

# ReMix: Training Generalized Person Re-identification on a Mixture of Data

Timur Mamedov<sup>1,2</sup> Anton Konushin<sup>3,2</sup> Vadim Konushin<sup>1</sup>

<sup>1</sup>Tevian, Moscow, Russia <sup>2</sup>Lomonosov Moscow State University <sup>3</sup>AIRI, Moscow, Russia

me@timmzak.com konushin@airi.net vadim@tevia.ai

## Abstract

Modern person re-identification (Re-ID) methods have a weak generalization ability and experience a major accuracy drop when capturing environments change. This is because existing multi-camera Re-ID datasets are limited in size and diversity, since such data is difficult to obtain. At the same time, enormous volumes of unlabeled single-camera records are available. Such data can be easily collected, and therefore, it is more diverse. Currently, single-camera data is used only for self-supervised pre-training of Re-ID methods. However, the diversity of single-camera data is suppressed by fine-tuning on limited multi-camera data after pre-training. In this paper, we propose ReMix, a generalized Re-ID method jointly trained on a mixture of limited labeled multi-camera and large unlabeled single-camera data. Effective training of our method is achieved through a novel data sampling strategy and new loss functions that are adapted for joint use with both types of data. Experiments show that ReMix has a high generalization ability and outperforms state-of-the-art methods in generalizable person Re-ID. To the best of our knowledge, this is the first work that explores joint training on a mixture of multi-camera and single-camera data in person Re-ID.

## 1. Introduction

Person re-identification (Re-ID) is the task of recognizing the same person in images taken by different cameras at different times. This task naturally arises in video surveillance and security systems, where it is necessary to track people across multiple cameras. The urgent need for robust and accurate Re-ID has stimulated scientific research over the years. However, modern Re-ID methods still have a weak generalization ability and experience a significant performance drop when capturing environments change, which limits their applicability in real-world scenarios.

The main reasons for the weak generalization ability of modern methods are the small amount of training data and the low diversity of capturing environments in this data. In person Re-ID, the same person may appear across multi-



Figure 1. Examples of multi-camera and single-camera data. As we can see, multi-camera data is much more complex in terms of Re-ID: background, lighting, capturing angle, etc., may differ significantly for one person in multi-camera data. In contrast, images of the same person are less complex in single-camera data.

Dataset	#images	#IDs	#scenes
CUHK03-NP [22]	14,096	1,467	2
Market-1501 [56]	32,668	1,501	6
DukeMTMC-reID [36]	36,411	1,812	8
MSMT17 [47]	126,441	4,101	15
LUPerson [10]	>4M	>200K	46,260

Table 1. Comparison between existing well-known multi-camera Re-ID datasets and the single-camera LUPerson dataset. As we can see, single-camera data is more voluminous and diverse.

ple cameras from different angles (multi-camera data), and such data is difficult to collect and label. Due to these difficulties, each of the existing Re-ID datasets is captured from a single location. In contrast, collecting images of people from one camera (single-camera data) is much easier; for example, these images can be automatically extracted from

YouTube videos [10], featuring numerous diverse identities in distinct locations and a high diversity of capturing environments (Tab. 1).

However, single-camera data is much simpler than multi-camera data in terms of the person Re-ID task: in single-camera data, the same person can appear on only one camera and from only one angle (Fig. 1). Directly adding such simple data to the training process degrades the quality of Re-ID. Therefore, single-camera data is currently used only for self-supervised pre-training [10, 27]. However, we hypothesize that this approach has a limited effect on improving the generalization ability of Re-ID methods because subsequent fine-tuning for the final task is performed on relatively small and non-diverse multi-camera data.

In this paper, we propose ReMix, a generalized Re-ID method jointly trained on a mixture of limited labeled multi-camera and large unlabeled single-camera data. ReMix achieves better generalization by training on diverse single-camera data, as confirmed by our experiments. We also experimentally validate our hypothesis regarding the limitations of self-supervised pre-training and show that our joint training on two types of data overcomes them. In our ReMix method, we propose:

- A novel data sampling strategy that allows for efficiently obtaining pseudo labels for large unlabeled single-camera data and for composing mini-batches from a mixture of images from labeled multi-camera and unlabeled single-camera datasets.
- A new Instance, Augmentation, and Centroids loss functions adapted for joint use with two types of data, making it possible to train ReMix. For example, the Instance and Centroids losses consider the different complexities of multi-camera and single-camera data, allowing for more efficient training of our method.
- Using self-supervised pre-training in combination with the proposed joint training procedure to improve pseudo labeling and the generalization ability of the algorithm.

Our experiments show that ReMix outperforms state-of-the-art methods in the cross-dataset and multi-source cross-dataset scenarios (when trained and tested on different datasets). To the best of our knowledge, this is the first work that explores joint training on a mixture of multi-camera and single-camera data in the person Re-ID task.

## 2. Related Work

### 2.1. Person Re-identification

Rapid progress in solving the person re-identification task over the past few years has been associated with the emergence of CNNs. Some Re-ID approaches used the entire image to extract features [31, 37, 57]. Other methods

divided the image of a person into parts, extracted features for each part, and aggregated them to obtain full-image features [38, 39, 42]. Recently, transformer-based Re-ID methods have emerged [14, 21, 40, 52], which also improve the quality of solving the problem.

Recent Re-ID methods perform well in the standard scenario, but their quality is significantly reduced when applied to datasets that differ from those used during training (when capturing environments change). In this paper, we explore the problem of weak generalization ability of existing Re-ID methods and show that it can be improved by properly using a mixture of two types of training data — multi-camera and single-camera.

### 2.2. Generalizable Person Re-identification

Generalizable person re-identification aims to learn a robust model that performs well across various datasets. To achieve this goal, improved normalizations adapted to generalizable person Re-ID were proposed in [6, 17, 18]. A new residual block, consisting of multiple convolutional streams, each detecting features at a specific scale, was proposed in [58] to create a specialized neural network architecture adapted to the person Re-ID task. In [59], the ideas from [58] were continued, and an updated architecture with normalization layers was proposed to improve the generalization ability of the algorithm. Transformer-based models were also used to solve the problem under consideration: in [29] it was shown that local parts of images are less susceptible to domain gap, making it more effective to compare two images by their local parts in addition to global visual information during training. In [26], a new effective method for composing mini-batches during training was suggested, which improved the generalization ability of the algorithm.

As we can see, in most existing approaches, improving the generalization ability of the Re-ID algorithm has been achieved through the use of complex architectures. In contrast, in this paper, we prove that generalization can be achieved by properly training an efficient model using a variety of data, which is important in practice.

### 2.3. Self-supervised Pre-training

Self-supervised pre-training is an approach for training neural networks using unlabeled data to learn high-quality primary features. Such pre-training is usually performed by defining relatively simple tasks that allow training data to be generated on the fly, for example: context prediction [12], solving a puzzle [32], predicting an image rotation angle [11]. In [3, 5, 8, 51], self-supervised approaches based on contrastive learning were proposed: there, the neural network was trained to bring images of the same class closer in space and push away negative instances. Self-supervised pre-training was also used in person re-identification [4, 10, 27].

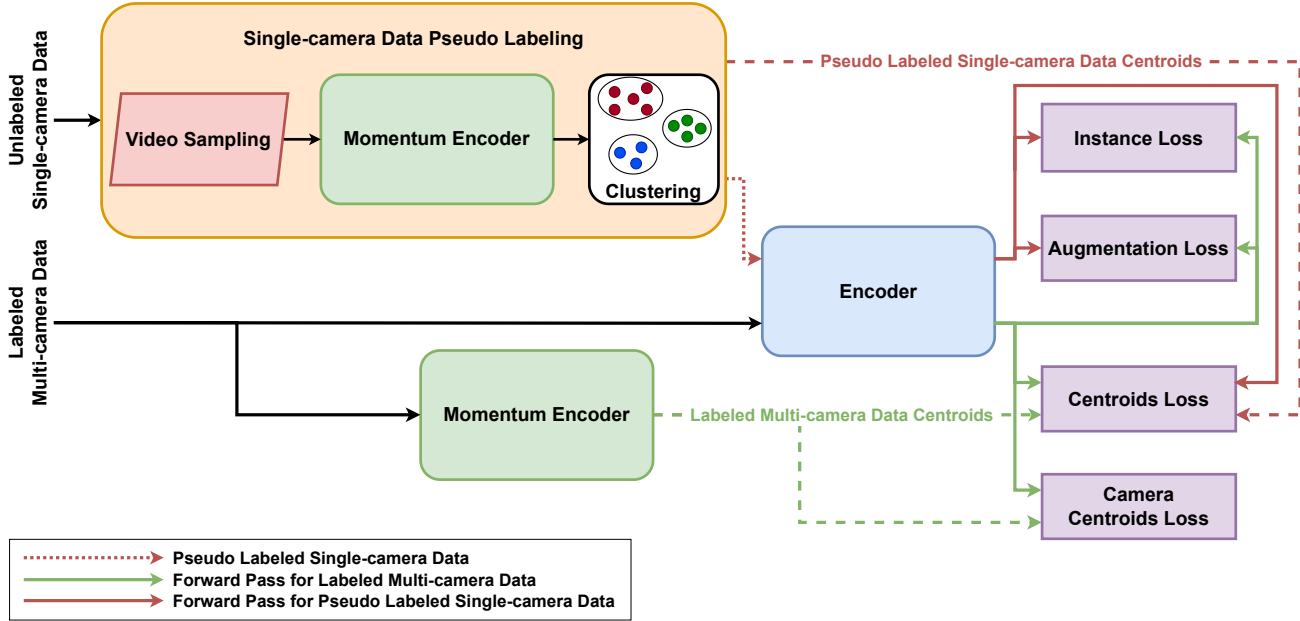


Figure 2. Scheme of ReMix. At the beginning of each epoch, all images from the person Re-ID dataset (multi-camera data) pass through the momentum encoder to obtain centroids for each identity (bottom part of the scheme). Simultaneously, videos are randomly sampled from the unlabeled single-camera dataset, and images from the selected videos are clustered using embeddings from the momentum encoder and pseudo labeled (top part of the scheme). After that, labeled multi-camera and pseudo labeled single-camera data are fed to the encoder as input. To train the encoder, the following new loss functions are used: the Instance Loss, the Augmentation Loss, and the Centroids Loss are calculated for both types of data, whereas the Camera Centroids Loss is calculated only for multi-camera data.

However, we suppose that this approach has a limited impact on improving the generalization ability of Re-ID methods, since subsequent fine-tuning for the final task is conducted on relatively small and non-diverse multi-camera data. In this paper, we show that the proposed joint training procedure in our ReMix method is more effective than pure self-supervised pre-training.

### 3. Proposed Method

#### 3.1. Overview

The scheme of ReMix is presented in Fig. 2. The proposed method consists of two neural networks with identical architectures — the encoder and the momentum encoder. The main idea of ReMix is to jointly train the Re-ID algorithm on a mixture of labeled multi-camera data for this task, and diverse unlabeled single-camera images of people. Therefore, during training, mini-batches consisting of these two types of data are used. The novel data sampling strategy is described in Sec. 3.2.

The encoder is trained using new loss functions that are adapted for joint use with two types of data: the Instance Loss  $\mathcal{L}_{ins}$  (Sec. 3.3.1), the Augmentation Loss  $\mathcal{L}_{aug}$  (Sec. 3.3.2), and the Centroids Loss  $\mathcal{L}_{cen}$  (Sec. 3.3.3) are calculated for both types of data, whereas the Camera Centroids Loss  $\mathcal{L}_{cc}$  (Sec. 3.3.4) is calculated only for multi-

camera data. The general loss function in ReMix has the following form:

$$\mathcal{L} = \mathcal{L}_{ins} + \mathcal{L}_{aug} + \mathcal{L}_{cen} + \gamma \mathcal{L}_{cc}. \quad (1)$$

The encoder is updated by backpropagation, and for the momentum encoder, the weights are updated using exponential moving averaging:

$$\theta_m^t = \lambda \theta_m^{t-1} + (1 - \lambda) \theta_e^t, \quad (2)$$

where  $\theta_e^t$  and  $\theta_m^t$  are the weights of the encoder and the momentum encoder at iteration  $t$ , respectively; and  $\lambda$  is the momentum coefficient.

The use of the encoder and the momentum encoder allows for more robust and noise-resistant training, which is important when using unlabeled single-camera data. During inference, only the momentum encoder is used to obtain embeddings. To train ReMix, loss functions involving centroids are applied. Therefore, to achieve training stability and frequent updating of centroids, only a portion of the images passes through the encoder in one epoch. Additionally, this approach reduces computational costs by generating pseudo labels only for a subset of single-camera data in one epoch, rather than for an entire large dataset (Sec. 3.2). ReMix is described in more detail in the supplementary material (see Algorithm 1).

## 3.2. Data Sampling

Let us formally describe the training datasets. Labeled multi-camera data (Re-ID datasets) consist of image-label-camera triples  $\mathcal{D}_m = \{(x_i, y_i, c_i)\}_{i=1}^{N_m^m}$ , where  $x_i \in \mathcal{X}$  is the image,  $y_i \in \mathcal{Y}_m = \{1, 2, \dots, M_m\}$  is the image’s identity label, and  $c_i \in \mathcal{C}_m = \{1, 2, \dots, K_m\}$  is the camera ID. As for unlabeled single-camera data  $\mathcal{D}_s$ , it is a set of videos  $\{\mathcal{V}_i\}_{i=1}^{N_s}$ , where each video  $\mathcal{V}_i$  is a set of unlabeled images  $\{\hat{x}_j^i\}_{j=1}^{N_s^i}$  of people. In single-camera data, each person appears on only one video.

**Single-camera data pseudo labeling.** Since the proposed method uses unlabeled single-camera data, pseudo labels are obtained at the beginning of each epoch. This is done according to the following algorithm: a video  $\mathcal{V}_i$  is randomly sampled from the set  $\mathcal{D}_s$ , and images from the selected video are clustered by DBSCAN [9] using embeddings from the momentum encoder and pseudo labeled. This procedure continues until pseudo labels are assigned to all images necessary for training in one epoch. As mentioned in Sec. 3.1, not all images are used for training in one epoch, so we know in advance how many images from unlabeled single-camera data should receive pseudo labels. Thus, our method iteratively obtains pseudo labels for almost all images from the large single-camera dataset. Additionally, it is worth noting that the pseudo labeling procedure uses embeddings from the momentum encoder with weights updated in the previous epoch, which leads to iterative improvements in the quality of pseudo labels. The proposed single-camera data pseudo labeling procedure is described in more detail in the supplementary material (see Algorithm 2).

**Mini-batch composition.** In our ReMix method, we compose a mini-batch from a mixture of images from multi-camera and single-camera datasets as follows:

- For multi-camera data,  $N_P^m$  labels are randomly sampled, and for each label,  $N_K^m$  corresponding images obtained from different cameras are selected.
- For single-camera data,  $N_P^s$  pseudo labels are randomly sampled, and for each pseudo label,  $N_K^s$  corresponding images are selected.

Thus, the mini-batch has a size of  $N_P^m \times N_K^m + N_P^s \times N_K^s$  images.

## 3.3. Loss Functions

### 3.3.1 The Instance Loss

The main idea of the proposed Instance Loss is to bring the anchor closer to all positive instances and push it away from all negative instances in a mini-batch. Thus, the Instance Loss forces the neural network to learn a more general solution.

Let us define  $\hat{\mathcal{Y}}_{m+s} = \hat{\mathcal{Y}}_m \cup \hat{\mathcal{Y}}_s$  as the set of all labels for multi-camera data and pseudo labels for single-camera data in a mini-batch.  $\hat{y}_i \in \hat{\mathcal{Y}}_{m+s}$  is either a label or pseudo label corresponding to the  $i$ -th image in a mini-batch.  $B_m = N_P^m \times N_K^m$  is the number of images from multi-camera data in a mini-batch. And  $B_s = N_P^s \times N_K^s$  is the number of images from single-camera data in a mini-batch. Then the Instance Loss is defined as follows:

$$\mathcal{L}_{ins} = \frac{1}{B_m + B_s} \left( \underbrace{\sum_{i=1}^{B_m} \mathcal{L}_{ins_m}^i}_{\text{multi-camera}} + \underbrace{\sum_{i=B_m+1}^{B_m+B_s} \mathcal{L}_{ins_s}^i}_{\text{single-camera}} \right), \quad (3)$$

$$\mathcal{L}_{ins_m}^i = \frac{-1}{N_K^m} \sum_{j:\hat{y}_i=\hat{y}_j} \log \frac{\exp(\langle f_i \cdot m_j \rangle / \tau_{ins_m})}{\sum_{k=1}^{N_m+1} \exp(\langle f_i \cdot m_k \rangle / \tau_{ins_m})}, \quad (4)$$

$$\mathcal{L}_{ins_s}^i = \frac{-1}{N_K^s} \sum_{j:\hat{y}_i=\hat{y}_j} \log \frac{\exp(\langle f_i \cdot m_j \rangle / \tau_{ins_s})}{\sum_{k=1}^{N_s+1} \exp(\langle f_i \cdot m_k \rangle / \tau_{ins_s})}, \quad (5)$$

where  $f_i, m_i$  are embeddings from the encoder and the momentum encoder for the anchor  $i$ -th image in a mini-batch, respectively;  $N_m$  and  $N_s$  are the numbers of negative instances for the anchor (for multi-camera and single-camera data, respectively); and  $\langle \cdot \rangle$  denotes cosine similarity. Since multi-camera and single-camera data have different complexities in terms of person Re-ID, we balance them by using temperature parameters in the Instance Loss:  $\tau_{ins_m}$  for multi-camera data and  $\tau_{ins_s}$  for single-camera data.

### 3.3.2 The Augmentation Loss

The distribution of inter-instance similarities produced by the algorithm can change under the influence of augmentations. After augmentations, an anchor image from the perspective of the neural network may become less similar to its positive pair, but at the same time, similarity to negative instances increases. Thus, current methods may be unstable to image changes and noise that may occur in practice.

To address this problem, we propose the new Augmentation Loss, which brings the augmented version of the image closer to its original and pushes it away from instances belonging to other identities in a mini-batch:

$$\mathcal{L}_{aug} = \mathbb{E} \left[ -\log \frac{\exp(\langle f_{aug}^i \cdot m_A^i \rangle / \tau_{aug})}{\sum_{j=1}^{N+1} \exp(\langle f_{aug}^i \cdot m_j \rangle / \tau_{aug})} \right], \quad (6)$$

where  $f_{aug}^i$  is the embedding from the encoder for the augmented  $i$ -th image in a mini-batch;  $m_A^i$  is the embedding from the momentum encoder for the original  $i$ -th image in a mini-batch; and  $N$  is the number of negative instances. It is important to note that in the Augmentation Loss, embeddings for the original images are obtained from the momentum encoder, as the momentum encoder is more stable.

### 3.3.3 The Centroids Loss

Let us define the concept of a centroid for a label or pseudo label  $\hat{y}_i \in \hat{\mathcal{Y}}_{m+s}$  as follows:

$$p_{\hat{y}_i} = \frac{1}{|M_{\hat{y}_i}|} \sum_{m \in M_{\hat{y}_i}} m, \quad (7)$$

where  $M_{\hat{y}_i}$  is the set of embeddings from the momentum encoder corresponding to the label or pseudo label  $\hat{y}_i$ , and  $m$  is an embedding from this set.

Then the new Centroids Loss can be defined as:

$$\mathcal{L}_{cen} = \frac{1}{B_m + B_s} \left( \underbrace{\sum_{i=1}^{B_m} \mathcal{L}_{cen}^i(\tau_{cen_m})}_{\text{multi-camera}} + \underbrace{\sum_{i=B_m+1}^{B_m+B_s} \mathcal{L}_{cen}^i(\tau_{cen_s})}_{\text{single-camera}} \right), \quad (8)$$

$$\mathcal{L}_{cen}^i(\tau) = -\log \frac{\exp(f_{\hat{y}_i} \cdot p_{\hat{y}_i} / \tau)}{\sum_{j=1}^{|\hat{\mathcal{Y}}_{m+s}|} \exp(f_{\hat{y}_j} \cdot p_{\hat{y}_j} / \tau)}, \quad (9)$$

where  $f_{\hat{y}_i}$  is the embedding from the encoder for the image with the label or pseudo label  $\hat{y}_i$ . Thus, this loss function brings instances closer to their corresponding centroids and pushes them away from other centroids. Like the Instance Loss, this loss function uses different temperature parameters for multi-camera and single-camera data.

### 3.3.4 The Camera Centroids Loss

Since the same person could be captured by different cameras in multi-camera data, it is useful to apply information about cameras for better feature generation. In our ReMix method, we use the Camera Centroids Loss [44]. This loss function brings instances closer to the centroids of instances with the same label, but captured by different cameras. Thus, the intra-class variance caused by stylistic differences between cameras is reduced.

## 4. Experiments

### 4.1. Datasets and Evaluation Metrics

**Multi-camera datasets.** We employ well-known datasets CUHK03-NP [22], Market-1501 [56], DukeMTMC-reID [36], and MSMT17 [47] as multi-camera data for evaluating our proposed method. The CUHK03-NP dataset consists of 14,096 images of 1,467 identities captured by two cameras. Market-1501 was gathered from six cameras and consists of 12,936 images of 751 identities for training and 19,732 images of 750 identities for testing. DukeMTMC-reID contains 16,522 training images of 702 identities and

19,889 images of 702 identities for testing, all of them collected from eight cameras. MSMT17, a large-scale Re-ID dataset, consists of 32,621 training images of 1,041 identities and 93,820 testing images of 3,060 identities captured by fifteen cameras. Additionally, we use MSMT17-merged, which combines training and test parts. We also employ a subset of the synthetic RandPerson [46] dataset, which contains 132,145 training images of 8,000 identities, for additional experiments. It is worth noting that DukeMTMC-reID was withdrawn by its creators due to ethical concerns, but this dataset is still used to evaluate other modern Re-ID methods. Therefore, we include it in our tests for fair and objective comparison.

**Single-camera dataset.** We use the LUPerson dataset [10] as unlabeled single-camera data. This dataset consists of over 4 million images of more than 200,000 people from 46,260 distinct locations. To collect it, YouTube videos were automatically processed. As we can see, this dataset is much larger than multi-camera datasets for person Re-ID and covers a much more diverse range of capturing environments (Tab. 1). Therefore, this kind of data is also useful for training Re-ID algorithms.

**Metrics.** In our experiments, we use Cumulative Matching Characteristics (CMC)  $Rank_1$ , as well as mean Average Precision ( $mAP$ ) to evaluate our method.

### 4.2. Implementation Details

In this paper, we use ResNet50 [13] with IBN-a [33] layers as the encoder and the momentum encoder. These encoders are self-supervised pre-trained on single-camera data from LUPerson using MoCo v2 [5]. Adam is used as an optimizer with a learning rate of 0.00035, a weight decay rate of 0.0005, and with a warm-up scheme in the first 10 epochs. As for the momentum coefficient  $\lambda$  in Eq. (2), we set  $\lambda = 0.999$ . ReMix is trained for 100 epochs. In our experiments, we set  $N_P^m = N_P^s = 8$  and  $N_K^m = N_K^s = 4$ , so the size of each mini-batch is 64. According to [44], we choose  $\gamma = 0.5$  in Eq. (1). In ReMix, all images are resized to  $256 \times 128$ , random crops, horizontal flipping, Gaussian blurring, and random grayscale are also applied to them.

### 4.3. Parameter Analysis

#### 4.3.1 Temperature Parameters

**Multi-camera parameters analysis.** First, we analyze the quality of our ReMix method for different values of parameters  $\tau_{ins_m}$  and  $\tau_{aug}$  in the Instance Loss (Sec. 3.3.1) and the Augmentation Loss (Sec. 3.3.2), respectively. Single-camera data is not used in these experiments. As can be seen from Tab. 2a, the best quality of cross-dataset Re-ID can be achieved with  $\tau_{ins_m} = \tau_{aug} = 0.1$ . According to [2], we choose  $\tau_{cen_m} = 0.5$  in Eq. (8).

**Single-camera parameters analysis.** Multi-camera and

$\tau_{ins_m} \backslash \tau_{aug}$	0.07	0.10	0.15
0.07	75.1/58.4	75.1/58.5	74.9/58.3
0.10	75.1/58.7	<b>75.8/58.7</b>	75.0/58.6
0.15	75.0/58.3	75.0/58.5	74.6/58.2

(a) Analysis of values for parameters  $\tau_{ins_m}$  and  $\tau_{aug}$  in Eq. (4) and Eq. (6). In this table, the first number is  $Rank_1$ , and the second is  $mAP$ .

$\tau_{ins_s}$	0.07	0.10	0.15	0.20	0.25
$Rank_1$	75.7	76.2	76.3	<b>76.3</b>	74.9
$mAP$	59.1	59.6	59.6	<b>59.9</b>	59.5

(b) Analysis of values for parameter  $\tau_{ins_s}$  in Eq. (5).

$\tau_{cen_s}$	0.40	0.50	0.60	0.65
$Rank_1$	76.3	76.3	<b>76.9</b>	75.8
$mAP$	60.0	60.4	<b>60.7</b>	60.4

(c) Analysis of values for parameter  $\tau_{cen_s}$  in Eq. (8).

Table 2. Temperature parameters analysis. In these experiments, we train the algorithm on MSMT17-merged and single-camera data from LUPerson, and test it on DukeMTMC-reID.

single-camera data have different complexities in terms of person Re-ID. So, in the Instance Loss (Sec. 3.3.1) and the Centroids Loss (Sec. 3.3.3) we propose to use special temperature parameters for single-camera data ( $\tau_{ins_s}$  and  $\tau_{cen_s}$ , respectively). According to Tab. 2b and Tab. 2c, the best results achieved when  $\tau_{ins_s} = 0.2$  and  $\tau_{cen_s} = 0.6$ .

**Conclusions from the analysis.** The temperature parameters  $\tau_{ins_m} = 0.1$  and  $\tau_{cen_m} = 0.5$  are selected for multi-camera data,  $\tau_{ins_s} = 0.2$  and  $\tau_{cen_s} = 0.6$  are selected for single-camera data. Higher temperature values make the probabilities closer together, which complicates training on simpler single-camera data. Accordingly, we confirm our hypothesis about the different complexities of multi-camera and single-camera data.

### 4.3.2 Epoch Duration

To achieve training stability and frequent updating of centroids, only a portion of the images is used during one epoch (Sec. 3.1). Also, this approach reduces computational costs by generating pseudo labels only for a subset of single-camera data in one epoch, rather than for an entire large dataset (Sec. 3.2). In this paper, one epoch consists of 400 iterations. As can be seen from the experimental results presented in Tab. 3, this number of iterations is a trade-off between the accuracy of our method and its training time.

### 4.4. Ablation Study

**Proof-of-concept.** We conduct a series of experiments to demonstrate the effectiveness of the proposed idea of joint

Iterations	300	400	600	800
$Rank_1$	76.4	<b>77.6</b>	77.1	77.2
$mAP$	61.1	61.6	<b>62.0</b>	61.6
Training Time*	~15h	~20h	~30h	~40h

\* Two Nvidia GTX 1080 Ti are used for training.

Table 3. Comparison of different numbers of iterations in one epoch. We train the algorithm on MSMT17-merged and single-camera data from LUPerson, and test it on DukeMTMC-reID.

Using s-cam. data		Market-1501		DukeMTMC-reID	
Pre-train	Joint	$Rank_1$	$mAP$	$Rank_1$	$mAP$
$\times$	$\times$	78.4	51.7	75.8	58.7
$\checkmark$	$\times$	81.7	54.9	75.1	59.2
$\times$	$\checkmark$	81.3	57.0	76.9	60.7
$\checkmark$	$\checkmark$	<b>84.0</b>	<b>61.0</b>	<b>77.6</b>	<b>61.6</b>

Table 4. Impact of using single-camera data in self-supervised pre-training and in our joint training procedure. In these experiments, we use MSMT17-merged and single-camera data from LUPerson (where applicable) for training.

Configuration	$Rank_1$	$mAP$
w/o single-camera data	75.8	58.7
+ in $\mathcal{L}_{aug}$	76.0	59.2
+ in $\mathcal{L}_{ins}$	76.3	59.9
+ in $\mathcal{L}_{cen}$ only as centroids	75.4	60.0
+ in $\mathcal{L}_{cen}$	<b>76.9</b>	<b>60.7</b>

Table 5. Step-by-step use of single-camera data in different loss functions. Here, "in  $\mathcal{L}_{cen}$  only as centroids" means that single-camera data is used only as centroids in the Centroids Loss. We train the algorithm on MSMT17-merged and single-camera data from LUPerson, and test it on DukeMTMC-reID.

training on multi-camera and single-camera data. The results of these experiments are presented in Tab. 4. As we can see, using single-camera data in addition to multi-camera data significantly improves the generalization ability of the algorithm and the quality of cross-dataset Re-ID. It is worth noting that the use of single-camera data most significantly affects the  $mAP$  metric. That is, our method produces higher similarity values for images of the same person and lower values for different ones. This is achieved due to a more diverse training data, which is primarily obtained from large amounts of single-camera data.

Moreover, the effectiveness of our approach is demonstrated in comparison with self-supervised pre-training: the model trained using the proposed joint training procedure achieves better accuracy than the self-supervised pre-trained model. In Sec. 1 we hypothesized that self-

Method	Reference	Market-1501	
		$Rank_1$	$mAP$
SNR [18]	CVPR20	66.7	33.9
MetaBIN [6]	CVPR21	69.2	35.9
MDA [30]	CVPR22	70.3	38.0
DTIN-Net [17]	ECCV22	69.8	37.4
ReMix (w/o s-cam.)	Ours	68.2	37.7
ReMix	Ours	<b>71.3</b>	<b>43.0</b>

(a) Training dataset: DukeMTMC-reID.

Method	Reference	DukeMTMC-reID	
		$Rank_1$	$mAP$
SNR [18]	CVPR20	55.1	33.6
MetaBIN [6]	CVPR21	55.2	33.1
MDA [30]	CVPR22	56.7	34.4
DTIN-Net [17]	ECCV22	57.0	36.1
ReMix (w/o s-cam.)	Ours	57.1	36.5
ReMix	Ours	<b>58.4</b>	<b>38.8</b>

(b) Training dataset: Market-1501.

Method	Reference	Training Dataset	CUHK03-NP		Market-1501		DukeMTMC-reID	
			$Rank_1$	$mAP$	$Rank_1$	$mAP$	$Rank_1$	$mAP$
SNR [18]	CVPR20	MSMT17	—	—	70.1	41.4	69.2	50.0
QAConv [24]	ECCV20		25.3	22.6	72.6	43.1	69.4	52.6
TransMatcher [25]	NeurIPS21		23.7	22.5	<b>80.1</b>	52.0	—	—
QAConv-GS [26]	CVPR22		20.9	20.6	79.1	49.5	67.3	49.4
PAT [29]	ICCV23		24.2	25.1	72.2	47.3	—	—
ReMix (w/o s-cam.)	Ours		24.1	24.5	73.0	42.5	68.9	49.2
ReMix	Ours		<b>27.3</b>	<b>27.4</b>	78.2	<b>52.4</b>	<b>71.6</b>	<b>52.8</b>
OSNet [58]	CVPR19	MSMT17-merged	—	—	66.5	37.2	—	—
OSNet-AIN [59]	TPAMI21		—	—	70.1	43.3	—	—
TransMatcher [25]	NeurIPS21		31.9	30.7	82.6	58.4	—	—
QAConv-GS [26]	CVPR22		27.6	28.0	80.6	55.6	71.3	53.5
ReMix (w/o s-cam.)	Ours		34.5	32.7	78.4	51.7	75.8	58.7
ReMix	Ours		<b>37.7</b>	<b>37.2</b>	<b>84.0</b>	<b>61.0</b>	<b>77.6</b>	<b>61.6</b>
RP Baseline [46]	ACMMM20	RandPerson	13.4	10.8	55.6	28.8	—	—
CBN [53]	ECCV20		—	—	64.7	39.3	—	—
QAConv-GS [26]	CVPR22		14.8	13.4	<b>74.0</b>	43.8	—	—
ReMix (w/o s-cam.)	Ours		17.1	15.7	71.1	42.4	61.2	39.0
ReMix	Ours		<b>19.3</b>	<b>18.4</b>	72.7	<b>45.4</b>	<b>63.2</b>	<b>42.8</b>

(c) Training datasets: MSMT17, MSMT17-merged, and RandPerson.

Table 6. Comparison of our ReMix method with others in the cross-dataset scenario. In this comparison, we use two versions of the proposed method: without using single-camera data and with using single-camera data during training. Here, we use the LUPerson dataset as single-camera data to train ReMix.

supervised pre-training has a limited effect, since subsequent fine-tuning for the final task is performed on relatively small multi-camera data. The results of our experiments validate this hypothesis. Indeed, by using our joint training procedure together with self-supervised pre-training, we can achieve the best quality. Thus, we experimentally confirm the importance of data volume at the fine-tuning stage. ReMix uses unlabeled single-camera data, and this result can also verify that self-supervised pre-training improves the quality of clustering and pseudo labeling.

**Using single-camera data in loss functions.** In addition to experiments showing the validity of our joint training procedure, we conduct an ablation study to demonstrate the ef-

fectiveness of adapting the proposed loss functions for joint use with two types of data. In this study, we gradually add single-camera data in losses and measure the final accuracy. As we can see from Tab. 5, each loss function added to a combination improves the performance, and using all losses with single-camera data jointly provides the highest quality. Thus, the proposed loss functions are successfully adapted for joint use with two types of training data — multi-camera and single-camera.

#### 4.5. Comparison with State-of-the-Art Methods

We compare our ReMix method with other state-of-the-art Re-ID approaches using two test protocols: the cross-

Method	Reference	M+D+MS $\rightarrow$ C3		D+C3+MS $\rightarrow$ M		M+C3+MS $\rightarrow$ D	
		$Rank_1$	$mAP$	$Rank_1$	$mAP$	$Rank_1$	$mAP$
MECL [50]	arXiv21	32.1	31.5	80.0	56.5	70.0	53.4
M <sup>3</sup> L [55]	ICCV21	36.4	35.2	81.5	59.6	71.8	54.5
RaMoE [7]	CVPR21	36.6	35.5	82.0	56.5	73.6	56.9
MetaBIN [6]	CVPR21	38.1	37.5	83.2	61.2	71.3	54.9
MixNorm [34]	TMM22	29.6	29.0	78.9	51.4	70.8	49.9
META [49]	ECCV22	46.0	45.9	85.3	65.7	76.9	59.9
IL [41]	TMM23	40.9	38.3	86.2	65.8	75.4	57.1
ReMix	Ours	<b>47.6</b>	<b>46.5</b>	<b>87.8</b>	<b>70.5</b>	<b>79.0</b>	<b>63.3</b>

Table 7. Comparison of our ReMix method with others in the multi-source cross-dataset scenario. Here, we use the LUPerson dataset as single-camera data to train ReMix. In this table, C3 is CUHK03-NP, M is Market-1501, D is DukeMTMC-reID, and MS is MSMT17.

dataset and multi-source cross-dataset scenarios. According to the first protocol, we train the algorithm on one multi-camera dataset and test it on another multi-camera dataset. In the multi-source cross-dataset scenario, we train the algorithm on several multi-camera datasets and test it on another multi-camera dataset. Thus, we evaluate the generalization ability of our method in comparison to other existing state-of-the-art Re-ID approaches. Also, we illustrate several complex examples in Fig. 3, where ReMix manages to notice important visual cues.

**The cross-dataset scenario.** As can be seen from Tab. 6, the proposed method demonstrates a high generalization ability and outperforms others in the cross-dataset scenario. In our ReMix method, the momentum encoder is trained to obtain embeddings for each query and gallery image, after which they are compared using cosine similarity. QAConv [24], TransMatcher [25], and QAConv-GS [26], which are among the most accurate methods in cross-dataset person Re-ID, use more complex architectures: in addition to the encoder, a separate neural network is used. This network compares features between the query and gallery images and predicts the probability that they belong to the same person. PAT [29] uses a transformer-based model, which is more computationally complex compared to ResNet50 with IBN-a layers in ReMix. Thus, most existing state-of-the-art approaches improve generalization ability by using complex architectures. In contrast, the high performance of our method is achieved through the training strategy that does not affect the computational complexity, so our method can seamlessly replace other methods used in real-world applications. It is also worth noting that in the comparison in Tab. 6, some methods use larger input images. In the supplementary material, we show that the accuracy of ReMix increases with the size of the input image (see Sec. 6.3).

**The multi-source cross-dataset scenario.** The comparison presented in Tab. 7 shows the effectiveness of our joint training procedure, even when using several multi-camera



Figure 3. Comparison of TOP-5 retrieved images on the Market-1501 dataset between ReMix and QAConv-GS [26]. Green boxes denote correct results, while red boxes denote incorrect results.

datasets and one single-camera dataset during training. This further proves the consistency and flexibility of ReMix.

## 5. Conclusion

In this paper, we proposed ReMix, a novel person Re-ID method that achieves generalization by jointly using limited labeled multi-camera and large unlabeled single-camera data for training. To the best of our knowledge, this is the first work that explores joint training on a mixture of multi-camera and single-camera data in person Re-ID. To provide effective training, we developed a novel data sampling strategy and new loss functions adapted for joint use with these two types of data. Through experiments, we showed that our method has a high generalization ability and outperforms state-of-the-art methods in the cross-dataset and multi-source cross-dataset scenarios. We believe our work will serve as a basis for future research dedicated to generalized, accurate, and reliable person Re-ID.



## References

- [1] Keni Bernardin and Rainer Stiefelwagen. Evaluating multiple object tracking performance: the clear mot metrics. *EURASIP Journal on Image and Video Processing*, 2008:1–10, 2008. 5
- [2] Hao Chen, Benoit Lagadec, and Francois Bremond. Ice: Inter-instance contrastive encoding for unsupervised person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14960–14969, 2021. 5
- [3] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 2
- [4] Weihua Chen, Xianzhe Xu, Jian Jia, Hao Luo, Yaohua Wang, Fan Wang, Rong Jin, and Xiuyu Sun. Beyond appearance: a semantic controllable self-supervised learning framework for human-centric visual tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15050–15061, 2023. 2, 4
- [5] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 2, 5
- [6] Seokeon Choi, Taekyung Kim, Minki Jeong, Hyoungseob Park, and Changick Kim. Meta batch-instance normalization for generalizable person re-identification. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 3425–3435, 2021. 2, 7, 8
- [7] Yongxing Dai, Xiaotong Li, Jun Liu, Zekun Tong, and Ling-Yu Duan. Generalizable person re-identification with relevance-aware mixture of experts. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16145–16154, 2021. 8
- [8] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. With a little help from my friends: Nearest-neighbor contrastive learning of visual representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9588–9597, 2021. 2
- [9] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, volume 96, pages 226–231, 1996. 4, 1
- [10] Dengpan Fu, Dongdong Chen, Jianmin Bao, Hao Yang, Lu Yuan, Lei Zhang, Houqiang Li, and Dong Chen. Unsupervised pre-training for person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14750–14759, 2021. 1, 2, 5
- [11] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018. 2
- [12] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. 2
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5, 3
- [14] Shuting He, Hao Luo, Pichao Wang, Fan Wang, Hao Li, and Wei Jiang. Transreid: Transformer-based object re-identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 15013–15022, 2021. 2
- [15] Jieru Jia, Qiuqi Ruan, and Timothy M Hospedales. Frustratingly easy person re-identification: Generalizing person re-id in practice. *arXiv preprint arXiv:1905.03422*, 2019. 3
- [16] Mengxi Jia, Xinhua Cheng, Shijian Lu, and Jian Zhang. Learning disentangled representation implicitly via transformer for occluded person re-identification. *IEEE Transactions on Multimedia*, 25:1294–1305, 2022. 4
- [17] Bingliang Jiao, Lingqiao Liu, Liying Gao, Guosheng Lin, Lu Yang, Shizhou Zhang, Peng Wang, and Yanning Zhang. Dynamically transformed instance normalization network for generalizable person re-identification. In *European Conference on Computer Vision*, pages 285–301. Springer, 2022. 2, 7
- [18] Xin Jin, Cuiling Lan, Wenjun Zeng, Zhibo Chen, and Li Zhang. Style normalization and restitution for generalizable person re-identification. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3143–3152, 2020. 2, 7
- [19] L Leal-Taixe. Motchallenge 2015: Towards a benchmark for multi-target tracking. *arXiv preprint arXiv:1504.01942*, 2015. 5
- [20] Hanjun Li, Gaojie Wu, and Wei-Shi Zheng. Combined depth space based architecture search for person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6729–6738, 2021. 4
- [21] Siyuan Li, Li Sun, and Qingli Li. 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, pages 1405–1413, 2023. 2, 4
- [22] Wei Li, Rui Zhao, Tong Xiao, and Xiaogang Wang. Deepreid: Deep filter pairing neural network for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 152–159, 2014. 1, 5
- [23] Yuan Li, Chang Huang, and Ram Nevatia. Learning to associate: Hybridboosted multi-target tracker for crowded scene. In *2009 IEEE conference on computer vision and pattern recognition*, pages 2953–2960. IEEE, 2009. 5
- [24] Shengcai Liao and Ling Shao. Interpretable and generalizable person re-identification with query-adaptive convolution and temporal lifting. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 456–474. Springer, 2020. 7, 8, 2
- [25] Shengcai Liao and Ling Shao. Transmatcher: Deep image matching through transformers for generalizable person re-identification. *Advances in Neural Information Processing Systems*, 34:1992–2003, 2021. 7, 8, 2
- [26] Shengcai Liao and Ling Shao. Graph sampling based deep metric learning for generalizable person re-identification. In

- Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7359–7368, 2022. 2, 7, 8
- [27] Timur Mamedov, Denis Kuplyakov, and Anton Konushin. Approaches to improve the quality of person re-identification for practical use. *Sensors*, 23(17):7382, 2023. 2
- [28] Anton Milan. Mot16: A benchmark for multi-object tracking. *arXiv preprint arXiv:1603.00831*, 2016. 5
- [29] Hao Ni, Yuke Li, Lianli Gao, Heng Tao Shen, and Jingkuan Song. Part-aware transformer for generalizable person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11280–11289, 2023. 2, 7, 8
- [30] Hao Ni, Jingkuan Song, Xiaopeng Luo, Feng Zheng, Wen Li, and Heng Tao Shen. Meta distribution alignment for generalizable person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2487–2496, 2022. 7
- [31] Xingyang Ni and Esa Rahtu. Flipreid: closing the gap between training and inference in person re-identification. In *2021 9th European Workshop on Visual Information Processing (EUVIP)*, pages 1–6. IEEE, 2021. 2, 4
- [32] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, pages 69–84. Springer, 2016. 2
- [33] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 464–479, 2018. 5, 3
- [34] Lei Qi, Lei Wang, Yinghuan Shi, and Xin Geng. A novel mix-normalization method for generalizable multi-source person re-identification. *IEEE Transactions on Multimedia*, 2022. 8
- [35] Yongming Rao, Guangyi Chen, Jiwen Lu, and Jie Zhou. Counterfactual attention learning for fine-grained visual categorization and re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1025–1034, 2021. 4
- [36] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *European conference on computer vision*, pages 17–35. Springer, 2016. 1, 5
- [37] Yantao Shen, Hongsheng Li, Shuai Yi, Dapeng Chen, and Xiaogang Wang. Person re-identification with deep similarity-guided graph neural network. In *Proceedings of the European conference on computer vision (ECCV)*, pages 486–504, 2018. 2
- [38] Yumin Suh, Jingdong Wang, Siyu Tang, Tao Mei, and Kyoung Mu Lee. Part-aligned bilinear representations for person re-identification. In *Proceedings of the European conference on computer vision (ECCV)*, pages 402–419, 2018. 2
- [39] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin Wang. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *Proceedings of the European conference on computer vision (ECCV)*, pages 480–496, 2018. 2
- [40] Lei Tan, Pingyang Dai, Rongrong Ji, and Yongjian Wu. Dynamic prototype mask for occluded person re-identification. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 531–540, 2022. 2
- [41] Wentao Tan, Changxing Ding, Pengfei Wang, Mingming Gong, and Kui Jia. Style interleaved learning for generalizable person re-identification. *IEEE Transactions on Multimedia*, 2023. 8
- [42] Guanshuo Wang, Yufeng Yuan, Xiong Chen, Jiwei Li, and Xi Zhou. Learning discriminative features with multiple granularities for person re-identification. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 274–282, 2018. 2
- [43] Haochen Wang, Jiayi Shen, Yongtuo Liu, Yan Gao, and Efstratios Gavves. Nformer: Robust person re-identification with neighbor transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7297–7307, 2022. 4
- [44] Menglin Wang, Baisheng Lai, Jianqiang Huang, Xiaojin Gong, and Xian-Sheng Hua. Camera-aware proxies for unsupervised person re-identification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 2764–2772, 2021. 5
- [45] Pingyu Wang, Zhicheng Zhao, Fei Su, and Hongying Meng. Ltireid: Factorizable feature generation with independent components for long-tailed person re-identification. *IEEE Transactions on Multimedia*, 25:4610–4622, 2022. 4
- [46] Yanan Wang, Shengcai Liao, and Ling Shao. Surpassing real-world source training data: Random 3d characters for generalizable person re-identification. In *Proceedings of the 28th ACM international conference on multimedia*, pages 3422–3430, 2020. 5, 7
- [47] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person transfer gan to bridge domain gap for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 79–88, 2018. 1, 5
- [48] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple online and realtime tracking with a deep association metric. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 3645–3649. IEEE, 2017. 5
- [49] Boqiang Xu, Jian Liang, Lingxiao He, and Zhenan Sun. Meta: Mimicking embedding via others’ aggregation for generalizable person re-identification. In *Proceedings of the European conference on computer vision (ECCV)*, 2022. 8
- [50] Shijie Yu, Feng Zhu, Dapeng Chen, Rui Zhao, Haobin Chen, Shixiang Tang, Jinguo Zhu, and Yu Qiao. Multiple domain experts collaborative learning: Multi-source domain generalization for person re-identification. *arXiv preprint arXiv:2105.12355*, 2021. 8
- [51] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR, 2021. 2
- [52] Guiwei Zhang, Yongfei Zhang, Tianyu Zhang, Bo Li, and Shiliang Pu. Pha: Patch-wise high-frequency augmentation for transformer-based person re-identification. In *Proceed-*

- ings of the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14133–14142, 2023. [2](#)
- [53] Tianyu Zhang, Lingxi Xie, Longhui Wei, Zijie Zhuang, Yongfei Zhang, Bo Li, and Qi Tian. Unrealperson: An adaptive pipeline towards costless person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11506–11515, 2021. [7](#)
- [54] Zhizheng Zhang, Cuiling Lan, Wenjun Zeng, Xin Jin, and Zhibo Chen. Relation-aware global attention for person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3186–3195, 2020. [4](#)
- [55] Yuyang Zhao, Zhun Zhong, Fengxiang Yang, Zhiming Luo, Yaojin Lin, Shaozi Li, and Nicu Sebe. Learning to generalize unseen domains via memory-based multi-source meta-learning for person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6277–6286, 2021. [8](#)
- [56] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *Proceedings of the IEEE international conference on computer vision*, pages 1116–1124, 2015. [1](#), [5](#)
- [57] Liang Zheng, Hengheng Zhang, Shaoyan Sun, Manmohan Chandraker, Yi Yang, and Qi Tian. Person re-identification in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1367–1376, 2017. [2](#)
- [58] Kaiyang Zhou, Yongxin Yang, Andrea Cavallaro, and Tao Xiang. Omni-scale feature learning for person re-identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3702–3712, 2019. [2](#), [7](#)
- [59] Kaiyang Zhou, Yongxin Yang, Andrea Cavallaro, and Tao Xiang. Learning generalisable omni-scale representations for person re-identification. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):5056–5069, 2021. [2](#), [7](#), [3](#)
- [60] Xiao Zhou, Yujie Zhong, Zhen Cheng, Fan Liang, and Lin Ma. Adaptive sparse pairwise loss for object re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19691–19701, 2023. [4](#)
- [61] Kuan Zhu, Haiyun Guo, Zhiwei Liu, Ming Tang, and Jinqiao Wang. Identity-guided human semantic parsing for person re-identification. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 346–363. Springer, 2020. [4](#)