

Deep-Person: Learning Discriminative Deep Features for Person Re-Identification

Xiang Bai, Mingkun Yang, Tengpeng Huang, Zhiyong Dou, Rui Yu, Yongchao Xu
Huazhong University of Science and Technology

{xbai, yangmingkun, tengtenghuang, zydou, yongchaoxu}@hust.edu.cn yurui.thu@gmail.com

Abstract

Recently, many methods of person re-identification (Re-ID) rely on part-based feature representation to learn a discriminative pedestrian descriptor. However, the spatial context between these parts is ignored for the independent extractor on each separate part. In this paper, we propose to apply Long Short-Term Memory (LSTM) in an end-to-end way to model the pedestrian, seen as a sequence of body parts from head to foot. Integrating the contextual information strengthens the discriminative ability of local representation. We also leverage the complementary information between local and global feature. Furthermore, we integrate both identification task and ranking task in one network, where a discriminative embedding and a similarity measurement are learned concurrently. This results in a novel three-branch framework named Deep-Person, which learns highly discriminative features for person Re-ID. Experimental results demonstrate that Deep-Person outperforms the state-of-the-art methods by a large margin on three challenging datasets including Market-1501, CUHK03, and DukeMTMC-reID. Specifically, combining with a re-ranking approach, we achieve a **90.84%** mAP on Market-1501 under single query setting.

1. Introduction

Person re-identification (Re-ID) refers the task of matching a specific person across multiple non-overlapping cameras. It has been receiving increasing attention in the computer vision community thanks to its various surveillance applications. Despite decades of study on person Re-ID task, it is still very challenging due to large variations in illumination, pose, background clutter, occlusion, and ambiguity in visual appearance.

Most early works in person Re-ID either focus on discriminative hand-craft feature representation or robust distance metric for similarity measurement. Benefiting by the development of deep learning and increasing large-scale



Figure 1. Part misalignment problems (green parts) caused by inaccurate pedestrian detection, occlusion, and pose variations when using predefined rigid body parts for local representation learning.

datasets [47, 43, 18], recent person Re-ID methods combine feature extraction and distance metric into an end-to-end deep convolution neural network (CNN). Nevertheless, most recent CNN-based methods endeavor to either design a better feature representation or develop a more robust feature learning, but rarely both aspects together.

The CNN-based methods focusing on better feature representations can be roughly divided into three categories: 1) Global full-body representation, which is adopted in many methods [18, 5]; 2) Local body-part representation, which has been exploited in many works with variant part partitions. A straightforward partition into predefined rigid body parts is used in many works [36, 31, 27]. In practice, as shown in Fig. 1, due to pose variations, imperfect pedestrian detectors, and occlusion, such trivial partition fails to correctly align two pedestrian images, leading to part-based features far from robust. Some recent works endeavor to develop better body partitions with some sophisticated methods [39, 17] or using extra pose annotation [41, 28]. All these methods ignore the contextual information between the body parts. In [33], the authors propose to first convert the original person image into sequential LOMO and Color Names features, then rely on Recurrent Neural Network (RNN) to model the spatial context. Yet, the separation of sequence feature extraction and spatial context modeling hinders the end-to-end training and optimization, resulting in degraded performance; 3) Combination of global and local representation [8, 17, 28, 36, 37], which concatenates the global and part-based feature as the final feature representation. This combined feature representation usually requires more computation and extra space in test phase

due to an extra branch compared to the single branch model, yielding slower runtime in practice.

The methods dedicated for robust feature learning usually consider the person Re-ID problem as either a classification task or a ranking task. Thanks to recently increasing large-scale Re-ID datasets, person Re-ID is regarded as a multi-class person identification task in many works [18, 45, 44]. The obtained ID-discriminative Embedding (IDE) given by the penultimate fully connected layer has shown great potentials for person Re-ID. Yet, the training objective of identification task is not totally consistent with the testing procedure. Such learned IDE is optimal for identification model on the training set, but may not be optimal to describe unseen person images in the test stage. Furthermore, the identification task does not explicitly learn a similarity measurement required for retrieval during the test stage. On the other hand, the ranking task aims to make the distance between the positive pairs closer than that of the negative pairs by a given margin. Therefore, a similarity measurement is explicitly learned. Yet, all identity information of the annotated Re-ID is not fully utilized. Recently, some works [5, 34, 22] leverage both classification and ranking tasks with triplet loss to learn more discriminative features for person Re-ID.

In this paper, we propose to model the pedestrian as a sequence of body parts from head to foot, and learn all part features together with the spatial contextual information rather than use independent branch for each separate part. For that, we apply Long Short-Term Memory (LSTM) [15] in an end-to-end way, enhancing the discriminative capacity of local feature with the prior knowledge of body structure and mitigating the misalignment caused by the separated part feature extractors. We also adopt global full-body representation in the proposed Deep-Person model. We feed the global and part-based features into two separate network branches for identification tasks. Different from the classical combination of global and local representation [8, 17, 28, 36, 37] concatenating global and part-based feature as the final feature representation for person Re-ID, we further add a ranking task branch using triplet loss to explicitly learn the similarity measurement. More specifically, we use the global average pooled feature f_m of the backbone feature f_b (shown in Fig. 2) as similarity estimation, also considered as the final pedestrian descriptor. Such three-branch Deep-Person model learns highly discriminative features for person Re-ID.

The main contributions of this paper are three folds: 1) We propose to regard the pedestrian as a sequence of body parts from head to foot, and apply LSTM in an end-to-end fashion to enhance the discriminative capacity of local representation; 2) We develop a novel three-branch framework which leverages two kinds of complementary advantages: local body-part and global full-body feature representation,

as well as identification task and ranking task for a better feature learning. The proposed Deep-Person yields highly discriminative features for person Re-ID; 3) In test phase, the proposed Deep-Person only performs a forward pass of the backbone network followed by a global average pooling. Consequently, compared to the single branch model, our model requires no additional runtime and space during testing, while still outperforms the state-of-the-art methods by large margins on three popular Re-ID datasets.

2. Related Work

We focus on two types of closely related deep learning methods for person Re-ID: those relying on part-based representations and those focusing on multi-loss learning. For a complete review of person Re-ID methods, the interested reader is referred to [44].

Part-based person Re-ID approaches. There are some methods use part-based representation to learn discriminative features for person Re-ID. Following the strategy of part partition, the part-based approaches can be roughly divided into two categories: 1) Rigid body part partition, which has been widely adopted in many methods [8, 36, 31, 27, 37, 3, 33], where the authors use predefined rigid grids as local parts. Each part is fed into an individual branch. All the individual part features are then concatenated together as the final part-based representation; 2) Flexible body part partition, which is more reasonable to localize appropriate parts. For example, Yao *et al.* [39] use an unsupervised method to generate a set of part boxes, and then employ the RoI pooling to produce part features. The Spatial Transformer Networks (STN) [16] with novel spatial constraints are applied to localize deformable person parts in [17]. With extra pose annotation, Zhao *et al.* [41] utilize the learned body joints to obtain subregion bounding boxes. Su *et al.* [28] further extend [41] by normalizing the pose-based parts into fixed size and orientation, and introducing the Pose Transformation Network (PTN) to eliminate the pose variations.

Although rigid body parts are simple to implement, due to the inaccurate pedestrian detection and occlusion, the misalignment of rigid body parts (see Fig. 1) is not beneficial for discriminative feature learning. The representation based on flexible body parts alleviates the misalignment problem to a certain extent by localizing appropriate parts. Yet, such methods usually have a much more complex pipeline or require extra prior knowledge (*e.g.*, human pose). Furthermore, the existing methods based on either rigid or flexible body parts all ignore the relationship between the body parts. Whereas, the pedestrian can always be decomposed into a sequence of body parts from head to foot. Based on this simple yet important property of pedestrian, we propose to learn all the parts together in a

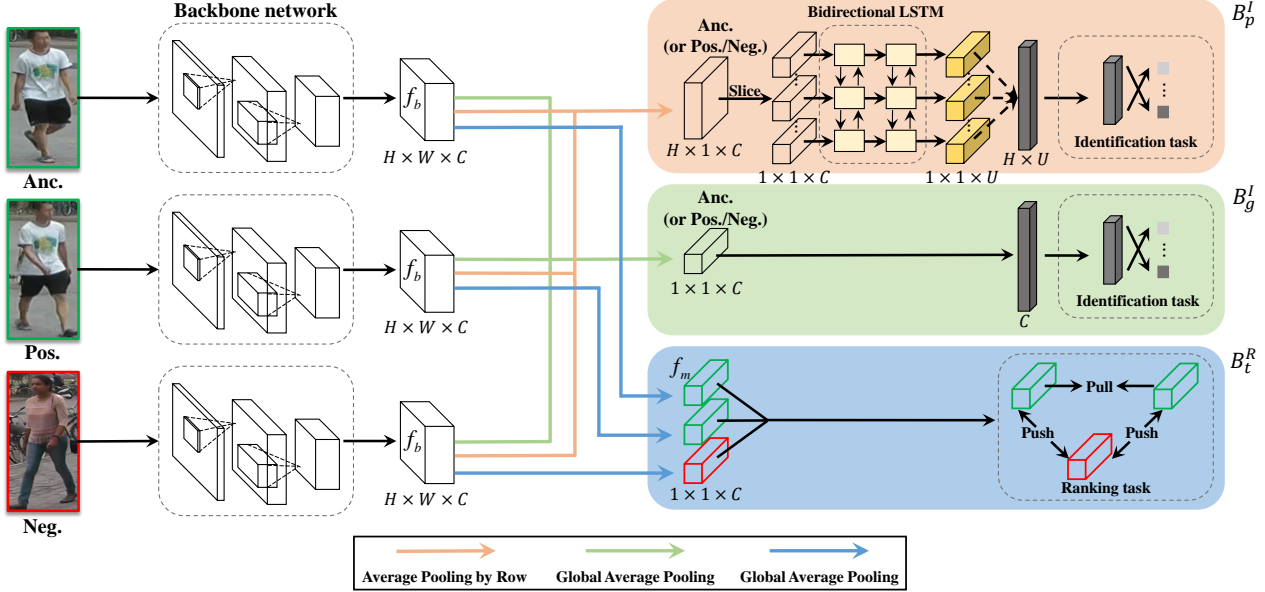


Figure 2. Illustration of the Deep-Person architecture. Given a triplet of images T , each image is fed into a part-based identification branch B_p^I and a global-based identification branch B_g^I . Meanwhile, a distance ranking branch B_t^R is applied on T using the triplet loss function. Note that each image in T is fed into the same backbone network. The duplication of the backbone is only for visualization purpose.

sequence way rather than discarding the spatial context by independent part feature extractors. For that, we naturally apply LSTM for the sequence-level modeling. The LSTM is also used in [33] to model the spatial contextual information. However, the authors first divide the input image into rigid parts and extract hand-craft feature for each individual part, then use LSTM to model the spatial relationship in a separate step. The proposed Deep-Person jointly integrates the deep feature extraction and sequence modeling in an end-to-end fashion, leading to more discriminative features for person Re-ID.

Approaches based on joint multi-loss learning. Recently, Zheng *et al.* [44] propose that person Re-ID lies in between instance retrieval task and image classification task. For the first point, the person Re-ID is regarded as a ranking task, where a ranking loss is adopted for feature learning. In [8], a new term is added to the original triplet loss to further pull the instances of the same person closer. Hermans *et al.* [14] introduce a variant of the standard triplet loss using hard mining within the batch. From the point of view of classification task, the person Re-ID problem is usually solved with a Softmax loss. There are two ways to perform person Re-ID as a classification task. The first one is known as verification network [18], which takes a pair of images as input and determines whether they belong to the same identity or not by a binary classification network. The second one is called identification network [44], namely multi-class recognition network, where each individual is regarded as an independent category.

The classification task and ranking task are complementary to each other. Some approaches optimize the network simultaneously with both type of loss functions, leveraging the complementary advantages of these two tasks. For example, In [5, 34], triplet loss and verification loss are trained together. The identification loss and verification loss are simultaneously optimized in [11, 24]. The combination of triplet loss and identification loss is adopted in [22] to optimize the comparative attention network. The proposed Deep-Person also rely on triplet loss and identification loss. Different from [22], the identification loss on the novel part-based representation is also adopted in addition to the identification loss on global representation, further boosting the discriminative ability of learned feature for person Re-ID.

3. Architecture of Deep-Person

3.1. Overview of Deep-Person

The proposed Deep-Person model focuses on both feature representation and feature learning. It is built upon two kinds of complementary designs detailed in the following: 1) Global representation and part-based local representation; 2) Softmax-based identification branch and ranking branch with the triplet loss.

Recent advances in person re-identification rely on deep learning to learn discriminative features from detected pedestrians. As pointed out in [17], the learned representation of the full body focus more on global information such as the pose and shape. Whereas, in some cases, only cer-




Figure 3. Each vector in the feature sequence describes the region of the corresponding receptive field in the raw image.

tain body parts such as the head, upper body, or lower body are discriminative for person re-identification [8, 27, 33]. In this sense, the part-based local representation of detected pedestrian is complementary to the global representation. We propose to use the LSTM-based RNN to naturally model the spatial dependency between each part of the pedestrian body, which can be seen as a sequence of body parts from head to foot. Hopefully, combining the complementary global representation and the LSTM-based local representation would strengthen the discriminative ability of learned features for person re-identification.

Recently, a common component of most deep models for person Re-ID is the Softmax-based identification branch that distinguishes different IDs based on the learned deep features. Yet, the training objective for identification branch is not totally consistent with the goal of person Re-ID, which aims to pair each probe with one of the gallery image. This is because that the identification branch does not explicitly learn a similarity measurement which is required for person Re-ID. Recently, as advocated in [22], a distance ranking branch with the triplet loss helps to learn a similarity measure between a pair of images. In this sense, the identification branch and the distance ranking branch with the triplet loss constitutes another complementarity.

Our proposed Deep-Person model leverages the above two kinds of complementarities. The overall architecture is depicted in Fig. 2. It is composed of two main components: (1) Backbone network for learning shared low-level feature f_b with size $H \times W \times C$; (2) Multi-branch network to learn a highly discriminative pedestrian representation thanks to the three complementary branches: part-based identification branch B_p^I , global-based identification branch B_g^I , and distance ranking branch B_t^R using the triplet loss. A joint learning strategy is adopted for simultaneously optimizing per-branch feature representation and discovering correlated complementary information.

3.2. Part-Sequence Learning Using LSTM

It has been shown in many methods [8, 27, 33] that part-based representation is useful for person Re-ID. Most part-based methods roughly decompose the extracted pedestrian into predefined rigid body parts (see Fig. 1) which approximately correspond to head, shoulder, upper-body, upper-leg and lower-leg, respectively. Each segmented part is then fed into an individual branch to learn the corresponding local feature. This may give interesting results in some cases. Yet, the individual process of each part ignores the spatial dependencies between different parts, which is useful to learn discriminative features for person Re-ID. Besides, such a partition exhibits a potential misalignment between different images of the same identity due to pose variations, inaccurate pedestrian detection and/or occlusion. Some examples are given in Fig. 1. The upper body of the left image in Fig. 1(a) is aligned with the head part of the right image.

We notice that pedestrian in images can be decomposed into a sequence of body parts from head to foot. Even though each part does not always lie on the same position in different images, all the pedestrian parts can be modeled in a sequence way thanks to the priori knowledge of body structure. The sequential representation of pedestrian naturally motivates us to resort to RNN with LSTM cells. Indeed, RNN with LSTM cells has been widely used in many sequence-based problem, such as image caption [38], machine translation [30], speech recognition [12], etc. LSTM has also shown a high potential in image classification [35] and object detection [2], where LSTM models the spatial dependencies and captures richer contextual information. Benefiting from the internal gating mechanisms, LSTM can control the information flow from the current state to the next state. Consequently, LSTM cell is competent to propagate certain relevant contexts and filter out some irrelevant parts. Based on the above insights, we propose to adopt LSTM to model the sequence of body parts for person Re-ID.

More specifically, to get the spatial contexts, we extract directly a sequence of feature vectors from the shared low-level features f_b without any explicit segmentation. As depicted in Fig. 2, each row of f_b undergoes an average pooling (*i.e.*, average pooling with $1 \times W$ kernel), which results in a corresponding feature vector of length C . As illustrated in Fig. 3, each feature vector describes a rectangle region in raw image given by the corresponding receptive field. A two-layer Bidirectional LSTM is then built upon the feature sequence. Thanks to the context modeling of LSTM, each resulting feature vector of length U may describe better its associated part. Finally, all the resulting feature vectors characterizing the underlying local parts are concatenated together as the final part-based person representation. This part-based feature is learned via a Softmax layer with N_c output neurons, where N_c denotes the number of pedestrian

identities.

3.3. Global Representation Learning

Part-based features focus more on discriminative pedestrian details. It is difficult to distinguish two different identities with very similar visual details (*e.g.*, wearing the same clothes) using only part-based representation. In this case, the pose and/or shape information is required to distinguish them. Indeed, the global feature is complementary to the part-based local feature, and contains more high-level semantics, such as pose and shape. Similar to many deep neural networks for person Re-ID, we straightforwardly extract a global representation by inserting a global average pooling and a fully-connected layer after the shared low-level feature f_b . A Softmax layer with N_c output neurons is then appended for global feature learning.

3.4. Deep Metric Learning with Triplet Loss

The part-based and global representation learning described in Sec. 3.2 and 3.3 do not explicitly learn a similarity measure which is required for person Re-ID during test. We propose a third branch in Deep-Person model that is responsible for distance ranking. For that, we apply another independent global average pooling to the shared low-level feature f_b , which results in a feature f_m in the metric space for similarity estimation. This feature f_m is also adopted as the final pedestrian descriptor for person Re-ID.

More specifically, we adopt the improved triplet loss for deep metric learning proposed in [14]. The main idea lies in forming batches named *PK sampling* by randomly sampling P classes (person identities), and then randomly sampling K images of each class (person), thus resulting in a batch of PK images. Given such a mini-batch $X = \{x_i\}_{i=1}^{PK}$, we can constitute a triplet $T_i = \{x_i^o, x_i^+, x_i^-\}$ for each selected anchor image x_i , where x_i^+ (*resp.* x_i^-) denotes a positive (*resp.* negative) sample in the mini-batch having the same (*resp.* different) person ID as x_i^o . Since hard triplet mining strategy is crucial for learning based on triplet loss, only the hardest positive and negative sample in the mini-batch are selected for each anchor sample to form the triplets for loss computation:

$$\mathcal{L}_{trp} = \frac{1}{PK} \sum_{i=1}^{PK} [m + \max_+ D(f_m(x_i^o), f_m(x_i^+)) - \min_- D(f_m(x_i^o), f_m(x_i^-))]_+, \quad (1)$$

where $[\cdot]_+ = \max(\cdot, 0)$, m is a margin that is enforced between positive and negative pairs, and $D(a, b)$ is referred as the distance function between feature vector a and b . In this paper, we use the Euclidean distance $D(f_m(x_i), f_m(x_j)) = \|f_m(x_i) - f_m(x_j)\|_2$ as the distance metric. This objective function encourages the features of positive pairs to be

closer in the learned feature space than the negative pairs by a predefined margin m .

3.5. Model Training

As discussed in Sec. 3.1, different branches have complementary strengths for learning discriminative pedestrian descriptors. To leverage these complementarities, we jointly train the whole network to predict person identity for both part-based and global feature learning while also satisfying the triplet objective. There are two identification subnets built with Softmax loss for multi-class person identification task, defined as:

$$\mathcal{L}_{cls} = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(W_{y_i}^T f_i + b_{y_i})}{\sum_{j=1}^{N_c} \exp(W_j^T f_i + b_j)}, \quad (2)$$

where f_i is the classification feature of i -th sample, y_i is the identity of i -th sample, N is the number of samples, W_j and b_j is respectively the weight and bias of the classifier for j -th identity. To simplify the notation and without any ambiguity, we replace f_i with f_p (*resp.* f_g) for the corresponding part-based (*resp.* global-based) classification loss function \mathcal{L}_{cls-p} (*resp.* \mathcal{L}_{cls-g}). Then the final loss is given by:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{trp} + \lambda_2 \mathcal{L}_{cls-p} + \lambda_3 \mathcal{L}_{cls-g}, \quad (3)$$

where λ_i ($i = 1, 2, 3$) denotes the loss weight for different branch.

4. Experiments

4.1. Implementation Details

The proposed Deep-Person model is built on the PyTorch framework. The backbone network is the ResNet-50 [13] model pre-trained on ImageNet, where the global average pooling and fully connected layers are discarded. The parameter C , number of channels of f_b (see Fig. 2) is set to 2048, and U , the number of hidden units in BLSTM is set to 256.

We follow common dataset augmentation strategies with different scales and aspect ratios to train the Deep-Person model. Concretely, all training images are first resized to 256×128 . Then, we randomly crop each resized image with scale in the interval $[0.64, 1.0]$ and aspect ratio in $[2, 3]$. Finally, the cropped patch is resized again to 256×128 . Randomly horizontal flip with a probability of 0.5 is also applied. During testing phase, images are simply resized to 256×128 . The globally pooled feature f_m , input of the distance ranking branch is used as the pedestrian descriptor for retrieval. Before feeding the input image to the network, we follow the image normalization procedure in original ResNet-50 model trained on ImageNet by subtracting the mean value and then dividing by the standard deviation.

As described in 3.4, *PKsampling* is used to form a batch, where P and K denote respectively the number of sampled classes and instances of each class. To train our Deep-Person model, we redefine the notion of “epoch” such that every $\lfloor N_c/P \rfloor$ trained batches form an “epoch”. This re-definition implies that an “epoch” covers approximately all the identities. Jointly training such a network is not trivial in the beginning of training stages. We apply gradient clipping to avoid gradient explosion. The Adam with the default hyper-parameters ($\epsilon = 10^{-8}$, $\beta_1 = 0.9$, $\beta_2 = 0.999$) is used to minimize the loss function \mathcal{L} given in Eq. (3), where λ_1 , λ_2 , and λ_3 are all set to 1. The margin m in Eq. (1) is set to 0.5. Following the common practice of learning rate decaying schedule in [14], we set the initial learning rate to 3×10^{-4} and apply the decay schedule at epoch 100. The total number of training epochs for all conducted experiments are set to 150. The mini-batch size PK is set to 128 for all experiments. Following different average number of samples for each identity in different datasets detailed in Sec. 4.2, K is set respectively to 8, 4, and 8 for Market-1501, CHUK03, and DukeMTMC-reID.

4.2. Datasets and Evaluation Protocol

Datasets. We evaluate our proposed method, Deep-Person, on three widely used large-scale datasets: Market-1501 [43], CUHK03 [18], and DukeMTMC-reID [47, 25] dataset. A brief description of them is given as follows:

Market-1501: This dataset consists of 32,668 images of 1,501 identities captured by six cameras in front of a super-market. The provided pedestrian bounding boxes are given by Deformable Part Model (DPM) [10] and further validated by manual annotations. The dataset is divided into training set consisting of 12,936 images of 751 identities and testing dataset containing the other 19,732 images of 750 identities. For each identity in the testing set, only one image from each camera is selected as a query, resulting in 3,368 query images in total.

CUHK03: This dataset contains 14,097 images of 1,467 identities shot by six cameras in the CUHK campus. It provides two types of annotations: manually labeled pedestrian bounding boxes and automatic detections given by DPM detector. We conduct experiments on both types of annotated datasets named as *labeled* dataset and *detected* dataset, respectively. This dataset also offers 20 random splits, each of which selects 100 test identities and use the other 1,367 persons for training. The average performance on all splits is reported for evaluation on this dataset.

DukeMTMC-reID: This dataset is a subset of DukeMTMC [25] for image-based re-identification, in the format of the Market-1501 dataset. It is composed of 36,411 images of 1,812 different identities taken by 8 high-resolution cameras, where 1,404 identities appear in more than two cameras and the other 408 identities are regarded as distrac-

| Methods | Single Query | | Multiple Query | |
|---------------------------|--------------|--------------|----------------|--------------|
| | rank1 | mAP | rank1 | mAP |
| OL-MANS [49] | 60.67 | - | 66.80 | - |
| DNS [40] | 61.02 | 35.68 | 71.56 | 46.03 |
| Gated S-CNN [32] | 65.88 | 39.55 | 76.04 | 48.45 |
| CRAFT-MFA [7] | 68.70 | 42.30 | 77.70 | 50.30 |
| P2S [50] | 70.72 | 44.27 | 85.78 | 55.73 |
| CADL [20] | 73.84 | 47.11 | 80.85 | 55.58 |
| Spindle [41] | 76.90 | - | - | - |
| MSCAN [17] | 80.31 | 57.53 | 86.79 | 66.70 |
| SVDNet [29] | 82.30 | 62.10 | - | - |
| Part-Aligned [42] | 81.00 | 63.40 | - | - |
| PDC [28] | 84.14 | 63.41 | - | - |
| JLML [36] | 85.10 | 65.50 | 89.70 | 74.50 |
| LSRO [47] | 83.97 | 66.07 | 88.42 | 76.10 |
| SSM [1] | 82.21 | 68.80 | 88.18 | 76.18 |
| TriNet [14] | 84.92 | 69.14 | 90.53 | 76.42 |
| DPFL [6] | 88.60 | 72.60 | 92.20 | 80.40 |
| Deep-Person (Ours) | 92.31 | 79.58 | 94.48 | 85.09 |

Table 1. Comparison with state-of-the-art results on Market-1501. The 1st/2nd best result is highlighted in red/blue.

tor IDs. Among the 1404 identities, 16,522 images of 702 identities are used for training, the other 702 identities are divided into 2,228 query images and 17,661 gallery images.

Evaluation Protocol. We follow the standard evaluation protocol. Concretely, we adopt the cumulative matching characteristics (CMC) at rank-1 and mean average precision (mAP) for performance evaluation on Market-1501 and DukeMTMC-reID. The precisions of rank-1, rank-5, and rank-10 are reported for CUHK03. Following most related works, the evaluation on CHUK03 and DukeMTMC-reID is performed under single query setting. Both single and multiple query settings are used for Market-1501 dataset.

4.3. Comparison with Related Methods

We compare the proposed method, Deep-Person, with the state-of-the-art approaches on Market-1501, CUHK03, and DukeMTMC-reID datasets. The proposed Deep-Person consistently outperforms the state-of-the-art methods on all three datasets. The details are given as follows:

Evaluation on Market-1501. As depicted in Table 1, the proposed Deep-Person achieves compelling results on Market-1501 dataset. Concretely, our Deep-Person significantly improves the state-of-the-art results by about 7% on mAP and 3.7% on rank-1 matching rate under single query mode. When combining with an effective re-ranking approach [48], the performance is further boosted, reaching **90.84%** mAP. To the best of our knowledge, this is the first time that an mAP higher than 90% is achieved on the Market-1501 dataset with single query setting. We also observe similar performance improvements using multiple query setting on this dataset, getting 4.7% and 2.3% improvement on mAP and rank-1, respectively.

| Methods | Labeled | | | Detected | | |
|---------------------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | r1 | r5 | r10 | r1 | r5 | r10 |
| EDM[27] | 61.3 | 88.9 | 96.4 | 52.1 | 82.9 | 91.8 |
| OL-MANS [49] | 61.7 | 88.4 | 95.2 | 62.7 | 87.6 | 93.8 |
| DNS [40] | 62.5 | 90.0 | 94.8 | 54.7 | 84.7 | 94.8 |
| Gated S-CNN [32] | - | - | - | 68.1 | 88.1 | 94.6 |
| GOG [23] | 67.3 | 91.0 | 96.0 | 65.5 | 88.4 | 93.7 |
| DictRW [9] | 71.1 | 91.7 | 94.7 | - | - | - |
| MSCAN [17] | 74.2 | 94.3 | 97.5 | 68.0 | 91.0 | 95.4 |
| MTDnet [5] | 74.7 | 96.0 | 97.5 | - | - | - |
| Quadruplet [4] | 75.5 | 95.2 | 99.2 | - | - | - |
| SSM [1] | 76.6 | 94.6 | 98.0 | 72.7 | 92.4 | 96.1 |
| MuDeep [24] | 76.9 | 96.1 | 98.4 | 75.6 | 94.4 | 97.5 |
| SVDNet [29] | - | - | - | 81.8 | 95.2 | 97.2 |
| CRAFT-MFA [7] | - | - | - | 84.3 | 97.1 | 98.3 |
| LSRO [47] | - | - | - | 84.6 | 97.6 | 98.9 |
| JLML [36] | 83.2 | 98.0 | 99.4 | 80.6 | 96.9 | 98.7 |
| Part-Aligned [42] | 85.4 | 97.6 | 99.4 | 81.6 | 97.3 | 98.4 |
| Spindle [41] | 88.5 | 97.8 | 98.6 | - | - | - |
| PDC [28] | 88.7 | 98.6 | 99.2 | 78.3 | 94.8 | 97.2 |
| Deep-Person (Ours) | 91.5 | 99.0 | 99.5 | 89.4 | 98.2 | 99.1 |

Table 2. Evaluation on CUHK03 in terms of rank-1 (r1), rank-5 (r5), and rank-10 (r10) matching rate, using manually labeled pedestrian bounding boxes and automatic detections by DPM.

Evaluation on CUHK03. The evaluation of the proposed Deep-Person on CUHK03 dataset in terms of rank-1, rank-5, and rank-10 matching rate is given in Table 2. Using the manually annotated pedestrian bounding boxes, our Deep-Person yields 2.8% rank-1 accuracy improvement. An improvement of 4.8% rank-1 accuracy is achieved when using an automatic method DPM to extract pedestrians. The later setting is coherent with practical application, which demonstrates the potential and robustness of Deep-Person in practice. Deep-Person also outperforms the state-of-the-art methods under both rank-5 and rank-10 matching rate.

Evaluation on DukeMTMC-reID. The comparison with state-of-the-art methods on DukeMTMC-reID dataset is depicted in Table 3. Our Deep-Person also outperforms all the state-of-the-art approaches, achieving an improvement of 1.7% rank-1 accuracy and 4.2% mAP. It is worth noting that the previous state-of-the-art DPFL [6] take multi-scale person images as input. With this design, the performance is expected to be further improved.

4.4. Ablation Study

We further evaluate several variants of Deep-Person to verify the effectiveness of each individual component. Without loss of generality, the ablation study is performed on Market-1501 dataset under single query setting, using the same settings as described in Sec. 4.1.

Effectiveness of LSTM-based Parts. The proposed Deep-Person leverages the complementary information be-

| Methods | rank1 | mAP |
|---------------------------|--------------|--------------|
| BoW+KISSME [43] | 25.13 | 12.17 |
| LOMO+XQDA [19] | 30.75 | 17.04 |
| LSRO [47] | 67.68 | 47.13 |
| AttIDNet [21] | 70.69 | 51.88 |
| PAN [46] | 71.59 | 51.51 |
| ACRN [26] | 72.58 | 51.96 |
| SVDNet [29] | 76.70 | 56.80 |
| DPFL [6] | 79.20 | 60.60 |
| Deep-Person (Ours) | 80.90 | 64.80 |

Table 3. Quantitative comparison with state-of-the-art methods on DukeMTMC-reID dataset.

| Feature type | Single Query | |
|--------------------------|--------------|--------------|
| | rank1 | mAP |
| B_g^I | 85.33 | 64.71 |
| B_p^I w/o LSTM | 82.39 | 60.75 |
| $B_g^I + B_p^I$ w/o LSTM | 86.49 | 66.74 |
| $B_g^I + B_p^I$ | 87.74 | 69.82 |

Table 4. Effectiveness of the complementary advantage between global features and the novel LSTM-based part features.

| Category | Methods | Single Query | |
|--------------------|-------------------|--------------|--------------|
| | | rank1 | mAP |
| Rigid Partition | SCSP [3] | 51.90 | 26.35 |
| | LSTM-Siamese [33] | 61.60 | 35.31 |
| | MR B-CNN [31] | 66.36 | 41.17 |
| | JLML [36] | 85.10 | 65.50 |
| Flexible Partition | Spindle [41] | 76.90 | - |
| | MSCAN [17] | 80.31 | 57.53 |
| | Part-Aligned [42] | 81.00 | 63.40 |
| | PDC [28] | 84.14 | 63.41 |
| | Part-Loss [39] | 85.60 | 65.91 |
| | $B_g^I + B_p^I$ | 87.74 | 69.82 |

Table 5. Comparison with state-of-the-art part-based models on Market-1501 dataset.

tween global representation and a novel LSTM-based part representation. We evaluate the contribution of this complementary advantage and the effect of the novel LSTM-based part representation. For that, we discard the ranking branch B_t^R in Deep-Person. As depicted in Table 4, the combination of global and part-based branches outperforms each individual branch alone. Furthermore, the adopted LSTM achieves 3.1% performance gain in mAP. To get some insights about the improvements, we analyze the learned feature maps f_b of three variants of Deep-Person: B_g^I branch alone, $B_g^I + B_p^I$ without LSTM, and $B_g^I + B_p^I$ with LSTM. Some illustrations are given in Fig. 4. Using global branch B_g^I alone only captures coarse areas with few part details. Combining with part-based representation without LSTM enriches some fine details. Adopting LSTM in the combi-




Figure 4. Visualization of feature maps extracted from three variants of Deep-Person. From left to right in (a-d): raw image, the feature map of using B_g^I branch alone, $B_g^I + B_p^I$ without LSTM, and $B_g^I + B_p^I$ with LSTM, respectively.

| Loss type | Single Query | |
|-------------------------------|--------------|--------------|
| | rank1 | mAP |
| Identification loss | 85.33 | 64.71 |
| Triplet loss | 81.29 | 63.51 |
| Identification + Triplet loss | 88.45 | 72.96 |

Table 6. Effectiveness of multi-loss on the variant of Deep-Person discarding part-based branch B_p^I .

| Feature type | Single Query | |
|--------------|--------------|--------------|
| | rank1 | mAP |
| f_c | 84.65 | 67.00 |
| f_m | 87.74 | 69.82 |

Table 7. Experimental results using different fused features as the final pedestrian descriptor on the $B_g^I + B_p^I$ two-branch model.

nation of global and part-based branches focuses more on the pedestrian with fine details, ignoring irrelevant regions (e.g., the top corners in Fig. 4 (d) and the motorbike in Fig. 4 (b)). As a result, the learned features of $B_g^I + B_p^I$ using LSTM are more robust for person Re-ID.

We also compare a variant of Deep-Person using global and LSTM-based part branches with some state-of-the-art part-based models. The comparison is depicted in Table 5. Thanks to the use of LSTM modeling spatial context information, an improvement of 3.9% mAP and 2.1% rank1 are achieved compared to the best part-based method. It is worth noting that Spindle [41] and PDC [28] use extra pose annotations for body-part detection. As compared to the method [33] using also LSTM to model spatial relationship between body parts, using LSTM in an end-to-end way in Deep-Person achieves a significant improvement.

Effectiveness of Multi-loss. Using the complementary advantage of identification and ranking loss is another important aspect of our Deep-Person. To simplify the evaluation of this complementary advantage, we discard the part-based branch B_p^I . As shown in Table 6, combining iden-

tification loss and ranking loss (e.g., triplet loss) performs much better than using each of them alone. This reveals that both the identification and ranking information are important to learn discriminative features for person Re-ID.

Choice of final pedestrian descriptor. As described in Sec. 3, Deep-Person uses the global average pooled feature f_m of the backbone feature f_b as the final pedestrian descriptor. Complementary information between global and part-based representation is implicitly integrated into f_m thanks to the combination of back-propagation error differentials of the global and part-based branch during training. Following recent works [8, 17, 28, 36, 37] using global and part-based branches, the concatenation of the penultimate fully connected layers of these two branches denoted as f_c seems to be a reasonable alternative. We use a two-branch $B_p^I + B_g^I$ variant of Deep-Person to evaluate these two choices. As shown in Table 7, the implicit fused feature f_m significantly outperforms the directly concatenated feature f_c . One possible reason for this is that the identification loss makes the features near the classification layer focus more on the difference of training identities. Such feature might be discriminative for identities in training images, but is not discriminative for the unseen identities during test. Whereas, the backbone feature may be more robust and generalize better to the unseen test categories. This comparison motivates us to append the ranking branch with triplet loss after f_m , which is considered as the final pedestrian descriptor in Deep-Person.

5. Conclusion

In this paper, we introduce a novel three-branch framework named Deep-Person to learn highly discriminative deep features for person Re-ID. Different from most existing methods which either focus on feature representation or feature learning alone, complementary advantages on both aspects are considered in Deep-Person. Concretely, local body-part and global full-body features are jointly employed. The identification loss and ranking loss are applied to simultaneously learn an ID-discriminative embedding and a similarity measurement. Furthermore, in contrast to the existing part-based methods which usually discard the spatial context of body structure, we use LSTM to enhance the discriminative ability of part representation with contextual information, mitigating the misalignment with the sequence-level pedestrian representation. Extensive evaluations on three popular and challenging datasets demonstrate the superiority of the proposed Deep-Person over the state-of-the-art methods. In the future, we would like to combine the spatial context modeling with attention mechanism to automatically select more discriminative parts for person Re-ID.

References

- [1] S. Bai, X. Bai, and Q. Tian. Scalable person re-identification on supervised smoothed manifold. In *CVPR*, 2017. 6, 7
- [2] S. Bell, C. Lawrence Zitnick, K. Bala, and R. Girshick. Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks. In *CVPR*, 2016. 4
- [3] D. Chen, Z. Yuan, B. Chen, and N. Zheng. Similarity learning with spatial constraints for person re-identification. In *CVPR*, 2016. 2, 7
- [4] W. Chen, X. Chen, J. Zhang, and K. Huang. Beyond triplet loss: A deep quadruplet network for person re-identification. In *CVPR*, 2017. 7
- [5] W. Chen, X. Chen, J. Zhang, and K. Huang. A multi-task deep network for person re-identification. In *AAAI*, 2017. 1, 2, 3, 7
- [6] Y. Chen, X. Zhu, and S. Gong. Person re-identification by deep learning multi-scale representations. In *ICCV*, 2017. 6, 7
- [7] Y.-C. Chen, X. Zhu, W.-S. Zheng, and J.-H. Lai. Person re-identification by camera correlation aware feature augmentation. *IEEE TPAMI*, PP(99):1–1, 2017. 6, 7
- [8] D. Cheng, Y. Gong, S. Zhou, J. Wang, and N. Zheng. Person re-identification by multi-channel parts-based cnn with improved triplet loss function. In *CVPR*, 2016. 1, 2, 3, 4, 8
- [9] L. L. A. G. H. Y. G. N. Z. De Cheng, Xiaojun Chang. Discriminative dictionary learning with ranking metric embedded for person re-identification. In *IJCAI*, 2017. 7
- [10] P. Felzenszwalb, D. McAllester, and D. Ramanan. A discriminatively trained, multiscale, deformable part model. In *CVPR*, 2008. 6
- [11] M. Geng, Y. Wang, T. Xiang, and Y. Tian. Deep transfer learning for person re-identification. *arXiv:1611.05244*, 2016. 3
- [12] A. Graves, A.-r. Mohamed, and G. Hinton. Speech recognition with deep recurrent neural networks. In *ICASSP*, 2013. 4
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [14] A. Hermans, L. Beyer, and B. Leibe. In defense of the triplet loss for person re-identification. *arXiv:1703.07737*, 2017. 3, 5, 6
- [15] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 2
- [16] M. Jaderberg, K. Simonyan, A. Zisserman, et al. Spatial transformer networks. In *NIPS*, 2015. 2
- [17] D. Li, X. Chen, Z. Zhang, and K. Huang. Learning deep context-aware features over body and latent parts for person re-identification. In *CVPR*, 2017. 1, 2, 3, 6, 7, 8
- [18] W. Li, R. Zhao, T. Xiao, and X. Wang. Deepreid: Deep filter pairing neural network for person re-identification. In *CVPR*, 2014. 1, 2, 3, 6
- [19] S. Liao, Y. Hu, X. Zhu, and S. Z. Li. Person re-identification by local maximal occurrence representation and metric learning. In *CVPR*, 2015. 7
- [20] J. Lin, L. Ren, J. Lu, J. Feng, and J. Zhou. Consistent-aware deep learning for person re-identification in a camera network. In *CVPR*, 2017. 6
- [21] Y. Lin, L. Zheng, Z. Zheng, Y. Wu, and Y. Yang. Improving person re-identification by attribute and identity learning. *arXiv:1703.07220*, 2017. 7
- [22] H. Liu, J. Feng, M. Qi, J. Jiang, and S. Yan. End-to-end comparative attention networks for person re-identification. *IEEE TIP*, 2017. 2, 3, 4
- [23] T. Matsukawa, T. Okabe, E. Suzuki, and Y. Sato. Hierarchical gaussian descriptor for person re-identification. In *CVPR*, 2016. 7
- [24] X. Qian, Y. Fu, Y.-G. Jiang, T. Xiang, and X. Xue. Multi-scale deep learning architectures for person re-identification. In *ICCV*, 2017. 3, 7
- [25] E. Ristani, F. Solera, R. Zou, R. Cucchiara, and C. Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *ECCV*, 2016. 6
- [26] A. Schumann and R. Stiefelhagen. Person re-identification by deep learning attribute-complementary information. In *CVPR Workshops*, 2017. 7
- [27] H. Shi, Y. Yang, X. Zhu, S. Liao, Z. Lei, W. Zheng, and S. Z. Li. Embedding deep metric for person re-identification: A study against large variations. In *ECCV*, 2016. 1, 2, 4, 7
- [28] C. Su, J. Li, S. Zhang, J. Xing, W. Gao, and Q. Tian. Pose-driven deep convolutional model for person re-identification. In *ICCV*, 2017. 1, 2, 6, 7, 8
- [29] Y. Sun, L. Zheng, W. Deng, and S. Wang. Svdnet for pedestrian retrieval. In *ICCV*, 2017. 6, 7
- [30] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *NIPS*, 2014. 4
- [31] E. Ustinova, Y. Ganin, and V. Lempitsky. Multi-region bilinear convolutional neural networks for person re-identification. In *AVSS*, 2017. 1, 2, 7
- [32] R. R. Variator, M. Haloi, and G. Wang. Gated siamese convolutional neural network architecture for human re-identification. In *ECCV*, 2016. 6, 7
- [33] R. R. Variator, B. Shuai, J. Lu, D. Xu, and G. Wang. A siamese long short-term memory architecture for human re-identification. In *ECCV*, 2016. 1, 2, 3, 4, 7, 8
- [34] F. Wang, W. Zuo, L. Lin, D. Zhang, and L. Zhang. Joint learning of single-image and cross-image representations for person re-identification. In *CVPR*, 2016. 2, 3
- [35] J. Wang, Y. Yang, J. Mao, Z. Huang, C. Huang, and W. Xu. Cnn-rnn: A unified framework for multi-label image classification. In *CVPR*, 2016. 4
- [36] S. G. Wei Li, Xiatian Zhu. Person re-identification by deep joint learning of multi-loss classification. In *IJCAI*, 2017. 1, 2, 6, 7, 8
- [37] S. Wu, Y.-C. Chen, X. Li, A.-C. Wu, J.-J. You, and W.-S. Zheng. An enhanced deep feature representation for person re-identification. In *WACV*, 2016. 1, 2, 8
- [38] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015. 4
- [39] H. Yao, S. Zhang, Y. Zhang, J. Li, and Q. Tian. Deep representation learning with part loss for person re-identification. *arXiv:1707.00798*, 2017. 1, 2, 7

- [40] L. Zhang, T. Xiang, and S. Gong. Learning a discriminative null space for person re-identification. In *CVPR*, 2016. 6, 7
- [41] H. Zhao, M. Tian, S. Sun, J. Shao, J. Yan, S. Yi, X. Wang, and X. Tang. Spindle net: Person re-identification with human body region guided feature decomposition and fusion. In *CVPR*, 2017. 1, 2, 6, 7, 8
- [42] L. Zhao, X. Li, Y. Zhuang, and J. Wang. Deeply-learned part-aligned representations for person re-identification. In *ICCV*, 2017. 6, 7
- [43] L. Zheng, L. Shen, L. Tian, S. Wang, J. Wang, and Q. Tian. Scalable person re-identification: A benchmark. In *ICCV*, 2015. 1, 6, 7
- [44] L. Zheng, Y. Yang, and A. G. Hauptmann. Person re-identification: Past, present and future. *arXiv:1610.02984*, 2016. 2, 3
- [45] L. Zheng, H. Zhang, S. Sun, M. Chandraker, Y. Yang, and Q. Tian. Person re-identification in the wild. In *CVPR*, 2017. 2
- [46] Z. Zheng, L. Zheng, and Y. Yang. Pedestrian alignment network for large-scale person re-identification. *arXiv:1707.00408*, 2017. 7
- [47] Z. Zheng, L. Zheng, and Y. Yang. Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In *ICCV*, 2017. 1, 6, 7
- [48] Z. Zhong, L. Zheng, D. Cao, and S. Li. Re-ranking person re-identification with k-reciprocal encoding. In *CVPR*, 2017. 6
- [49] J. Zhou, P. Yu, W. Tang, and Y. Wu. Efficient online local metric adaptation via negative samples for person re-identification. In *CVPR*, 2017. 6, 7
- [50] S. Zhou, J. Wang, J. Wang, Y. Gong, and N. Zheng. Point to set similarity based deep feature learning for person re-identification. In *CVPR*, 2017. 6