

Person Re-identification via Recurrent Feature Aggregation

Yichao Yan^{1(✉)}, Bingbing Ni¹, Zhichao Song¹, Chao Ma¹, Yan Yan²,
and Xiaokang Yang¹

¹ Shanghai Jiao Tong University, Shanghai, China

{yanyichao, nibingbing, 5110309394, chaoma, xkyang}@sjtu.edu.cn

² University of Michigan, Ann Arbor, USA

tom.yan.555@gmail.com

Abstract. We address the person re-identification problem by effectively exploiting a globally discriminative feature representation from a sequence of tracked human regions/patches. This is in contrast to previous person re-id works, which rely on either single frame based person to person patch matching, or graph based sequence to sequence matching. We show that a progressive/sequential fusion framework based on long short term memory (LSTM) network aggregates the frame-wise human region representation at each time stamp and yields a sequence level human feature representation. Since LSTM nodes can *remember and propagate* previously accumulated good features and *forget* newly input inferior ones, even with simple hand-crafted features, the proposed recurrent feature aggregation network (**RFA-Net**) is effective in generating highly discriminative sequence level human representations. Extensive experimental results on two person re-identification benchmarks demonstrate that the proposed method performs favorably against state-of-the-art person re-identification methods.

Keywords: Person re-identification · Feature fusion · Long short term memory networks

1 Introduction

Person re-identification (re-id) deals with the problem of re-associating a specific person across non-overlapping cameras. It has been receiving increasing popularity [1] due to its important applications in intelligent video surveillance.

Existing methods mainly focus on addressing the single-shot person re-id problem. Given a probe image of one person taken from one camera, a typical scenario for single-shot person re-id is to identify this person in a set of gallery images taken from another camera. Usually, the identification results are based on ranking the similarities of the probe-gallery pairs. The performance of person re-id is measured by the rank- k matching rate if the correct pair hits the retrieved top- k ranking list. To increase the matching rate, state-of-the-art approaches either employ discriminative features in representing persons

or apply distance metric learning methods to increase the similarity between matched image pairs. Numerous types of features have been explored to represent persons, including global features like color and texture histograms [2,3], local features such as SIFT [4] and LBP [5], and deep convolutional neural network (CNN) features [6,7]. In the meantime, a large number of metric learning approaches have been applied to person re-id task, such as LMNN [8], Mahalanobis distance metric [9], and RankSVM [10]. Despite the significant progress in recent years, the performance achieved by these methods do not fulfill the real-application requirement due to the following reasons. First, images captured by different cameras undergo large amount of appearance variations caused by illumination changes, heavy occlusion, background clutter, or human pose changes. More importantly, for real surveillance, persons always appear in a video rather than in a single-shot image. These single-shot based methods fail to make full use of the temporal sequence information in surveillance videos.

Several algorithms have been proposed to tackle the multi-shot person re-id problem, *i.e.*, to match human instances in the sequence level (these human patch sequences are usually obtained by visual detection and tracking). To exploit the richer sequence information, existing methods mainly resort to: (1) key frame/fragment representation [11]; (2) feature fusion/encoding [12–14] and (3) spatio-temporal appearance model [15]. Despite their favorable performance on recent benchmark datasets [11,16], we have the following observations on their limitations. First, key frame/fragment representation based algorithms [11] often assume the discriminative fragments to be located in the local minima and maxima of Flow Energy Profile [11], which may not be accurate. During the final matching, since only one fragment is selected to represent the whole sequence, richer information contained by the rest of the sequences is not fully utilized. Second, feature fusion/encoding methods [12–14] take bag-of-words approach to encode a set of frame-wise features into a global vector, but ignore the informative spatio-temporal information of human sequence. To overcome these shortages, the recently proposed method [15] employs the spatially and temporally aligned appearance of the person in a walking cycle for matching. However, such approach is extremely computationally inefficient and thus inappropriate for real applications. It is therefore of great importance to explore a more effective and efficient scheme to make full use of the richer sequence information for person re-id.

The fundamental challenge of multi-shot person re-id is how to systematically aggregate both the frame-wise appearance information as well as temporal dynamics information along the human sequence to generate a more discriminative sequence level human representation. To this end, we propose a recurrent feature aggregation network (**RFA-Net**) that builds a sequence level representation from a temporally ordered sequence of frame-wise features, based on a typical version of recurrent neural network, namely, Long Short-Term Memory network [17]. Figure 1 shows an overview of the difference between our method and previous approaches. The proposed feature aggregation framework possesses various advantages. First, it allows discriminative information of frame-wise data

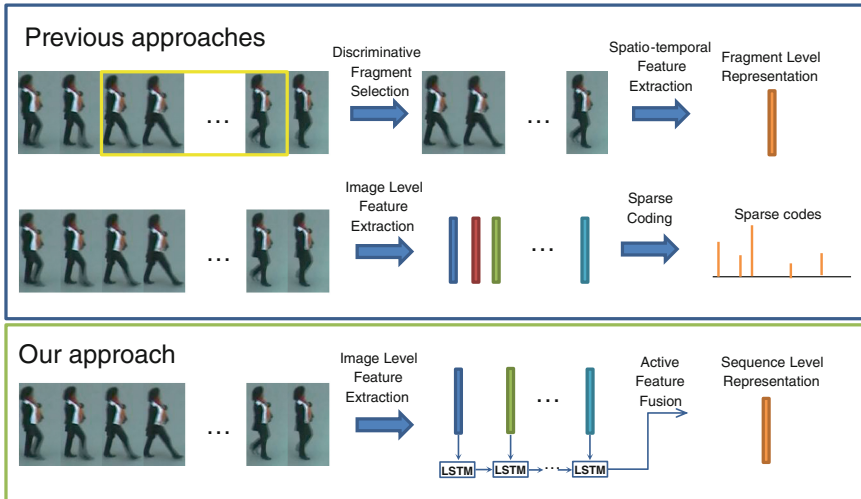


Fig. 1. Pipeline of the proposed recurrent feature aggregation network for multi-shot person re-id with comparison to previous methods. The top part depicts two common used methods for multi-shot person re-id, *i.e.*, to choose a most discriminative fragment for matching, or to use sparse coding method to enhance feature representation. The bottom part gives an illustration of our approach, frame-wise features are input into a feature aggregation (LSTM) network, and the output vector of the network builds the sequence level representation

to propagate along the temporal direction, and discriminative information could be accumulated from the first LSTM node to the deepest one, thus yielding a highly discriminative sequence level human representation. Second, during feature propagation, this framework can *prevent* non-informative information from reaching the deep nodes, therefore it is robust to noisy features (due to occlusion, tracking/detection failure, or background clutter etc.). Third, the proposed fusion network is simple yet efficient, which is able to deal with sequences with variable length.

The main contributions of this work lie in that we present a recurrent feature aggregation network to address the multi-shot person re-id problem. The proposed network jointly takes the feature fusion and spatio-temporal appearance model into account and makes full use of the temporal sequence information for multi-shot person re-id. With the use of the proposed network, hand-crafted low-level features can be augmented with temporal cues and significantly improve the accuracy of person re-id. Extensive experiments on two publicly available benchmark datasets demonstrate that the proposed person re-id method performs favorably against state-of-the-arts algorithms in terms of effectiveness and efficiency.

2 Related Work

Person re-identification has been explored for several years. A large body of works mainly focus on solving two subproblems: feature representation and distance metric learning.

Low-level features such as texture, color and gradient are most commonly used for person representation, however, these features are not powerful enough to discriminate a person from similar ones. Some methods address this problem by combining features together to build a more discriminative representation [18] or by selecting the most discriminative features to represent a person [2].

When multiple images are available for one person, image level features are accumulated or averaged [14] to build a more robust representation. When image sequences or videos are available, the space-time information can be used to build space-time representations. These kind of features such as 3D SIFT [19] and 3D HOG [20] are simply extensions of 2D features, which are not powerful enough as in the single-shot case. Other methods try to build more powerful descriptors, such as in [21], interest points descriptors are accumulated along the time space to capture appearance variability. Different color-based features combined with a graph based approach was proposed in [22]. This method tries to learn the global geometric structure of the people and to realize comparison of the video sequences. Many other methods also exploit image sequences to enhance feature description [23,24]. Wang *et al.* [11] presented a model to automatically select the most discriminative video fragments and simultaneously to learn a video ranking function for person re-id. Karanam *et al.* [14] proposed a Local Discriminant Analysis approach that is capable of encoding a set of frame-wise features representing a person into a global vector. Liu *et al.* recently proposed a spatio-temporal appearance representation method [15], and feature vectors that encode the spatially and temporally aligned appearance of the person in a walking cycle are extracted. These methods are usually feature-sensitive, designing complex/good features or combining several features together may achieve good performance, while only using simple features would lead to significant performance drop [11]. In contrast to these methods, the method we proposed is simple and effective, *i.e.*, it could learn to aggregate very simple features (LBP and color information) into a highly discriminative representation.

A large number of metric learning as well as ranking algorithms have been proposed to solve the person re-id problem. Mahalanobis distance metric [9] is most commonly used in person re-id. Boosting was applied to Mahalanobis distance learning by Shen *et al.* [25]. Person re-id was formulated as a learning-to-rank problem in [26]. Zheng *et al.* have also formulated person re-id as a relative distance comparison (RDC) problem [27]. Although an efficient distance metric learning algorithm can model the transition in feature space between two camera views and thus enhance the re-identification performance, it will be no longer necessary if the features possess high discriminativeness. As demonstrated by our experiments, our feature representation will still generate promising results when only a simple metric (cosine distance) is applied.

Deep learning based methods [6, 7] have also been applied to person re-id to simultaneously solve feature representation and metric learning problem. A typical structure of this approach consists of two parts. The first part is a feature learning network, usually a convolutional neural network or other kinds of networks. The feature pairs extracted from the first step are then fed into several metric learning layers to predict whether a pair of images belong to the same person. Usually the networks have simple structure and can be trained end-to-end. However, this kind of approaches use image pairs to learn feature representation and distance metric, which is unable to consider the space-time information in video sequences, this makes it more suitable for single-shot problem. Moreover, deep learning based methods are easy to overfit on small database. In contrast, our method adopts a simple LSTM network to learn space-time information, which is more efficient than the deep learning based models. We also realize two contemporary works [28, 29].

3 Sequential Feature Aggregation for Person Re-identification

3.1 Overview

As mentioned above, previous single-shot based re-id methods rely on the information extracted from a single image, which is insufficient to identify the person in a gallery of images. For multi-shot person re-id, on one hand, methods based on graph matching to calculate the similarity between sequences are computationally complex and sensitive to outliers (e.g., false human detections). On the other hand, methods based on pooling frame-level features into a sequence level representation usually cannot well model the temporal change of human appearance, e.g., human dynamics. Therefore, in this work, we aim to aggregate both image level features and human dynamics information into a discriminative sequence level representation. More specifically, our sequential feature aggregation scheme is based on a recurrent neural network (e.g., LSTMs). Frame-wise human features are sequentially fed into this network and discriminative/descriptive information is propagated and accumulated through the LSTM nodes, and yielding final sequence level human representation. Accordingly, both the frame-wise human appearance information and human dynamics information are properly encoded.

3.2 Sequential Feature Aggregation Network

Network Architecture. Motivated by the success of recurrent neural networks (RNNs) in temporal sequence analysis, we employ a LSTM network as our feature aggregation network prototype. A LSTM network maintains the information of the previous time stamp and uses it to update the current state sequentially, therefore it is a powerful tool to deal with sequential data. With a LSTM network, information could be propagated from the first LSTM node to the last one, and this information aggregation mechanism is very useful for our given

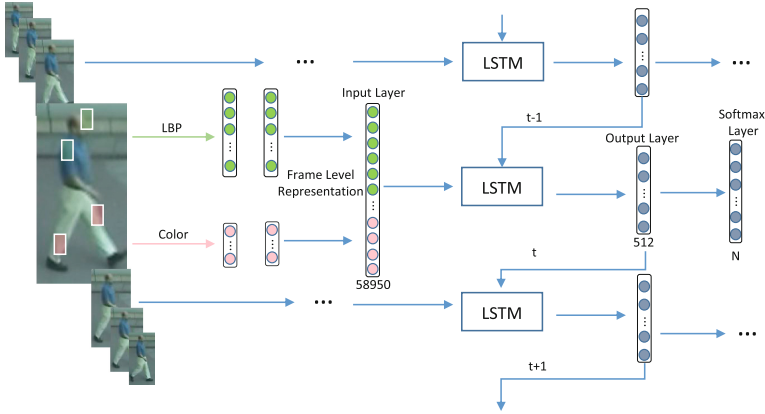


Fig. 2. Details of the network structure. LBP and color features are first extracted from overlapping rectangular image patches, and then concatenated as frame level representation. This representation as well as the previous LSTM outputs are input to the current LSTM node, the output of the LSTM is further connected to a softmax loss layer for error calculation and gradient backpropagation (Color figure online)

task, *i.e.*, to fuse the frame-wise features to a sequence level feature representation. In other words, the output of the last LSTM nodes contains aggregated discriminative information of the sequence. Each LSTM node of the feature aggregation network contains input nodes, hidden nodes and output nodes, as illustrated in Fig. 2. The input features, denoted as $\{\mathbf{x}_t\}_{t=1:T}$, could be either deep CNN features or simply hand crafted features. In this work, we use simple hand crafted features. This is because although deep CNN features are more prevailing recently, to train a valid CNN model requires large amount of data. Unfortunately, existing person re-id databases are mostly with small size, and the learned deep CNN features inevitably suffer from over-fitting. The dimensionality of the input feature vector of each time stamp is 58950. Although in many other applications, stack of LSTM nodes are used for better network representation capability, our proposed network employs only one layer of LSTM nodes (*i.e.*, no stack) with 512 hidden units, as our experimental results show that it is sufficient for our re-id task and deeper models or more hidden units do not bring too much performance gain.

Frame Level Representation. For each human bounding box, we use both color and texture features to represent human appearance features, which have proven to be effective for the task of person re-id [18, 30]. These features are naturally complementary to each other, because texture features are usually extracted from gray-scale images, while color features could capture the chromatic patterns within an image. Particularly, we use Local Binary Patterns (LBP) [5], HSV and Lab color channels to build an image frame level person representation, as have been used by many previous works [31, 32]. All images are resized to the size of

128 × 64 before feature extraction. Features are extracted from 16 × 8 rectangular patches, with a spatial overlap of 8 and 4 pixels, vertically and horizontally. For each patch, we compute the histogram of LBP codes from gray scale representation. The mean values of HSV and Lab color channels are also computed. All these patch level features are concatenated to build an image frame level representation, *i.e.*, for each rectangular patch, the dimension is 262 (256 + 6), and the dimensionality of the input feature is 58950 (262*225) for each time stamp. The output of the LSTM unit at each time stamp is a 512-dimensional vector.

Network Training. The sequence of image level feature vectors are input to a LSTM network for progressive/sequential feature fusion. The network is trained as a classification problem of N classes (N is the number of persons), *i.e.*, we aim to classify the feature vectors which belong to the same person into the same class. In particular, each LSTM node includes three gates, (*i.e.* the input gate \mathbf{i} , the output gate \mathbf{o} and the forget gate \mathbf{f}) as well as a memory cell. At each time stamp t (in our case, t indicates the order in the sequence), given the input \mathbf{x}_t and the previous hidden state \mathbf{h}_{t-1} , we update the LSTM network as follows:

$$\mathbf{i}_t = \sigma(\mathbf{W}_i \mathbf{x}_t + \mathbf{U}_i \mathbf{h}_{t-1} + \mathbf{V}_i \mathbf{c}_{t-1} + \mathbf{b}_i) \tag{1}$$

$$\mathbf{f}_t = \sigma(\mathbf{W}_f \mathbf{x}_t + \mathbf{U}_f \mathbf{h}_{t-1} + \mathbf{V}_f \mathbf{c}_{t-1} + \mathbf{b}_f) \tag{2}$$

$$\mathbf{c}_t = \mathbf{f}_t \cdot \mathbf{c}_{t-1} + \mathbf{i}_t \cdot \tanh(\mathbf{W}_c \mathbf{x}_t + \mathbf{U}_c \mathbf{h}_{t-1} + \mathbf{b}_c) \tag{3}$$

$$\mathbf{o}_t = \sigma(\mathbf{W}_o \mathbf{x}_t + \mathbf{U}_o \mathbf{h}_{t-1} + \mathbf{V}_o \mathbf{c}_t + \mathbf{b}_o) \tag{4}$$

$$\mathbf{h}_t = \mathbf{o}_t \cdot \tanh(\mathbf{c}_t) \tag{5}$$

where σ is the sigmoid function and \cdot denotes the element-wise multiplication operator. \mathbf{W}_* , \mathbf{U}_* and \mathbf{V}_* are the weight matrices, and \mathbf{b}_* are the bias vectors. The memory cell \mathbf{c}_t is a weighted sum of the previous memory cell \mathbf{c}_{t-1} and a function of the current input. The weights are the activations of forget gate and input gate respectively. Therefore, intuitively, the network can learn to propagate/accumulate the good features to the deeper nodes of the network and forget noisy features in the meantime. The output of the LSTM hidden state \mathbf{h}_t serves as the fused image level feature representation vector, which is further connected to a softmax layer. The output of the N -way softmax is the prediction of the probability distribution over N different identities.

$$y_i = \frac{\exp(y'_i)}{\sum_{k=1}^N \exp(y'_k)}, \tag{6}$$

where $y'_j = \mathbf{h}_t \cdot \mathbf{w}_j + b_j$ linearly combines the LSTM outputs. The network is learned by minimizing $-\log y_k$, where k is the index of the true label for a given input. Stochastic gradient descent is used with gradients calculated by back-propagation.

As only two sequences are available for a single person, and the length of the sequences are variable. To increase training instances and to make the model applicable for sequences of variable length, we randomly extract subsequences

of fixed length L for training. L can be determined according to the nature of the dataset, we fix $L = 10$ in our experiments, considering the tradeoff between the amount of training instances and the information contained within each subsequence.

Network Implementation Details. In the proposed framework, the LSTM networks are trained/tested using the implementation of Donahue *et al.* [33] based on *caffe* [34]. The LSTM layer contains 512 output units, and a drop-out layer is placed after it to avoid over-fitting. The weights of the LSTM layer is initialized from a uniform distribution in $[-0.01, 0.01]$. The training phase lasts for 400 epoches. As mentioned above, there is only a single LSTM layer to learn, therefore, we start with small learning rate 0.001 and after 200 epoches it decreases to 0.0001. The training procedure will terminate after another 200 epoches. We train the network on a Titan X GPU card, and a whole training duration is about 8 h for a dataset of 400 image sequences.

3.3 Person Re-identification

Once the feature aggregation network is trained, the sequence level representation can be obtained as follows. First, the testing sequence of a human patches is input to the aggregation network, then the LSTM outputs of each time stamp are concatenated to form the sequence level human feature representation for re-id purpose. The LSTM network we trained contains 10 time stamps, and from each time stamp we can extract a 512-dimensional feature vector. Therefore the total length of the sequence level representation is 5120(512×10), denoted as \mathbf{s}_i for the i -th human sequence sample. In our test phase, since the lengths of the human sequences are varying, the representations of K randomly selected subsequences (we assume the length of the sequence is larger than 10) are averaged to build the final sequence level representation. There are two benefits from this scheme. First, noisy information can be further diluted while discriminative information can be retained. Second, influences of different poses and walking circles are decreased.

After we get the person representation, we can apply a simple metric or distance metric learning methods to measure the similarity of two input instances. In this paper, we employ the RankSVM [10] method to measure the similarity as in [26]. In the meantime, a non supervised method (cosine distance metric) is also applied as a baseline metric. Experiments show that this simple combination works well in our re-id task. Consider the sequence level features $\{(\mathbf{s}_i^a, \mathbf{s}_i^b)\}_{i=1}^N$, where \mathbf{s}_i^a and \mathbf{s}_i^b denote the sequence level features of the i -th person from camera a and b . Suppose images from camera a are probe images and images from camera b are galleries. The cosine distance of a probe-gallery pair is represented as

$$d_{ij} = \frac{\mathbf{s}_i^a \cdot \mathbf{s}_j^b}{\|\mathbf{s}_i^a\| \|\mathbf{s}_j^b\|}, \quad (7)$$

where the $\|\cdot\|$ denotes the L_2 norm of the vector. Note that we do not use any complex metric learning technique. For a probe \mathbf{s}_i^a , the distances of all the pairs $\{d_{ij}\}_{i=1}^N$ are ranked in ascending order. If the distance of the true probe-gallery pair is ranked within the k top matched samples, it is counted in the top- k matching rate.

To train the RankSVM model, probe-gallery pairs with ground-truth labels are used for training. There is a single positive instance for a person, denoted as $\mathbf{s}_i^+ = |\mathbf{s}_i^a - \mathbf{s}_i^b|$. The negative instances consist of non-matching pairs: $\mathbf{s}_{ij}^- = |\mathbf{s}_i^a - \mathbf{s}_j^b|$, for all $j \neq i$. Consider a linear function $F(\mathbf{s}) = \mathbf{w} \cdot \mathbf{s}$, we wish to learn a weight vector such that:

$$\forall j \neq i : F(\mathbf{s}_i^+) > F(\mathbf{s}_{ij}^-). \quad (8)$$

The solution can be obtained using the SVM solver by minimizing following objective function:

$$\begin{aligned} & \frac{1}{2} \mathbf{w} \cdot \mathbf{w} + C \sum \xi_{ij} \\ \text{s.t. } & \forall j \neq i : \mathbf{w} \cdot \mathbf{s}_i^+ \geq \mathbf{w} \cdot \mathbf{s}_{ij}^- + 1 - \xi_{ij} \\ & \forall (i, j) : \xi_{ij} \geq 0. \end{aligned} \quad (9)$$

Once the RankSVM model is trained, it can be used to give the similarity rankings out of a set of probe-gallery pairs, *i.e.*, to fulfill the re-id task.

4 Experiments

In this section, we present extensive experimental evaluations and in-depth analysis of the proposed method in terms of effectiveness and robustness. We also quantitatively compare our method with the state-of-the-art multi-shot re-id methods. Algorithmic evaluations are performed on the following two person re-id benchmarks:

iLIDS-VID dataset. The iLIDS-VID dataset contains of 600 image sequences for 300 people in two non-overlapping camera views. Each image sequence consists of frames of variable length from 23 to 192, with an average length of 73. This dataset was created at an airport arrival hall under a multi-camera CCTV network, and the images were captured with significant background clutter, occlusions, and viewpoint/illumination variations, which makes the dataset very challenging.

PRID 2011 dataset. The PRID 2011 dataset includes 400 image sequences for 200 people from two camera. Each image sequence has a variable length consisting of 5 to 675 image frames, with an average number of 100. The images were captured in an uncrowded outdoor environment with relatively simple and clean background and rare occlusion; however, there are significant viewpoint and illumination variations as well as color inconsistency between the two views.




Fig. 3. Rank-1,5,10 and 20 person re-id matching rates based on different *fusion depth* on (a) iLIDS-VID and (b) PRID 2011 datasets

Following [11], the sequence pairs with more than 21 frames are used in our experiments.

Evaluation settings. Following [11], the whole set of human sequence pairs of both datasets is randomly split into two subsets with equal size, *i.e.*, one for training and the other for testing. The sequence of the first camera is used as probe set while the gallery set is from the other one. For both datasets, we report the performance of the average Cumulative Matching Characteristics (CMC) curves over 10 trials.

4.1 Effectiveness of Recurrent Feature Aggregation Network

To evaluate the effectiveness of the proposed recurrent feature aggregation network (**RFA-Net**), we extract the learned feature representation from each LSTM node of our network and evaluate the person re-id matching rate on the testing set. The results on both datasets are illustrated in Fig. 3. From Fig. 3, we have two observations. First, the matching rate increases consistently when the features are extracted from *deeper* LSTM nodes of our network. And the accumulated sequence level person representation (last node output) outperforms single frame level representation (first node output) for about 10% on both datasets. This demonstrates that the proposed network is able to aggregate discriminative information from frame level towards a more discriminative sequence level feature representation. Second, we note that the rank-1 matching rate increases (with respect to the network depth) faster than others, which shows that the rank-1 matching performance relies more on the feature representation.

Based on the same type of visual features (*i.e.*, Color&LBP), we further compare our method with state-of-the-art matching method for person re-id. The compared methods are as follows. (1) Baseline method: the averagely pooled feature representation (*i.e.*, Color&LBP) of each frame in the person sequence combined with RankSVM [10] based matching method; (2) Sequence matching method: the two person sequences are matched by the Dynamic Time Wrapping

Table 1. Performance of different methods based on Color&LBP feature

Dataset	iLIDS-VID				PRID 2011				
	Rank R	R=1	R=5	R=10	R=20	R=1	R=5	R=10	R=20
Color&LBP [31]+RSVM		23.2	44.2	54.1	68.8	34.3	56.0	65.5	77.3
Color&LBP+DTW [35]		9.3	21.7	29.5	43	14.6	33	42.6	47.8
Color&LBP+DVR [11]		34.5	56.7	67.5	77.5	37.6	63.9	75.3	89.4
Color&LBP+ RFA-Net +Cosine		44.5	71.9	82.0	90.1	54.9	84.2	93.7	98.4
Color&LBP+ RFA-Net +RSVM		49.3	76.8	85.3	90.0	58.2	85.8	93.4	97.9

(DTW) [35] method; and (3) Discriminative feature selection method: the Discriminative Video Ranking (DVR) [11] method which selects the most discriminative fragments for person matching. The comparison results on both datasets are summarized in Table 1. For our method, we perform experiment by using the simple cosine distance metric or the learned distance based on RankSVM and report both results. Note that for the rank-1 matching rate, our method outperforms the baseline method and the Discriminative Video Ranking (DVR) based method for more than 60%. We further notice that even without metric learning, our method still achieves a rather good performance, which well demonstrates the discriminativeness of our learned sequence level representation, *i.e.*, our propose network is capable of fusing frame level human features into a more discriminative human sequence level feature representation. Also, based on recurrent scheme, our network implicitly encodes human dynamics. In contrast, the comparing methods which are only based on frame level feature matching (or even discriminative feature selection) do not possess good discriminative capability.

4.2 Robustness of Recurrent Feature Aggregation Network

As there exist a large amount of color/illumination change as well as occlusion and background clutter, a good re-id method should be robust to these noise. Here, we evaluate the robustness of our method by adding noise to human patch sequences to be matched. For the test image sequences in each dataset, part of the images are replaced by noise images, which are randomly selected from the other dataset (e.g., the noise for PRID 2011 dataset are images from iLIDS-VID dataset). Then we extract feature vectors from the noise contaminated images for person re-id. As shown in Table 2, although the performance suffers from the noise, our method still achieves 29.8% and 44.7% rank-1 matching rate on the two datasets when up to 50% images (frames) are contaminated, which is remarkable. We also notice that the rank-10 and rank-20 matching rates just decrease slightly compared to the corresponding noise-free sequences. This demonstrates that despite large amount of noises, our feature representation still remains good discriminativeness. This is because that LSTM network is capable of *propagating* discriminative features to deeper LSTM nodes as well as *forgetting* irrelevant information (*i.e.*, noises) during propagation. Therefore the aggregated sequence level representation is highly discriminative.

Table 2. Performance of our method in existence of noises

Dataset	iLIDS-VID				PRID 2011			
Rank R	R=1	R=5	R=10	R=20	R=1	R=5	R=10	R=20
Noise Level: 0 %	49.3	76.8	85.3	90.0	58.2	85.8	93.4	97.9
Noise Level: 10 %	43.4	70.6	81.5	88.9	52.3	83.2	91.4	97.5
Noise Level: 30 %	40.0	67.4	77.5	87.0	51.4	81.1	90.5	96.9
Noise Level: 50 %	29.8	60.5	71.9	81.5	44.7	75.2	85.6	95.5

Table 3. Performance of our method choosing different number of subsequences

Dataset	iLIDS-VID				PRID 2011			
Rank R	R=1	R=5	R=10	R=20	R=1	R=5	R=10	R=20
No. of subsequences: 1	33.1	60.9	73.0	82.9	45.6	73.0	84.7	94.3
No. of subsequences: 5	43.9	70.7	80.7	89.2	52.4	81.6	91.7	97.3
No. of subsequences: 10	49.3	76.8	85.3	90.0	58.2	85.8	93.4	97.9
No. of subsequences: 15	47.8	72.2	82.4	89.7	57.9	83.4	92.8	98.0

4.3 Effectiveness of Feature Averaging

We also evaluate the performance of our method when different numbers of subsequences are selected in the test phase. Choosing a single subsequence is similar to single-shot case, *i.e.*, the feature of a randomly selected fragment is used for matching. Therefore to be more discriminative, we first randomly select several subsequences from the input human sequence and average the subsequence level human features (based on our recurrent feature aggregation network) over several subsequences. The results are shown in Table 3. As expected, averaging several subsequences significantly enhance the performance because it makes the feature vector more robust to pose/illumination changes for representing a person, *i.e.*, noise is also diluted by this scheme. We also notice that the performance does not have much difference when the number of subsequences is 10 and 15. This is because 10 subsequences are already sufficient to well represent human dynamics.

4.4 Comparison to the State-of-the-Art

In this section, our method is compared with the state-of-the-art multi-shot person re-id approaches. We choose three methods for comparison and the results are shown in Table 4. From the results, we see that our method has outperformed the Discriminative Video Ranking (DVR) [36] for more than 20% on rank-1 matching rate. This is due to the fact that instead of selecting just one discriminative video fragment, we randomly select several subsequences in the video, which makes our feature more discriminative. Our method also achieves

Table 4. Performance of our method compared against state-of-the-art methods

Dataset	iLIDS-VID				PRID 2011				
	Rank R	R=1	R=5	R=10	R=20	R=1	R=5	R=10	R=20
3D Hog&Color+DVR [36]		39.5	61.1	71.7	81.0	40.0	71.7	84.5	92.2
DVDL [14]		25.9	48.2	57.3	68.9	40.6	69.7	77.8	85.6
STFV3D [15]		37.0	64.3	77.0	86.9	21.6	46.4	58.3	73.8
STFV3D+KISSME [37]		44.3	71.7	83.7	91.7	64.1	87.3	89.9	92.0
Color&LBP+ RFA-Net +Cosine		44.5	71.9	82.0	90.1	54.9	84.2	93.7	98.4
Color&LBP+ RFA-Net +RSVM		49.3	76.8	85.3	90.0	58.2	85.8	93.4	97.9

better performance than the dictionary learning method DVDL [14], which discriminatively trained viewpoint invariant dictionaries. However, DVDL is unable to encode temporal information into the dictionaries. In our case, LSTM effectively solves this problem. Spatio-temporal Fisher vector (STFV3D) [15] generates a spatio-temporal body-action model that consists of a series of body-action units, which are represented by Fisher vectors built upon low-level features that combines color, texture and gradient information. This representation aligns the spatio-temporal appearance of a person globally and thus displays high discriminativeness when combines with a metric learning method (KISSME [37]). Compared to this method, our approach is simple and does not need to consider the temporal alignment problem. Even though, we still achieves comparable results to STFV3D+KISSME. In addition, the impressive performance of STFV3D+KISSME is largely due to the effectiveness of the metric learning method. Thus, using no metric learning leads to a large performance drop for this method. On the other hand, our method does not necessarily require a metric learning method to enhance performance. Even using a simple cosine distance metric, it still achieves impressive results on both datasets. This demonstrates that the feature we learned is very discriminative itself.

4.5 Matching Examples

In this section, we discuss about some matching examples in our experiments, as shown in Fig. 4. Here, four typical failure re-identification examples and a successful example are illustrated, where each row displays one. The first column in the figure consists of probe sequences, the second corresponds to their correct matching pairs and the third one contains the top matching sequence generated by our method. The top two rows show the failure examples from iLIDS-VID datasets, where the first failure is largely due to color inconstancy and occlusion. In the second case, the persons in last two sequences nearly have the same pose and they wear clothes of similar color, which make it difficult for a texture and color based feature to discriminate. In fact, these two sequence have very close similarity scores. The next two failure cases are similar, which are illustrated in the third and fourth row. The matching failure is caused by similar local

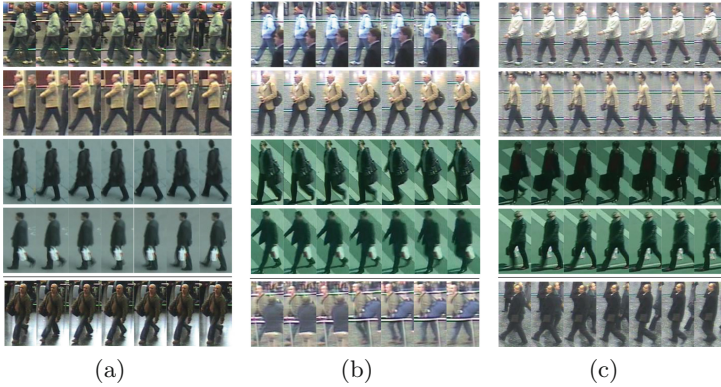


Fig. 4. Matching examples. (a) probe sequence, (b) Correct match, (c) Rank-1 matching sequence. The first four rows correspond to failure examples using features accumulated to the 10th LSTM node. The bottom row illustrates an example that fails to match the correct sequence when using the features from the first LSTM node, but correctly matches the probe sequence when using the features accumulated to the 10th LSTM node

information such as black briefcase in and white bag in hand, which makes it difficult to identify the correct match. These failure examples demonstrate that both color and local features have their limitations, using or combing more features has great opportunity to enhance the performance of our system.

An interesting example is displayed in the bottom row, where the third column corresponds to the rank-1 matching sequence using the features extracted from the first LSTM node. The failure is due to the occlusion, pose change and color inconsistency in the gallery sequence. However, when features are accumulated to the deeper (10th) LSTM nodes, our system correctly matches the probe sequence with the corresponding gallery sequence, which well demonstrate the effectiveness of our feature aggregation method. In addition, notice that all the examples displayed here are also difficult for humans to identify. The performance of our method using simple texture and color features is remarkable.

5 Conclusions

In this paper, we proposed a novel recurrent feature aggregation framework for person re-identification. In contrast to existing multi-shot person re-id methods that use complex feature descriptors or design complex matching metric, our method is capable of learning discriminative sequence level representation from simple frame-wise features. Experimental results show that features learned by this recurrent network not only possess high discriminativeness, and also show great robustness to noise. Even using simple texture and color features as input, the performance of the proposed method is better/comparable to the state-of-the-art methods.

Acknowledgements. The work was supported by State Key Research and Development Program (2016YFB1001003), NSFC (61527804, 61521062, 61502301), STCSM (14XD1402100), the 111 Program (B07022) and China's Thousand Youth Talents Plan.

References

1. Bedagkar-Gala, A., Shah, S.K.: A survey of approaches and trends in person re-identification. *Image Vis. Comput.* **32**(4), 270–286 (2014)
2. Farenzena, M., Bazzani, L., Perina, A., Murino, V., Cristani, M.: Person re-identification by symmetry-driven accumulation of local features. In: *CVPR*, pp. 2360–2367 (2010)
3. Gray, D., Tao, H.: Viewpoint invariant pedestrian recognition with an ensemble of localized features. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) *ECCV 2008*, Part I. LNCS, vol. 5302, pp. 262–275. Springer, Heidelberg (2008)
4. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *IJCV* **60**(2), 91–110 (2004)
5. Ojala, T., Pietikäinen, M., Mäenpää, T.: Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *TPAMI* **24**(7), 971–987 (2002)
6. Li, W., Zhao, R., Xiao, T., Wang, X.: Deepreid: deep filter pairing neural network for person re-identification. In: *CVPR*, pp. 152–159 (2014)
7. Ahmed, E., Jones, M.J., Marks, T.K.: An improved deep learning architecture for person re-identification. In: *CVPR*, pp. 3908–3916 (2015)
8. Weinberger, K.Q., Blitzer, J., Saul, L.K.: Distance metric learning for large margin nearest neighbor classification. In: *NIPS*, 1473–1480 (2005)
9. Roth, P.M., Hirzer, M., Köstinger, M., Beleznai, C., Bischof, H.: Mahalanobis distance learning for person re-identification. In: Gong, S., Cristani, M., Yan, S., Loy, C.C. (eds.) *Person Re-Identification*. *ACVPR*, pp. 247–267. Springer, Heidelberg (2014)
10. Chapelle, O., Keerthi, S.S.: Efficient algorithms for ranking with svms. *Inf. Retrieval* **13**(3), 201–215 (2010)
11. Wang, T., Gong, S., Zhu, X., Wang, S.: Person re-identification by video ranking. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014*, Part IV. LNCS, vol. 8692, pp. 688–703. Springer, Heidelberg (2014)
12. Zheng, L., Shen, L., Tian, L., Wang, S., Wang, J., Tian, Q.: Scalable person re-identification: A benchmark. In: *ICCV*, pp. 1116–1124 (2015)
13. Zheng, L., Shen, L., Tian, L., Wang, S., Bu, J., Tian, Q.: Person re-identification meets image search. *CoRR* abs/1502.02171 (2015)
14. Karanam, S., Li, Y., Radke, R.J.: Person re-identification with discriminatively trained viewpoint invariant dictionaries. In: *ICCV*, pp. 4516–4524 (2015)
15. Liu, K., Ma, B., Zhang, W., Huang, R.: A spatio-temporal appearance representation for video-based pedestrian re-identification. In: *ICCV*, pp. 3810–3818 (2015)
16. Hirzer, M., Beleznai, C., Roth, P.M., Bischof, H.: Person re-identification by descriptive and discriminative classification. In: Heyden, A., Kahl, F. (eds.) *SCIA 2011*. LNCS, vol. 6688, pp. 91–102. Springer, Heidelberg (2011)
17. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
18. Ma, B., Su, Y., Jurie, F.: Bicov: a novel image representation for person re-identification and face verification. In: *BMVC*, pp. 1–11 (2012)

19. Scovanner, P., Ali, S., Shah, M.: A 3-dimensional sift descriptor and its application to action recognition. In: ICME, pp. 357–360 (2007)
20. Kläser, A., Marszalek, M., Schmid, C.: A spatio-temporal descriptor based on 3d-gradients. In: BMVC, pp. 1–10 (2008)
21. Hamdoun, O., Moutarde, F., Stanculescu, B., Steux, B.: Person re-identification in multi-camera system by signature based on interest point descriptors collected on short video sequences. In: ICDSC, pp. 1–6 (2008)
22. Truong Cong, D.N., Achard, C., Khoudour, L., Douadi, L.: Video sequences association for people re-identification across multiple non-overlapping cameras. In: Foggia, P., Sansone, C., Vento, M. (eds.) ICIAP 2009. LNCS, vol. 5716, pp. 179–189. Springer, Heidelberg (2009)
23. Gheissari, N., Sebastian, T.B., Hartley, R.I.: Person reidentification using spatiotemporal appearance. In: CVPR, pp. 1528–1535 (2006)
24. Xu, Y., Lin, L., Zheng, W., Liu, X.: Human re-identification by matching compositional template with cluster sampling. In: ICCV, pp. 3152–3159 (2013)
25. Shen, C., Kim, J., Wang, L., van den Hengel, A.: Positive semidefinite metric learning using boosting-like algorithms. *J. Mach. Learn. Res.* **13**, 1007–1036 (2012)
26. Prosser, B., Zheng, W., Gong, S., Xiang, T.: Person re-identification by support vector ranking. In: BMVC, pp. 1–11 (2010)
27. Zheng, W., Gong, S., Xiang, T.: Reidentification by relative distance comparison. *TPAMI* **35**(3), 653–668 (2013)
28. McLaughlin, N., Martinez del Rincon, J., Miller, P.: Recurrent convolutional network for video-based person re-identification. In: CVPR (June 2016)
29. Haque, A., Alahi, A., Fei-Fei, L.: Recurrent attention models for depth-based person identification. In: CVPR (June 2016)
30. Zhao, R., Ouyang, W., Wang, X.: Unsupervised salience learning for person re-identification. In: CVPR, pp. 3586–3593 (2013)
31. Hirzer, M., Roth, P.M., Köstinger, M., Bischof, H.: Relaxed pairwise learned metric for person re-identification. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012, Part VI. LNCS, vol. 7577, pp. 780–793. Springer, Heidelberg (2012)
32. Jing, X., Zhu, X., Wu, F., You, X., Liu, Q., Yue, D., Hu, R., Xu, B.: Super-resolution person re-identification with semi-coupled low-rank discriminant dictionary learning. In: CVPR, pp. 695–704 (2015)
33. Donahue, J., Hendricks, L.A., Guadarrama, S., Rohrbach, M., Venugopalan, S., Saenko, K., Darrell, T.: Long-term recurrent convolutional networks for visual recognition and description. In: CVPR (2015)
34. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: Convolutional architecture for fast feature embedding (2014). arXiv preprint [arXiv:1408.5093](https://arxiv.org/abs/1408.5093)
35. Simonnet, D., Lewandowski, M., Velastin, S.A., Orwell, J., Turkbeyler, E.: Re-identification of pedestrians in crowds using dynamic time warping. In: Fusiello, A., Murino, V., Cucchiara, R. (eds.) ECCV 2012. LNCS, vol. 7583, pp. 423–432. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-33863-2_42](https://doi.org/10.1007/978-3-642-33863-2_42)
36. Wang, T., Gong, S., Zhu, X., Wang, S.: Person re-identification by discriminative selection in video ranking. *CoRR* abs/1601.06260 (2016)
37. Köstinger, M., Hirzer, M., Wohlhart, P., Roth, P.M., Bischof, H.: Large scale metric learning from equivalence constraints. In: CVPR, pp. 2288–2295 (2012)