

Image-Text Knowledge Modeling for Unsupervised Multi-Scenario Person Re-Identification

Zhiqi Pang¹, Lingling Zhao¹, Yang Liu¹, Chunyu Wang^{1*}, Gaurav Sharma²

¹Faculty of Computing, Harbin Institute of Technology, Harbin, China

²Department of Electrical and Computer Engineering, University of Rochester, Rochester, USA
zqpang98@gmail.com, zhaoll@hit.edu.cn, liuyang@hit.edu.cn, chunyu@hit.edu.cn, g.sharma@ieee.org

Abstract

We propose unsupervised multi-scenario (UMS) person re-identification (ReID) as a new task that expands ReID across diverse scenarios (cross-resolution, clothing change, etc.) within a single coherent framework. To tackle UMS-ReID, we introduce image-text knowledge modeling (ITKM) – a three-stage framework that effectively exploits the representational power of vision-language models. We start with a pre-trained CLIP model with an image encoder and a text encoder. In Stage I, we introduce a scenario embedding in the image encoder and fine-tune the encoder to adaptively leverage knowledge from multiple scenarios. In Stage II, we optimize a set of learned text embeddings to associate with pseudo-labels from Stage I and introduce a multi-scenario separation loss to increase the divergence between inter-scenario text representations. In Stage III, we first introduce cluster-level and instance-level heterogeneous matching modules to obtain reliable heterogeneous positive pairs (e.g., a visible image and an infrared image of the same person) within each scenario. Next, we propose a dynamic text representation update strategy to maintain consistency between text and image supervision signals. Experimental results across multiple scenarios demonstrate the superiority and generalizability of ITKM; it not only outperforms existing scenario-specific methods but also enhances overall performance by integrating knowledge from multiple scenarios.

1 Introduction

Given a query image of an individual, person re-identification (ReID) (Ye et al. 2021b,a; Li et al. 2024; Gong et al. 2024; Tan et al. 2024) seeks to identify images of the same person from a gallery containing a large number of person images. To reduce human labor and tedium, researchers have explored various unsupervised traditional ReID (UT-ReID) methods (Wang et al. 2021; Xuan and Zhang 2021; Chen, Lagadec, and Bremond 2021; Dai et al. 2022; He et al. 2024), which typically utilize pseudo-labels generated by clustering algorithms in place of manually annotated identity labels to guide model optimization. Although advanced UT-ReID methods have achieved promising performance in simple scenarios, they often struggle to handle more challenging scenarios. Therefore, researchers have explored var-

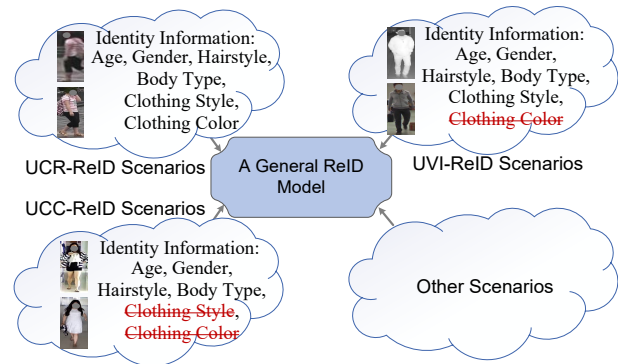


Figure 1: Identity information is scenario-dependent.

ious challenging scenarios, such as unsupervised visible-infrared ReID (UVI-ReID) (Pang et al. 2023; Shi et al. 2024; Yang, Chen, and Ye 2023; Wang et al. 2025), unsupervised clothing change ReID (UCC-ReID) (Pang, Zhao, and Wang 2024a; Pang et al. 2025), and unsupervised cross-resolution ReID (UCR-ReID) (Pang, Zhao, and Wang 2024b), and proposed methods tailored to each scenario.

The nature of identity-related information is typically different across various scenarios. As depicted in Figure 1, in UCR-ReID, both the style and color of clothing are considered identity-related information (Pang, Zhao, and Wang 2024b). However, color is irrelevant for identity in UVI-ReID (Yang et al. 2022) because of the lack of color information in infrared images. Furthermore, UCC-ReID typically assumes that the same individual may wear different clothing, thus the style and color of clothing might not be identity-related (Pang et al. 2025). Because the afore-mentioned unsupervised scenario-specific ReID (USS-ReID) methods focus on identity-related features unique to each scenario, limiting their generalizability across multiple scenarios. This results in two key limitations. First, practical ReID applications often span multiple scenarios—for instance, a person may change clothing or appear under low-light conditions. Existing USS-ReID methods typically require separate models for each scenario, involving distinct training strategies and multiple sets of weights, which increases both system complexity and deployment cost. Second, existing USS-ReID methods are

*Corresponding author: Chunyu Wang

designed for individual scenarios and do not support joint training across multiple scenarios, thus failing to leverage the performance benefits of increased data diversity (Zheng et al. 2024; Shi et al. 2023).

To address the aforementioned issues, we propose a novel image-text knowledge modeling (ITKM) framework that effectively leverages the representational power of vision-language models (VLMs) (Radford et al. 2021) for addressing the challenging unsupervised multi-scenario person re-identification (UMS-ReID). UMS-ReID is a novel task that aims to train a single general model in an unsupervised manner to handle multiple diverse scenarios—unlike prior tasks (Yang et al. 2022; Pang et al. 2025), which target individual scenarios. To tackle the challenges of UMS-ReID, ITKM introduces advances in both data processing and model optimization. In terms of data, based on the relationship between data heterogeneity (e.g., modality gaps, clothing variations, and resolution disparities) and encoder structure (Wu and Ye 2023; Shi et al. 2024), we consistently divide the heterogeneous data in each scenario into two homogeneous groups, ensuring that the image encoder in ITKM can accommodate inputs from all scenarios. In terms of optimization, given the powerful representation and transfer capabilities of VLMs, we design ITKM as a three-stage framework based on CLIP (Radford et al. 2021; Zhou et al. 2022). In Stage I, we introduce a scenario embedding in the CLIP image encoder and fine-tune the encoder to adapt and effectively leverage knowledge from multiple scenarios. In Stage II, to obtain scenario-specific text representations, we: (a) optimize a set of learned text embeddings to associate, through the CLIP, with identity pseudo-labels from Stage I and (b) introduce a multi-scenario separation loss to increase the divergence between inter-scenario text representations. Stage III first introduces cluster-level heterogeneous matching (CHM) and instance-level heterogeneous matching (IHM) modules to obtain reliable heterogeneous positive pairs (e.g., a visible image and an infrared image of the same person, images of the same person wearing different clothing, or images of the same person with different resolutions) within each scenario. Next, Stage III introduces a dynamic text representation update (DRU) strategy that is guided by the latest pseudo-labels and helps maintain consistency between text and image supervision signals in unsupervised setting.

The main contributions are summarized as follows:

- We propose a novel task, UMS-ReID, along with a tailored framework, ITKM, to address it. To the best of our knowledge, this is the first study to explore unsupervised person re-identification across multiple scenarios.
- To adapt to different scenarios and encourage scenario-specific text representations, we introduce a novel scenario embedding for the CLIP image encoder and propose a multi-scenario separation loss.
- We construct two new modules, CHM and IHM, for obtaining reliable heterogeneous positive pairs that are critical for learning in our unsupervised setting. Additionally, we introduce the DRU strategy to ensure consistency between text and image supervision signals.
- Experimental results on datasets from different scenar-

ios validate the superiority and generalization of ITKM, which not only surpasses existing USS-ReID methods within each scenario but also enhances overall performance by incorporating multi-scenario knowledge.

2 Related Work

UT-ReID methods (Cho et al. 2022; Lan et al. 2023) typically assume that different images of the same person exhibit high similarity. However, real-world applications often involve multiple challenging scenarios, where different images of the same person may exhibit significant heterogeneity. These real-world complexities are partly addressed in prior research by establishing specific tasks and tailored methods. For example, given a visible (infrared) image of a person, UVI-ReID (Yang, Chen, and Ye 2024) aims to find images with the same identity in a gallery composed of infrared (visible) images. UCC-ReID (Li et al. 2023) typically assumes that people may change clothes and seeks to match images of the same person wearing different clothing. Additionally, UCR-ReID (Pang, Zhao, and Wang 2024b) aims to match low-resolution images with high-resolution images of the same person. Building on prior work, this paper proposes a novel UMS-ReID method that supports multiple scenarios simultaneously, aiming to promote broader real-world applications of ReID.

Pre-trained VLMs (Radford et al. 2021; Zhou et al. 2022; Lin et al. 2024) have demonstrated promising transferability in downstream tasks. A pioneer in this area is the CLIP (Radford et al. 2021) model that trains a pair of image and text encoders by maximizing the representation similarity between matched image-text pairs. Subsequent research, such as CoOp (Zhou et al. 2022), leverages learnable text embeddings to further enhance the flexibility of CLIP. Researchers in the ReID field have also explored CLIP-based methods (Li, Sun, and Li 2023; Chen et al. 2023; Yang et al. 2024). For example, CLIP-ReID (Li, Sun, and Li 2023) first learns a set of text embeddings for each identity, then uses the learned text embeddings to assist in optimizing the image encoder. CCLNet (Chen et al. 2023) optimizes learnable text embeddings based on pseudo-labels in the UVI-ReID task and uses the text representations as additional supervisory signals. The proposed ITKM framework also benefits from CLIP’s representation capabilities. Unlike prior works, we focus on the UMS-ReID task and explicitly introduce DRU to update text representations, addressing the limitations of outdated offline representations used in previous methods.

3 Proposed Framework

To illustrate the UMS-ReID task, we use a multi-scenario setting that includes UVI-ReID (Yang, Chen, and Ye 2024), UCC-ReID (Pang, Zhao, and Wang 2024a), and UCR-ReID (Pang, Zhao, and Wang 2024b). During the training phase, we define the combination of the training sets from these three scenarios as $\{X^s\}_{s=1}^S$, where s indexes the $S=3$ scenarios. This paper aims to demonstrate the potential of UMS-ReID. Accordingly, we do not prioritize large-scale datasets or extensive model parameters, leaving these aspects for future work. Given the data distribution, we first

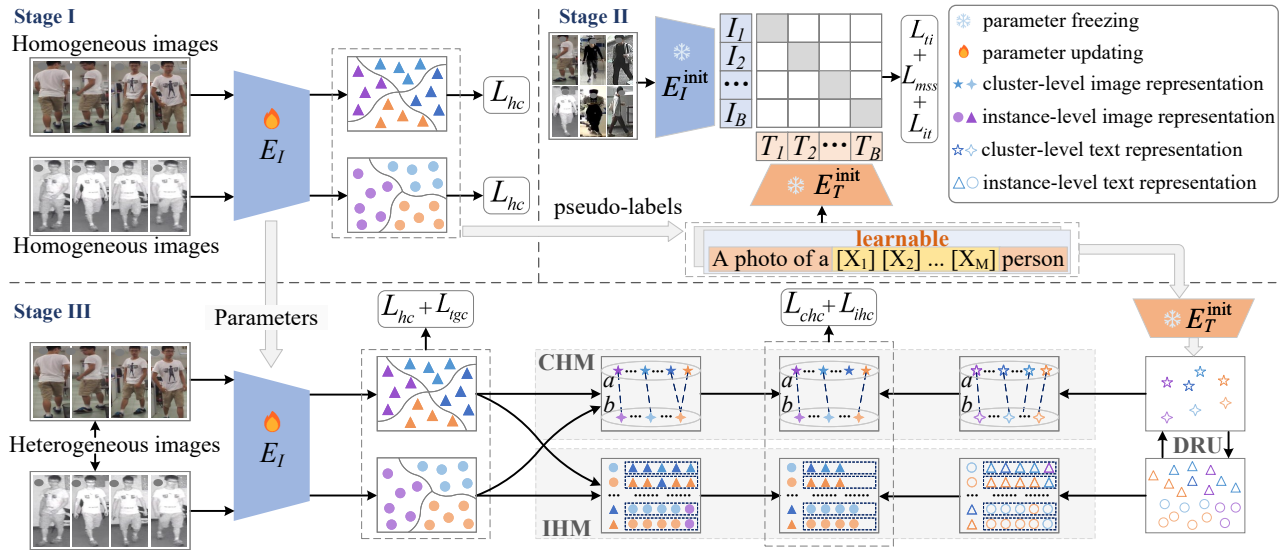


Figure 2: The proposed ITKM framework for UMS-ReID consists of three stages. Stage I performs unsupervised homogeneous learning to generate pseudo-labels. Stage II learns text embeddings in the sentence “A photo of a $[X_1][X_2]\dots[X_M]$ person” to associate it with the images for a homogeneous pseudo-label via CLIP. Stage III conducts unsupervised heterogeneous learning using CHM and IHM. In Stages I and III, colors indicate pseudo-labels and shapes indicate homogeneous groups.

divide the heterogeneous images within each scenario into two homogeneous groups, labeled as a and b . As shown in Figure 2, our proposed ITKM framework consists of three stages: *unsupervised homogeneous learning* (Stage I), *text representation learning* (Stage II), and *unsupervised heterogeneous learning* (Stage III).

3.1 Unsupervised Homogeneous Learning

In Stage I, we adapt the image encoder from a pre-trained CLIP model to transform an input image into a scenario-specific representation that preserves identity information. Specifically, we adapt the vision transformer (ViT) (Dosovitskiy et al. 2021) E_I^{init} , which constitutes the CLIP image encoder, to introduce a dual-branch frontend, where each branch is used for one homogeneous group, and, additionally, introduce a scenario embedding in the image encoder. The process is illustrated in Figure 3 using an input image $x_{a,m}^s$ from the group a of s -th scenario. The image $x_{a,m}^s$ is partitioned into N patches which are flattened into vectors x_1, x_2, \dots, x_N and a learnable embedding matrix P maps each of the patches to a $1 \times D$ vector. The class embedding is concatenated at the head of the sequence of vectors and position and scenario are then additively added into the resulting vector to obtain an $(N+1)$ -length sequence of $1 \times D$ vectors

$$z_0 = [z_{a,m}^{cls} + e_s; [x_1; x_2; \dots; x_N]P] + [z_p^0; \dots; z_p^N] \quad (1)$$

where $[z_p^0; \dots; z_p^N]$ represent the position embeddings and e_s is the learnable scenario embedding. The flattening and projection of the patches, concatenation of $z_{a,m}^{cls}$, and the position embedding are based on the standard ViT approach (Dosovitskiy et al. 2021). The scenario embedding is introduced in this work to address our UMS-ReID setting.




Figure 3: Architecture of image encoder used to obtain identity representations $f_{a,m}^{s,i}$ from an input image $x_{a,m}^s$.

As shown in Figure 3, z_0 is processed through a series of transformer layers of the ViT, producing (at the head of the final layer) the identity representation $f_{a,m}^{s,i}$, which is referred to as the instance-level image representation in the following description. The instance-level image representation $f_{b,m}^{s,i}$ for group b is similarly obtained. The instance-level image representations are then clustered within each homogeneous group to obtain homogeneous pseudo-labels.

We then calculate cluster-level image representations. For example, the representation of the u -th cluster in group a in the s -th scenario is defined as:

$$c_{a,u}^{s,i} = \frac{1}{N_{a,u}^s} \sum_{m=1}^{N_{a,u}^s} f_{a,m}^{s,i} \quad (2)$$

where $N_{a,u}^s$ represents the number of images in this cluster and $f_{a,m}^{s,i}$ denotes the instance-level image representation in this cluster. Subsequently, we refine the image encoder E_I^{init} based on homogeneous learning, and the optimized image encoder is referred to as E_I . Specifically, the homogeneous contrastive loss is introduced to optimize E_I based

on the pseudo-labels. For an instance-level image representation $f_{a,m}^{s,i}$, the homogeneous contrastive loss is defined as:

$$L_{hc}^{s,a,m} = -\log \frac{\exp(f_{a,m}^{s,i} \cdot c_{a,u}^{s,iT} / \tau)}{\sum_{v=1}^{C_a^{s,i}} \exp(f_{a,m}^{s,i} \cdot c_{a,v}^{s,iT} / \tau)}, \quad (3)$$

where $c_{a,u}^{s,i}$ is from the same cluster as $f_{a,m}^{s,i}$, $C_a^{s,i}$ represents the number of clusters in group a in the s -th scenario, and τ is the temperature hyperparameter. For an instance-level representation $f_{b,m}^{s,i}$ in group b , the corresponding loss $L_{hc}^{s,b,m}$ is obtained analogously. The overall homogeneous contrastive loss is then obtained as:

$$L_{hc} = \frac{1}{B \cdot S} \sum_{s=1}^S \sum_{m=1}^B (L_{hc}^{s,a,m} + L_{hc}^{s,b,m}), \quad (4)$$

where B is the batch size. Clustering for assigning homogeneous pseudo-labels and optimizing E_I (to minimize the loss in Eq. 4) are performed iteratively. The final pseudo-labels are then used in Stage II.

3.2 Text Representation Learning

Stage II, which is illustrated in Figure 2, learns identity related text representations by using a (second) pre-trained CLIP model with an image encoder E_I^{init} and a text encoder E_T^{init} . Specifically, for each homogeneous pseudo-label saved during Stage I, a set of text embeddings $[X]_1[X_2]\dots[X_M]$ are learned to optimize the association between the *pseudo-label tagging sentence* ‘‘A photo of a $[X]_1[X_2]\dots[X_M]$ person’’ and the images with the pseudo-label by minimizing an overall loss:

$$L_{s2} = \frac{1}{B \cdot S} \sum_{s=1}^S \sum_{m=1}^B (L_{it}^{s,m} + L_{ti}^{s,m}) + \lambda_{mss} L_{mss}, \quad (5)$$

where $L_{it}^{s,m}$ and $L_{ti}^{s,m}$ are image-to-text and text-to-image contrastive losses, respectively, L_{mss} is a multi-scenario separation loss designed to encourage scenario-specific text representations and enable the model to effectively adapt to different scenarios, and λ_{mss} is a weight hyperparameter that controls the relative significance of L_{mss} . These individual components of the overall loss are defined as:

$$L_{it}^{s,m} = -\log \frac{\exp(f_m^{s,i} \cdot f_m^{s,tT})}{\sum_{v=1}^B \exp(f_m^{s,i} \cdot f_v^{s,tT})}, \quad (6)$$

$$L_{ti}^{s,m} = \frac{-1}{|\varphi_m^{s,t}|} \sum_{f_u^{s,i} \in \varphi_m^{s,t}} \log \frac{\exp(f_m^{s,t} \cdot f_u^{s,iT})}{\sum_{v=1}^B \exp(f_m^{s,t} \cdot f_v^{s,iT})}, \quad (7)$$

$$L_{mss} = \sum_{g=1}^S \sum_{\substack{h=1 \\ h \neq g}}^S \left[\kappa - \left\| \frac{1}{B} \sum_{m=1}^B (f_m^{g,t} - f_m^{h,t}) \right\|_2 \right]_+, \quad (8)$$

where $f_v^{s,i}$ denotes the representation of the v -th image in the batch, $f_v^{s,t}$ is the representation of the sentence for the pseudo-label assigned to the v -th image, $\varphi_m^{s,t}$ denotes the set of image representations (in the batch) that share the

same pseudo-label as the $f_m^{s,t}$, κ is a margin hyperparameter that ensures adequate separation between inter-scenario text representations, and $[\zeta]_+$ denotes the positive part of (the real-number) ζ . Note that the learnable text embeddings $[X]_1[X_2]\dots[X_M]$ are optimized by minimizing the loss in Eq. 5 while keeping the parameters of the encoders E_I^{init} and E_T^{init} frozen. Since the same operations are performed on all images and do not involve interactions between heterogeneous images, we do not differentiate between them in the equations.

We refer to the text representation output by E_T^{init} for a sentence as the cluster-level text representation. The final cluster-level text representations for the sentences are saved, and for each image within a cluster, an instance-level text representation is initialized as the corresponding cluster-level text representation. The saved text representations are referred to as *offline text representations*.

3.3 Unsupervised Heterogeneous Learning

Stage III aims to optimize the dual front-end image encoder from Stage I to ensure its identity representations are effective for UMS-ReID. To achieve this, we employ an iterative approach that alternates between updating the image encoder and leveraging the current text representations. Specifically, each iteration in Stage III begins by using the current image encoder E_I to obtain instance-level image representations for $\{X^s\}_{s=1}^S$. These representations are then clustered within each homogeneous group to generate homogeneous pseudo-labels and cluster-level image representations are computed similarly to Eq. 2. DRU is utilized to obtain instance-level *online text representations* guided by the latest pseudo-labels. CHM and IHM are the utilized to obtain cluster-level heterogeneous image pairs and instance-level heterogeneous positive sets, respectively. An overall contrastive loss—comprising homogeneous, cluster-level heterogeneous, instance-level heterogeneous, and text-guided contrastive losses—is used to update E_I . Details of DRU, CHM and IHM are provided next.

Dynamic text representation update. The offline text representations saved in Stage II may have pseudo-labels that are inconsistent with the newly obtained pseudo-labels in Stage III, potentially misleading the optimization process. To address this, we develop a dynamic text representation update (DRU) strategy to obtain *online text representations*. Mutual updating between cluster-level and instance-level text representations is accomplished through DRU. Specifically, after each clustering iteration in Stage III, we first calculate the cluster-level text pseudo-representation:

$$\hat{c}_u^t = \frac{1}{N_u^s} \sum_{m=1}^{N_u^s} f_m^{s,t}, \quad (9)$$

where N_u^s is the number of images in the u -th cluster of the s -th scenario, and $f_m^{s,t}$ is the instance-level text representation in the corresponding cluster. Subsequently, we define the top η proportion of instance-level text representations within the cluster that are closest to \hat{c}_u^t as the neighboring

representation set $\xi(\hat{c}_u^t)$, and then calculate the new cluster-level text representation:

$$c_u^t = \frac{1}{|\xi(\hat{c}_u^t)|} \sum_{f_m^{s,t} \in \xi(\hat{c}_u^t)} f_m^{s,t}. \quad (10)$$

Then, based on the new cluster-level text representation, we update each instance-level text representation:

$$f_m^{s,t} \leftarrow (1 - \alpha) f_m^{s,t} + \alpha c_u^t, \quad (11)$$

where α is the update rate hyperparameter. From the above process, it is evident that the DRU allows text representations to be updated alongside pseudo-labels.

DRU effectively prevents inconsistencies between the text and image supervision signals. Moreover, this strategy has the potential to ensure inter-cluster separability and intra-cluster compactness. On one hand, based on Eq. 9 and Eq. 10, DRU only utilizes the inner instance-level text representations to compute the new cluster-level text representation, thereby ensuring inter-cluster separability. On the other hand, DRU updates all instance-level text representations within the cluster using the new cluster-level text representation, enhancing intra-cluster compactness.

Finally, we design a text-guided contrastive loss to provide supervisory signals from text semantics for E_I . For any $f_m^{s,i}$, the text-guided contrastive loss is defined as:

$$L_{tgc} = -\frac{1}{B \cdot S} \sum_{s=1}^S \sum_{m=1}^B \log \frac{\exp(f_m^{s,i} \cdot f_m^{s,tT})}{\sum_{v=1}^B \exp(f_m^{s,i} \cdot f_v^{s,tT})}, \quad (12)$$

where $f_m^{s,i}$ and $f_m^{s,t}$ are the instance-level image and text representations of the same image, respectively.

Cluster-level heterogeneous matching. CHM aims to identify cluster-level heterogeneous image pairs of the same identity, with each pair consisting of one cluster from each of the two homogeneous groups. After each clustering, based on cluster-level image representations, we match cluster-level heterogeneous image pairs and cluster-level heterogeneous text pairs using the graph matching strategy (Wu and Ye 2023). For example, in UVI-ReID, a cluster-level heterogeneous image pair consists of one cluster composed of infrared images and another composed of visible images, both presumed to contain images of the same person. We next assess consistency between cluster-level heterogeneous image and text pairs. For consistent pairs, we directly retain the corresponding cluster-level heterogeneous image pairs; for each inconsistent pair, we stochastically retain the corresponding cluster-level heterogeneous image pair with a probability $\beta \in [0, 1]$. Each retained cluster-level heterogeneous image pair corresponds to two clusters that are (putatively) mutually heterogeneous positive clusters. Finally, we construct a cluster-level heterogeneous contrastive loss to mitigate intra-scenario heterogeneity. For any $f_{a,m}^{s,i}$, the cluster-level heterogeneous contrastive loss is defined as:

$$L_{chc}^{s,ab,m} = -\log \frac{\exp(f_{a,m}^{s,i} \cdot c_{b,u}^{s,iT} / \tau)}{\sum_{v=1}^{C_b^{s,i}} \exp(f_{a,m}^{s,i} \cdot c_{b,v}^{s,iT} / \tau)}, \quad (13)$$

where $c_{b,u}^{s,i}$ represents the cluster-level image representation of the heterogeneous positive cluster corresponding to $f_{a,m}^{s,i}$, $C_b^{s,i}$ represents the number of clusters in group b in the s -th scenario. For any $f_{b,m}^{s,i}$, the cluster-level heterogeneous contrastive loss $L_{chc}^{s,ba,m}$ is similarly defined. The overall cluster-level heterogeneous contrastive loss is then defined as:

$$L_{chc} = \frac{1}{B \cdot S} \sum_{s=1}^S \sum_{m=1}^B (L_{chc}^{s,ab,m} + L_{chc}^{s,ba,m}). \quad (14)$$

Instance-level heterogeneous matching. In addition to matching heterogeneous positive clusters, we also focus on obtaining (putative) heterogeneous positive sets at the instance-level. Specifically, we first search for heterogeneous neighbor sets for each image in both image and text representation spaces based on cosine similarity. For the instance-level image representation $f_{a,m}^{s,i}$ of any given image $x_{a,m}^s$, we define the heterogeneous neighbor set $\psi_i(x_{a,m}^s)$ as the set of images in group b corresponding to the top- k instance-level image representations that have the highest cosine similarity with $f_{a,m}^{s,i}$. Similarly, we define the heterogeneous neighbor set $\psi_t(x_{a,m}^s)$ as the set of images in group b corresponding to the top- k instance-level text representations that have the highest cosine similarity with the instance-level text representation $f_{a,m}^{s,t}$ of the image $x_{a,m}^s$. Subsequently, we take the set of the instance-level image representations corresponding to all images in the intersection $\psi_i(x_{a,m}^s) \cap \psi_t(x_{a,m}^s)$ as the instance-level heterogeneous positive set $U_{a,m}^{s,i} = \{f_{b,u}^{s,i}\}$ for $f_{a,m}^{s,i}$. Similarly, we also identify a instance-level heterogeneous positive set $U_{b,m}^{s,i} = \{f_{a,v}^{s,i}\}$ for $f_{b,m}^{s,i}$. Finally, we construct a instance-level heterogeneous contrastive loss to further bridge intra-scenario heterogeneity. For $f_{a,m}^{s,i}$, the instance-level heterogeneous contrastive loss is defined as:

$$L_{ihc}^{s,ab,m} = \frac{-1}{|U_{a,m}^{s,i}|} \sum_{f_{b,u}^{s,i} \in U_{a,m}^{s,i}} \log \frac{\exp(f_{a,m}^{s,i} \cdot f_{b,u}^{s,iT} / \tau)}{\sum_{v=1}^{N_b^s} \exp(f_{a,m}^{s,i} \cdot f_{b,v}^{s,iT} / \tau)}, \quad (15)$$

where N_b^s represents the number of images in group b within the s -th scenario. For $f_{b,m}^{s,i}$, the loss $L_{ihc}^{s,ba,m}$ is similarly defined. The overall instance-level heterogeneous contrastive loss is then obtained as:

$$L_{ihc} = \frac{1}{B \cdot S} \sum_{s=1}^S \sum_{m=1}^B (L_{ihc}^{s,ab,m} + L_{ihc}^{s,ba,m}). \quad (16)$$

For any given image, we refer to the combination of the image with any image from the heterogeneous positive cluster obtained by CHM, as well as the combination with any image from the instance-level heterogeneous positive set obtained by IHM, as a *heterogeneous positive pair*.

In summary, the overall objective function of Stage III is defined as:

$$L_{s3} = L_{hc} + L_{chc} + L_{ihc} + \lambda_{tgc} L_{tgc}, \quad (17)$$

where λ_{tgc} is a weight hyperparameter for L_{tgc} . The final trained image encoder E_I for ITKM is obtained by optimizing L_{s3} .

4 Experiments

4.1 Datasets and Evaluation Metrics

We evaluate the proposed method on the commonly used SYSU-MM01 (Wu et al. 2017), LTCC (Qian et al. 2020), and MLR-CUHK03 (Pang, Zhao, and Wang 2024b), using mean average precision (mAP) and cumulative matching characteristic (CMC) as evaluation metrics.

4.2 Comparison with State-of-the-Art

In Table 1, we evaluate ITKM trained in a single scenario (denoted as ITKM(S)) and ITKM jointly trained across three scenarios (denoted as ITKM(M)) on SYSU-MM01, LTCC, and MLR-CUHK03, respectively.

First, we find that ITKM(S) outperforms existing unsupervised traditional (UT) methods, including ICE (Chen, Lagadee, and Bremond 2021), CC (Dai et al. 2022), Purification (Lan et al. 2023), and DCCC (He et al. 2024). This is because UT-ReID does not account for the heterogeneity within scenarios, making it difficult to obtain sufficient heterogeneous positive pairs, whereas ITKM utilizes CHM and IHM to acquire ample heterogeneous positive pairs.

Next, we compare ITKM with unsupervised scenario-specific (USS) ReID methods for each individual scenario, including the UVI-ReID methods SDCL (Yang, Chen, and Ye 2024) and TokenMatcher (Wang et al. 2025), the UCC-ReID methods CICL (Pang, Zhao, and Wang 2024a) and JAPL (Pang et al. 2025), and the UCR-ReID method DRFM (Pang, Zhao, and Wang 2024b). We find that ITKM(S) is competitive with state-of-the-art USS-ReID methods on SYSU-MM01 and LTCC, and significantly outperforms the USS-ReID method DRFM on MLR-CUHK03. These results not only validate the superiority of ITKM but also confirm its generalizability across multiple scenarios.

Finally, we train and test two existing methods, SDCL (Yang, Chen, and Ye 2024) and TokenMatcher (Wang et al. 2025), on the unsupervised multi-scenario (UMS) setting that combines the three datasets. These methods are denoted as TokenMatcher(M) and SDCL(M), respectively. As shown in Table 1, joint training across multiple scenarios does not improve the performance of TokenMatcher(M) and SDCL(M); instead, they perform worse compared to single-scenario training. This indicates that these methods struggle to adapt to multiple scenarios. In contrast, ITKM(M), which is trained across multiple scenarios, shows slight improvements in performance over ITKM(S) across all three scenarios and significantly outperforms TokenMatcher(M) and SDCL(M). This is because the scenario embedding and multi-scenario separation loss in ITKM enable the image encoder to adapt to different scenarios, thereby improving the model’s generalization ability.

4.3 Ablation Study

In Table 2, we assess the effectiveness of the five components within ITKM, specifically: scenario embedding e^s , multi-scenario separation loss L_{mss} , DRU, CHM, and IHM. M1 uses PGM (Wu and Ye 2023) to obtain heterogeneous positive pairs and employs ACCL (Wu and Ye 2023) along




Figure 4: Two-dimensional t-SNE (Van der Maaten and Hinton 2008) visualizations of the text representations for ablated methods M2 and M3. Different colors representing text representations from different scenarios.

with offline text representations to jointly optimize the image encoder. Other ablation methods sequentially embed the above components or alternatives into M1.

Effectiveness of e^s , L_{mss} and DRU. As shown in Table 2, M2 achieves a significant performance improvement over M1 after introducing scenario embedding e^s . Specifically, the Rank-1 accuracy increases by 2.3%, 2.4%, and 1.2% on SYSU-MM01, LTCC, and MLR-CUHK03, respectively. This validates the effectiveness of scenario embedding in the UMS-ReID task. Additionally, M3 further improves recognition performance by introducing L_{mss} on top of M2. To further evaluate the effectiveness of L_{mss} , we use t-SNE (Van der Maaten and Hinton 2008) visualizations of the text representations corresponding to the text embeddings learned by M2 and M3 during Stage II. To avoid the influence of data heterogeneity, we only visualize the text representations corresponding to high-resolution visible images, excluding those corresponding to infrared or low-resolution images. As shown in Figure 4, compared to M2, the text representations of M3 exhibit greater inter-scenario separability. This confirms that L_{mss} enhances model performance by encouraging scenario-specific text representations. As shown in Table 2, M4 achieves Rank-1 accuracy improvements of 2.3%, 1.7%, and 0.7% over M3 on SYSU-MM01, LTCC, and MLR-CUHK03, respectively. These results validate the effectiveness of DRU.

Effectiveness of CHM. As shown in Table 2, M5 achieves significant performance improvements over M4. Additionally, to assess the effectiveness of CHM in obtaining reliable heterogeneous positive pairs, we calculate the F-score (Goutte and Gaussier 2005) of the training pseudo-labels for M4 and M5. A higher F-score indicates greater accuracy of the pseudo-labels. As illustrated in Figure 5, M5 consistently achieves a higher F-score than M4 throughout the training process. This confirms that CHM enhances model performance by improving the accuracy of heterogeneous positive pairs.

Effectiveness of IHM. In Table 2, M6 and M7 introduces CNL (Yang, Chen, and Ye 2024) and IHM, respectively, to obtain instance-level heterogeneous positive sets based on M5, and both achieve performance improvements. This confirms the feasibility of obtaining heterogeneous positive

Methods	SYSU-MM01				LTCC				MLR-CUHK03		
	All Search		Indoor Search		Clothing Change		General		Rank-1	Rank-5	Rank-10
	Rank-1	mAP	Rank-1	mAP	Rank-1	mAP	Rank-1	mAP			
ICE	20.5	20.4	29.8	38.4	14.5	7.1	61.1	28.4	27.6	65.3	80.6
CC	20.2	22.0	23.3	34.0	7.4	6.0	17.0	11.0	31.6	66.8	80.8
Purification	20.8	21.3	30.6	39.0	15.6	8.7	52.8	23.9	30.5	65.8	78.9
DCCC	18.8	18.6	21.3	33.1	13.7	7.2	47.6	21.3	27.6	60.3	73.6
SDCL	64.5	63.2	71.4	76.9	-	-	-	-	-	-	-
TokenMatcher	65.1	62.8	69.0	74.9	-	-	-	-	-	-	-
CICL	-	-	-	-	25.7	13.0	60.3	29.6	-	-	-
JAPL	-	-	-	-	26.1	12.9	63.0	31.8	-	-	-
DRFM	-	-	-	-	-	-	-	-	35.8	72.2	83.6
ITKM(S)	64.6	63.2	72.0	76.7	26.2	13.2	62.8	31.6	62.5	85.1	90.3
SDCL(M)	63.0	61.8	69.2	75.0	18.8	9.3	55.2	25.7	35.3	71.6	82.1
TokenMatcher(M)	62.5	60.9	68.7	73.6	16.9	9.0	55.3	25.5	37.7	72.5	83.0
ITKM(M)	64.9	63.3	72.3	77.1	27.3	14.4	63.3	32.5	63.6	85.7	91.6

Table 1: Comparison of ITKM with existing methods on SYSU-MM01, LTCC and MLR-CUHK03.

Methods	e^s	L_{mss}	DRU	Cluster	Instance	SYSU-MM01(All Search)		LTCC(Clothing Change)		MLR-CUHK03	
						Rank-1	mAP	Rank-1	mAP	Rank-1	Rank-5
M1				PGM		53.8	49.0	15.2	9.7	33.2	68.6
M2	✓			PGM		56.1	52.2	17.6	10.7	34.4	71.2
M3	✓	✓		PGM		57.0	52.7	18.5	10.9	36.3	72.5
M4	✓	✓	✓	PGM		59.3	54.9	20.2	11.3	37.0	72.9
M5	✓	✓	✓	CHM		62.6	59.1	23.8	12.3	50.7	79.8
M6	✓	✓	✓	CHM	CNL	63.2	61.5	25.6	13.5	56.5	82.1
M7/ITKM	✓	✓	✓	CHM	IHM	64.9	63.3	27.3	14.4	63.6	85.7

Table 2: Results of the ablation study. ‘‘Cluster’’ and ‘‘Instance’’ refer to the methods used for obtaining heterogeneous positive clusters and instance-level heterogeneous positive sets, respectively.




Figure 5: F-score for M4 and M5.



Figure 6: Visualization of instance-level heterogeneous positive sets obtained by M6 and M7. True and false positive samples are marked with green and red boxes, respectively.

pairs at the instance-level. Furthermore, we find that M7 outperforms M6. To explore this further, we visualize some of the instance-level heterogeneous positive sets obtained by

M6 and M7 in Figure 6. Compared to M6, M7 effectively excludes some false positive samples that are visually similar. This may be due to the different cognitive focuses of the text encoder and image encoder, allowing IHM to avoid certain misjudgments of the image encoder by relying on consistency. In contrast, although CNL also relies on consistency, specifically between shallow and deep representations of the image encoder, it is inherently dependent on the knowledge within the image encoder, making it difficult to overcome the image encoder’s inherent limitations. This validates the effectiveness of IHM.

5 Conclusion

We introduce the UMS-ReID as a new task and propose the ITKM framework to tackle the challenges associated with UMS-ReID. Ablation experiments validate the effectiveness of various components in ITKM, namely the scenario embeddings and multi-scenario separation loss effectively encourage the model to adapt to multiple scenarios, while CHM and IHM obtain reliable heterogeneous positive pairs. Additionally, DRU effectively enhances model performance by maintaining consistency between text and image supervision signals. Experimental results validate the superiority and generalizability of ITKM, demonstrating that it not only competes effectively with advanced scenario-specific methods within each scenario, but also leverages knowledge from multiple scenarios to further enhance performance.

Acknowledgments

This work is supported by National Key Research and Development Program of China (Grant no. 2023YFC3305003) and the National Natural Science Foundation of China (NSFC, Grant no. 62231013).

References

- Chen, H.; Lagadec, B.; and Bremond, F. 2021. ICE: Inter-instance contrastive encoding for unsupervised person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14960–14969.
- Chen, Z.; Zhang, Z.; Tan, X.; Qu, Y.; and Xie, Y. 2023. Unveiling the power of clip in unsupervised visible-infrared person re-identification. In *Proceedings of the ACM International Conference on Multimedia*, 3667–3675.
- Cho, Y.; Kim, W. J.; Hong, S.; and Yoon, S.-E. 2022. Part-based pseudo label refinement for unsupervised person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7308–7318.
- Dai, Z.; Wang, G.; Yuan, W.; Zhu, S.; and Tan, P. 2022. Cluster contrast for unsupervised person re-identification. In *Proceedings of the Asian Conference on Computer Vision*, 1142–1160.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houshby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *Proceedings of the International Conference on Learning Representations*.
- Gong, Y.; Zhong, Z.; Qu, Y.; Luo, Z.; Ji, R.; and Jiang, M. 2024. Cross-modality perturbation synergy attack for person re-identification. *Advances in Neural Information Processing Systems*, 37: 23352–23377.
- Goutte, C.; and Gaussier, E. 2005. A probabilistic interpretation of precision, recall and F-score, with implication for evaluation. In *Proceedings of the European Conference on Information Retrieval*, 345–359.
- He, Z.; Xue, M.; Du, Y.; Zhao, Z.; and Su, F. 2024. Dynamic Clustering and Cluster Contrastive Learning for Unsupervised Person Re-Id With Feature Distribution Alignment. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, 3610–3614.
- Lan, L.; Teng, X.; Zhang, J.; Zhang, X.; and Tao, D. 2023. Learning to purification for unsupervised person re-identification. *IEEE Transactions on Image Processing*, 32: 3338–3353.
- Li, M.; Xu, P.; Li, C.-G.; and Guo, J. 2023. SiCL: Silhouette-Driven Contrastive Learning for Unsupervised Person Re-Identification with Clothes Change. arXiv:2305.13600.
- Li, S.; Sun, L.; and Li, Q. 2023. CLIP-ReID: exploiting vision-language model for image re-identification without concrete text labels. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1405–1413.
- Li, Z.; Liu, H.; Peng, X.; and Jiang, W. 2024. Inter-Intra Modality Knowledge Learning and Clustering Noise Alleviation for Unsupervised Visible-Infrared Person Re-Identification. *IEEE Transactions on Knowledge and Data Engineering*.
- Lin, G.; Xu, Y.; Lai, H.; and Yin, J. 2024. Revisiting few-shot learning from a causal perspective. *IEEE Transactions on Knowledge and Data Engineering*.
- Pang, Z.; Wang, C.; Zhao, L.; Liu, Y.; and Sharma, G. 2023. Cross-modality hierarchical clustering and refinement for unsupervised visible-infrared person re-identification. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Pang, Z.; Zhao, L.; Liu, Y.; Sharma, G.; and Wang, C. 2025. Joint Augmentation and Part Learning for Unsupervised Clothing Change Person Re-Identification. *IEEE Transactions on Information Forensics and Security*.
- Pang, Z.; Zhao, L.; and Wang, C. 2024a. Clothing-invariant contrastive learning for unsupervised person re-identification. *Neural Networks*, 106477.
- Pang, Z.; Zhao, L.; and Wang, C. 2024b. Dual-Resolution Fusion Modeling for Unsupervised Cross-Resolution Person Re-Identification. In *Proceedings of the ACM International Conference on Multimedia*.
- Qian, X.; Wang, W.; Zhang, L.; Zhu, F.; Fu, Y.; Xiang, T.; Jiang, Y.-G.; and Xue, X. 2020. Long-term cloth-changing person re-identification. In *Proceedings of the Asian Conference on Computer Vision*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning*, 8748–8763.
- Shi, J.; Yin, X.; Chen, Y.; Zhang, Y.; Zhang, Z.; Xie, Y.; and Qu, Y. 2024. Multi-Memory Matching for Unsupervised Visible-Infrared Person Re-Identification. arXiv:2401.06825.
- Shi, J.; Zhang, Y.; Yin, X.; Xie, Y.; Zhang, Z.; Fan, J.; Shi, Z.; and Qu, Y. 2023. Dual pseudo-labels interactive self-training for semi-supervised visible-infrared person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11218–11228.
- Tan, L.; Zhang, Y.; Han, K.; Dai, P.; Zhang, Y.; Wu, Y.; and Ji, R. 2024. RLE: A unified perspective of data augmentation for cross-spectral re-identification. *Advances in Neural Information Processing Systems*, 37: 126977–126996.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(11).
- Wang, M.; Lai, B.; Huang, J.; Gong, X.; and Hua, X.-S. 2021. Camera-aware proxies for unsupervised person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 2764–2772.
- Wang, X.; Liu, L.; Yang, B.; Ye, M.; Wang, Z.; and Xu, X. 2025. TokenMatcher: Diverse Tokens Matching for Unsupervised Visible-Infrared Person Re-Identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 7934–7942.
- Wu, A.; Zheng, W.-S.; Yu, H.-X.; Gong, S.; and Lai, J. 2017. RGB-infrared cross-modality person re-identification. In

Proceedings of the International Conference on Computer Vision, 5380–5389.

Wu, Z.; and Ye, M. 2023. Unsupervised visible-infrared person re-identification via progressive graph matching and alternate learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9548–9558.

Xuan, S.; and Zhang, S. 2021. Intra-inter camera similarity for unsupervised person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11926–11935.

Yang, B.; Chen, J.; and Ye, M. 2023. Towards grand unified representation learning for unsupervised visible-infrared person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11069–11079.

Yang, B.; Chen, J.; and Ye, M. 2024. Shallow-Deep Collaborative Learning for Unsupervised Visible-Infrared Person Re-Identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16870–16879.

Yang, B.; Ye, M.; Chen, J.; and Wu, Z. 2022. Augmented dual-contrastive aggregation learning for unsupervised visible-infrared person re-identification. In *Proceedings of the ACM International Conference on Multimedia*, 2843–2851.

Yang, Z.; Wu, D.; Wu, C.; Lin, Z.; Gu, J.; and Wang, W. 2024. A Pedestrian is Worth One Prompt: Towards Language Guidance Person Re-Identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17343–17353.

Ye, M.; Ruan, W.; Du, B.; and Shou, M. Z. 2021a. Channel augmented joint learning for visible-infrared recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13567–13576.

Ye, M.; Shen, J.; Lin, G.; Xiang, T.; Shao, L.; and Hoi, S. C. 2021b. Deep learning for person re-identification: A survey and outlook. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6): 2872–2893.

Zheng, X.; Zhang, Y.; Lu, Y.; and Wang, H. 2024. Semi-supervised Visible-Infrared Person Re-identification via Modality Unification and Confidence Guidance. In *Proceedings of the ACM International Conference on Multimedia*, 5761–5770.

Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9): 2337–2348.