

STA: Spatial-Temporal Attention for Large-Scale Video-Based Person Re-Identification

Yang Fu,^{1,2} Xiaoyang Wang,² Yunchao Wei,¹ Thomas Huang¹

¹IFP, Beckman, UIUC, IL, ²Nokia Bell Labs, Murray Hill, NJ
 {yangfu2, yunchao, t-huang1}@illinois.edu
 xiaoyang.wang@nokia-bell-labs.com

Abstract

In this work, we propose a novel Spatial-Temporal Attention (STA) approach to tackle the large-scale person re-identification task in videos. Different from the most existing methods, which simply compute representations of video clips using frame-level aggregation (*e.g.* average pooling), the proposed STA adopts a more effective way for producing robust clip-level feature representation. Concretely, our STA fully exploits those discriminative parts of one target person in both spatial and temporal dimensions, which results in a 2-D attention score matrix via inter-frame regularization to measure the importances of spatial parts across different frames. Thus, a more robust clip-level feature representation can be generated according to a weighted sum operation guided by the mined 2-D attention score matrix. In this way, the challenging cases for video-based person re-identification such as pose variation and partial occlusion can be well tackled by the STA. We conduct extensive experiments on two large-scale benchmarks, *i.e.* MARS and DukeMTMC-VideoReID. In particular, the mAP reaches **87.7%** on MARS, which significantly outperforms the state-of-the-arts with a large margin of more than **11.6%**.

Introduction

Person re-identification (Re-ID) aims at matching images of a person in one camera with the images of this person from another different camera. In recent years, person Re-ID under the image sequence (video) setting has drawn significant attention in literature since it is critical for applications like intelligent video surveillance and multimedia. (McLaughlin, Martinez del Rincon, and Miller 2016; Wang et al. 2016; Zhou et al. 2017; Liu, Yan, and Ouyang 2017; Jimin Xiao 2019).

Most of the existing video-based person Re-ID work (Liu, Yan, and Ouyang 2017; Song et al. 2017; Li et al. 2018) focuses on very small datasets, *e.g.* PRID-2010 (Hirzer et al. 2011) and iLIDS-VID (Wang et al. 2014) that only contains about 300 person identities with 600 tracklets in total. Although the existing approaches achieve good performances on PRID-2010 and iLIDS-VID, their accuracies in the state-of-the-art large-scale datasets such as

the MARS dataset (Zheng et al. 2016) and DukeMTMC-VideoReID dataset (Wu et al. 2018; Ristani et al. 2016) are far from satisfying. Their performances are largely limited by the tremendous variations in camera viewpoints, human poses, illumination, occlusions, and background clutter in the large-scale video-based Re-ID dataset.

For the video-based Re-ID task, the key is to learn a mapping function that converts videos into a low-dimensional feature space in which each video can be represented by a single vector. Most existing methods represent a frame of the video as a feature vector, and then employ average or maximum pooling across the frames to obtain the representation of the input video (Hirzer et al. 2011; You et al. 2016). However, this approach usually fails when occlusions appear frequently in the video. In addition, basic operations like maximum or average pooling across video frames cannot handle the spatial misalignment caused by the variation of human poses among frames. In order to distill the relevant information from a video and weaken the influence of noisy samples (*e.g.* occlusion), recent studies (Zhou et al. 2017; Xu et al. 2017) introduce the attention mechanism and achieve improved results. However, these existing attention-based methods only assign an attention weight to each frame, thus lack the capabilities of discovering either the discriminative frames in a video sequence or the discriminative body parts in each frame. Also, the attention mechanisms in most existing attention-based approaches are parametric, *e.g.* fully connected layers, which require the input video sequence to be in fixed length.

In order to address the issues mentioned above, we propose an effective yet easy-to-implement **Spatial-Temporal Attention (STA)** framework to address the large-scale video-based person re-identification problem, as shown in Fig. 1. Instead of simply encoding a sequence of images by pooling or assigning weights to each frame by the parametric model, our STA framework jointly incorporates multiple novel components including frame selection, discriminative parts mining, and feature aggregation without using any additional parameters. In all, our major contributions in this paper can be summarized as follows:

- We propose a simple yet effective STA model that assigns attention score for each spatial region to achieve discriminative parts mining and frame selection without using any additional parameters.

Yang Fu did this work during the intern at Nokia Bell Labs. Xiaoyang Wang and Yunchao Wei are joint corresponding authors. Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

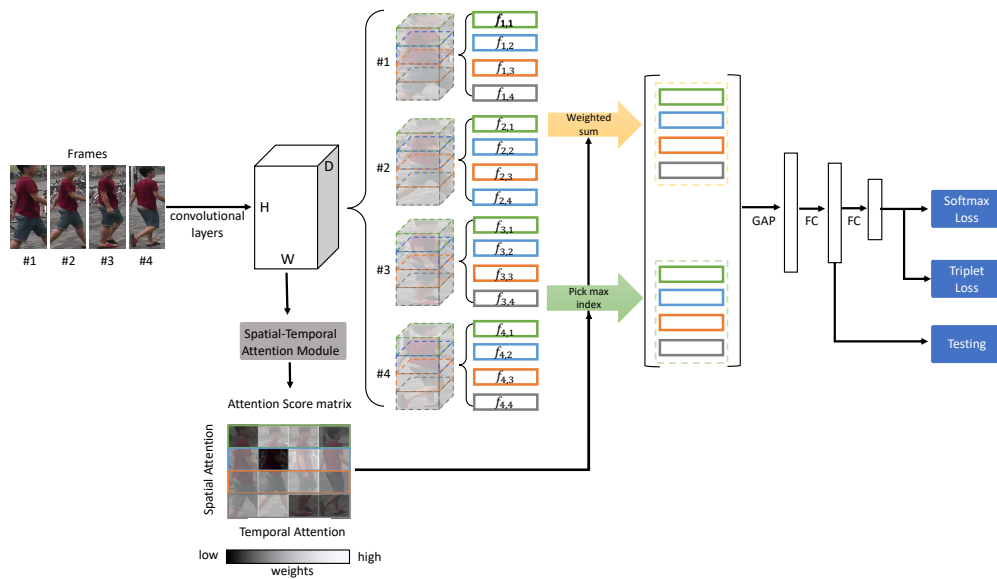


Figure 1: **Architecture of STA framework.** The input video tracklet is first reduced to N frames by random sampling. (1) Each selected frame is fed into the backbone network to be transformed into feature maps. (2) Then, the feature maps are sent to our proposed spatial-temporal attention model to assign an attention score to each spatial region of different frames and then generate a 2-D attention score matrix. An Inter-Frame Regularization is employed to restrict the difference among different frames (detailed in Fig. 2). (3) Using the attention score, we extract the spatial region feature map with the highest attention score among all frames, and operate the weighted sum of the spatial region feature maps based on the assigned attention score. (4) Then, we adopt feature fusion strategy to concatenate the spatial feature maps from different spatial regions to generate two sets of feature maps of the whole person body as the global representation and discriminative representation. (5) Finally, a global pooling layer and a fully connected layer are used to transform the feature maps to a vector for person re-identification. During training, we combine both the triplet loss and the softmax loss. During testing, we choose the feature vector after the first fully connected layer as the representation for the input video tracklet.

- We introduce the inter-frame regularization term to restrict the dissimilarity among different frames and ensure that each frame shares the same identity.
- We design a novel feature fusion strategy that combines both global information and discriminative information from the video sequence for better feature aggregation.
- We conduct extensive experiments and ablation study to demonstrate the effectiveness of each component. The final results achieve the state-of-the-art on two mainstream large-scale datasets: MARS and DukeMTMC-VideoReID.

Related Work

Image-based person re-identification is extensively explored in the literature, and the existing studies can generally be divided into two categories: discriminative learning (Xiao et al. 2016; Zheng, Zheng, and Yang 2017; Sun et al. 2017; Fu et al. 2019) and metric learning (Ahmed, Jones, and Marks 2015; Hermans, Beyer, and Leibe 2017; Ding et al. 2015). In (Hermans, Beyer, and Leibe 2017), Hermans *et al.* propose a variant of triplet loss to perform end-to-end deep metric learning, and their model can outperform many other published methods by a large margin. In (Zheng, Zheng, and Yang 2017), Zheng *et al.* employ siamese net-

work and combine both verification loss and classification losses to learn a discriminative embedding and a similarity measurement at the same time.

Video-based person re-identification is an extension of the image-based person re-identification, and is widely studied recently (McLaughlin, Martinez del Rincon, and Miller 2016; Liu et al. 2017; Wang et al. 2016; Zheng et al. 2016). For example, McLaughlin *et al.* (McLaughlin, Martinez del Rincon, and Miller 2016) adopt Recurrent Neural Network (RNN) to pass the message of each frame extracted from Convolution Network (CNN). Liu *et al.* (Liu et al. 2017) focus on learning a long-range motion context features from adjacent frames for a robust identification.

Attention models in person re-identification. Since Xu *et al.* (Xu et al. 2015) propose the attention mechanism, it has been applied to lots of person re-identification work (Zhou et al. 2017; Xu et al. 2017; Liu, Yan, and Ouyang 2017; Li et al. 2018). In (Liu, Yan, and Ouyang 2017), Liu *et al.* propose a method to estimate quality score of each frame automatically and weaken the influence of noisy samples. Xu *et al.* (Xu et al. 2017) introduce the joint Spatial and Temporal Attention Pooling Network that can extract the discriminative frames from probe and gallery videos, and obtain temporal attention weights for one sequence guided by the features of the other sequences. Li *et*

al. (Li et al. 2018) employ multiple spatial attention models with the temporal attention model to learn latent representations of different body parts of each person. Comparing to these existing attention models, our proposed STA model has two major differences that further boost our person Re-ID performance: First, the simple yet effective STA model has no additional parameters, which means that the length of the input sequence does not have to be fixed. Second, the STA model can learn an attention score for each region in different frames, which can achieve the discriminative region mining and frame selection jointly.

Proposed Method

Given a tracklet of person sequence, we propose the STA framework (Fig. 1) to better handle the video-based person re-identification problem through spatial-temporal attention model with inter-frame regularization. We first randomly select a constant number of frames from the input video query, and feed them into a backbone network to extract features from each frame. Then, we send the obtained feature maps into the proposed STA model to generate a 2-D attention score matrix which assigns an attention weight for each spatial region of each frame. In order to restrict the difference among the frames in the single video tracklet, we propose the inter-frame regularization to estimate the inter-frame similarity. Next, we use both the spatial region with the maximum corresponding weights in each frame, and the weighted sum among all attention weights to obtain two sets of feature maps of the whole person body. Finally, we concatenate them together as the global representation and discriminative representation, and employ a global average pooling followed by a fully connected layer to represent the video query. For the objective function, we combine the softmax loss and batch-hard triplet loss (Hermans, Beyer, and Leibe 2017).

Spatial-Temporal Attention (STA) Framework

Backbone Network. Various network architectures, like VGG (Simonyan and Zisserman 2014), Resnet (He et al. 2016), and Google Inception (Szegedy et al. 2016), can be employed as the backbone network to extract feature maps for each frame. We choose ResNet50 (He et al. 2016) as the backbone network, which was adopted by most previous works. In particular, ResNet50 has one convolutional block named *conv1*, and followed by four residual block named *conv2*, 3, 4, 5 respectively. We further make two modifications on the original ResNet50: 1) the stride of the first residual block *conv5* is set to 1; 2) the average pooling layer and fully connected layer are removed. The input video is first reduced to N frames by random sampling, and each selected frame is fed into the backbone network. As a result, each video $V = \{I_1, \dots, I_n, \dots, I_N\}$ is represented by a set of 16×8 feature maps $\{f_n\}_{n=1:N}$, and each feature map has $D = 2048$ channels.

Spatial-Temporal Attention Model. We propose the spatial-temporal attention model to automatically learn from each image frame the discriminative regions that are useful for re-identification. Previous video-based person re-identification methods (Liu, Yan, and Ouyang 2017; Zhou

et al. 2017) consider each frame as a whole image and assign one weight for each frame. However, different regions of a person body should have different influences on the re-identification task. Thus, our approach aims to discover the discriminative representation of these region for each frame. Li et al. (Li et al. 2018) also employ a spatial-temporal attention model, in which they use different convolutional layers to extract salient region of person body and adopt traditional temporal attention model for frame selection. There are three major drawbacks of this method. First, it involves more computation because of more convolutional layers, and its input sequence length has to be fixed due to the temporal attention model. Second, multiple spatial attention models used in their approach are independent from each other, without utilizing the spatial relationships that exist between human body parts. As a result, the extracted spatial attentions could be scattered and do not reflect the complete human body in the foreground. Third, the spatial attention information and temporal attention information are obtained by two different models which would cause error accumulation. Different from the existing methods, our spatial-temporal attention assigns attention weights, which contain both of spatial attention information and temporal attention information, to each spatial region in different frames automatically without any additional parameters. Experiments in Table 4 demonstrate the advantages of our method compared to (Li et al. 2018). To the best of our knowledge, our model is the first video-based person Re-ID model that can discover the discriminative parts but reserve the spatial relationship, and achieve frame selection at the same time.

The illustration of the spatial-temporal attention model is shown in Fig. 2. Given the feature maps of an input video $\{f_n\}_{n=1:N}$, we first generate the corresponding attention map g_n by operating the ℓ_2 normalization on the square sum through the depth channel. Specifically,

$$g_n(h, w) = \frac{\|\sum_{d=1}^{d=D} f_n(h, w, d)^2\|_2}{\sum_{h,w} \|\sum_{d=1}^{d=D} f_n(h, w, d)^2\|_2} \quad (1)$$

where H, W are the height and width of feature maps. Thus, each frame has one corresponding attention map. Both the feature maps and the attention maps of the N frames are then divided into K blocks horizontally:

$$\begin{cases} g_n = [g_{n,1}, \dots, g_{n,k}, \dots, g_{n,K}] \\ f_n = [f_{n,1}, \dots, f_{n,k}, \dots, f_{n,K}] \end{cases} \quad (2)$$

Here, $g_{n,k}$ represents the spatial attention map on k th regions of n th frame. After that, we employ the ℓ_1 normalization on all values in each block to obtain one spatial attention score for that region.

$$s_{n,k} = \sum_{i,j} \|g_{n,k}(i, j)\|_1 \quad (3)$$

Since the feature maps are after ReLU activation and all the values are greater than or equal to zero, the higher response of an attention map means the better representation of the person for re-identification task. The same procedure is operated on all selected frames of the input video to obtain the $N \times K$ matrix S of spatial attention scores.

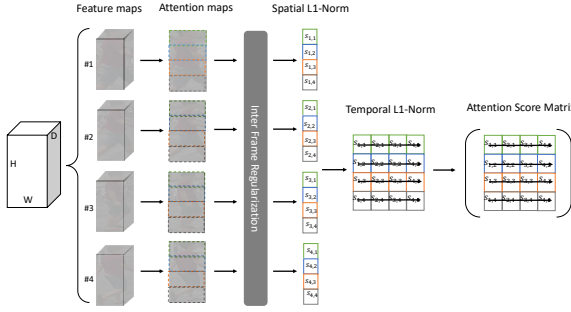


Figure 2: **Details of the spatial-temporal attention model with inter-frame regularization.** Given a set of feature maps from the input video, we generate corresponding attention maps for each frame. The inter-frame regularization is used to restrict the difference among frames in the same video tracklet. Then, attention maps are horizontally split into four equal spatial region, and the spatial region from the same spatial region but different frames are used to calculate the 2-D attention score matrix

Instead of using multiple convolutional layers to formulate the temporal attention model as in (Li et al. 2018; Zhou et al. 2017), we directly compare the attention scores that are from different frames but on the same spatial region, and compute each attention over the ℓ_1 normalization among them to obtain the normalized spatial-temporal attention scores. Specifically,

$$S(n, k) = \frac{s_{n,k}}{\sum_n \|s_{n,k}\|_1} \quad (4)$$

As a result, each spatial region from different frames is assigned with a specific attention score based on the spatial-temporal attention information.

Inter-Frame Regularization

For the video-based person Re-ID, the images from the same video tracklet of a person should represent the appearance of the same person. Such information is further exploited by our approach as the inter-frame regularization to restrict the difference of the learned attention maps among frames. This inter-frame regularization helps to avoid the cases in which the learned attention scores of each spatial region focus on one specific frame and largely ignore the other frames.

Specifically, since each frame has a corresponding feature map f_n , which is used to classify the person identification during training. One possible way is to add a classification loss to all frames to make sure they share the same identification. However, there maybe some noisy samples which are hard to classify and hence make the training processing unstable. An alternative solution is to use Kullback-Leibler (KL) divergence to evaluate the similarity of each frame, but lots of close-to-zero elements exist in attention maps. These elements will drop dramatically when employing the \log operation in the KL divergence, and make the training processing unstable as well (Lin et al. 2017). Thus, in order to

encourage the spatial-temporal attention model to preserve the similarity and meanwhile avoid focusing on one frame, we design the inter-frame regularization which measures the difference among input image frames. For convenience, we define G as the collection of attention maps generated from the input image frames,

$$G = [g_1, \dots, g_N] \quad (5)$$

Assume g_i, g_j are attention maps calculated by Eqn. (1) for two frames i and j . We employ the square Frobenius Norm (Meyer 2000) of the difference between g_i and g_j . Specifically,

$$\begin{aligned} Reg &= \|g_i - g_j\|_F \\ &= \sqrt{\sum_{h=1}^H \sum_{w=1}^W |g_i(h, w) - g_j(h, w)|^2} \end{aligned} \quad (6)$$

Note that we randomly choose two frames i and j from the N frames of each video for this regularization term. In order to restrict the difference between two frames, we minimize this regularization term Reg by adding it to the original objective function L_{total} defined in Eqn. (11) after multiplied by a coefficient λ .

$$\min(L_{total} + \lambda Reg) \quad (7)$$

Feature Fusion Strategy

After STA model with inter-frame regularization, we obtain an $N \times K$ matrix S that assigns an attention score $s_{n,k}$ for the feature map $f_{n,k}$ of each spatial region and each frame. Inspired by (Fu et al. 2019), we propose a strategy for feature fusion by combining the global and discriminative information of each tracklet as described in Alg. 1.

Given the attention score matrix and a set of feature maps, we first divide feature maps into several spatial regions just like what we operate on attention map, and pick the spatial region that has the highest corresponding score compared to other frames. Then, we repeat this operation for every spatial region and concatenate those regions together to obtain a feature map that contains the most discriminative regions of input frames. Next, we use every attention score as a weight and employ the element-wise multiplication on every split feature map to generate another feature map with the global information of input frames. Finally, we concatenate these two feature maps together and employ the global average pooling followed by a fully connected layer to generate the representation vector X for the Re-ID task.

$$X = [x_1, x_2, \dots, x_n] \quad (8)$$

Loss Function

In this paper, we utilize both the batch-hard triplet loss proposed in (Hermans, Beyer, and Leibe 2017) and the softmax loss jointly to train the STA model through the combination the metric learning and discriminative learning.

The triplet loss with hard mining is first proposed in (Hermans, Beyer, and Leibe 2017) as an improved version of the original semi-hard triplet loss (Schroff, Kalenichenko,

Algorithm 1: Algorithm of Feature Fusion Strategy

Input : A set of feature maps $\{f_n\}_{n=1:N}$;
An attention score matrix S ;
Size of Feature map, H, W, D ;
Length of Sequence N ;
Number of Spatial blocks K ;

Output: Feature map after fusion F_{fuse} ;

Initialization : $F_1 = \text{Zeros}(H, W, D)$,
 $F_2 = \text{Zeros}(H, W, D)$; $H_s = \lfloor H/4 \rfloor$;

for $n = 1 : N$ **do**
 for $k = 1 : K$ **do**
 $f_{n,k} \leftarrow f_n(H_s * (k - 1) : H_s * k, :, :)$;
 end
end
for $k = 1 : K$ **do**
 $m \leftarrow$ the index of the maximum value in $S(:, k)$;
 $F_1(H_s : H_s, :, :) \leftarrow f_{mk}$;
 for $n = 1:N$ **do**
 $F_2(H_s * (k - 1) : H_s * k, :, :) += f_{nk} * s_{nk}$;
 end
end
 $F_{fuse} \leftarrow [F_1, F_2]$;
return F_{fuse} ;

and Philbin 2015). We randomly sample P identities and K tracklets for each mini-batch to meet the requirement of the batch-hard triplet loss. Typically, the loss function is formulated as follows:

$$L_{triplet} = \sum_{i=1}^P \sum_{a=1}^K [\alpha + \underbrace{\max_{p=1\dots K} \|x_a^{(i)} - x_p^{(i)}\|_2}_{\text{hardest positive}} - \underbrace{\min_{\substack{n=1\dots K \\ j=1\dots P \\ j \neq i}} \|x_a^{(i)} - x_p^{(i)}\|_2}_{\text{hardest negative}}]_+ \quad (9)$$

where $x_a^{(i)}, x_p^{(i)}, x_n^{(i)}$ are features extracted from the anchor, positive and negative samples respectively, and α is the margin hyperparameter to control the differences of intra and inter distances. Here, positive and negative samples refer to the person with same or different identity from the anchor.

Besides batch-hard triplet loss, we employ softmax cross entropy loss for discriminative learning as well. The original softmax cross entropy loss can be formulated as follows:

$$L_{softmax} = - \sum_{i=1}^P \sum_{a=1}^K \log \frac{e^{W_{y_{i,a}}^T x_{a,i}}}{\sum_{k=1}^C e^{W_k^T x_{a,i}}} \quad (10)$$

where $y_{i,a}$ is the ground truth identity of the sample $\{a, i\}$, and C is number of classes. Our loss function for optimization is the combination of softmax loss and batch-hard triplet loss as follows:

$$L_{total} = L_{softmax} + L_{triplet} \quad (11)$$

Experiments

Datasets and Evaluation Protocol

Mars dataset (Zheng et al. 2016) is one of the largest video-based person re-identification dataset. It contains 17,503 tracklets from 1,261 identities, and additional 3,248 tracklets serving as distractors. These video tracklets are captured by six cameras in a university campus. The total 1,261 identities are split into 625 identities for training and 636 identities for testing. Every identity in the training set has 13 video tracklets on average, and each tracklet has 59 frames on average. The ground truth labels are detected and tracked using the Deformable Part Model (DPM) (Felzenszwalb, McAllester, and Ramanan 2008) and GMCP tracker (Zamir, Dehghan, and Shah 2012).

DukeMTMC-VideoReID dataset (Wu et al. 2018) is another large-scale benchmark dataset for video-based person Re-ID. It is derived from the DukeMTMC dataset (Ristani et al. 2016). The DukeMTMC-VideoReID dataset contains 4,832 tracklets from 1,812 identities, and it is split into 702, 702 and 408 identities for training, testing and distraction respectively. In total, it has 369,656 frames of 2,196 tracklets for training, and 445,764 frames of 2,636 tracklets for testing and distraction. Each tracklet has 168 frames on average. The bounding boxes are annotated manually.

Evaluation Protocol. In our experiments, we use the Cumulative Matching Characteristic (CMC) curve and the mean average precision (mAP) to evaluate the performance of the STA model. For each query, CMC represents the accuracy of the person retrieval. We report the Rank-1, Rank-5, Rank-20 scores to represent the CMC curve. The CMC metric is effective when each query corresponds to only one ground truth clip in the gallery. However, when multiple ground truth clips exist in the gallery, and the objective is to return to the user as many correct matches as possible, CMC would not effectively measure the performances of models on this objective. Comparatively, mAP is a comprehensive metric that is well-suited for both single-match and multiple-match objectives.

Implementation Details

As discussed in ‘‘Proposed Method’’, we first randomly select $N = 4$ frames from the input tracklet, and use the modified ResNet50 initialized on the ImageNet (Deng et al. 2009) dataset as the backbone network. The number of spatial regions is set to $K = 4$. And, each frame is augmented by random horizontal flipping and normalization. Each mini-batch is sampled with randomly selected P identities and randomly sampled K images for each identity from the training set. In our experiment, we set $P = 16$ and $K = 4$ so that the mini-batch size is 64. And, we recommend to set the margin parameter in triplet loss to 0.3. During training, we use the Adam (Kingma and Ba 2014) weight decay 0.0005 to optimize the parameters for 70 epochs. The overall learning rate is initialized to 0.0003 and decay to 3×10^{-5} and 3×10^{-6} after training for 200 and 400 epochs respectively. The total training process lasts for 800 epochs. For evaluation, we extract the feature vector after the first fully connected layer as the representation of query tracklet. Our model is imple-

Model	MARS				DukeMTMC-VideoReID			
	R1	R5	R10	mAP	R1	R5	R10	mAP
Baseline	74.5	88.8	91.8	64.0	79.1	93.9	96.0	76.8
Baseline + TL	80.8	92.1	94.3	74.0	90.6	95.8	96.7	89.7
Baseline + TL + Avg	82.5	92.9	94.9	75.0	91.8	97.4	98.0	91.0
Baseline + TL + STA	84.8	94.6	96.2	78.0	93.3	98.1	98.6	92.7
Baseline + TL + STA + Fusion	85.3	95.1	96.4	79.1	95.3	98.1	99.1	93.9
Baseline + TL + STA + Fusion + Reg	86.3	95.7	97.1	80.8	96.2	99.3	99.6	94.9

Table 1: Comparison of different proposed components, where TL, Avg, STA, Fusion, and Reg represent the triplet loss, average pooling, spatial-temporal attention module, feature fusion strategy, and inter-frame regularization respectively. R-1, R-5, R-10 accuracies (%) and mAP (%) are reported. Baseline model corresponds to ResNet50 trained with softmax loss on video datasets MARS or DukeMTMC-VideoReID respectively.

Sequence Length	MARS				DukeMTMC-VideoReID			
	R1	R5	R10	mAP	R1	R5	R10	mAP
N=2	81.7	93.8	95.7	75.7	90.3	97.6	98.6	89.0
N=4	86.3	95.7	97.1	80.8	96.2	99.3	99.6	94.9
N=6	86.2	95.7	96.9	81.0	96.0	99.4	99.7	95.0
N=8	86.2	95.7	97.1	81.2	96.0	99.3	99.6	95.0

Table 2: Performance comparison of STA model with different sequence lengths during testing on MARS dataset and Duke-VideoReID dataset. Here, we use the model trained with the sequence length of 4 and the spatial regions number of 4.

mented on Pytorch platform and trained with two NVIDIA TITAN X GPUs. All our experiments on different datasets follow the same settings as above.

Ablation Study

To verify the effectiveness of each component in STA model, we conduct several analytic experiments including w/ or w/o triplet loss, w/ or w/o spatial-temporal attention model, w/ or w/o feature fusion, and w/ or w/o inter-frame regularization. In addition, we carry out experiments to investigate the effect of varying the sequence length N and the number K of spatial regions. Note that all the remaining settings are the same as those discussed in ‘‘Implementation Details’’.

Effectiveness of Components. In Table 1, we list the results of each component in our STA framework. **Baseline** represents the ResNet50 model trained with softmax loss on MARS/DukeMTMC-VideoReID dataset. **TL** corresponds to the hard-batch triplet loss, and ‘‘+ TL’’ means the hard-batch triplet loss is combined with the softmax loss in Baseline model. Note that, the **Baseline** and **Baseline + TL** both treat each tracklet frame by frame, i.e., the image-based models. **STA** is our proposed spatial-temporal attention model in which the number of spatial regions is set to $K = 4$ and the input sequence length is set to $N = 4$ as well. It generates a 4×4 attention score matrix, and uses the score of the same region but from a different frame to calculate the weighted sum feature maps of each region. Compared to **Baseline + TL**, **STA** improves Rank-1 and mAP accuracy by 4.0% and 4.0% on MARS, as well as 2.7% and 2.4% on DukeMTMC-VideoReID respectively. These results show that the spatial-temporal attention model is very effective at discovering discriminative image regions which are useful for boosting re-identification performance. **Fusion** means aggregating feature representation by the proposed fusion strategy

described in Alg. 1. It is obvious that the proposed fusion strategy can further improve the performance by combining the most discriminative information and global information together. **Reg** refers to the proposed inter-frame regularization term. From the comparison of w/ and w/o **Reg**, we can find the Rank-1 accuracy and mAP improve by 1.0% and 1.7% on Mars, as well as 0.9% and 1.0% on DukeMTMC-VideoReID respectively. This improvement shows the proposed inter-frame regularization term can balance the frame diversity and thus further improve the performance.

Consistence of Sequence Length. In Table 2, we show the robustness of the trained model to different sequence length of input tracklet. For fair comparison, we use the model trained with sequence length of 4 and spatial region number of 4, and evaluate the performance with different sequence length: 2, 4, 6, and 8. As we can see, the performances are very consistent when input sequence length is 4, 6 or 8. (e.g. all the Rank-1 accuracies are above 85% and the mAP are above 80% on MARS.) This is because the proposed STA model does not involve more parameters, so there is no restriction on the sequence length. In addition, the performances with $N = 4$ surpass most of the state-of-the-art methods which usually need more frames for good representation of the input video.

Influence of Spatial Region Number. To investigate how the number of spatial regions influences the final performance, we conduct experiments with three different spatial regions: 2, 4, and 8 with the same sequence length 4. The results are listed in Table 3. In these experiments, the STA framework always achieves the best result with 4 spatial regions. Since the size of feature maps is 16×8 , with 2 or 8 spatial regions, it could be too coarse to contain enough information or too small to contain enough information for the re-identification task.

Number of Spatial Regions	MARS				DukeMTMC-VideoReID			
	R1	R5	R10	mAP	R1	R5	R10	mAP
K=2	85.3	95.1	96.6	80.3	94.7	99.0	99.6	93.8
K=4	86.3	95.7	97.1	80.8	96.2	99.3	99.6	94.9
K=8	85.5	95.3	96.9	80.4	95.2	99.1	99.4	93.8

Table 3: Performance comparison of the STA model trained with different number of spatial regions on MARS dataset and Duke-VideoReID dataset. Here, we keep the sequence length as a constant of 4.

Model	R1	R5	R20	mAP
CNN+Kiss.+MQ (Zheng et al. 2016)	68.3	82.6	89.4	49.3
SeeForest (Zhou et al. 2017)	70.6	90.0	97.6	50.7
Latent Parts (Li et al. 2017)	71.8	86.6	93.0	56.1
QAN (Liu, Yan, and Ouyang 2017)	73.7	84.9	91.6	51.7
K-reciprocal (Zhong et al. 2017)	73.9	–	–	68.5
TriNet (Hermans, Beyer, and Leibe 2017)	79.8	91.4	–	67.7
RQEN (Song et al. 2017)	77.8	88.8	94.3	71.1
CSACSE (Chen et al. 2018)	81.2	92.1	–	69.4
STAN (Li et al. 2018)	82.3	–	–	65.8
CSACSE (Chen et al. 2018) + Optical Flow	86.3	94.7	98.2	76.1
STA	86.3	95.7	98.1	80.8
STA + ReRank	87.2	96.2	98.6	87.7

Table 4: Comparison of the STA model with the state-of-the-arts on MARS dataset. Here, we show the results tested with sequence length of 4 and spatial region number of 4 as well.

Model	R1	R5	R20	mAP
ETAP-Net(supervised) (Wu et al. 2018)	83.6	94.6	97.6	78.3
STA	96.2	99.3	99.6	94.9

Table 5: Comparison with state-of-the-art on DukeMTMC-VideoReID. Here, we show the results tested with sequence length of 4 and spatial region number of 4 as well.

Comparison with the State-of-the-arts

Table 4 and Table 5 report the comparison of our proposed STA model with the state-of-the-art techniques. On each dataset, our approach achieves the best performance, especially on mAP. We attain R1/mAP: 86.3/80.8(87.2/87.7) on MARS before and after re-ranking. In addition, we achieve R1/mAP: 96.2/94.9 on DukeMTMC-VideoReID.

Results on MARS. Comparisons between our approach and the state-of-the-art approaches on MARS are shown in Table 4. The results show that our approach achieves **80.8%** in mAP, which surpasses all existing work by more than **4.0%**. Even for the Rank-1 and Rank-5, our approach achieves competitive results compared to the most recent work listed in Table 4. It’s noted that the CSACSE + Optical Flow method (Chen et al. 2018) incorporates optical flow as the extra information, which not only brings in more computation, but also causes the drawback that its whole network cannot be trained end-to-end. Comparing to other related work that also does not use optical flow, our approach improves the Rank-1 accuracy and mAP by **4.0%** and **15.0%** respectively. After implementing re-ranking, the Rank-1 accuracy and mAP can be improved to **87.2%** and **87.7%**, which outperforms the CSACSE + Optical Flow method (Chen et al. 2018) by **11.6%** on mAP.

Results on DukeMTMC-VideoReID. Table 5 shows the video-base Re-ID performance on DukeMTMC-VideoReID dataset. This dataset is new to the field, and there is only one published baseline (Wu et al. 2018) on this dataset. Comparing to this baseline, our approach improves more than **10%** on Rank-1 accuracy and mAP. Although there are few results reported on this dataset, we have good reason to believe that our approach works well because it achieves **96.2%** and **94.9%** on Rank-1 accuracy and mAP respectively.

Conclusion

This paper addresses the large-scale video-based person re-identification (Re-ID) problem with the proposed Spatial-Temporal Attention (STA) module. In STA, instead of directly extracting video representation through frame-level feature aggregation (*e.g.* average pooling), the 2-D spatial-temporal map is used to calculate the more robust clip-level feature representation without using any additional parameters. The inter-frame regularization and feature fusion strategy are proposed to further improve the clip-level representation. Extensive experiments conducted on two benchmarks including MARS and DukeMTMC-VideoReID demonstrate the effectiveness of the proposed module for large-scale video-based person Re-ID problem comparing to the existing methods. Our current work focuses on person Re-ID with person tracklets as inputs. A worthy future study would be applying the module in the multi-camera video systems.

Acknowledgements. Yang Fu is supported in part by CloudWalk at UIUC. Yunchao Wei is supported in part by IBM-ILLINOIS Center for Cognitive Computing Systems Research (C3SR) - a research collaboration as part of the IBM AI Horizons Network. Nokia owns the intellectual property rights of this work.

References

- Ahmed, E.; Jones, M.; and Marks, T. K. 2015. An improved deep learning architecture for person re-identification. In *IEEE ICCV*.
- Chen, D.; Li, H.; Xiao, T.; Yi, S.; and Wang, X. 2018. Video person re-identification with competitive snippet-similarity aggregation and co-attentive snippet embedding. In *IEEE CVPR*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *IEEE CVPR*.
- Ding, S.; Lin, L.; Wang, G.; and Chao, H. 2015. Deep feature learning with relative distance comparison for person re-identification. *Pattern Recognition*.
- Felzenszwalb, P.; McAllester, D.; and Ramanan, D. 2008. A discriminatively trained, multiscale, deformable part model. In *IEEE CVPR*.
- Fu, Y.; Wei, Y.; Zhou, Y.; Shi, H.; Huang, G.; Wang, X.; Yao, Z.; and Huang, T. 2019. Horizontal pyramid matching for person re-identification. In *AAAI*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *IEEE CVPR*.
- Hermans, A.; Beyer, L.; and Leibe, B. 2017. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*.
- Hirzer, M.; Beleznai, C.; Roth, P. M.; and Bischof, H. 2011. Person re-identification by descriptive and discriminative classification. In *SCIA*.
- Jimin Xiao, Yanchun Xie, T. T. K. H. Y. W. J. F. 2019. Ian: The individual aggregation network for person search. *Pattern Recognition*.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Li, D.; Chen, X.; Zhang, Z.; and Huang, K. 2017. Learning deep context-aware features over body and latent parts for person re-identification. In *IEEE CVPR*.
- Li, S.; Bak, S.; Carr, P.; and Wang, X. 2018. Diversity regularized spatiotemporal attention for video-based person re-identification. In *IEEE CVPR*.
- Lin, Z.; Feng, M.; Santos, C. N. d.; Yu, M.; Xiang, B.; Zhou, B.; and Bengio, Y. 2017. A structured self-attentive sentence embedding. *arXiv preprint arXiv:1703.03130*.
- Liu, H.; Jie, Z.; Jayashree, K.; Qi, M.; Jiang, J.; Yan, S.; and Feng, J. 2017. Video-based person re-identification with accumulative motion context. *IEEE TCSVT*.
- Liu, Y.; Yan, J.; and Ouyang, W. 2017. Quality aware network for set to set recognition. In *IEEE CVPR*.
- McLaughlin, N.; Martinez del Rincon, J.; and Miller, P. 2016. Recurrent convolutional network for video-based person re-identification. In *IEEE CVPR*.
- Meyer, C. D. 2000. *Matrix analysis and applied linear algebra*. Siam.
- Ristani, E.; Solera, F.; Zou, R.; Cucchiara, R.; and Tomasi, C. 2016. Performance measures and a data set for multi-target, multi-camera tracking. In *ECCV*.
- Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. Facenet: A unified embedding for face recognition and clustering. In *IEEE CVPR*.
- Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Song, G.; Leng, B.; Liu, Y.; Hetang, C.; and Cai, S. 2017. Region-based quality estimation network for large-scale person re-identification. *arXiv preprint arXiv:1711.08766*.
- Sun, Y.; Zheng, L.; Yang, Y.; Tian, Q.; and Wang, S. 2017. Beyond part models: Person retrieval with refined part pooling. *arXiv preprint arXiv:1711.09349*.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the inception architecture for computer vision. In *IEEE CVPR*.
- Wang, T.; Gong, S.; Zhu, X.; and Wang, S. 2014. Person re-identification by video ranking. In *ECCV*.
- Wang, T.; Gong, S.; Zhu, X.; and Wang, S. 2016. Person re-identification by discriminative selection in video ranking. *IEEE TPAMI*.
- Wu, Y.; Lin, Y.; Dong, X.; Yan, Y.; Ouyang, W.; and Yang, Y. 2018. Exploit the unknown gradually: One-shot video-based person re-identification by stepwise learning. In *IEEE CVPR*.
- Xiao, T.; Li, H.; Ouyang, W.; and Wang, X. 2016. Learning deep feature representations with domain guided dropout for person re-identification. In *IEEE CVPR*.
- Xu, K.; Ba, J.; Kiros, R.; Cho, K.; Courville, A.; Salakhudinov, R.; Zemel, R.; and Bengio, Y. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*.
- Xu, S.; Cheng, Y.; Gu, K.; Yang, Y.; Chang, S.; and Zhou, P. 2017. Jointly attentive spatial-temporal pooling networks for video-based person re-identification. *arXiv preprint arXiv:1708.02286*.
- You, J.; Wu, A.; Li, X.; and Zheng, W.-S. 2016. Top-push video-based person re-identification. In *IEEE CVPR*.
- Zamir, A. R.; Dehghan, A.; and Shah, M. 2012. Gmcp-tracker: Global multi-object tracking using generalized minimum clique graphs. In *ECCV*.
- Zheng, L.; Bie, Z.; Sun, Y.; Wang, J.; Su, C.; Wang, S.; and Tian, Q. 2016. Mars: A video benchmark for large-scale person re-identification. In *ECCV*.
- Zheng, Z.; Zheng, L.; and Yang, Y. 2017. A discriminatively learned cnn embedding for person reidentification. *ACM TOMM*.
- Zhong, Z.; Zheng, L.; Cao, D.; and Li, S. 2017. Re-ranking person re-identification with k-reciprocal encoding. In *IEEE CVPR*.
- Zhou, Z.; Huang, Y.; Wang, W.; Wang, L.; and Tan, T. 2017. See the forest for the trees: Joint spatial and temporal recurrent neural networks for video-based person re-identification. In *IEEE CVPR*.