

Jointly Attentive Spatial-Temporal Pooling Networks for Video-based Person Re-Identification

Shuangjie Xu^{1*} Yu Cheng^{2*} Kang Gu¹ Yang Yang³ Shiyu Chang⁴ Pan Zhou¹

¹Huazhong University of Science and Technology ²AI Foundations, IBM Research
⁴Northwestern University ³IBM T.J. Watson Research Center

Abstract

Person Re-Identification (person re-id) is a crucial task as its applications in visual surveillance and human-computer interaction. In this work, we present a novel joint Spatial and Temporal Attention Pooling Network (ASTPN) for video-based person re-identification, which enables the feature extractor to be aware of the current input video sequences, in a way that interdependency from the matching items can directly influence the computation of each other's representation. Specifically, the spatial pooling layer is able to select regions from each frame, while the attention temporal pooling performed can select informative frames over the sequence, both pooling guided by the information from distance matching. Experiments are conducted on the iLIDS-VID, PRID-2011 and MARS datasets and the results demonstrate that this approach outperforms existing state-of-art methods. We also analyze how the joint pooling in both dimensions can boost the person re-id performance more effectively than using either of them separately¹.

1. Introduction

Person Re-Identification has been viewed as one of the key subproblems of the generic object recognition task. It is also important due to its applications in surveillance, and human-computer interaction communities. Given a query image, the task is to identify a set of matching person images from a pool, usually captured from the same/different cameras, from different viewpoints, at the same/different time points. It is a very challenging task due to the large variations of lighting conditions, viewing angles, body poses and occlusions.

Methods for re-identification in still images setting have been extensively investigated, including feature represen-

*indicates equal contributions.

¹the code is available at <https://github.com/shuangjiexu/Spatial-Temporal-Pooling-Networks-ReID>




Figure 1. Sample video frames from one person captured by three cameras a, b and c, simulating how human compare different video pairs. The regions under the circles are the parts which visual attentions are drawn on.

tation learning [11, 19, 15, 38], distance metric learning [14, 41, 30, 36, 37, 21] and CNN-based schemes [33, 32, 24, 25]. Very recently, researchers began to explore solving this problem in video-based setting, which is a more natural way to perform re-identification. The intuition of this kind of methods is that temporal information related to person motion can be captured from video. Moreover, sequences of images provide rich samples of persons' appearances, helping boosting the re-identification performance with more discriminative features. In [20], a temporal deep neural network architecture combines optical flow, recurrent layers and mean-pooling and achieves reasonable success. The work in [32] exploited a novel recurrent feature aggregation framework, which is capable of learning discriminative sequence level representation from simple frame-wise features.

The main idea of video-based methods is first to extract useful representations from video images with RNN (or CNN-RNN) models. Then exploit a distance function to judge their extent of matching. However, most of these ap-

proaches derive each sequence’s representation separately, rarely considering the impact of the others, which neglect the mutual influence of the two video sequences in the context of the matching task. Let’s think about how human visual processing works when comparing video sequences. For example, the pair-wise case described in Figure 1, when comparing video frames a with two other b and c separately, as b and c are different, it is natural for our brain to draw different focuses on different frames of a . On the other hand, the interaction of compared sequences should also have effect on the spatial dimension, which guides human to pay attentions on different regions of the input a . This is extremely important for the scenario with large viewpoint changes or fast moving object. The example demonstrates why we should draw different attention when comparing different pairs of video frames.

Motivated by recent success of attention models [1, 31, 34, 5], we proposed jointly Attentive Spatial-Temporal Pooling Networks (ASTPN), a powerful mechanism for learning the representation of video sequences by taking into account the interdependence among them. Specifically, ASTPN first learns a similarity measure over the features extracted from recurrent-convolutional networks of the two input items, and uses the similarity scores between the features to compute attention vectors in both spatial (regions in each frame) and temporal (frames over sequences) dimensions. Next, the attention vectors are used to perform pooling. Finally, a Siamese network architecture is deployed over the attention vectors. The proposed architecture can be trained efficiently with the end-to-end training schema.

We perform extensive experiments on three datasets, iLIDS-VID, PRID-2011 and MARS. The results clearly demonstrate that our proposed method for person re-identification outperforms well established baselines significantly and offers new state-of-the-art performance. The cross dataset test also derives the same conclusion. ASTPN is also a general component that can handle a wide variety of person re-identification tasks.

2. Related Work

Person re-id, a challenging task which has been explored for several years, still remains to be further focused on to overcome the problems of viewpoint difference, illumination change, occlusions and even similar appearance of different people. A majority of recent works mainly develop their solutions from two aspects: extracting reliable feature representations [28, 6, 11, 19, 15, 38] or learning a robust distance metric [14, 41, 2, 30, 36, 37, 21, 13]. To be specific, features including color histograms [37, 30], texture histograms [6], Local Binary Patterns [30], Color Names [40] and so on are widely utilized for person re-id to address identity information in the existence of challenges like lighting change. In the meantime, metric learning methods

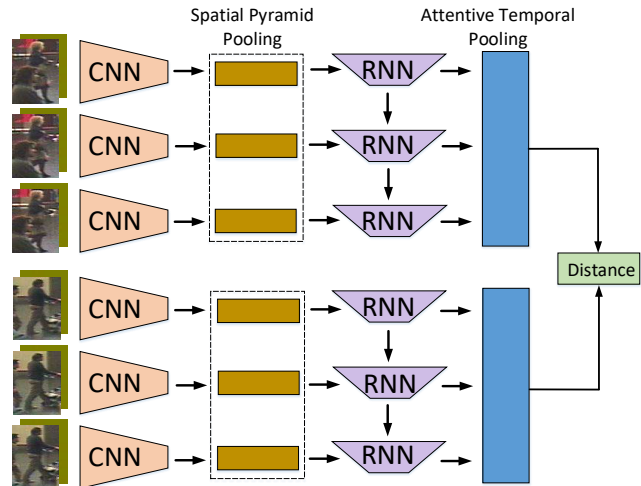


Figure 2. Our video-based person re-identification system. We adopt Siamese network architecture for spatial-temporal feature extraction, and jointly attentive spatial-temporal pooling for inter-dependence information learning.

such as large margin nearest neighbor (LMNN) [29], Mahalanobis distance metric (RCA) [2], Locally Adaptive Decision Function (LADF) [13] and RankSVM [38] have also been applied to person re-id task. Despite the prominent progress in recent years, most of these works are still based on image-to-image level. Video setting is intuitively more close to the practical scenario as video is the first-hand material captured by surveillance camera [4, 3]. Besides, temporal information relevant to a person’s motion, gait for instance, may help to discriminate similar pedestrians. Moreover, video provides abundant samples of the target for us with the cost of increasing computation.

Gradually, more and more works began to explore video-to-video matching problem in person re-id. Discriminative Video Ranking model [27] used discriminative video fragments selection to capture more accurate space-time information, while simultaneously learning a video ranking function for person re-id. Bag-of-words [40] method aimed to encode frame-wise features into a global vector. However, neither of these models could be considered effective for ignoring the rich temporal information contained in the videos. However, video-based person re-id raises new challenges: some inter-class difference of video-based representation can be much more ambiguous compared with the one when using image-based representation, since it’s likely that different people could not only have similar appearance but also similar motions, making alignment tough to achieve. Therefore, space-time information must be fully utilized to solve those extra problems. Besides, a top-push distance learning (TDL) model has been proposed to effectively make use of space-time information, with a top-push constraint to quantify ambiguous video representation [35].

Deep learning offers an approach to solve feature representation and metric learning problem at the same time. The typical architecture is composed of two parts: a feature extracting network, usually a CNN or RNN, and multiple metric learning layers to make final prediction. The first Siamese-CNN (SCNN) structure [33] proposed for person re-id leveraged a set of 3 SCNNs to the three overlapped parts of the image. [32] exploited a novel recurrent feature aggregation framework, which is capable of learning discriminative sequence level representation for frame-wise features. A recent work [20] used CNN to obtain feature representation from the multiple frames of the video, then applied RNN to learn the interaction between them. Temporal pooling layer followed the recurrent layer, aiming to capture sequential interdependence (the pooling might be max-pooling or mean-pooling). Those layers were jointly trained to function as a feature extractor. However, the max-pooling and mean-pooling adapted by them were not robust enough to compress and produce the person’s appearance over a period of time, since max-pooling only employed the most active feature map at one temporal step of whole sequence and mean-pooling, which produced a representation averaged over all time steps, thus couldn’t preclude the impact of ineffective features well.

More importantly, the re-id frameworks usually take the form of similarity measure with other inputs. Most of prior works ignored the mutual influence of other items when performing representation learning. Thus we would like to fill this gap by introducing the attention mechanism, which already achieved great success in image caption generation [31], machine translation [1], question-answering [34] as well as action recognition [23]. [16] presented a comparative attention architecture and addressed the problem in the spatial dimension. [5] proposed a two-way attention mechanism to matching the text sequence, which is exploited in our framework as the temporal pooling component.

3. The Proposed Model Architecture

This work builds a recurrent-convolutional network with jointly attentive spatial-temporal pooling (ASTPN) for video-based person re-identification. Our ASTPN architecture works by passing a pair of video sequences through a Siamese networks to obtain two representations and producing the Euclidean distance between them. As shown in Figure 2, each input (one frame from a video with optic flow involved) is passed through a CNN network to extract feature maps from the last convolutional layer. Then those feature maps are fed into our spatial pooling layer to obtain image-level representation at one time step. After that, we take temporal information into consideration by utilizing a recurrent network to generate the feature set of a video sequence. Finally, all time steps resulting from recurrent network are combined by attentive temporal pooling to form

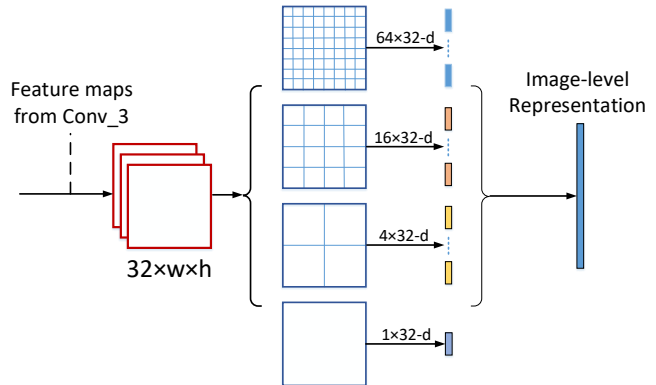


Figure 3. Jointly attentive spatial pooling architecture. Here *Conv_3* is the last convolutional layer. In the spatial pooling layer, we use a spatial pyramid pooling structure with multi-level spatial bins (8×8, 4×4, 2×2 and 1×1). The image-level representation is then generated by joining all pooling outputs with those spatial bin.

the sequence-level representation.

The crucial part of ASTPN relies on the jointly attentive spatial-temporal pooling (ASTP) layers. Instead of using general max (mean) pooling and over-time temporal pooling layers, this pooling mechanism could take information to form the distance at each step, allowing our model to be more attentive both on region of interests in image level and on effective time step in sequence level. Moreover, attentive spatial-temporal pooling also makes our model adaptive to image sequence of arbitrary resolution/length. Detailed techniques about the attentive spatial and temporal pooling will be presented in following subsections.

3.1. Spatial Pooling Layer

In person re-identification, due to the overlooking angle of most surveillance equipment, pedestrians only take a part in whole spatial images. Therefore, local spatial attention is necessary for deep networks. The design of such layer should 1) generate multi-scales region patches of each image and feed them into RNN/attention pooling layer; 2) make the model robust to image sequence of arbitrary resolution/length. In this work, we use spatial pyramid pooling (SPP) layer [7] as the component attentive spatial pooling to concentrate our model on important region in spatial dimension. Shown in Figure 3, the SPP layer has multi-level spatial bins to generate multi-level spatial representations, and then those representations are combined into a fixed-length image-level representation. Since the image-level representations involve pedestrian position and multi-scale spatial information, our joint attentive spatial pooling mechanism is able to select regions from each frame.

Given the input sequence $\mathbf{v} = \{v^1, \dots, v^T\}$, we obtain the feature maps set $\mathbf{C} = \{C^1, \dots, C^T\}$ by utilizing the convolutional network shown in Table 1. Each

$C^i \in \mathbb{R}^{c \times w \times h}$ is then fed into spatial pooling layer to get image-level representation r^i . Assuming that the size set of spatial bins is $\{(m_w^j, m_h^j) | j = 1, \dots, n\}$, the window size $win^j = \left(\left\lceil \frac{w}{m_w^j} \right\rceil, \left\lceil \frac{h}{m_h^j} \right\rceil\right)$ and pooling stride $str^j = \left(\left\lfloor \frac{w}{m_w^j} \right\rfloor, \left\lfloor \frac{w}{m_w^j} \right\rfloor\right)$ for the j -th spatial bin are determined. Then the result vector r^i is obtained by formula:

$$\begin{aligned} b^{i,j} &= f_R \{f_p(C^i; win^j, str^j)\} \\ r^i &= b^{i,1} \oplus b^{i,2} \oplus \dots \oplus b^{i,n} \end{aligned} \quad (1)$$

where f_p represents the max pooling function with window size win and stride str . $\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$ denote ceiling and floor operations. f_R means the reshape operation which reshapes a matrix to a vector. Besides, \oplus denotes vector connection operation. Let $\mathbf{r} = \{r^i \in \mathbb{R}^L | i = 1, \dots, T\}$ be a sequence representation, where $L = \sum_j m_w^j m_h^j$, we then pass \mathbf{r} forward to the recurrent network to extract information between time-steps. The hidden layer is formulized by:

$$o^t = U r^t + W s^{t-1}, \quad s^t = \tanh(o^t) \quad (2)$$

where $s^{t-1} \in \mathbb{R}^N$ is the hidden state containing information from previous time step, and o^t is the output as time t . Fully-connected weight $U \in \mathbb{R}^{L \times N}$ projects the recurrent layer input r^t from \mathbb{R}^L to \mathbb{R}^N , and $W \in \mathbb{R}^{N \times N}$ projects hidden state s^{t-1} from \mathbb{R}^N to \mathbb{R}^N . Notice that the recurrent layer embeds the feature vector into a lower-dimensional feature by matrix U . The hidden state s^0 is initialized to zero at first time step, and between time steps the hidden state is passed through \tanh activation function.

3.2. Attentive Temporal Pooling Layer

Although the recurrent layer is able to capture temporal information with hidden states, those raw temporarily contains much redundant information. For instance, there are only minor changes in a series of continuous frames, shown in Figure 1, thus features learned from these sequence input involve a lot of redundant information such as the ambiguous background and clothing. In order to avoid the "Bad money drives out good" issue, we propose an attentive temporal pooling architecture to enable our model to concentrate on effective information. Attentive temporal pooling reinforces the pooling layer to perceive the input probe and gallery data pair, and allows the probe input sequence I_P to directly influence the computation of gallery sequence representation v_g . We put attentive temporal pooling layer between the recurrent layer and the distance computation layer. In the training phase, attentive temporal pooling is jointly learning with recurrent-convolutional network and the spatial pooling layer, guiding our model for effective information extraction in temporal dimension.

In Figure 4, we illustrate our attentive temporal pooling architecture, which follows the design in [5, 34]. Given the

matrices $P \in \mathbb{R}^{T \times N}$ and $G \in \mathbb{R}^{T \times N}$, whose i -th row represents the output of the recurrent layer in the i -th time step with probe data and gallery data respectively, we compute the attention matrix $A \in \mathbb{R}^{T \times T}$ as follows:

$$A = \tanh(PUG^T) \quad (3)$$

where $U \in \mathbb{R}^{N \times N}$ is a intent information sharing matrix to be learned by networks. When the convolution and recurrent layer are employed to obtain matrix P and G , the attention matrix is able to have a sight on both probe and gallery sequence features, and computes weight scores in temporal dimension. In the gradient descent phase, U is updated by back propagation and influences parameters of convolution and hidden state to guide our model to focus on effective information.

Next, we apply column-wise and row-wise max pooling on A respectively to obtain temporal weight vector $t_p \in \mathbb{R}^T$ and $t_g \in \mathbb{R}^T$. The i -th element of t_p represents importance score for the i -th frame in the probe sequence, which is the same with t_g . Due to the participation of P in the computation of t_g , the vector t_g can capture the attentive scores of gallery features related to probe data.

After that, we apply softmax function on temporal weight vectors t_p and t_g to generate attention vectors $a_p \in \mathbb{R}^T$ and $a_g \in \mathbb{R}^T$. The softmax function transforms the i -th weight $[t_p]_i$ and $[t_g]_i$ to the attention ratio $[a_p]_i$ and $[a_g]_i$. For instance, the i -th element in a_g is computed as follows:

$$[a_g]_i = e^{[t_g]_i} / \sum_{j=1}^T e^{[t_g]_j} \quad (4)$$

Finally, we apply dot product between the feature matrices P , G and attention vectors a_p , a_g to obtain the sequence-level representation $v_p \in \mathbb{R}^N$ and $v_g \in \mathbb{R}^N$, respectively:

$$v_p = P^T a_p, \quad v_g = G^T a_g \quad (5)$$

3.3. Model Details

The main thought of our work is to construct a feature extracting network which is able to map the sequence data into feature vector in a low dimensional space, where the feature vectors from the sequences of the same person are close, and the feature vectors from the sequences of different persons are separated by a margin. More details on the components of our proposed network will be explained in follows.

Input: The input to our network consists of three color channels and two optical flow. The color channels provide spatial information such as clothing and background, while optical flow channels provide the temporal motion information. Compared with only use color channels as input, there

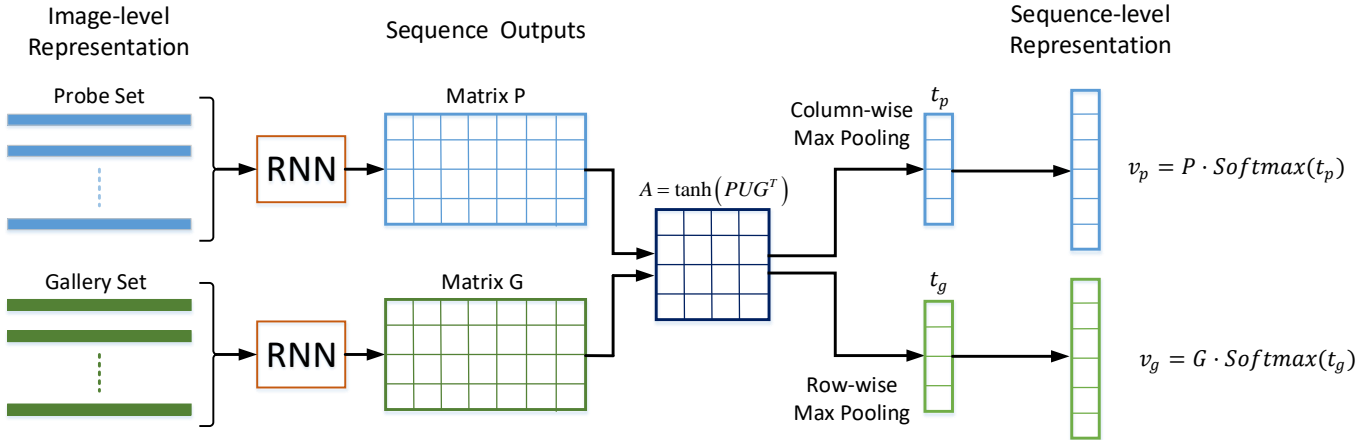


Figure 4. Attentive temporal pooling architecture. With the RNNs output matrices P and G , we compute attention matrix by introduce a parameter matrix U to capture attentive score in temporal dimension. With column/row-wise max pooling operation and softmax function, the attention vector is obtained, which contains the attentive weight for each time step. The sequence-level representation is computed by dot product between the feature matrices P , G and attention vectors a_p , a_g .

should be a promotion for person re-id when utilizing both of color channels and optical flow channels .

The Siamese Network: We use a Siamese network architecture as shown in Figure 2. As mentioned above, our network architecture is grouped by four functional parts: convolutional layers, the attentive spatial pooling layer, the recurrent layer and the attentive temporal pooling layer. As for convolutional layers, we use a convolutional architecture with parameters shown in Table 1, and the pooling layer in the final layer is replaced by the attentive spatial pooling layer. Notice that convolutional layers are unrolled along with the recurrent layer, and these layers share their parameters in all time steps, which means all frames are passed through the same spatial feature extractor. Similarly, the two recurrent layers also share their parameters to process a pair sequences input.

The Training Objective: Given a pair of sequences (I_p, I_g) of persons p and g , the sequence-level representations (v_p, v_g) are obtained by our Siamese network. After that, we use the Euclidean distance Hinge loss to train our model as follows:

$$E(v_p, v_g) = \begin{cases} \|v_p - v_g\|^2 & p = g \\ \max(0, m - \|v_p - v_g\|^2) & p \neq g \end{cases} \quad (6)$$

where m denotes the margin to separate features of different persons in Hinge loss. In the training phase, the network is shown positive and negative input pairs alternately. While in the testing phase for a new sequence input, we copy the sequence to form a new pair and pass the pair through our Siamese network to obtain identity feature. By computing the distance between the identity feature with previously saved features of other identities, the most similar identity is indicated with the lowest distance. In addition, we also

Table 1. layer parameter of the CNN network

| Layer | Type | Conv(size, channel, pad, stride) | Max Pooling |
|--------|-------|----------------------------------|-------------|
| Conv_1 | c+t+p | 5×5, 16, 4, 1 | 2×2 |
| Conv_2 | c+t+p | 5×5, 32, 4, 1 | 2×2 |
| Conv_3 | c+t | 5×5, 32, 4, 1 | N/A |

c: Convolutional layer; t: Tanh layer; p: Pooling layer

take identity classification loss into consideration, following the work [20]. We apply softmax regression on the final features (v_p, v_g) to predict the identity of persons. By using the cross-entropy loss, we obtain the identity loss $I(v_p)$ and $I(v_g)$. Since that the joint learning of Siamese loss and identity loss brings about a great promotion, the final training objective is the combination of the Siamese loss and the identity loss $L(v_p, v_g) = E(v_p, v_g) + I(v_p) + I(v_g)$.

4. Experimental Results

We evaluate our model for video-based person re-id on three different datasets: iLIDS-VID [27], PRID-2011 [8] and MARS [39]. We also investigate how the joint pooling strategy can bring benefit to the proposed network, the difference between adapting attentive temporal pooling and other common temporal pooling strategies, and the use of attentive spatial pooling.

4.1. iLIDS-VID & PRID-2011

The iLIDS-VID dataset [27] contains 300 people in total, where each person is represented by two image sequences recorded by a pair of non-overlapping cameras. The length

of frames forming each image sequence ranges from 23 to 192, with an average length of 73. The challenging dataset was created at an airport arrival hall under a multi-camera CCTV network, whose image sequences were accompanied by clothing similarities among people, lighting and view-point variations, cluttered background and occlusions.

The PRID-2011 re-id dataset [8] consists of 400 image sequences for 200 people captured by two cameras that are adjacent to each other. Each image sequence is composed of frames of length from 5 to 675, with an average number of 100. It's captured in relatively simple environments with rare occlusions, compared with the iLIDS-VID dataset.

4.1.1 Experiment Settings

Following [20], we split the whole set of human sequence pairs of iLIDS-VID and PRID-2011 randomly into two subsets with equal size. One is used for training, and the other is used for testing. We report the performance of the average Cumulative Matching Characteristics (CMC) curves over 10 trials with different train/test splits. Data augmentation was done in several forms. Firstly, since the probe and gallery sequences are of variable-length, sub-sequences of $k = 16$ consecutive frames were chosen randomly at each epoch during training process. Yet we considered the first camera as probe and the second camera as gallery during testing. Secondly, positive pair was composed of a sub-sequence from camera 1 and a sub-sequence from camera 2 containing the same person A, and negative pair was composed of a sub-sequence from camera 1 of person A and a sub-sequence from camera 2 of person B, who was selected arbitrarily from the rest of people in training set. Positive and negative sequence pairs were sent to our system successively so that the model is capable of distinguishing correct match and wrong match. Lastly, the image level augmentation was performed by cropping and mirroring. Sub-image of both width and length 8 pixels less than its progenitor was produced after cropping, and then we fixed the cropping area within the same sequence. Mirroring operation was randomly applied to a whole sequence together with a probability of $p = 0.5$. Test data also underwent the augmentation to eliminate bias.

Preprocessing steps included the following actions [20]: Images were converted to YUV color space firstly, and each color channel was normalized to have zero mean and unit variance; Optical flow, both vertical and horizontal, were extracted between each pair of adjoining images using the Lucas-Kanade method [18], and then optical flow channels were normalized to the range -1 to 1; The learning rate, when network was trained with stochastic gradient descent, was 0.001 at the beginning, with batch size set as one.

The initialization of hyper-parameters of convolutional network was performed based on [20], optimized already on

Table 2. Comparison of our model with other state-of-the-art methods on iLIDS-VID and PRID-2011 according to CMC curves (%).

| Dataset | iLIDS-VID | | | | PRID-2011 | | | |
|--------------|-----------|----|----|----|-----------|----|----|----|
| | 1 | 5 | 10 | 20 | 1 | 5 | 10 | 20 |
| ASTPN | 62 | 86 | 94 | 98 | 77 | 95 | 99 | 99 |
| RNN-CNN [20] | 58 | 84 | 91 | 96 | 70 | 90 | 95 | 97 |
| RFA [32] | 49 | 77 | 85 | 92 | 64 | 86 | 93 | 98 |
| STA [17] | 44 | 72 | 84 | 92 | 64 | 87 | 90 | 92 |
| VR [27] | 35 | 57 | 68 | 78 | 42 | 65 | 78 | 89 |
| AFDA [12] | 38 | 63 | 73 | 82 | 43 | 73 | 85 | 92 |

the challenging VIPeR person re-identification dataset [22]. Besides, the margin in the Siamese cost function was set to 3, and the dimension of feature space was set to 128. We alternately showed our Siamese network positive and negative sequence pairs, and a full epoch consisted of the equal number of both. As the training set contains 150 people with a maximum sequence length of 192, it takes approximately 3 hours to train for 700 epochs, using the Nvidia GTX-1080 GPU.

4.1.2 Results

We display the results on iLIDS-VID and PRID-2011 in Table 2. The competitor methods are introduced as follows:

- RNN-CNN: A recurrent convolutional network (RCN) [20] with temporal pooling.
- RFA: Recurrent feature aggregation network [32] based on LSTM, which aggregates the frame-wise human region representation at each time stamp and produces a sequence-level representation.
- STA: Spatio-temporal body-action model that takes the video of a walking person as input and builds a spatio-temporal appearance representation for pedestrian re-identification.
- VR: A DVR framework presented in [27] for person re-id uses discriminative space-time feature selection to automatically discover and exploit the most reliable video fragments.
- AFDA: An algorithm [12] that hierarchically clusters image sequences and uses the representative data samples to learn a feature subspace maximizing the Fisher criterion.

Comparing the CMC results of our proposed architecture with the RNN-CNN method and other systems on iLIDS-VID, we can conclude that the attentive mechanism enables our network to outperform all mentioned networks by a

large margin. Note that even for the rank-1 matching rate, our method also achieves 62%, exceeding the RNN-CNN method by more than 4%. We further notice that even without attentive spatial pooling layer, the utilization of attentive temporal pooling can still lead to a fairly good performance. Thus our proposed network is capable of capturing frame-level human features and then fusing them into a discriminative representation. Another point is that the performances of DNNs seem to apparently surpass the existing state-of-the-art algorithms [10, 12, 9], proving the power of DNNs when sufficient training data is available.

Less challenging than iLIDS-VID as PRID-2011 is, we can observe overall increments in matching rate. Our model still outperforms other methods prominently in terms of Table 2, with rank-1 accuracy achieving 77%—transcending the RNN-CNN method by 7%. Besides, our system is more efficient and robust since its CMC rank rate reaches 95% at level of rank 5 and further goes up to the summit of 99% quickly at the level of rank 10. The tendency of accuracy demonstrates that our system is an effective space-time feature extractor, able to obtain more discriminative sequence-level representation through learning process. DNNs [20, 32] still exhibit distinctive capability of capturing human features on the whole, with accuracy converging at earlier point.

4.2. MARS

This is a dataset introduced in [39], which is also claimed to be the largest video re-id dataset to date. MARS consists of 1261 different pedestrians, each of whom was captured by at least two cameras. Compared with iLIDS-VID and PRID-2011, MARS is 4 times larger in the number of identities and 30 times larger in total tracklets. The tracklets of MARS are generated automatically by DPM detector and GMMCP tracker, whose error makes MARS more realistic and of course more challenging than previous dataset. Each identity has 13.2 tracklets on average. For instance, most identities are captured by 2-4 cameras, and most identities have 5-15 tracklets, most of which contain 25-50 frames.

To perform our experiments on MARS, simplification should be done in two steps. Firstly, as pedestrians were recorded by at least 2 cameras, we randomly chose 2 camera viewpoints of the same person out of the ensemble. Then one of them was set as probe set and the other was set as gallery set. Here the case was reduced to our previous experiences with iLIDS-VID and PRID-2011.

The performances of our models are displayed in Table 3, compared with the baseline RNN-CNN. ASTPN still achieves the best accuracy while general results dropping obviously in contrast with Table 2. Compared with iLIDS-VID and PRID-2011, the improvement is larger (around 4% in all ranks). The reasons may be attributed to that a considerable part of image sequences of MARS are accompanied

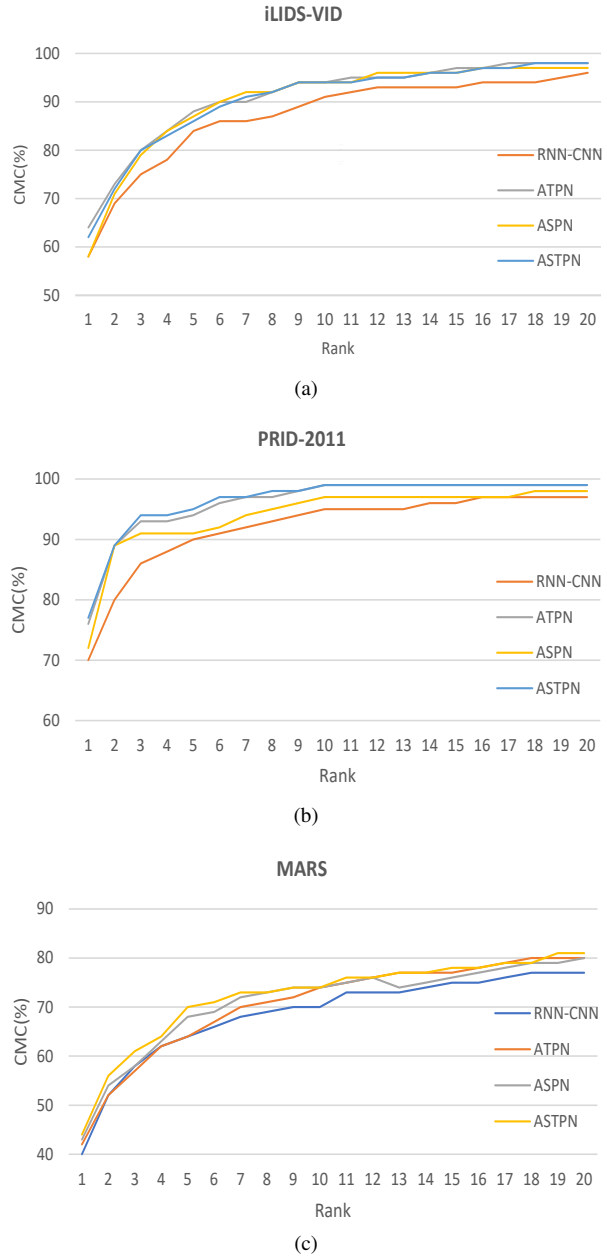


Figure 5. The variants of our model are tested on three datasets respectively. ATPN refers to attentive temporal pooling network and ASP refers to attentive spatial pooling network. Finally ASTPN stands for the combination of ATPN and ASPN.

by cluttered backgrounds, ambiguity in visual appearance, or drastic viewpoint changes between sequence pairs.

4.3. Control Experiments with Different Pooling Strategies

We investigate the effects of attentive temporal pooling (ATPN), attentive spatial pooling (ASPN), and the coexistence of them (ASTPN). The related CMC curves on iLIDS-VID, PRID-2011 and MARS are presented in Figure 5 re-

Table 3. Performance comparison with CMC Rank accuracy on MARS (%).

| Dataset | MARS | | | |
|----------|------|----|----|----|
| | 1 | 5 | 10 | 20 |
| CMC Rank | | | | |
| RNN-CNN | 40 | 64 | 70 | 77 |
| ASTPN | 44 | 70 | 74 | 81 |

spectively.

ATPN: the overall performance of ATPN curve is obviously better than the RNN-CNN method in Figure 5(a) and 5(b). For example, ATPN curve exceeds the RNN-CNN method by almost 10% at the rank 2 accuracy on PRID-2011. Meanwhile, on iLIDS-VID ATPN curve also outperforms the RNN-CNN method by 5% at the rank 3 accuracy. We can safely conclude that ATPN can efficiently utilize temporal human appearance to form powerful sequence-level representation, which is more subtle and discriminative than the output of simple pooling strategies (max-pooling and mean-pooling).

ASPEN exhibits equally prominent capability of matching compared with ATPN on iLIDS-VID. Although it is less robust than ATPN on PRID-2011, distinct margin still exists between ASPEN curve and the RNN-CNN method. We may reason that ASPEN, attentive spatial pooling network, mainly leverages relevant contextual information to enhance the discriminative power of the final representation. However, as we have mentioned about these two datasets, iLIDS-VID was created in a rather complicated environment, which means the contextual information could be more valuable clue due to the ambiguity of human appearance. On the contrary, ASPEN thus doesn't perform as competitively as ATPN in Figure 5(b).

ASTPN: combining ATPN and ASPEN together, ASTPN can capitalize on frame-wise interactions effectively as well as selectively propagate additional contextual information through the network. Based on Figure 5(c), where apparent distinction between ATPN curve and ASTPN curve can be observed with overall accuracy decreasing caused by MARS, ASTPN exceeds ATPN by about 5% at rank 3 point. It's proven that ASTPN is a more robust joint method especially on dataset as challenging as MARS.

4.4. Cross-Dataset Testing

Data bias is inevitable since a particular dataset only represents a small fragment of data of whole real world. The machine-learning model trained on A dataset would perform much worse when tested on B dataset. It can be regarded as over-fitting to the particular scenario, thus reducing the generality of the model. Cross-data testing is designed to evaluate the model's potentials in practical application.

The settings are introduced as follows: Both ASTPN and

Table 4. Cross-dataset testing matching rate on PRID-2011 (%). * indicates that both probe set and gallery set are composed of single image during test.

| Model | Trained on | 1 | 5 | 10 | 20 |
|----------|------------|----|----|----|----|
| ASTPN | iLIDS-VID | 30 | 58 | 71 | 85 |
| ASTPN* | iLIDS-VID | 15 | 33 | 46 | 63 |
| RNN-CNN | iLIDS-VID | 28 | 57 | 69 | 81 |
| RNN-CNN* | iLIDS-VID | 14 | 31 | 45 | 61 |

RNN-CNN are trained on diverse iLIDS-VID dataset, and then are tested on 50% of the PRID-2011 dataset. Apart from distinction brought out by cross-dataset training, the contrast of single-shot method and multi-shot method is also shown in Table 4. Although results are much worse than Table 2, ASTPN still achieves 30% on rank 1 accuracy, close to SRID [10] which is trained on PRID-2011 with the rank 1 accuracy of 35%. Moreover, using video-based re-id seems to improve the scores of both models by 100% than using single-shot re-id in terms of rank 1 score. It can be concluded that the valuable temporal information provided by video-based re-id really enhance the generalization performance of re-id networks greatly.

5. Conclusion

We proposed ASTPN, a novel deep architecture with jointly attentive spatial-temporal pooling for video-based person re-identification, enabling a joint learning of the representations of the inputs as well as their similarity measurement. ASTPN extends the standard RNN-CNNs by decomposing pooling into two steps: a spatial-pooling on feature map from CNN and an attentive temporal-pooling on the output of RNN. In effect, explicit or implicitly attention is performed at each pooling stage which select key regions or frames over the sequences for the feature representation learning.

Extensive experiments on iLIDS-VID, PRID-2011 and MARS have demonstrated that ASTPN significantly outperforms standard max and temporal pooling approaches. In particular, by executing control experiments, we show the joint pooling power than either of spatial/temporal pooling separately. Additionally, ASTPN is simple to implement and introduces little computational overhead compared to general max pooling, which makes it a desirable design choice for deep RNN-CNNs used in person re-identification in future. We would also consider to apply the current method into target tracking/detection systems [26].

6. Acknowledgment

This research is supported by NSFC No. 61401169 with principal investigator Pan Zhou.

References

- [1] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473, 2014. 2, 3
- [2] A. Barhillel, T. Hertz, N. Shental, and D. Weinshall. Learning a mahalanobis metric from equivalence constraints. *JMLR*, pages 937–965, 2005. 2
- [3] Y. Cheng, L. M. Brown, Q. Fan, R. S. Feris, S. Pankanti, and T. Zhang. Riskwheel: Interactive visual analytics for surveillance event detection. In *IEEE International Conference on Multimedia and Expo, ICME 2014, Chengdu, China, July 14-18, 2014*, pages 1–6, 2014. 2
- [4] Y. Cheng, Q. Fan, S. Pankanti, and A. Choudhary. Temporal sequence modeling for video event detection. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014. 2
- [5] C. N. dos Santos, M. Tan, B. Xiang, and B. Zhou. Attentive pooling networks. *CoRR*, abs/1602.03609, 2016. 2, 3, 4
- [6] M. Farenzena, L. Bazzani, A. Perina, V. Murino, and M. Cristani. Person re-identification by symmetry-driven accumulation of local features. In *IEEE CVPR*, pages 2360–2367, 2010. 2
- [7] K. He, X. Zhang, S. Ren, and J. Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *CoRR*, abs/1406.4729, 2014. 3
- [8] M. Hirzer, C. Beleznai, P. M. Roth, and H. Bischof. Person re-identification by descriptive and discriminative classification. In *scandinavian conference on image analysis*, pages 91–102, 2011. 5, 6
- [9] S. Karanam, Y. Li, and R. Radke. Person re-identification with discriminatively trained viewpoint invariant dictionaries. In *IEEE ICCV*, pages 4516–4524, 2015. 7
- [10] S. Karanam, Y. Li, and R. Radke. Sparse re-id: Block sparsity for person re-identification. In *IEEE CVPR Workshops*, pages 33–40, 2015. 7, 8
- [11] I. Kviatkovsky, A. Adam, and E. Rivlin. Color invariants for person reidentification. *IEEE TPAMI*, 35(7):1622–1634, 2013. 1, 2
- [12] Y. Li, Z. Wu, S. Karanam, and R. Radke. Multi-shot human re-identification using adaptive fisher discriminant analysis. In *BMVC*, 2015. 6, 7
- [13] Z. Li, S. Chang, F. Liang, T. S. Huang, L. Cao, and J. R. Smith. Learning locally-adaptive decision functions for person verification. In *IEEE CVPR*, pages 3610–3617, 2013. 2
- [14] S. Liao and S. Z. Li. Efficient psd constrained asymmetric metric learning for person re-identification. In *IEEE ICCV*, pages 3685–3693, 2015. 1, 2
- [15] C. Liu, S. Gong, and C. C. Loy. Person re-identification: what features are important? In *IEEE ICCV*, pages 391–401, 2012. 1, 2
- [16] H. Liu, J. Feng, M. Qi, J. Jiang, and S. Yan. End-to-end comparative attention networks for person re-identification. *CoRR*, abs/1606.04404, 2016. 3
- [17] K. Liu, B. Ma, W. Zhang, and R. Huang. A spatio-temporal appearance representation for video-based pedestrian re-identification. In *IEEE CVPR*, pages 3810–3818, 2015. 6
- [18] B. D. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *IJCAI*, pages 674–679, 1981. 6
- [19] B. Ma, Y. Su, and F. Jurie. Local descriptors encoded by fisher vectors for person re-identification. In *IEEE ICCV*, pages 413–422, 2012. 1, 2
- [20] N. McLaughlin, J. M. Rincon, and P. Miller. Recurrent convolutional network for video-based person re-identification. In *IEEE CVPR*, pages 1325–1334, 2016. 1, 3, 5, 6, 7
- [21] S. Paisitkriangkrai, C. Shen, and A. V. D. Hengel. Learning to rank in person re-identification with metric ensembles. In *IEEE CVPR*, pages 1846–1855, 2015. 1, 2
- [22] D. G. S, S. Brennan, and H. Tao. Evaluating appearance models for recognition, reacquisition, and tracking. *PETS*, 3, 2007. 6
- [23] S. Sharma, R. Kiros, and R. Salakhutdinov. Action recognition using visual attention. *CoRR*, abs/1511.04119, 2015. 3
- [24] A. Subramaniam, M. Chatterjee, and A. Mittal. Deep neural networks with inexact matching for person re-identification. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in NIPS*, pages 2667–2675. Curran Associates, Inc., 2016. 1
- [25] R. R. Varior, B. Shuai, J. Lu, D. Xu, and G. Wang. A siamese long short-term memory architecture for human re-identification. In *ECCV*, pages 135–153, 2016. 1
- [26] J. Wang, Y. Cheng, and R. Schmidt Feris. Walk and learn: Facial attribute representation learning from egocentric video and contextual data. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 8
- [27] T. Wang, S. Gong, X. Zhu, and S. Wang. Person re-identification by video ranking. In *ECCV*, pages 688–703, 2014. 2, 5, 6
- [28] X. Wang, G. Doretto, T. Sebastian, J. Rittscher, and P. Tu. Shape and appearance context modeling. In *IEEE ICCV*, pages 1–8, 2007. 2
- [29] K. Q. Weinberger and L. K. Saul. Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research*, pages 207–244, 2009. 2
- [30] F. Xiong, M. Gou, O. Camps, and M. Szaier. Person re-identification using kernel-based metric learning methods. In *ECCV*, pages 1–16, 2014. 1, 2
- [31] K. Xu, J. Ba, R. Kiros, K. Cho, A. C. Courville, R. Salakhutdinov, R. S. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, pages 2048–2057, 2015. 2, 3
- [32] Y. Yan, B. Ni, Z. Song, C. Ma, Y. Yan, and X. Yang. Person re-identification via recurrent feature aggregation. In *ECCV*, pages 701–716, 2016. 1, 3, 6, 7
- [33] D. Yi, Z. Lei, S. Liao, and S. Z. Li. Deep metric learning for person re-identification. In *IEEE CVPR*, pages 24–39, 2014. 1, 3
- [34] W. Yin, H. Schütze, B. Xiang, and B. Zhou. ABCNN: attention-based convolutional neural network for modeling sentence pairs. *TACL*, 4:259–272, 2016. 2, 3, 4

- [35] J. You, A. Wu, X. Li, and W. Zheng. Top-push video-based person re-identification. In *:IEEE CVPR*, pages 1345–1253, 2016. [2](#)
- [36] Z. Zhang, Y. chen, and V. Saligrama. Group membership prediction. In *:IEEE ICCV*, pages 3916–3924, 2015. [1](#), [2](#)
- [37] R. Zhao, W. Ouyang, and X. Wang. Person re-identification by salience matching. In *:IEEE ICCV*, pages 2528–2535, 2013. [1](#), [2](#)
- [38] R. Zhao, W. Ouyang, and X. Wang. Learning mid-level filters for person re-identification. In *:IEEE CVPR*, pages 144–151, 2014. [1](#), [2](#)
- [39] L. Zheng, Z. Bie, Y. Sun, J. Wang, C. Su, S. Wang, and Q. Tian. Mars: A video benchmark for large-scale person re-identification. In *ECCV*, pages 868–884, 2016. [5](#), [7](#)
- [40] L. Zheng, L. Shen, L. Tian, S. Wang, J. Wang, J. Bu, and Q. Tian. Scalable person re-identification: a benchmark. In *:IEEE ICCV*, pages 1116–1124, 2015. [2](#)
- [41] W. S. Zheng, S. Gong, and T. Xiang. Reidentification by relative distance comparison. *IEEE TPAMI*, 35(3):653–668, 2013. [1](#), [2](#)