Averaging Method:

$$\mu_{sa}(y_i^{\sim t1}, y_i^{\sim t2}) = \frac{y_i^{\sim t1} + y_i^{\sim t2}}{2}$$

Weighted Averaging Method (with w_1 and w_2 as weights):

$$\mu_w(y_i^{\sim t1}, y_i^{\sim t2}) = \frac{w_1 \cdot y_i^{\sim t1} + w_2 \cdot y_i^{\sim t2}}{w_1 + w_2} \tag{6}$$

Meanwhile, a comparatively smaller classifier C^t is fed with the target sample image x_i^t to obtain the feature vector $y_i^{\sim t}$. The loss function \mathcal{L}_{st} uses both $y_i^{\sim t}$ and $Y_i^{\sim t}$ to fine-tune the trainable parameters θ_3 of C^t . In place of the triplet loss \mathcal{L}_{trip}^s function used in the previous phase (3), a modified soft-triplet loss \mathcal{L}_{st} is used. As the classifier C^3 , deals with soft labels $Y_i^{\sim t}$, which is a collection of probabilities. Therefore, it cannot be handled by a triplet loss. In subsection 5.2, the proposed method is also tested with the triplet loss by computing $\max(Y_i^{\sim t})$.

The proposed soft triplet loss function \mathcal{L}_{st} is an adaptation of the original soft triplet loss \mathcal{L}_{trip} introduced in [13]. This adaptation addresses the complex demands of UDKD, which involves multiple teacher networks. The refined loss function incorporates a softmax function, ensuring its smooth and differentiable nature for effective backpropagation during model training.

$$\begin{aligned} \tau_{i,p} &= \exp(\chi - \|y_i^{\sim t} - \mu_{w=m}(y_{i,p}^{\sim t1}, y_{i,p}^{\sim t2})\|) \\ \tau_{i,n} &= \exp(\|y_i^{\sim t} - \mu_{w=m}(y_{i,n}^{\sim t1}, y_{i,n}^{\sim t2})\| - \chi) \\ \mathcal{L}_{st}(\theta_3) &= \frac{1}{n} \sum_{i=1}^n \log(1 + \tau_{i,p}) + \log(1 + \tau_{i,n}) \end{aligned} \tag{7}$$

In the UDKD context, the soft triplet loss serves the dual purpose of minimizing the discrepancy between positive and negative samples while maximizing the distance between the positive and the negative samples. This nuanced training objective is essential for acquiring

a feature representation capable of effectively discriminating between similar and dissimilar samples, aligning seamlessly with the unique architecture of UDKD. Appendix B provides comprehensive details on the modification of the original triplet loss, offering an in-depth exploration of the adjustments made to suit the distinctive features and requirements of UDKD.

To fine-tune the network parameters of C^t (students) in real time, the above procedures will be carried out concurrently. For this particular task, online learning or knowledge distillation is preferred to offline learning for several reasons. Firstly, online learning facilitates continuous improvement by perpetually refining and advancing the model, as it remains updated with the incorporation of fresh data. Secondly, it enhances scalability and data flexibility by enabling the model to dynamically adjust to shifts in the data distribution as they occur. Additionally, online learning proves advantageous in handling larger and more dynamic datasets, as it does not require storing all the data at once. Moreover, it offers the benefit of real-time performance by updating the model on the fly with new data, which can be critical in time-sensitive applications. Furthermore, online learning tends to be more computationally efficient, as it does not necessitate storing all the data simultaneously, leading to reduced computational costs. Importantly, it effectively addresses challenges such as concept drift, where the underlying distribution of the data changes over time, as well as non-stationary data, whose statistical properties evolve over time.

3.4 Model evaluation

In order to assess the effectiveness of the proposed UDKD approach, two prominent person reID datasets, namely DukeMTMC [30] and Market1501 [8], are utilized for evaluation purposes. Details of these datasets are provided in Table 1.

To evaluate and compare the performance of the UDKD model, two distinct combinations of source-to-target datasets are employed: Market-to-Duke and Duke-to-Market. Various evaluation metrics are utilized to gauge the model's performance, including mean average precision (mAP), as well as the Cumulative Matching Characteristic (CMC), ranks of 1, 5, and 10, respectively. These metrics are commonly employed in the field of person reID research and serve as benchmarks for evaluating the efficacy of different models.

4 Implementation details

4.1 Training and optimization details

In the experimental setup, we employ a Windows server machine equipped with an Intel Xeon Silver 4114 CPU operating at a frequency of 2.20 GHz, accompanied by a generous 64 GB of RAM. Complementing this configuration is a single NVIDIA Quadro RTX 8000 GPU endowed with an impressive 48 GB of dedicated Graphics-RAM. This computational infrastructure enables us to undertake both supervised pre-training on the source domain and unsupervised fine-tuning on the target domain.

During the training process, a mini-batch approach is adopted, where each batch consists of 64 images belonging to 16 distinct individuals. To ensure uniformity and compatibility within our network, all input images undergo a preprocessing step, wherein they are uniformly resized to dimensions of 256×128 pixels.

Table 1 Description of Market-1501 and DukeMTMC datasets

Characteristic	Market-1501	DukeMTMC
Number of Identities	1,501	1,404
Number of Images	32,668	36,411
Number of Cameras	6	8
Training Set Size	12,936 images of 751 identities	16,522 images of 702 identities
Testing Set Size	19,732 images of 750 identities	17,661 images of 702 identities
Query Set Size	3,368 images	2,228 images
Image Resolution	Varies, typically 128x64 pixels	Varies, typically 128x64 pixels
Annotation	Bounding boxes	Bounding boxes
Environment	Outdoor	Both indoor and outdoor
Challenges	Occlusions, varying lighting, background clutter	Occlusions, varying lighting, background clutter
Collection Method	Collected in front of a supermarket	Collected on Duke University campus

4.2 Hyperparameter tuning

The hyperparameter values were determined through a pragmatic "heat and trial" method, considering hardware limitations. Due to computational constraints, a manual iterative approach was adopted, allowing for adjustments based on intuition and observed model performance. This heuristic method proved effective in achieving satisfactory results within the given limitations. Further details regarding specific hyperparameter values used in the experiments are provided in Table 2.

4.3 Dual-teacher training

The teacher networks C_1^s and C_2^s employ the ResNet152 [31] and DenseNet169 [32] CNNs, respectively, which are pretrained on ImageNet [33]. These networks are trained independently using the image samples x^s and corresponding labels y^s from the source domain D_s . The training procedure, as outlined in [29], involves fine-tuning the weights θ_1 and θ_2 over 100 iterations, each consisting of 30 epochs. An initial learning rate of 0.00035 is set, which decreases to 1/10 of its previous value every 10 epochs.

4.4 Clustering

After the pre-training process, the images x^t of the target domains are used as input for the teacher networks to generate n number of clusters of identities by deploying the k-means clustering algorithm. With 0.0004 learning rate, 100 iterations of training consisting of 80 epochs are executed to re-tune the weights θ_1 and θ_2 . At the end, each cluster is considered as soft-pseudo-labels $y^{\sim t1}$ and $y^{\sim t2}$.

4.5 Knowledge distillation

Finally, C_1^s , C_2^s and C^t fed with the image samples x^t of D_t to generate refined pseudo-labels $y_i^{\sim t1}$ and $y_i^{\sim t2}$. Further, the $\mu_{w=m}$ is calculated for each sample by using the (6) to fine-tune the weights θ_3 of the student network (ResNet50) in an online manner. The training involves 400 iterations of 40 epochs with a 0.00035 learning rate.

5 Experimental setup and results

This section aims to describe, analyze, and signify the different components of the proposed scheme, such as the type of loss, averaging methods, and the number of clusters. The proposed

Table 2 Training Parameters for Each Stage

Stage	Iterations	Epochs per iteration	Learning rate
Dual-Teacher Training	100	30	0.00035
Clustering	100	80	0.0004
Knowledge Distillation	400	40	0.00035

Table 3 Experimental results of the proposed UDKD on Duke-to-Market and Market-to-Duke datasets with two kinds of loss functions (triplet loss \mathcal{L}_{trip}^t and soft-triplet loss \mathcal{L}_{st}^t), three types averaging methods (simple averaging μ_{sa} , R1-weighted averaging $\mu_{w=r}$, and mAP-weighted averaging $\mu_{w=m}$), and five different number of clusters ($n = 500, 750, 1000, 1250, 1500$)

Loss	Averaging	Clusters	Duke-to-Market			Market-to-Duke				
			mAP	Rank 1	Rank 5	Rank 10	mAP	Rank 1	Rank 5	Rank 10
\mathcal{L}_{trip}^t	μ_{sa}	500	73.09	85.16	91.58	93.36	58.41	79.65	86.74	93.74
\mathcal{L}_{trip}^t	μ_{sa}	750	68.92	78.35	87.74	93.21	52.54	68.18	86.44	92.85
\mathcal{L}_{trip}^t	μ_{sa}	1000	64.97	75.32	88.57	91.64	47.57	63.86	83.42	90.32
\mathcal{L}_{trip}^t	μ_{sa}	1250	63.02	70.47	83.16	89.36	47.03	58.17	82.49	91.22
\mathcal{L}_{trip}^t	μ_{sa}	1500	61.52	69.81	84.01	88.64	41.36	49.06	78.31	89.79
\mathcal{L}_{trip}^t	$\mu_{w=r}$	500	78.78	88.75	95.00	96.61	66.36	82.81	90.66	95.52
\mathcal{L}_{trip}^t	$\mu_{w=r}$	750	75.61	83.85	92.66	94.78	59.90	75.22	89.76	94.41
\mathcal{L}_{trip}^t	$\mu_{w=r}$	1000	72.46	80.29	90.93	93.55	56.22	69.74	88.28	93.68
\mathcal{L}_{trip}^t	$\mu_{w=r}$	1250	69.72	77.09	88.35	91.93	52.86	64.75	85.37	92.68
\mathcal{L}_{trip}^t	$\mu_{w=r}$	1500	65.26	75.87	86.94	91.18	48.28	56.38	82.49	91.60
\mathcal{L}_{trip}^t	$\mu_{w=m}$	500	80.74	89.66	95.44	97.17	67.17	83.62	91.75	95.73
\mathcal{L}_{trip}^t	$\mu_{w=m}$	750	77.35	84.84	93.42	95.20	61.70	76.78	90.67	95.02
\mathcal{L}_{trip}^t	$\mu_{w=m}$	1000	74.54	80.63	91.64	94.10	58.10	71.15	89.10	94.35
\mathcal{L}_{trip}^t	$\mu_{w=m}$	1250	70.35	77.99	89.31	92.56	54.57	65.60	86.03	92.90
\mathcal{L}_{trip}^t	$\mu_{w=m}$	1500	66.03	76.56	87.39	91.77	48.90	58.04	82.86	92.07
\mathcal{L}_{st}	μ_{sa}	500	77.50	87.03	90.82	96.45	62.95	79.57	91.32	95.47
\mathcal{L}_{st}	μ_{sa}	750	72.84	87.89	90.71	93.75	63.15	78.57	89.84	92.47
\mathcal{L}_{st}	μ_{sa}	1000	70.40	84.24	88.33	96.04	61.24	79.22	88.77	93.91
\mathcal{L}_{st}	μ_{sa}	1250	71.04	78.83	83.40	95.53	55.33	77.24	86.04	92.69
\mathcal{L}_{st}	μ_{sa}	1500	68.18	77.08	84.96	94.26	55.60	76.03	82.80	90.46

Table 3 continued

Loss	Averaging	Clusters	Duke-to-Market			Market-to-Duke				
			mAP	Rank 1	Rank 5	Rank 10	mAP	Rank 1	Rank 5	Rank 10
\mathcal{L}_{st}	$\mu_{w=r}$	500	84.03	93.80	96.82	98.74	72.74	87.77	94.79	97.13
\mathcal{L}_{st}	$\mu_{w=r}$	750	79.77	92.92	94.78	97.77	69.71	85.44	93.88	95.39
\mathcal{L}_{st}	$\mu_{w=r}$	1000	78.55	89.99	91.97	97.69	67.69	85.21	91.44	94.50
\mathcal{L}_{st}	$\mu_{w=r}$	1250	76.62	86.24	89.43	97.48	63.15	84.40	90.01	93.63
\mathcal{L}_{st}	$\mu_{w=r}$	1500	74.35	84.43	88.23	97.25	60.75	81.08	87.12	92.37
\mathcal{L}_{st}	$\mu_{w=m}$	500	<u>84.57</u>	<u>94.34</u>	<u>97.31</u>	<u>98.83</u>	<u>73.32</u>	<u>88.26</u>	<u>95.1</u>	<u>97.14</u>
\mathcal{L}_{st}	$\mu_{w=m}$	750	80.36	93.57	95.1	98	70.3	85.92	94.27	95.45
\mathcal{L}_{st}	$\mu_{w=m}$	1000	79.22	90.4	92.33	97.71	68.2	85.63	91.5	94.49
\mathcal{L}_{st}	$\mu_{w=m}$	1250	77.27	86.73	89.8	97.57	63.63	85.01	90.5	93.62
\mathcal{L}_{st}	$\mu_{w=m}$	1500	75.02	85	88.59	97.34	61.26	81.52	87.51	92.36

UDKD is evaluated using all possible combinations of two loss functions (triplet loss \mathcal{L}_{trip}^t and soft-triplet loss \mathcal{L}_{st}), three averaging methods (simple averaging μ_{sa} , R1-weighted averaging $\mu_{w=r}$, and mAP-weighted averaging $\mu_{w=m}$), and five different numbers of clusters ($n = 500, 750, 1000, 1250, 1500$).

To facilitate a clearer comprehension, Table 3 illustrates the significance of each component. This table presents the experimental results of the proposed UDKD on Duke-to-Market and Market-to-Duke datasets with various configurations. The detailed breakdown includes performance metrics when employing the triplet loss \mathcal{L}_{trip}^t and soft-triplet loss \mathcal{L}_{st} , alongside comparisons of the three averaging methods (simple averaging μ_{sa} , R1-weighted averaging $\mu_{w=r}$, and mAP-weighted averaging $\mu_{w=m}$). Additionally, it explores the impact of different numbers of clusters, ranging from 500 to 1500, on the overall performance. These evaluations underscore the relative importance and contribution of each component to the efficacy of the UDKD scheme, highlighting the nuanced improvements brought by specific configurations.

5.1 Soft triplet loss analysis

To investigate the necessity of the soft-triplet loss \mathcal{L}_{st} , we have compared the model's performance with or without using it. In Fig. 4, which illustrates the performance improvement achieved using soft-triplet loss compared to triplet loss in the context of a fixed number of clusters ($n = 500$) and using the mAP-weighted averaging method ($\mu_{w=m}$), the number of clusters $n = 500$ and pseudo-label averaging method $\mu_{w=m}$ are kept constant. In the context of the Market-to-Duke dataset, we observed an increase of 6.15% in mAP and a significant 4.64% in the rank 1 score. Similarly, in the Duke-to-Market dataset, there is a 3.83% and 4.68% increase in mAP and rank 1 score. Based on these observations, we adopted a soft triplet loss during the knowledge distillation process.

5.2 mAP-weighted averaging evaluation

Selecting an effective method for combining pseudo labels becomes paramount when dealing with diverse teacher networks exhibiting varying performance on the target domain. A

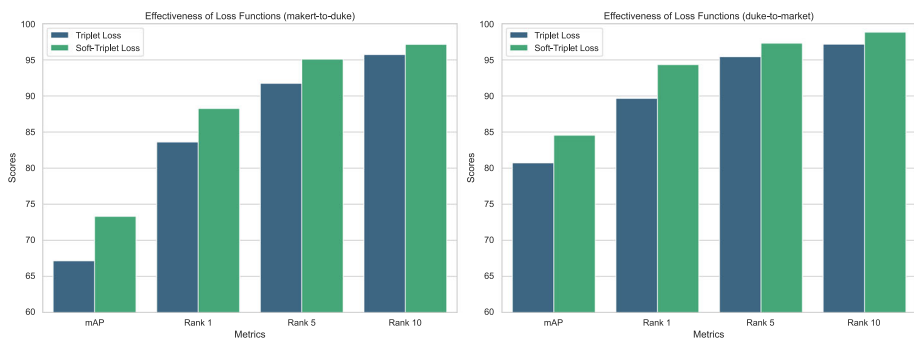


Fig. 4 Comparison of the significance of soft-triplet loss over triplet loss with fixed $n = 500$ and $\mu_{w=m}$. This figure illustrates the performance improvement achieved using soft-triplet loss compared to triplet loss in the context of a fixed number of clusters ($n = 500$) and using the mAP-weighted averaging method ($\mu_{w=m}$)

simplistic equal-weights (μ_{sa}) approach falls short, prompting the use of weighted averaging. Here, each teacher’s output is assigned a specific weight, considering its individual performance.

Our mAP-weighted averaging method, denoted as $\mu_{w=m}$, plays a pivotal role in refining pseudo labels within our UDKD learning scheme. Notably, we incorporate the mAP ($\mu_{w=m}$) and R1 ($\mu_{w=r}$) scores of each teacher network as their respective weights in the averaging process.

To underscore the significance of $\mu_{w=m}$, we conducted comparative evaluations with alternative averaging methods ($\mu_{w=r}$ and μ_{sa}), utilizing \mathcal{L}_{st} as the loss function and maintaining a constant cluster number ($n = 500$). Results, based on the Duke-to-Market dataset, reveal a 7.07% and 0.54% increase in mAP scores compared to the simple average and R1-weighted average, respectively. Similarly, in the Market-to-Duke dataset, a remarkable increase of 10.37% and 3.61% in mAP is observed with the mAP-weighted average over the other two methods. These findings underscore the efficacy of our chosen $\mu_{w=m}$ method. Refer to Fig. 5 for a detailed barplot comparison.

5.3 Impact of cluster number

Selecting a large number of clusters in clustering algorithms can engender overfitting, whereby the algorithm models the intrinsic noise within the data rather than discerning the underlying patterns. This can yield clusters that are excessively specific and lack generalizability to novel data instances. Conversely, opting for a low number of clusters may result in underfitting, where the algorithm fails to capture all pertinent patterns present in the data. Consequently, the produced clusters may prove excessively broad and inadequately specific. Thus, it is imperative to ascertain the appropriate number of clusters or employ methodologies such as the elbow method and silhouette score to determine the optimal number of clusters. By leveraging these techniques, one can navigate the challenge of cluster selection, thereby enhancing the fidelity and generalizability of the clustering results.

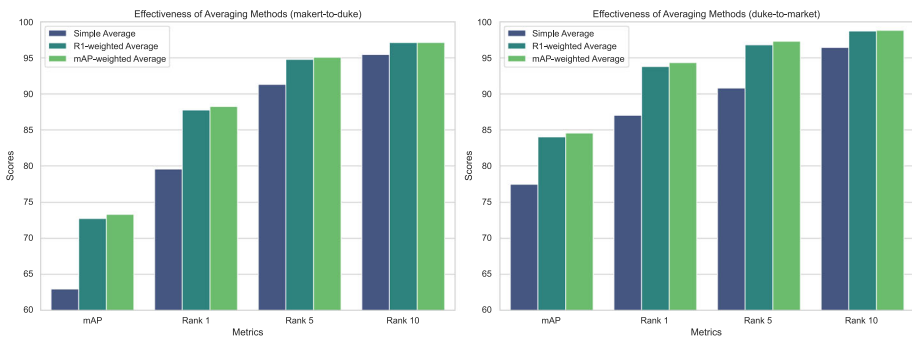


Fig. 5 Comparison of the significance of $\mu_{w=m}$ over $\mu_{w=r}$ and μ_{sa} with soft-triplet loss and $n = 500$. This figure demonstrates the superiority of mAP-weighted averaging ($\mu_{w=m}$) over other averaging methods ($\mu_{w=r}$ and μ_{sa}) in conjunction with soft-triplet loss, particularly in the scenario where the number of clusters is fixed at $n = 500$

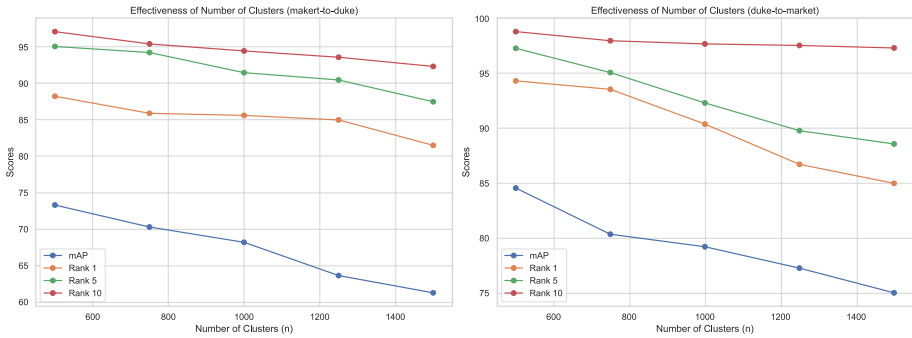


Fig. 6 Comparison of the significance of $n = 500$ over 750, 1000, 1200, and 1500 with soft-triplet loss and $\mu_w = m$. x -axis represents evaluation metrics, and y -axis represents the percentage (60-100%). This figure illustrates the impact of varying the number of clusters on performance metrics when using soft-triplet loss and the mAP-weighted averaging method

To establish a baseline, [13] opted for 500 clusters in their research, guiding our initial evaluation and ensuring a meaningful comparison. Our model underwent systematic testing across different cluster counts, including $n = 500, 750, 1000, 1250,$ and 1500 , utilizing the soft-triplet loss with mAP-weighted averaging. On the Market-to-Duke dataset, notable distinctions of 3.02% and 2.34% were observed in mAP and R1 scores, respectively, between $n = 500$ and $n = 750$. The Duke-to-Market dataset revealed an even more pronounced contrast, with differences of 4.21% and 6.77% for mAP and R1 scores, respectively. Given these insights, our model settles on $n = 500$ clusters. Figure 6 visually summarizes the detailed comparison through a barplot.

5.4 Comparative analysis with state-of-the-art methods

The evaluation of the proposed method is conducted by comparing its performance with state-of-the-art UDA methods for person re-ID using the Duke-to-Market and Market-to-Duke datasets. The assessment results, including the mAP, Rank 1, Rank 5, and Rank 10 scores for both datasets, are provided in Table 4. Notably, the experimental findings clearly

Table 4 The performance of the proposed UDKD compared with current state-of-the-art methods for UDA for person reID. Bold values indicate the highest score

Methods	Duke to Market				Market to Duke			
	mAP	Rank 1	Rank 5	Rank 10	mAP	Rank 1	Rank 5	Rank 10
MMT [13]	76.5	90.9	96.4	97.9	65.7	79.3	89.1	92.4
NRMT [16]	71.7	87.8	94.6	96.5	62.2	77.8	86.9	89.5
ANL [17]	77.4	91.4	96.6	97.7	64.1	79.2	84.9	91.3
NRNet [22]	78.6	91.1	69.4	98	67.6	79.8	90.1	93
MDL [34]	73.4	88.9	95.5	97.6	65.4	79.3	89.2	92.9
Dual-Refinement [14]	78	90.9	96.4	97.7	67.7	82.1	90.1	92.5
INCLR [35]	82.2	92.6	97.6	98.3	70.9	82.6	91.8	93.2
DUCL+PAFR [36]	84.0	93.9	97.5	98.6	72.7	84.9	92.0	93.9
UDKD (ours)	84.57	94.34	97.31	98.83	73.32	88.26	95.1	97.14

demonstrate the superior performance of the proposed method across all evaluation metrics, surpassing all existing approaches in a significant manner.

6 Discussion

The results from our experiments show that the proposed Unsupervised Domain Knowledge Distillation (UDKD) approach improves the performance of person re-identification (reID) systems across different domains. By using dual-teacher networks pre-trained on a source domain and incorporating soft-triplet loss and mAP-weighted averaging, we observed significant improvements in both mean Average Precision (mAP) and Cumulative Matching Characteristics (CMC) metrics. The dual-teacher framework, involving ResNet152 and DenseNet169, allows the student network, ResNet50, to learn diverse and robust features, leading to better generalisation of the target domain.

A key finding of this study is the effectiveness of the soft-triplet loss function over the traditional triplet loss. The soft-triplet loss provides a softer margin between positive and negative pairs, leading to more stable and effective training, which is particularly beneficial in an unsupervised setting where ground truth labels are absent. Additionally, the mAP-weighted averaging strategy for combining teacher network outputs significantly enhances the quality of pseudo-labels used for training the student network. This method outperforms other averaging techniques by prioritising the most relevant features, thereby improving the model's adaptability to new domains.

The number of clusters used for pseudo-label generation also plays a crucial role. Our experiments indicate that setting the number of clusters to 500 strikes a balance between intra-cluster compactness and inter-cluster separation, resulting in more accurate and meaningful pseudo-labels. When compared to other state-of-the-art unsupervised domain adaptation methods for person reID, the UDKD approach shows superior performance across various evaluation metrics. This highlights the robustness and effectiveness of our method, demonstrating its potential for practical applications in real-world reID systems.

One of the strengths of the UDKD approach is its ability to operate effectively in an unsupervised manner. By leveraging knowledge from pre-trained networks and generating high-quality pseudo-labels, the method reduces the reliance on labelled data, which is often scarce and expensive to obtain. Furthermore, the use of dual-teacher networks ensures that the student network benefits from diverse perspectives, leading to more comprehensive feature learning.

However, there are limitations to this study. Hardware constraints limited the extent of hyperparameter tuning, which may have affected the model's performance. Future work could explore more extensive hyperparameter optimisation and apply UDKD to a broader range of datasets and domains. Additionally, while this implementation uses two specific teacher networks, examining different combinations of teacher architectures could provide further insights into optimising the dual-teacher framework. Exploring other advanced clustering algorithms and loss functions could also enhance the model's performance.

7 Conclusion and future work

The suggested scheme presented an efficient UDKD learning scheme designed specifically for UDA in person reID tasks. The proposed method addresses the challenge of noisy pseudo

labels by leveraging the outputs of two large networks to minimize misclassifications and train a smaller classifier using the knowledge distilled from both networks.

Through an extensive experimental investigation, we have demonstrated the significant contributions of two key components: the mAP-weighted average method and the soft-triplet loss method. The mAP-weighted average method effectively combines the predictions of the dual teacher networks, providing a robust and reliable output for UDKD. Additionally, the modified soft-triplet loss facilitates improved discrimination between classes and enhances the discriminative capabilities of the smaller classifier.

The UDKD has exhibited superior performance compared to the current state-of-the-art UDA methods in Market-to-Duke and Duke-to-Market domain adaptation tasks. The comprehensive evaluation of the proposed approach demonstrates its effectiveness in achieving higher accuracy in person identification results. The proposed method achieves an mAP of **84.57** and **73.32**, and Rank 1 scores of **94.34** and **88.26** for Duke to Market and Market to Duke scenarios, respectively. These improvements underscore the efficacy of UDKD in advancing UDA techniques for person reID, highlighting its potential to enhance performance and robustness in real-world applications.

However, to further enhance the performance of UDKD, future research should prioritize addressing the challenges associated with noisy labels generated by clustering techniques. By mitigating the impact of noisy labels, the performance of UDKD can be elevated to a level comparable to that of fully supervised models.

Moreover, as an additional consideration, exploring the potential and effects of incorporating more than two teacher networks into the UDKD could provide valuable insights. Investigating the benefits and trade-offs of increasing the number of teacher networks could lead to improved knowledge distillation techniques and further advancements in UDA for person identification.

Appendix A Modifications of the baseline

The Stronger Baseline serves as the foundation for the proposed method, constructed based on Luo et al.'s work [28]. It employs a robust architecture forming the baseline for teacher training and clustering phases. Enhancements have been made for training a resilient Person ReID model, incorporating modifications introducing identity (ID) loss and triplet loss. Additionally, a clustering phase refines the model's capabilities, aiming to capture fine-grained features, enhance discrimination between identities, and optimize feature representations.

In evaluations, ResNet152 and DenseNet169 are utilized as backbone architectures. Superior performance is observed with the modified approach compared to the original baseline, as demonstrated in Table 5. Thus, the modified baseline is chosen as the teacher network due to its more accurate and robust feature representations.

The overall loss function is defined as:

$$\mathcal{L}(\theta) = \mathcal{L}_{\text{id}}(\theta) + (\lambda \times \mathcal{L}_{\text{trip}}(\theta))$$

Here, $\mathcal{L}_{\text{id}}(\theta)$ represents the identity loss function, $\mathcal{L}_{\text{trip}}(\theta)$ is the triplet loss function, and λ is the weighting parameter for the triplet loss.

The introduction of triplet loss enhances the model's ability to distinguish between individuals by creating anchor-positive and anchor-negative pairs. This encourages mapping similar images closer together in the feature space, thereby improving discriminative power.

Table 5 Comparison of the Baseline (BL) and Modified Baseline on various backbones for UDA Person reID

Models	Backbone	Duke to Market		Market to Duke	
		mAP	R1	mAP	R1
BL [28]	ResNet50	79.5	92.7	66.7	80.0
BL [28]	ResNet152	82.53	93.02	68.19	82.18
Our Baseline	ResNet152	83.09	94.85	69.54	82.25
BL [28]	DenseNet169	84.13	95.44	70.29	83.05
Our Baseline	DenseNet169	84.79	95.49	70.86	83.51

Identity loss is refined to align predicted probabilities closely with ground truth labels. This modification penalizes deviations from ground truth, emphasizing accurate identity predictions and improving overall classification accuracy.

The weighting parameter (λ) for triplet loss provides flexibility to fine-tune its contribution, balancing between identity preservation and feature discrimination as per specific task requirements.

$$\mathcal{L}_{\text{cluster}} = \sum_{j=1}^k \sum_{i=1}^n \|x_i - c_j\|^2$$

A clustering phase is introduced to further refine the model, employing K-means clustering to group similar instances. This guides the model to generate more cohesive and well-defined clusters in the feature space, enhancing both intra-class compactness and inter-class separability.

The clustering phase serves as a supplementary step to refine feature representations learned by the model. K-means clustering groups similar instances, contributing to the formation of compact and distinguishable clusters, thereby enhancing overall performance.

Appendix B Modification of the soft triplet loss

This appendix delves into the derivation and analysis of the modifications applied to the soft triplet loss, as outlined in the main manuscript. These modifications are crucial for adapting the loss function to the unique challenges of Unsupervised Dual-Teacher Knowledge Distillation (UDKD). We will dissect each step of the derivation, providing a clear understanding of how the original soft triplet loss is transformed into a version specifically tailored for UDKD.

The foundation of our analysis lies in the original soft triplet loss, denoted as $L_t^{\text{tri}}(\theta_1)$, which is rooted in the methodology presented in the cited study. This loss function is formally expressed as:

$$L_t^{\text{tri}}(\theta_1) = \frac{1}{N_t} \sum_{i=1}^{N_t} L_{\text{bce}} \left(T_i(\theta_1), T_i \left(E^{(T)}[\theta_2] \right) \right)$$

Here, $T_i(\theta_1)$ denotes the prediction made by the student network with parameters θ_1 for the i th sample in the target domain. Additionally, $T_i \left(E^{(T)}[\theta_2] \right)$ represents the temporal average prediction for the same sample, derived from the teacher network with parameters θ_2 through the knowledge distillation process. The variable N_t signifies the total number of samples present in the target domain. Our loss function, denoted by L_{bce} , quantifies the

disparity between the network's prediction and the temporal average prediction, facilitating the optimization process. This is mathematically represented as:

The original soft triplet loss encourages the student network to produce predictions that are closer to the temporal average predictions from the teacher network, guiding the student towards better knowledge acquisition.

However, directly applying the original soft triplet loss to UDKD poses challenges. In UDKD, we don't have access to ground-truth labels in the target domain. Instead, we rely on pseudo-labels generated through clustering, which can be noisy and imprecise. This noise can negatively impact the loss function, hindering the training process. To overcome these challenges, we propose modifications by incorporating Weighted Averaging and a Modified Loss Function with Exponential Terms.

Instead of solely relying on the current prediction y_t^i for positive samples, we introduce a weighted average of past predictions, denoted by $\mu_w(y_{t1}^p, y_{t2}^p)$. This mechanism leverages the valuable information contained in past predictions to provide a more stable and reliable representation for positive samples. The specific weighting scheme used in μ_w can be customized to suit the specific requirements of the UDKD task.

We replace the binary cross-entropy loss in the original formulation with a modified version that utilizes exponential terms:

$$L_{st}(\theta_1) = \frac{1}{N_t} \sum_{i=1}^{N_t} \log(1 + \tau_{i,p}) + \log(1 + \tau_{i,n})$$

Here, $\tau_{i,p}$ and $\tau_{i,n}$ represent the "softness" scores for positive and negative samples, respectively. These scores are calculated using exponential terms that capture the relative distances between the current prediction y_t^i and the weighted averages of positive and negative samples:

$$\tau_{i,p} = \exp(\chi - \|y_t^i - \mu_w(y_{t1}^p, y_{t2}^p)\|)$$

$$\tau_{i,n} = \exp(\|y_t^i - \mu_w(y_{t1}^n, y_{t2}^n)\| - \chi)$$

The parameter χ controls the steepness of the exponential functions, influencing how quickly the scores decrease with increasing distance. Higher values of χ lead to more focused relationships, where only very close samples are considered positive or negative, while lower values create softer relationships, allowing for some tolerance to noise in the pseudo-labels.

Appendix C List of abbreviations

Abbreviation	Full Form
AdaMG	Adaptive Memorization with Group labels
ADMP	Adversarial Double Mask based Pruning
ANL	Anti-Noise Learning
CCTV	Close Circuit Television
CMC	Cumulative Matching Characteristics
CNN	Convolutional Neural Network
CoP	Color Prompting
DUCL	Dual-Uncertainty Guided Curriculum Learning
GLT	Group-aware Label Transfer
INCLR	Intensifying the Consistency of Pseudo Label Refinement
MDL	Multi-Domain Learning
MMT	Mutual Mean Teaching

mAP	mean Average Precision
MTMC	Multi Tracking Multi Camera
NRMT	Noise Resistible Mutual-Training
NRNet	Noise Resistible Network
P ² LR	Probabilistic Uncertainty-guided Progressive Label Refinery
PAFR	Part-Aware Feature Refinement
reID	person re-identification
RPCA	Robust Principal Component Analysis
R1, R5 and R10	Rank 1, Rank 5 and Rank 10
UDA	Unsupervised Domain Adaptation
UDKD	Unsupervised Dual-teacher Knowledge Distillation

Acknowledgements The authors gratefully credit the Image and Video Processing Lab of IIIT Bhubaneswar for providing invaluable computational resources. The support and access to computing facilities have been instrumental in the successful execution of this research work.

Author Contributions All authors contributed equally to the conception, design, execution, and interpretation of the research study presented in this manuscript. Each author played a pivotal role in various aspects of the study, including but not limited to data collection, analysis, manuscript preparation, and critical revision. The collaborative effort of all authors was essential in ensuring the integrity and validity of the research findings reported herein.

Data Availability and Access The authors confirm that there is no new or modified dataset directly related to the findings reported in this manuscript. The study relied on publicly available datasets, which are appropriately referenced and cited within the manuscript. Detailed descriptions of the methods and materials utilized in the study are comprehensively provided in the manuscript to facilitate reproducibility and transparency.

Declarations

Competing Interests The authors declare that they have no competing interests that could potentially bias the interpretation or presentation of the research findings disclosed in this manuscript. There are no financial, personal, or professional relationships that could be perceived as influencing the conduct or outcomes of this study.

Ethical and Informed Consent for Data Used In our study, we utilized publicly available datasets for person re-identification. Appropriate measures were taken to uphold ethical standards and respect the privacy of individuals depicted in the images. Specifically, the facial features of the individuals were intentionally erased to prevent potential identification.

References

1. Ali AA, El-Hafeez TA, Mohany YK (2019) An accurate system for face detection and recognition. *J Advan Math Comput Sci* 33(3):1–19
2. Ali AA, El-Hafeez TA, Mohany YK (2019) A robust and efficient system to detect human faces based on facial features. *Asian J Res Comput Sci* 2(4):1–12
3. Eman M, Mahmoud TM, Ibrahim MM, Abd El-Hafeez T (2023) Innovative hybrid approach for masked face recognition using pretrained mask detection and segmentation, robust pca, and knn classifier. *Sensors* 23(15):6727
4. Taha ME, Mostafa T, El-Rahman A, Abd El-Hafeez T (2023) A novel hybrid approach to masked face recognition using robust pca and goa optimizer. *Scientific J Damietta Faculty Sci* 13(3):25–35
5. Saabia AA-B, El-Hafeez T, Zaki AM (2019) Face recognition based on grey wolf optimization for feature selection. In: *Proceedings of the international conference on advanced intelligent systems and informatics 2018 4*, Springer, pp 273–283
6. Liu H, Guo F, Xia D (2021) Domain adaptation with structural knowledge transfer learning for person re-identification. *Multimed Tool Appl* 80(19):29321–29337

7. Zhang Z, Zhang H, Liu S (2021) Person re-identification using heterogeneous local graph attention networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 12136–12145
8. Zheng L, Shen L, Tian L, Wang S, Wang J, Tian Q (2015) Scalable person re-identification: a benchmark. In: Proceedings of the IEEE international conference on computer vision, pp 1116–1124
9. Xiang S, Fu Y, Xie M, Yu Z, Liu T (2020) Unsupervised person re-identification by hierarchical cluster and domain transfer. *Multimedia Tool Appl* 79:19769–19786
10. Chen S, Fan Z, Yin J (2020) Pseudo label based on multiple clustering for unsupervised cross-domain person re-identification. *IEEE Signal Process Lett* 27:1460–1464
11. Dubourvieux F, Audigier R, Loesch A, Ainouz S, Canu S (2021) Unsupervised domain adaptation for person re-identification through source-guided pseudo-labeling. In: 2020 25th International Conference on Pattern Recognition (ICPR), IEEE, pp 4957–4964
12. Feng H, Chen M, Hu J, Shen D, Liu H, Cai D (2021) Complementary pseudo labels for unsupervised domain adaptation on person re-identification. *IEEE Trans Image Process* 30:2898–2907
13. Ge Y, Chen D, Li H (2020) Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. [arXiv:2001.01526](https://arxiv.org/abs/2001.01526)
14. Dai Y, Liu J, Bai Y, Tong Z, Duan L-Y (2021) Dual-refinement: joint label and feature refinement for unsupervised domain adaptive person re-identification. *IEEE Trans Image Process* 30:7815–7829
15. Huang T, Chen L (2021) Unsupervised domain adaptation person re-identification method based on softened pseudo labeling. In: 2021 2nd international conference on artificial intelligence and information systems, pp 1–5
16. Zhao F, Liao S, Xie G-S, Zhao J, Zhang K, Shao L (2020) Unsupervised domain adaptation with noise resistible mutual-training for person re-identification. In: European conference on computer vision, Springer, pp 526–544
17. Zhang H, Han S, Pan X, Zhao J (2020) Anl: anti-noise learning for cross-domain person re-identification. [arXiv:2012.13853](https://arxiv.org/abs/2012.13853)
18. Sekh AA, Dogra DP, Choi H, Chae S, Kim I-J (2020) Person re-identification in videos by analyzing spatio-temporal tubes. *Multimedia Tool Appl* 79:24537–24551
19. Zeng K, Ning M, Wang Y, Guo Y (2020) Hierarchical clustering with hard-batch triplet loss for person re-identification. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 13657–13665
20. Zheng K, Liu W, He L, Mei T, Luo J, Zha Z-J (2021) Group-aware label transfer for domain adaptive person re-identification. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 5310–5319
21. Han J, Li Y-L, Wang S (2022) Delving into probabilistic uncertainty for unsupervised domain adaptive person re-identification. In: Proceedings of the AAAI conference on artificial intelligence, vol 36, pp 790–798
22. Zhang S, Zeng Y, Hu H, Liu S (2021) Noise resistible network for unsupervised domain adaptation on person re-identification. *IEEE Access* 9:60740–60752
23. Feng X, Yuan Z, Wang G, Liu Y (2020) Admp: an adversarial double masks based pruning framework for unsupervised cross-domain compression. [arXiv:2006.04127](https://arxiv.org/abs/2006.04127)
24. Gu J, Luo H, Wang K, Jiang W, You Y, Zhao J (2023) Color prompting for data-free continual unsupervised domain adaptive person re-identification. [arXiv:2308.10716](https://arxiv.org/abs/2308.10716)
25. Peng J, Jiang G, Wang H (2023) Adaptive memorization with group labels for unsupervised person re-identification. *IEEE Trans Circuits Syst Video Technol* 33(10):5802–5813
26. Deng J, Yao H, Shi P (2023) Enhanced 3d pose estimation in multi-person, multi-view scenarios through unsupervised domain adaptation with dropout discriminator. *Sensors* 23(20):8406
27. Xie M-S, Chen S-B (2023) Global-local transformer for unsupervised person re-identification. In: Fourth International Conference on Artificial Intelligence and Electromechanical Automation (AIEA 2023), SPIE, vol 12709, pp 241–246
28. Luo H, Jiang W, Gu Y, Liu F, Liao X, Lai S, Gu J (2019) A strong baseline and batch normalization neck for deep person re-identification. *IEEE Trans Multimedia* 22(10):2597–2609
29. Zheng K, Liu W, He L, Mei T, Luo J, Zha Z-J (2021) Group-aware label transfer for domain adaptive person re-identification. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp 5310–5319
30. Gou M, Karanam S, Liu W, Camps O, Radke RJ (2017) Dukemtmc4reid: A large-scale multi-camera person re-identification dataset. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops, pp 10–19
31. He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp 770–778

32. Huang G, Liu Z, Van Der Maaten L, Weinberger KQ (2017) Densely connected convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp 4700–4708
33. Deng J, Dong W, Socher R, Li L-J, Li K, Fei-Fei L (2009) Imagenet: a large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition, IEEE, pp 248–255
34. Xu X, Zhang L (2021) Rectifying pseudo label by mutual disagreement learning for unsupervised domain adaptation person re-identification. In: 2021 IEEE International Conference on Multimedia and Expo (ICME), IEEE, pp 1–6
35. Zha L, Chen Y, Zhou P, Zhang Y (2023) Intensifying the consistency of pseudo label refinement for unsupervised domain adaptation person re-identification. In: 2023 IEEE International Conference on Multimedia and Expo (ICME), IEEE, pp 1547–1552
36. Liu Z, Liu B, Zhao Z, Chu Q, Yu N (2023) Dual-uncertainty guided curriculum learning and part-aware feature refinement for domain adaptive person re-identification. ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 1–5

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.