

Temporal Knowledge Propagation for Image-to-Video Person Re-identification

Xinqian Gu^{1,2}, Bingpeng Ma², Hong Chang^{1,2}, Shiguang Shan^{1,2,3}, Xilin Chen^{1,2}

¹Key Laboratory of Intelligent Information Processing of Chinese Academy of Sciences (CAS),
Institute of Computing Technology, CAS, Beijing, 100190, China

²University of Chinese Academy of Sciences, Beijing, 100049, China

³CAS Center for Excellence in Brain Science and Intelligence Technology, Shanghai, 200031, China

xinqian.gu@vip1.ict.ac.cn, bpbma@ucas.ac.cn, {changhong, sgshan, xlchen}@ict.ac.cn

Abstract

In many scenarios of Person Re-identification (Re-ID), the gallery set consists of lots of surveillance videos and the query is just an image, thus Re-ID has to be conducted between image and videos. Compared with videos, still person images lack temporal information. Besides, the information asymmetry between image and video features increases the difficulty in matching images and videos. To solve this problem, we propose a novel Temporal Knowledge Propagation (TKP) method which propagates the temporal knowledge learned by the video representation network to the image representation network. Specifically, given the input videos, we enforce the image representation network to fit the outputs of video representation network in a shared feature space. With back propagation, temporal knowledge can be transferred to enhance the image features and the information asymmetry problem can be alleviated. With additional classification and integrated triplet losses, our model can learn expressive and discriminative image and video features for image-to-video re-identification. Extensive experiments demonstrate the effectiveness of our method and the overall results on two widely used datasets surpass the state-of-the-art methods by a large margin.

1. Introduction

Person Re-identification (Re-ID) aims to find the sample among the gallery set which has the same identity with a given query. In general, Re-ID problems fall into two categories: image-based Re-ID [1, 36, 12, 8] and video-based Re-ID [34, 32, 13]. The main difference is that the query and gallery in image-based Re-ID are both images, while the query and gallery in video-based Re-ID are both videos.

However, in many real-world scenarios, the gallery set is usually constituted by lots of surveillance videos, while the query only consists of one image, thus Re-ID has to

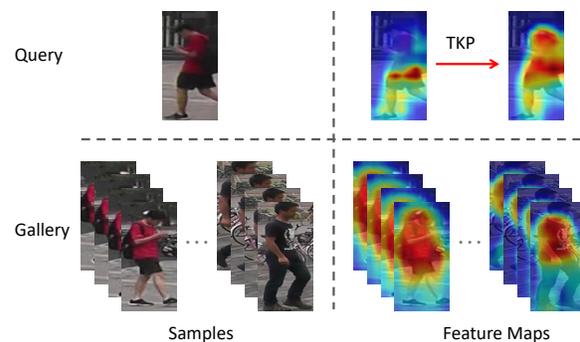


Figure 1: Visualisation of the feature maps. Due to lack of temporal information, compared with video features, the query image feature only pays attention to some local regions. With the proposed TKP method, temporal knowledge can be transferred to enhance the image feature and the learned image feature focuses on more foreground of person image.

be conducted between image and videos. One instance is rapidly locating and tracking a criminal suspect among a mass of surveillance videos according to one photo of the suspect (e.g., Boston marathon bombings event). Due to its crucial role in video surveillance system, image-to-video (I2V) Re-ID [44, 45] has attracted increasing attention in recent years.

In I2V Re-ID, the query is a still image, while the gallery of videos contains additional temporal information. Some research [37, 44] indicates that modeling temporal relations between video frames makes the gallery video features robust to disturbing conditions. However, these methods ignore that the query of I2V Re-ID only consists of one still image and lacks temporal information. As a result, on one hand, the image feature cannot benefit from the advantages of modeling temporal relations (see Figure 1). On the other hand, the information asymmetry between image and video features increases the difficulty in measuring the image-to-video similarity. Hence, it is essential and desirable to develop a method to supplement temporal information to im-

age feature representation.

In this paper, we propose a novel Temporal Knowledge Propagation (TKP) method to address the problems of ignored image temporal representation and information asymmetry in I2V Re-ID. This is inspired by knowledge distillation [11], which transfers dark knowledge from a large and powerful teacher network to a smaller and faster student network. In our TKP method, the temporal knowledge learned by video representation network is propagated to image representation network. During training, given the same input videos, we enforce the frame features extracted by the image representation network to match the outputs of video representation network in a shared feature space. After training with back propagation, the temporal knowledge can be naturally transferred from video representation network to image representation network. In the test stage, we use the trained image representation network to extract the query image features. Thanks to the transferred temporal knowledge, the extracted image features manifest robustness to disturbing conditions just like the video frame features (see Figure 1). Meanwhile, the information asymmetry problem between image and video features is addressed, thus it is much easier to measure the similarity between images and videos.

Extensive experiments validate the effectiveness of the proposed method. For instance, on MARS dataset, our method improves the performance from 67.1% to 75.6% (+8.5%) w.r.t. top-1 accuracy, surpassing the state-of-the-art methods by a large margin.

2. Related Work

I2V Re-ID. Recently, several relevant methods [31, 37, 44, 45] are proposed for I2V Re-ID task. Among them, Zhu *et al.* [44, 45] firstly investigate this problem and propose a heterogeneous dictionary pair learning framework to map image features and video features to a shared feature space. Wang *et al.* [31] try to use deep learning based methods to solve this problem and Zhang *et al.* [37] use LSTM to model the temporal information of gallery videos to enhance the robustness of the video features. However, these methods [44, 45, 37] neglect that there is no temporal information in query image features. In contrast, our proposed TKP method transfers temporal knowledge learned by video representation network to image representation network, which can effectively reduce the information asymmetry between image and video features.

Modeling Temporal Relations. Dealing with temporal relations between video frames is of central importance in video feature extraction. A natural solution is to apply Recurrent Neural Networks (RNN) to model sequences [5, 23, 24]. There are also some methods [7, 15, 6] that use 3D convolution (C3D) to process temporal neighborhoods.

However, both RNN and C3D only process one local neighborhood at a time. Recently, Wang *et al.* [33] propose to use non-local operation to capture long-range dependencies, which achieves higher results on video classification task. In this paper, we also attempt to utilize non-local operation to model temporal relations in person videos.

Knowledge Distillation. Knowledge distillation [2, 11, 25, 4] is a widely used technique to transfer knowledge from a teacher network to a student network. Generally, it is used to transfer from a powerful and large network to a faster and small network. In contrast, our TKP method is the first time to transfer temporal knowledge from video representation network to image representation network. Besides, instead of using a well-trained teacher, our image-to-video representation learning and temporal knowledge transferring are trained simultaneously.

As for the distillation forms, Hinton *et al.* [11, 40] minimize the Kullback-Leibler divergence of the final classification probability of teacher network and student network to transfer knowledge. In contrast, Bengio *et al.* [25] directly minimize the Mean Square Error (MSE) of the middle outputs of these two networks. For deep metric learning tasks, Chen *et al.* [4, 39] transfer knowledge via cross sample similarities. In this paper, we transfer temporal knowledge by minimizing the MSE of the image features and the corresponding video frame features in a shared feature space, which is similar to [25] in loss design but different in models. Besides, we also formulate the TKP loss based on cross sample distances in the shared feature space.

3. The Proposed Method

In this section, we first introduce the overall network architecture for the proposed TKP method. Secondly, the details of image representation network and video representation network are illustrated. Then, our TKP method is presented, followed by the final objective function and sampling strategy. Finally, these two learned networks are used to perform I2V Re-ID testing.

The framework for our proposed TKP method in I2V Re-ID training is depicted in Figure 2. Given input video clips, the image representation network extracts visual information of single-frame images, while the video representation network extracts visual information and deals with temporal relations between video frames simultaneously. Temporal knowledge propagation from the video representation network to image representation network is formulated by constructing the TKP loss. By minimizing the TKP loss, together with the classification and triplet losses, the image features and video features are mapped to a shared feature space. Details of the proposed method are given as follows.

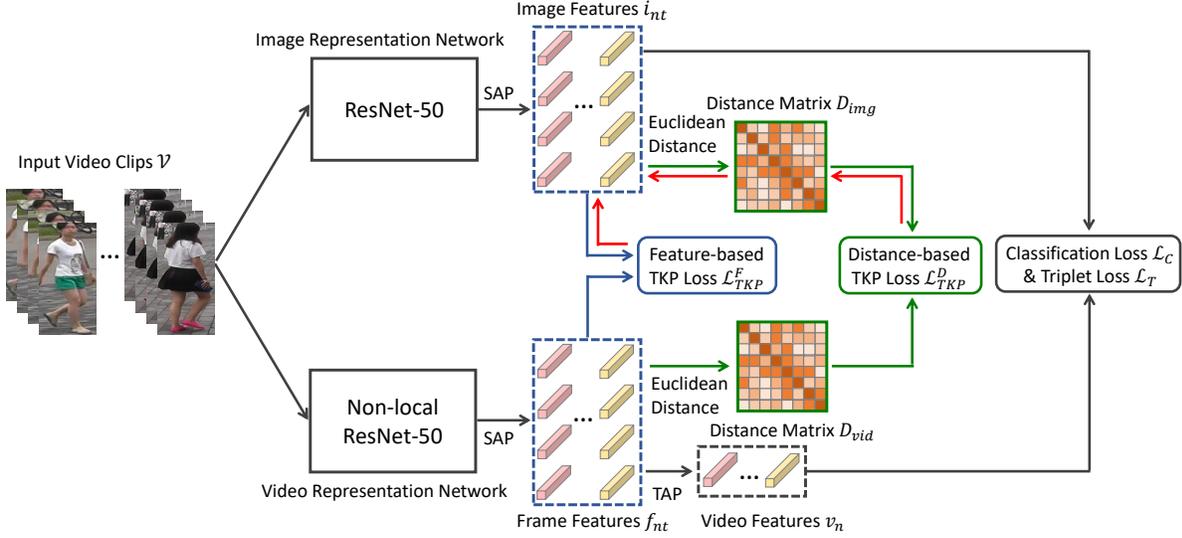


Figure 2: The framework of TKP method in I2V Re-ID training. SAP and TAP represent spatial average pooling and temporal average pooling respectively. The classification loss and triplet loss are used to guide image-to-video representation learning. The blue arrow represents the process of TKP via features, while the green arrow represents the process of TKP via cross sample distances. And the red arrow denotes the back propagation process of TKP loss. The representation learning and temporal knowledge transferring are trained simultaneously. Best viewed in color.

3.1. Image Representation Network

We use ResNet-50 [9] without final fully-connected layer as image representation network for visual feature learning. To enrich the granularity of image features, we remove the last down-sampling operation of ResNet-50 following [30].

Given N person video clips $\mathcal{V} = \{V_n\}_{n=1}^N$, each V_n contains T frames $V_n = \{F_{n1}, F_{n2}, \dots, F_{nT}\}$ (Unless specified, we set $T = 4$). If we discard the temporal relations between video frames, these video clips \mathcal{V} can be considered as a set of individual images $\{F_{nt}\}_{n=1, t=1}^{N, T}$. Hence, we can use the image representation network $\mathcal{F}_{img}(\cdot)$ to extract features of these images, for all n, t ,

$$i_{nt} = \mathcal{F}_{img}(F_{nt}), \quad (1)$$

where $i_{nt} \in \mathbb{R}^D$ is the corresponding image feature of video frame F_{nt} . As for ResNet-50, D is 2048.

3.2. Video Representation Network

To model the visual and temporal information of video sequences simultaneously, we combine CNN with non-local neural network [33] as video representation network. The non-local block computes the response at a position as a weighted sum of the features at all positions in the input feature map. It can naturally handle the temporal relations between video frames.

Table 1 shows the model structure of our video representation network on the backbone of ResNet-50. Specifically, we add two non-local blocks to res_3 and three non-local

Table 1: The architecture of video representation network. Each input video clip contains 4 frames, each frame with 256×128 pixels.

	layer	output size
conv ₁	7×7 , stride 2, 2	$4 \times 128 \times 64 \times 64$
pool	3×3 max, stride 2, 2	$4 \times 64 \times 32 \times 64$
res ₂	residual block $\times 3$	$4 \times 64 \times 32 \times 256$
res ₃	$\left[\begin{array}{l} \text{residual block} \times 2 \\ \text{non-local block} \times 1 \end{array} \right] \times 2$	$4 \times 32 \times 16 \times 512$
res ₄	$\left[\begin{array}{l} \text{residual block} \times 2 \\ \text{non-local block} \times 1 \end{array} \right] \times 3$	$4 \times 16 \times 8 \times 1024$
res ₅	residual block $\times 3$	$4 \times 16 \times 8 \times 2048$
	spatial average pool	4×2048
	temporal average pool	2048

blocks to res_4 and remove the last down-sampling operation in res_5 to enrich granularity. Given an input video clip $V_n = \{F_{n1}, F_{n2}, \dots, F_{nT}\}$ with T frames, the video representation network $\mathcal{F}_{vid}(\cdot)$ is defined as:

$$\{f_{n1}, f_{n2}, \dots, f_{nT}\} = \mathcal{F}_{vid}(F_{n1}, F_{n2}, \dots, F_{nT}), \quad (2)$$

where $f_{nt} \in \mathbb{R}^D$, $t = 1, \dots, T$ is the video frame feature of F_{nt} . With temporal average pooling, multiple video frame features of a video clip can be integrated to a video feature $v_n \in \mathbb{R}^D$.

3.3. Temporal Knowledge Propagation

In general, the performance of Re-ID highly depends on the robustness of feature representation. It has been proved that modeling temporal relations between video frames makes the person appearance representations robust

to large variations [35]. However, image representation network takes in still images and cannot process temporal relations, thus the output image features cannot benefit from temporal knowledge. To solve this problem, we propose TKP method which enforces the outputs of image representation network to fit the robust outputs of video representation network in the shared feature space. By back propagation algorithm, the image representation network can learn temporal knowledge from the video frame features. Consequently, the features extracted by image representation network are assigned, though not directly, some video temporal information.

Specifically, given the input video clips \mathcal{V} , we can use Eq. (1) and Eq. (2) to extract image features i_{nt} and video frame features f_{nt} for all $n = 1, \dots, N, t = 1, \dots, T$. Since $\mathcal{F}_{vid}(\cdot)$ extracts visual information and deals with temporal relations between video frames simultaneously, f_{nt} not only contains the visual information of video frame F_{nt} , but also involves temporal relations with other frames. In order to use video frame feature f_{nt} to propagate the temporal knowledge to image representation network, we formulate the TKP method as an optimization problem from the following two ways.

Propagation via Features. The first way is enforcing image representation network to fit the robust video frame features in a shared feature space. In this case, TKP method can be formulated to minimize the MSE between the image features and the corresponding video frame features:

$$\mathcal{L}_{TKP}^F = \frac{1}{NT} \sum_{n=1}^N \sum_{t=1}^T \|i_{nt} - f_{nt}\|_2^2, \quad (3)$$

where $\|\cdot\|_2$ denotes l_2 distance.

Eq. (3) can be considered as simplified *moving least squares* [19, 27], which is capable of reconstructing continuous functions from a set of labeled samples. Here, our target is to reconstruct the image representation function $\mathcal{F}_{img}(\cdot)$ from the video frame representations $(F_{nt}, f_{nt})_{n=1, t=1}^{N, T}$. This formulation is similar to FitNets [25], except that the outputs of teacher network and student network in [25] are mapped to the same dimension via an extra convolutional regressor. In contrast, the outputs of the image and video representation networks in our framework have the same dimension and the network construction is similar, thus we do not need additional convolutional regressor.

Propagation via Cross Sample Distances. Another way to propagate temporal knowledge from video representations to image representations may resort to neural network embedding. The structure of the target embedding space is characterized by cross sample distances. For all video frame features $\{f_{nt}\}_{n=1, t=1}^{N, T}$, we compute the cross sample

Euclidean distances matrix as $D_{vid} \in \mathbb{R}^{NT \times NT}$. To estimate the image representation in the embedding space, we constrain that the cross image distances $D_{img} \in \mathbb{R}^{NT \times NT}$ are consistent with the cross video frame distances D_{vid} . In this way, the temporal information, as well as the sample distribution, are propagated to the image representation network. The TKP loss is formulated as:

$$\mathcal{L}_{TKP}^D = \frac{1}{NT} \|D_{img} - D_{vid}\|_F^2, \quad (4)$$

where $\|\cdot\|_F$ denotes Frobenius norm. This formulation is similar to *multidimensional scaling* [17], except that we use a deep network to model the embedding function $\mathcal{F}_{img}(\cdot)$ instead of directly computing the embedded features via eigen-decomposition.

Eq. (3) and Eq. (4) transfer knowledge from different levels and are complementary to each other. The empirical comparison of the two ways is provided in Section 4.3.

Note that both image and video networks use ResNet-50 as backbone. The only difference is that video network add extra non-local blocks to model temporal information. Given the same inputs, TKP loss enforces these two networks to output similar features. Obviously, the weights of additional non-local blocks being 0 is the optimal solution of minimizing TKP loss. In that case, the non-local blocks can not capture any temporal information. So updating video network by TKP deteriorates modeling temporal knowledge. Unless specified, in our experiments, \mathcal{L}_{TKP}^F and \mathcal{L}_{TKP}^D are not back-propagated through the video representation network during model training.

3.4. Objective Function

Besides the TKP loss, additional identification losses are also needed to learn discriminative features for image-to-video re-identification. In this paper, we utilize the widely used classification loss and integrated triplet loss. In fact, other identification losses are also applicable.

Classification Loss. Considering person identities as category-level annotations, we build two shared weights classifiers to map the image features and video features to a shared identity space. The classifiers are implemented as a linear layer following by a softmax operation and the output channel is the number of identities of training set. The classification loss \mathcal{L}_C can be formulated as the cross entropy error between the predicted identities and the correct labels.

Integrated Triplet Loss. We also use triplet loss with hard sample mining [10] to constrain the relative sample distances in the shared feature space. Specifically, we integrate four kinds of triplet losses, image-to-video (I2V), video-to-image (V2I), image-to-image (I2I) and video-to-video (V2V) triplet losses. The final triplet loss \mathcal{L}_T is defined as:

$$\mathcal{L}_T = \mathcal{L}_{I2V} + \mathcal{L}_{V2I} + \mathcal{L}_{I2I} + \mathcal{L}_{V2V}, \quad (5)$$

where

$$\mathcal{L}_{I2V} = \left[m + \max_{v_p \in S_a^+} d(i_a, v_p) - \min_{v_n \in S_a^-} d(i_a, v_n) \right]_+, \quad (6)$$

$$\mathcal{L}_{V2I} = \left[m + \max_{i_p \in S_a^+} d(v_a, i_p) - \min_{i_n \in S_a^-} d(v_a, i_n) \right]_+, \quad (7)$$

$$\mathcal{L}_{I2I} = \left[m + \max_{i_p \in S_a^+} d(i_a, i_p) - \min_{i_n \in S_a^-} d(i_a, i_n) \right]_+, \quad (8)$$

$$\mathcal{L}_{V2V} = \left[m + \max_{v_p \in S_a^+} d(v_a, v_p) - \min_{v_n \in S_a^-} d(v_a, v_n) \right]_+. \quad (9)$$

Here m is a pre-defined margin, $d(\cdot, \cdot)$ denotes the Euclidean distance, and $[\cdot]_+ = \max(0, \cdot)$. S_a^+ and S_a^- are the positive and negative sample sets of the anchor sample (i_a or v_a) respectively.

Among the four losses, Eq. (6) and Eq. (7) constrain the distance between image feature and video feature, which can improve the between-modality feature discriminativity. In contrast, Eq. (8) and Eq. (9) constrain within-modality relative distances, which makes our model distinguish the fine-grained differences between different identities within the same modality. The between-modality and within-modality losses are complementary and their integration can improve image-to-video representation learning.

Objective Function. The image-to-video representation learning and temporal knowledge transferring are trained simultaneously. The final objective function is formulated as the combination of classification loss, integrated triplet loss and the proposed TKP loss:

$$\mathcal{L} = \mathcal{L}_C + \mathcal{L}_T + \mathcal{L}_{TKP}^F + \mathcal{L}_{TKP}^D. \quad (10)$$

3.5. Sampling Strategy

To better train the model with multiple losses, we design a particular sampling strategy. For each batch, we randomly select P persons. For each person, we randomly select K video clips, each with T frames. All the $P \times K = N$ video clips are fed into video representation network. Meanwhile, all $N \times T$ frames form an image batch and are fed into image representation network. In this way, all the samples in a mini batch can be reused to compute these three losses in Eq. (10), which can reduce the computational cost.

3.6. Image-to-video Re-ID Testing

In the test stage, each query is a still image and the gallery set consists of masses of person videos. The process of image-to-video Re-ID testing is shown in Figure 3. Specifically, we use the learned image representation network after TKP to extract image feature of the query and the gallery video features are extracted by video representation network. After feature extraction, we compute the distances between the query feature and each gallery video feature and then conduct image-to-video retrieval according to the distances.

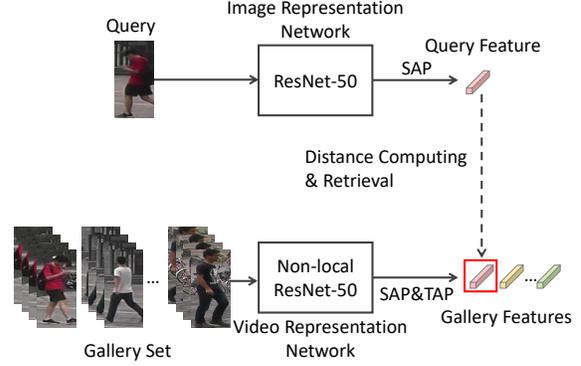


Figure 3: The pipeline of I2V Re-ID testing.

4. Experiments

4.1. Datasets and Evaluation Protocol

Datasets. We evaluate our method on MARS [41], DukeMTMC-VideoReID (Duke) [34] and iLIDS-VID [32] datasets. Among them, MARS and Duke are multi-camera datasets, while iLIDS-VID is captured by only two cameras. The amounts of person videos on MARS, Duke and iLIDS-VID are 20478, 5534 and 600 respectively, and the average lengths of person videos on these three datasets are 58, 168 and 71 respectively.

Evaluation Protocol. All the three datasets above are video Re-ID datasets. For multi-camera datasets (MARS and Duke), we just use the first frame of every query video as query image to perform I2V Re-ID testing following [31]. For iLIDS-VID, we use the first frames of all person videos captured by the first camera for both training and testing in order to be consistent with [44, 31, 37].

We use the Cumulative Matching Characteristics (CMC) to evaluate the performance of each approach. For iLIDS-VID, the experiment is repeated 10 times and the average result is presented. For multi-camera datasets, we also report the mean Average Precision (mAP) [42] as a complement to CMC.

The comparison experiments are mainly conducted on MARS and Duke, since these two datasets have fixed training/testing splits, which is convenient for extensive evaluation. We also present the final results on iLIDS-VID as well as MARS to compare with the state-of-the-art methods.

4.2. Implementation Details

We pre-train ResNet-50 on ImageNet [26] and adopt the method in [33] to initialize the non-local blocks. During training, we randomly sample 4 frames with a stride of 8 frames from the original full-length video to form an input video clip. For the original video less than 32 frames, we duplicate it to meet the length. The parameters P and K in Section 3.5 are both set to 4. The input video frames are resized to 256×128 pixels. Only horizontal flip is used for data augmentation. We adopt Adaptive Moment Estimation

Table 2: The results of I2V, I2I and V2V Re-ID on the MARS and Duke datasets. In I2I setting, only the first frames of the query and gallery samples are used. In V2V setting, the full-length query videos and gallery videos are used. All the image features in I2I and I2V Re-ID are extracted by the learned image representation network. All the video features in I2V and V2V Re-ID are extracted by the learned video representation network.

Models	Losses				MARS						Duke					
	\mathcal{L}_C	\mathcal{L}_T	\mathcal{L}_{TKP}^F	\mathcal{L}_{TKP}^D	I2V Re-ID		I2I Re-ID		V2V Re-ID		I2V Re-ID		I2I Re-ID		V2V Re-ID	
					top-1	mAP										
baseline	✓	✓			67.1	55.5	65.9	49.2	83.4	72.6	67.5	65.6	60.4	52.8	93.2	91.3
TKP-F	✓	✓	✓		75.0	64.2	71.0	54.7	83.2	72.6	76.8	74.2	63.0	54.5	93.6	91.5
TKP-D	✓	✓		✓	75.0	63.1	70.3	55.0	84.1	72.9	76.5	74.9	62.0	53.5	93.3	91.4
TKP	✓	✓	✓	✓	75.6	65.1	71.0	55.0	84.0	73.3	77.9	75.9	63.4	54.8	94.0	91.7

(Adam) [16] with weight decay 0.0005 to optimize the parameters. The model is trained for 150 epochs in total. The learning rate is initialized to 0.0003 and divided by 10 after every 60 epochs. For iLIDS-VID, we first pre-train the model on large-scale dataset and then fine-tune it on iLIDS-VID following [31].

In the test phase, the query image features are extracted by image representation model. For each gallery video, we first split it into several 32-frame clips. For each clip, we utilize video representation model to extract video representation. The final video feature is the averaged representation of all clips.

4.3. Results on I2V Re-ID

To validate the effectiveness of the proposed TKP method for I2V Re-ID, we implement and test a baseline and several variants of our model. The configurations of these approaches are presented in Table 2. Among them, *baseline* only adopts classification loss and triplet loss for image-to-video representation learning. *TKP-F* and *TKP-D* extra use Eq. (3) and Eq. (4) to transfer temporal knowledge respectively. *TKP* combines these two transfer ways during training. The results of I2V Re-ID on the MARS and Duke dataset can be seen in Table 2.

Compared with *baseline*, *TKP-F* and *TKP-D* consistently improve the performance by a large margin. Specifically, *TKP-F* increases the mAP by 8.7% and 8.6% on MARS and Duke respectively. For *TKP-D*, the improvement of mAP is 7.6% on MARS, and 9.3% on Duke. This comparison shows that temporal knowledge transfer is essential for image-to-video representation learning. We also compare the combination method *TKP* with *TKP-F* and *TKP-D*. It can be seen that a farther improvement is achieved. This result demonstrates that these two transfer ways transfer temporal knowledge from different perspectives and are complementary to each other.

4.4. How does TKP Work?

To investigate how TKP works, we extra use the image representation networks trained in Section 4.3 to conduct image-to-image (I2I) Re-ID tests, where only the first frames of the original query and gallery videos are utilized.

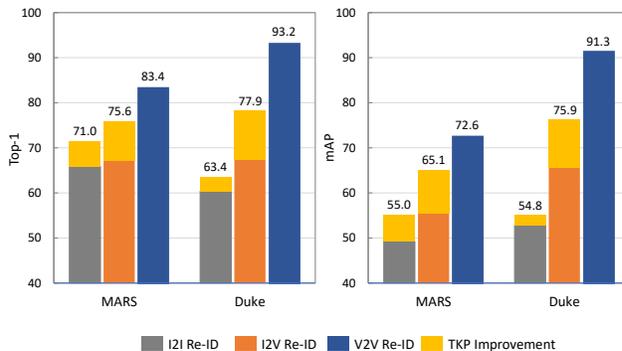


Figure 4: Comparison among I2I, I2V and V2V Re-ID on the MARS and Duke datasets. With the proposed TKP method, the performance gap can be significantly reduced.

Moreover, we also use the trained video representation networks to perform video-to-video (V2V) Re-ID experiments, where the original full-length query and gallery videos are used. The experiment results are also presented in Table 2.

Compared with *baseline*, different transfer methods consistently improve the I2I Re-ID performance on both MARS and Duke dataset. Especially, *TKP* increases the mAP from 49.2% to 55.0% (+5.8%) on the MARS dataset. Moreover, the V2V Re-ID performance of different transfer methods is close to *baseline*. The comparisons demonstrate that the proposed TKP method can improve the robustness of learned image features, meanwhile, does not reduce the discriminativity of video features. In addition, with the transferred temporal knowledge, the information asymmetry between image and video features can also be alleviated, thus the I2V Re-ID performance gains more improvement.

4.5. Comparison among I2I, I2V and V2V Re-ID

I2I (image-based) Re-ID is a task where the query and each gallery are both images, while the query and each gallery in V2V (video-based) Re-ID are both videos. In I2V setting, the query is an image while each gallery is a video. We compare the three different tasks under the same configurations and the comparisons on the MARS and Duke dataset are shown in Figure 4. Due to the lack of additional visual and temporal information, the performance of I2V Re-ID is lower than that of V2V Re-ID, and I2I Re-ID is

Table 3: Comparison with state-of-the-art I2V Re-ID methods on the iLIDS-VID dataset.

Models	top-1	top-5	top-10	top-20
PSDML [43]	13.5	33.8	45.6	56.3
LERM [14]	15.3	37.1	49.7	62.0
XQDA [21]	16.8	38.6	52.3	63.6
KISSME [18]	17.6	41.7	55.3	68.7
PHDL [44]	28.2	50.4	65.9	80.4
TMSL [37]	39.5	66.9	79.6	86.6
P2SNet [31]	40.0	68.5	78.1	90.0
TKP	54.6	79.4	86.9	93.5

Table 4: Comparison with state-of-the-art I2V Re-ID methods on the MARS dataset.

Models	top-1	top-5	top-10	mAP
P2SNet [31]	55.3	72.9	78.7	-
ResNet-50 [9]+XQDA [21]	67.2	81.9	86.1	54.9
TKP	75.6	87.6	90.9	65.1

Table 5: Comparison with state-of-the-art V2V Re-ID methods on the MARS dataset.

Models	top-1	top-5	top-10	mAP
SDM [38]	71.2	85.7	91.8	-
MGCAM [29]	77.2	-	-	71.2
DuATM [28]	78.7	90.9	-	62.3
multi-snippets [3]	81.2	92.1	-	69.4
DRSA [20]	82.3	-	-	65.8
TKP	84.0	93.7	95.7	73.3

lower than I2V Re-ID. Especially, the performance gap on Duke is much larger. The reason can be attributed that the average length of the videos on Duke is longer than that on MARS. When we only use one frame of the original video to conduct I2I and I2V tests, the information loss is more serious. But our proposed TKP method can transfer temporal knowledge to image features, thus the performance gap can be greatly reduced on these two datasets.

4.6. Comparison with State-of-the-art Methods

We compare the proposed approach with state-of-the-art I2V Re-ID methods on the iLIDS-VID and MARS datasets. The results are presented in Table 3 and Table 4 respectively. Among them, *PSDML* [43], *LERM* [14], *XQDA* [21], *KISSME* [18] and *PHDL* [44] are handcrafted feature based methods, while *ResNet-50* [9]+*XQDA* [21], *TMSL* [37] and *P2SNet* [31] are deep learning based methods. It can be seen that deep learning based methods significantly outperform the traditional methods with handcrafted features, while our method further surpasses the existing deep learning based methods by a large margin. Since Duke is a newly released dataset, existing methods have not conducted I2V Re-ID experiments on it. Therefore, we do not compare with state-of-the-art methods on this dataset. Anyway, the results of our method can be seen in Table 2.

Note that the V2V Re-ID performance decides the upper bound of the I2V Re-ID performance. We also compare the

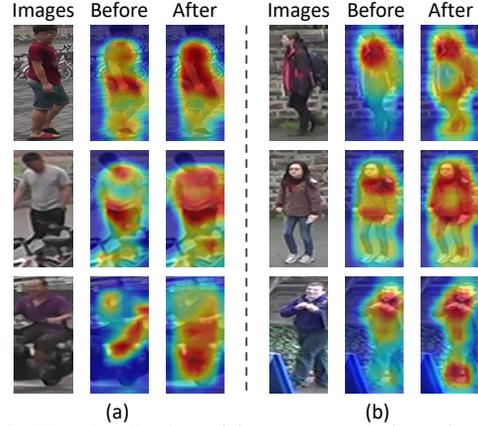


Figure 5: The visualisation of feature maps before/after TKP on the (a) MARS and (b) Duke datasets. Best viewed in color.

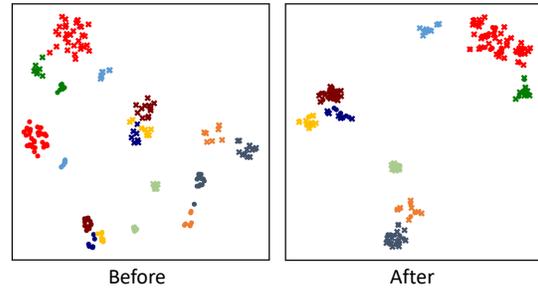


Figure 6: The visualisation of feature distribution before/after TKP on the MARS dataset. Circle represents video feature, while cross represents image feature. Different colors denote different identities. Best viewed in color.

proposed approach with state-of-the-art V2V Re-ID methods on the MARS dataset. To compare fairly, for *multi-snippets* [3], we use the results without optical flow. As shown in Table 5, our TKP consistently outperforms these methods. As for the iLIDS-VID dataset, since we do not use all of the training set (only use the first frames of all videos captured by the first camera), we do not compare V2V Re-ID results with these methods on this dataset.

4.7. Visualization

The visualization of feature maps. We visualise the feature maps of the image features before/after TKP in Figure 5. It can be seen that the original image features only pay attention to some local discriminative regions. After TKP transferring temporal knowledge, the learned image representations can focus on more foreground and manifest robustness to occlusion and blur just like video features in Figure 1, which benefits I2V matching. So the final I2V performance can be improved significantly.

The visualization of feature distribution. We also visualize the distribution of the learned features before/after TKP using t-SNE [22] as shown in Figure 6. Before temporal knowledge transferring, the image features and video features with the same identity are incompact. After TKP

Table 6: The I2V Re-ID results with/without non-local blocks on the MARS dataset. w/ NL? denotes whether the model contains non-local blocks or not. The performance improvement is provided in brackets.

Models	w/ NL?	top-1	mAP
baseline	-	66.1	51.8
TKP-F	-	68.9(+2.8)	57.8(+6.0)
baseline	✓	67.1	55.5
TKP-F	✓	75.0(+7.9)	64.2(+8.7)

Table 7: The results with/without the TKP loss propagating gradient to video representation network on the MARS dataset. BP2v? denotes whether the gradient of TKP loss is propagated to video representation network or not.

Models	BP2v?	I2I Re-ID		I2V Re-ID		V2V Re-ID	
		top-1	mAP	top-1	mAP	top-1	mAP
baseline		65.9	49.2	67.1	55.5	83.4	72.6
TKP-F	✓	66.6	51.0	72.7	60.3	78.5	66.6
	-	71.0	54.7	75.0	64.2	83.2	72.6
TKP-D	✓	66.3	50.3	74.2	61.7	79.3	66.1
	-	70.3	55.0	75.0	63.1	84.1	72.9

Table 8: Comparison with the method using pre-trained video model on the MARS dataset.

Models	baseline	pre-trained	TKP
top-1	67.1	73.2	75.6
mAP	55.5	61.5	65.1

Table 9: Comparing the methods with different identification losses on the MARS dataset.

Models	I2V tri.	Integrated tri.	baseline
top-1	54.4	59.1	67.1
mAP	42.6	47.3	55.5

transferring temporal knowledge, the feature distributions of these two modalities become more consistent. Therefore, it is easier to measure the similarity between image and video features.

4.8. Ablation Study

Whether the non-local blocks are required? In our framework, we use non-local blocks to model temporal relations between video frames. To verify whether the non-local blocks are required or not, we remove the non-local blocks from the methods *baseline* and *TKP-F*. And the video frame feature f_{nt} in Eq. (3) is replaced by the video feature v_n after temporal average pooling. As shown in Table 6, when the non-local blocks are removed, *TKP-F* still surpasses *baseline* by a reasonable margin. But the performance and improvement without non-local blocks are both lower than those with non-local blocks. We argue that, compared with simple temporal average pooling, non-local blocks can model temporal information much better, which makes temporal knowledge propagation more effective as well.

Whether TKP loss should propagate gradient to video

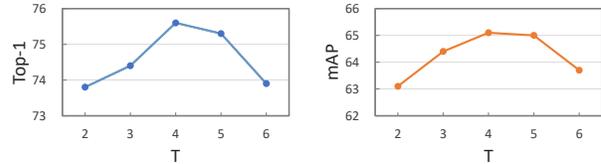


Figure 7: The results with different T on the MARS dataset.

representation network? As discussed in Section 3.3, enforcing TKP loss to propagate gradient to video representation network will degenerate video representations w.r.t. temporal knowledge. To verify this, we add two additional experiments and the results are reported in Table 7. It can be seen that, when the gradient of TKP loss is propagated to video network, *TKP-F* and *TKP-D* can still increase the performance of I2I and I2V Re-ID by a considerable margin. But these two methods consistently gain lower V2V performance. If the back propagation to video network is banned, all I2I, I2V and V2V results can be further improved.

Whether using a pre-trained video model is beneficial for the convergence of networks? Our method aims to solve I2V matching, but the pre-trained video model supervised by V2V loss function may not be optimal for I2V matching. To verify this, we add an experiment which uses a pre-trained video model to perform knowledge propagation. As shown in Table 8, though the method *pre-trained* outperforms *baseline*, it is inferior to *TKP* which learns two networks simultaneously.

The influence of different triplet losses. To explore this, we perform the experiments with different kinds of triplet losses. As shown in Table 9, the method *Integrated tri.* using integrated triplet loss surpasses *I2V tri.* which only uses I2V triplet loss. With additional classification loss, *baseline* outperforms these two methods.

The influence of the video clip size T . By varying T , we show the experimental results in Figure 7. It can be seen that the best top-1 and mAP are consistently achieved when T is 4.

5. Conclusion

In this paper, we propose a novel TKP method for I2V Re-ID. TKP can transfer temporal knowledge from video representation network to image representation network. With the transferred temporal knowledge, the robustness of the image features can be improved and the information asymmetry between image and video features can also be alleviated. Extensive experiments demonstrate the effectiveness of our method and the results on two widely used datasets significantly surpass start-of-the-art performance.

Acknowledgement This work is partially supported by National Key R&D Program of China (No.2017YFA0700800), Natural Science Foundation of China (NSFC): 61876171 and 61572465.

References

- [1] Song Bai, Xiang Bai, and Qi Tian. Scalable person re-identification on supervised smoothed manifold. In *CVPR*, 2017.
- [2] Cristian Bucilu, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *KDD*, 2006.
- [3] Dapeng Chen, Hongsheng Li, Tong Xiao, Shuai Yi, and Xiaogang Wang. Video person re-identification with competitive snippet-similarity aggregation and co-attentive snippet embedding. In *CVPR*, 2018.
- [4] Yuntao Chen, Naiyan Wang, and Zhaoxiang Zhang. Dark-rank: Accelerating deep metric learning via cross sample similarities transfer. In *AAAI*, 2018.
- [5] Dahjung Chung, Khalid Tahboub, and Edward J. Delp. A two stream siamese convolutional neural network for person re-identification. In *ICCV*, 2017.
- [6] Tran Du, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015.
- [7] Christoph Feichtenhofer, Axel Pinz, and Richard P. Wildes. Spatiotemporal multiplier networks for video action recognition. In *CVPR*, 2017.
- [8] Yongxin Ge, Xinqian Gu, Min Chen, Hongxing Wang, and Dan Yang. Deep multi-metric learning for person re-identification. In *ICME*, 2018.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [10] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *ArXiv:1703.07737*, 2017.
- [11] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In *NIPS workshop*, 2014.
- [12] Ruibing Hou, Bingpeng Ma, Hong Chang, Xinqian Gu, Shiguang Shan, and Xilin Chen. Interaction-and-aggregation network for person re-identification. In *CVPR*, 2019.
- [13] Ruibing Hou, Bingpeng Ma, Hong Chang, Xinqian Gu, Shiguang Shan, and Xilin Chen. Vrstc: Occlusion-free video person re-identification. In *CVPR*, 2019.
- [14] Zhiwu Huang, Ruiping Wang, Shiguang Shan, and Xilin Chen. Learning euclidean-to-riemannian metric for point-to-set classification. In *CVPR*, 2014.
- [15] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3d convolutional neural networks for human action recognition. In *ICML*, 2010.
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [17] Joseph B. Kruskal. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 1964.
- [18] Martin Kstinger, Martin Hirzer, Paul Wohlhart, Peter M Roth, and Horst Bischof. Large scale metric learning from equivalence constraints. In *CVPR*, 2012.
- [19] David Levin. The approximation power of moving least-squares. *Mathematics of Computation*, 1998.
- [20] Shuang Li, Slawomir Bak, Peter Carr, and Xiaogang Wang. Diversity regularized spatiotemporal attention for video-based person re-identification. In *CVPR*, 2018.
- [21] Shengcai Liao, Yang Hu, Xiangyu Zhu, and Stan Z. Li. Person re-identification by local maximal occurrence representation and metric learning. In *CVPR*, 2015.
- [22] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *JMLR*, 2008.
- [23] Niall McLaughlin, Jesus Martinez Del Rincon, and Paul Miller. Recurrent convolutional network for video-based person re-identification. In *CVPR*, 2016.
- [24] Yue Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *CVPR*, 2015.
- [25] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In *ICLR*, 2015.
- [26] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, and Michael Bernstein. Imagenet large scale visual recognition challenge. *IJCV*, 2015.
- [27] Scott Schaefer, Travis McPhail, and Joe Warren. Image deformation using moving least squares. In *SIGGRAPH*, 2006.
- [28] Jianlou Si, Honggang Zhang, Chun-Guang Li, Jason Kuen, Xiangfei Kong, Alex C. Kot, and Gang Wang. Dual attention matching network for context-aware feature sequence based person re-identification. In *CVPR*, 2018.
- [29] Chunfeng Song, Yan Huang, Wanli Ouyang, and Liang Wang. Mask-guided contrastive attention model for person re-identification. In *CVPR*, 2018.
- [30] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin Wang. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *ECCV*, 2018.
- [31] Guangcong Wang, Jianhuang Lai, and Xiaohua Xie. P2snet : Can an image match a video for person re-identification in an end-to-end way? *TCSVT*, 2017.
- [32] Taiqing Wang, Shaogang Gong, Xiatian Zhu, and Shengjin Wang. Person re-identification by video ranking. In *ECCV*, 2014.
- [33] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018.
- [34] Yu Wu, Yutian Lin, Xuanyi Dong, Yan Yan, Wanli Ouyang, and Yi Yang. Exploit the unknown gradually: One-shot video-based person re-identification by stepwise learning. In *CVPR*, 2018.
- [35] Jinjie You, Ancong Wu, Xiang Li, and Wei-Shi Zheng. Top-push video-based person re-identification. In *CVPR*, 2016.
- [36] Rui Yu, Zhiyong Dou, Song Bai, Zhaoxiang Zhang, Yongchao Xu, and Xiang Bai. Hard-aware point-to-set deep metric for person re-identification. In *ECCV*, 2018.
- [37] Dongyu Zhang, Wenxi Wu, Hui Cheng, Ruimao Zhang, Zhenjiang Dong, and Zhaoquan Cai. Image-to-video person re-identification with temporally memorized similarity learning. *TCSVT*, 2017.

- [38] Jianfu Zhang, Naiyan Wang, and Liqing Zhang. Multi-shot pedestrian re-identification via sequential decision making. In *CVPR*, 2018.
- [39] Xuan Zhang, Hao Luo, Xing Fan, Weilai Xiang, Yixiao Sun, Qiqi Xiao, Wei Jiang, Chi Zhang, and Jian Sun. Aligned-reid: Surpassing human-level performance in person re-identification. *ArXiv:1711.08184*, 2017.
- [40] Ying Zhang, Tao Xiang, Timothy M. Hospedales, and Huchuan Lu. Deep mutual learning. In *CVPR*, 2018.
- [41] Liang Zheng, Zhi Bie, Yifan Sun, Jingdong Wang, Chi Su, Shengjin Wang, and Qi Tian. Mars: A video benchmark for large-scale person re-identification. In *ECCV*, 2016.
- [42] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *ICCV*, 2015.
- [43] Pengfei Zhu, Lei Zhang, Wangmeng Zuo, and David Zhang. From point to set: Extend the learning of distance metrics. In *ICCV*, 2013.
- [44] Xiaoke Zhu, Xiao-Yuan Jing, Fei Wu, Yunhong Wang, Wangmeng Zuo, and Wei-Shi Zheng. Learning heterogeneous dictionary pair with feature projection matrix for pedestrian video retrieval via single query image. In *AAAI*, 2017.
- [45] Xiaoke Zhu, Xiao-Yuan Jing, Xinge You, Wangmeng Zuo, Shiguang Shan, and Wei-Shi Zheng. Image to video person re-identification by learning heterogeneous dictionary pair with feature projection matrix. *TIFS*, 2018.