Multi-shot Person Re-identification through Set Distance with Visual Distributional Representation

Ting-Yao Hu
Carnegie Mellon University
tingyao@cs.cmu.edu

Alexander G. Hauptmann
Carnegie Mellon University
alex@cs.cmu.edu

ABSTRACT

Person re-identification aims to identify a specific person at distinct times and locations. It is challenging because of occlusion, illumination, and viewpoint change in camera views. Recently, multi-shot person re-id task receives more attention since it is closer to real-world application. A key point of a good algorithm for multi-shot person re-id is the temporal aggregation of the person appearance features. While most of the current approaches apply pooling strategies and obtain a fixed-size vector representation, these may lose the matching evidence between examples. In this work, we propose the idea of *visual distributional representation*, which interprets an image set as samples drawn from an unknown distribution in appearance feature space. Based on the supervision signals from a downstream task of interest, the method reshapes the appearance feature space and further learns the unknown distribution of each image set. In the context of multi-shot person re-id, we apply this novel concept along with Wasserstein distance and jointly learn a distributional set distance function between two image sets. In this way, the proper alignment between two image sets can be discovered naturally in a non-parametric manner. Our experiment results on three public datasets show the advantages of our proposed method compared to other state-of-the-art approaches.

CCS CONCEPTS

• **Computing methodologies** → **Object identification; Tracking; Neural networks; Image representations.**

KEYWORDS

person re-identification, representation learning, Wasserstein distance

ACM Reference Format:

Ting-Yao Hu and Alexander G. Hauptmann. 2019. Multi-shot Person Re-identification through Set Distance with Visual Distributional Representation. In *International Conference on Multimedia Retrieval (ICMR '19)*, June 10–13, 2019, Ottawa, ON, Canada. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3323873.3325030>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICMR '19, June 10–13, 2019, Ottawa, ON, Canada

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6765-3/19/06...\$15.00

<https://doi.org/10.1145/3323873.3325030>

1 INTRODUCTION

Person re-identification (person re-id) aims to identify a specific person at distinct times and locations. It is an essential task for several applications, such as long-term person tracking across camera views [25]. Re-id task is still challenging due to the appearance variations of person. These variations usually come from occlusion, illumination, and viewpoint change in camera views. Recent approaches usually treat re-id task as a retrieval problem: given a query based on a single image or a set of images, and a gallery set of candidate person images, we need to rank these candidates according to some similarity metrics. Researchers have considered two scenarios, single-shot and multi-shot, for person re-id task. A lot of previous works have focused mainly on the single-shot scenario, while only a few lie in the latter. However, multi-shot person re-id is more suitable for practical surveillance applications, since person tracklets are available by applying object detection and tracking algorithms. In this paper, we investigate person re-id task in the multi-shot scenario.

Multi-shot person re-id methods require the comparison between two sets of images. Thanks to the rapid development of deep learning techniques, recent approaches adopt convolutional neural network (CNN) to extract the appearance feature of each image, and take temporal pooling strategies to aggregate an appearance feature sequence and form a fixed-size vector representation. Common temporal pooling strategies include mean and max pooling, recurrent neural network (RNN), and attention models. Finally, the dissimilarity between two vectors is calculated based on a distance function, such as the Euclidean and cosine distance. However, using a fixed-size vector as the representation of a set of images, previous re-id algorithms may ignore the matching evidence between two person tracklets. For example, when the system compares the image set of identity a to the image sets of a' and b from the gallery set, the evidence showing that a and a' are identical, and that showing a and b are not, may come from different images in set a . Hence, instead of aggregating appearance features by conventional pooling strategies, a better way should be to represent an image set by the whole appearance feature set, and discover the alignment/attention between two sets, as shown in Figure 1.

In this paper, we propose the idea of *visual distributional representation* and use it to solve multi-shot person re-id task. Our approach treats a set of images, or a person tracklet as samples drawn from a probability distribution in appearance feature space. Based on this concept, our multi-shot person re-id algorithm consists of three parts: an appearance feature extractor, a probability estimator, and a distributional distance function. Specifically, we choose the *Wasserstein distance* between distributions as the function calculating the dissimilarity between two image sets. By doing so, we can model the diversity and uncertainty of each image set, and embed the

Figure 1: In multi-shot re-id task, a good algorithm is supposed to focus on different images of the query while comparing to different examples in the gallery set. Our algorithm aims to discover the most proper alignment between the query and gallery examples.

alignment/attention mechanism into our re-id algorithm without introducing additional parameters. Furthermore, the dissimilarity score also serves as a supervision signal to reshape the appearance feature extractor and learn the distributions of all the image sets.

We conduct experiments on iLIDS-VID [32], MARS [41] and DukeMTMC-VideoReID [34] datasets. The results demonstrate that the proposed person re-id algorithm based on visual distributional representation outperforms some conventional pooling strategies, and achieve state-of-the-art performance.

In summary, the contributions of this work are three folds: (1) We outline the concept of visual distributional representation, which aims to preserve the diversity and uncertainty of a set of observations, and conduct deep feature learning. (2) Based on this concept, we utilize Wasserstein distance to design an image set distance function and build a network architecture to solve multi-shot person re-id problem. This architecture is simple yet effective in the alignment between two image sets. (3) We show that our proposed method outperforms several baselines and provides competitive/superior performance comparing to state-of-the-art approaches.

2 RELATED WORKS

Person re-identification task has been investigated for a long time. Earlier works in this area can be categorized into two major parts: feature extraction and metric learning. For feature extraction, a variety of visual features have been proposed to capture the appearance of a person, such as color histogram [40], Local Binary Pattern (LBP) [36], and Local Maximal Occurrence (LOMO) [19]. Some other works leveraged metric learning techniques to discover a discriminative distance measure in appearance feature space, such as KISSME [16], LFDA [22], and LMNN [4]. On the other hand, deep learning based methods [1, 5, 29, 35] have been proposed to reduce the demand of hand-crafted feature design, and jointly conduct feature extraction and metric learning.

Recently, more and more works focused on multi-shot person re-id task, since it is closer to the requirement of a real system. The re-id algorithms receive a person tracklet as a input and should take the advantage of the availability of multi-shot images. McLaughlin et al. [21] proposed an RNN model to encode temporal information, and adopted mean/max pooling to aggregate features over each dimension. Xu et al. [37] built a network incorporating spatial pyramid attention and temporal RNN jointly. In [21] and [37], the authors utilized both RGB and optical flow channel as network input. Most of the recent work considered RGB channels to be the input, and focused on the development of attention models. Liu et al. [20] used temporal attention to estimate the quality score of each image, and re-weighted the appearance features according to quality scores. Li et al. [18] proposed an attention network that automatically discovered a diverse set of discriminative body parts. The features from these local body regions were aggregated by temporal attention. Fu et al. [8] designed an attention mechanism conducting region and frame selection without additional network branch. Zhou et al. [42] used RNN to model spatial and temporal attention simultaneously. Si et al. [27] proposed a dual attention model to perform context-aware feature sequence comparison.

Some previous methods for video face recognition task [10, 33] also modeled an image set as a probability distribution, and incorporated supervision signal. In [10], authors used kernel density estimation to obtain the distribution of an image set, and considered supervised dimensionality reduction. In [33], authors exploited the supervision signal to conduct kernel metric learning. In comparison to these works, we adopt distributional distance and supervised training to finetune the parameters of appearance feature extractor, which is usually a CNN. Hence, our approach can be treated as a new type of deep feature learning framework.

To calculate the distance between two probability distributions, we select Wasserstein distance, which has been successfully applied to many different applications. Arjovsky et al. [2] improved the training process of generative adversarial network (GAN). Frogner et al. [7] utilized it to design a new loss function for multi-label classification. Shen et al. [26] measured the distance between source and target domains in domain adaptation task. Another set of works [3, 38] focused on reducing the computation cost of Wasserstein distance calculation, which is crucial to the inference process of our framework.

3 PROPOSED METHOD: VISUAL DISTRIBUTIONAL REPRESENTATION

We propose the concept of visual distributional representation, and apply it to solve multi-shot person re-id task. Our algorithm learns how to represent an image set as a probability distribution, and conducts set-to-set comparison by calculating the distance between two distributions. We claim that our method can effectively preserve the diversity of person appearance, and discover a proper alignment between two image sets. In the mean time, it does not increase the number of parameters comparing to models with naive mean pooling. In this section, we will elaborate on the details.

3.1 Problem Definition: Multi-shot Person Re-ID

In multi-shot person re-id task, multiple images of one identity are available. A key point of a good algorithm is its set-to-set comparison. Given two sets of images $S_i = \{I_1^i, I_2^i, \dots, I_{n_i}^i\}$ and $S_j = \{I_1^j, I_2^j, \dots, I_{n_j}^j\}$, we have to define a proper set distance function $D(S_i, S_j)$ to calculate the dissimilarity between them. In most of the previous works, this function can be decomposed into three parts: an appearance feature extractor, a pooling method, and a distance metric in feature vector space. The appearance feature extractor is usually implemented as a CNN. And the pooling method aggregates the appearance features into a fixed-size vector. For example, if we adopt mean pooling and the Euclidean distance, the set distance function would become:

$$D(S_i, S_j) = \left\| \frac{1}{n_i} \sum_{k=1}^{n_i} f(I_k^i) - \frac{1}{n_j} \sum_{k=1}^{n_j} f(I_k^j) \right\| \quad (1)$$

where f represents the appearance feature extractor. Other possible choices of pooling method include max pooling, RNN and attention models.

However, the conventional set distance functions like Equation (1) fail to align the evidence appearing in two sets, and thus deteriorate the performance. This is because previous pooling methods aggregate appearance features and generate a fixed-size vector before doing the comparison between two image sets. To resolve this issue, we propose the concept of visual distributional representation, which naturally handles the diversity and uncertainty in a set of images.

3.2 Visual Distributional Representation

We assume that an image set can be treated as samples of a probability distribution in appearance feature space. The purpose of visual distributional representation is to find a proper distribution to describe an image set, and optimize the function for appearance feature extraction, simultaneously. Hence, the basic components of visual distributional representation learning framework include an appearance feature extractor f , a probability estimation method pe , and a downstream task, which provides the supervision signal for the optimization of f .

In the context of multi-shot person re-id task, we use visual distributional representation to describe the appearance of a person, and design a new type of set distance function, which can be formulated as followed:

$$D(S_i, S_j) = d(v_i, v_j) \quad (2)$$

$$v_i = pe(f(I_1^i), f(I_2^i), \dots, f(I_{n_i}^i))$$

where v_i and v_j are the visual distributional representations of S_i and S_j , respectively, and d is a distance measure between two distributions. It is important to define all three basic components: f , pe , and d properly. In this study, we mainly focus on the investigation of probability estimator pe and distributional distance measure d .

The goal of the probability estimator pe is to summarize a set of appearance features by estimating its distribution. It is in parallel to the pooling strategies in conventional person re-id approaches. In this paper, we consider two types of probability estimator: (1) kernel density estimation (KDE) with Dirac delta function $\delta: \hat{v}_i =$

Figure 2: The whole process of visual distributional representation learning.

$\frac{1}{n_i} \sum_{k=1}^{n_i} \delta(f(I_k^i))$ (2) Gaussian estimation (GE): assuming that the samples in appearance feature space follow a multi-variate Gaussian distribution $\hat{v}_i \sim N(m_i, \Sigma_i)$, while mean m_i and covariance Σ_i can be obtained from the appearance features easily.

Since each person tracklet is represented as a distribution, we need a distributional distance measure to calculate the dissimilarity between two tracklets. In this study, we specifically choose Wasserstein distance because it discovers the optimal alignment between two tracklets. More details are discussed in the following.

To sum up, visual distributional representation is proposed to replace the conventional pooling strategies. Comparing to a fixed-size vector representation, it effectively captures the diversity and uncertainty of a set of images. Figure 2 indicates the whole process of learning framework for visual distributional representation

3.3 Visual Distributional Representation Learning with Wasserstein Distance

We design our multi-shot person re-id algorithm based on the concept of visual distributional representation with Wasserstein distance. In this section, we first revisit the basic idea of Wasserstein distance, discuss practical issues of visual distributional representation learning, and describe the details of our model.

3.3.1 Wasserstein Distance Revisit. Given a metric space (χ, d_χ) , where $d_\chi(x, y)$ is a distance function between two elements $x \in \chi$ and $y \in \chi$, the p -th Wasserstein distance between two distributions v_x and v_y can be defined as:

$$W^p(v_x, v_y) = \left(\inf_{\pi \in \Pi(v_x, v_y)} \int_{\chi \times \chi} d_\chi(x, y) d\pi(x, y) \right)^{\frac{1}{p}} \quad (3)$$

where $\Pi(v_x, v_y)$ is the set of all joint probability π on $\chi \times \chi$ with marginals v_x and v_y . In this study, we choose to set $p = 2$ so d_χ is the Euclidean distance.

In the discrete case, we need to compare two empirical measures $\hat{v}_x = \frac{1}{n_x} \sum_{k=1}^{n_x} \delta(x^{(k)})$ and $\hat{v}_y = \frac{1}{n_y} \sum_{l=1}^{n_y} \delta(y^{(l)})$ represented by the uniformly weighted sums of n_x and n_y Diracs delta with mass at positions $x^{(k)}$ and $y^{(l)}$. $\Pi(\hat{v}_x, \hat{v}_y) = \{\sum_{k,l} w_{k,l} \delta(x^{(k)}) \delta(y^{(l)}) \mid \sum_k w_{k,l} = \frac{1}{n_y}, \sum_l w_{k,l} = \frac{1}{n_x}\}$ is the valid set of joint probability distributions, whose elements can be noted as a $n_x \times n_y$ matrix P . Hence, 2-Wasserstein distance becomes:

$$W^2(\hat{v}_x, \hat{v}_y) = \min_{P \in \Pi(\hat{v}_x, \hat{v}_y)} \langle P, M \rangle_F \quad (4)$$

where M is a $n_x \times n_y$ distance matrix and M_{kl} stores the Euclidean distance between $x^{(k)}$ and $y^{(l)}$. Finding the discrete version Wasserstein distance is usually known as optimal transport problem [30].

The discrete version of Wasserstein distance provides a natural way to conduct set-to-set comparison. For each element in a set, it considers only the distances with the nearest neighbors in the other set. When the appearance evidence of re-id is not dense, this approach helps the system automatically discover a proper alignment, and makes set-to-set comparison more robust. In our visual distributional representation learning framework, if we use kernel density estimation (KDE) with Dirac delta function to represent each person tracklet, then the final set distance function would become:

$$D_E(S_i, S_j) = W^2(\hat{v}_i, \hat{v}_j) \quad (5)$$

where $\hat{v}_i = \frac{1}{n_i} \sum_{k=1}^{n_i} \delta(f(I_k^i))$ and $\hat{v}_j = \frac{1}{n_j} \sum_{k=1}^{n_j} \delta(f(I_k^j))$.

3.3.2 Approximation of Wasserstein Distance. The efficiency of Wasserstein distance calculation is an important issue to our proposed model. Finding the exact solution of Equation (4) is equivalent to solving a linear programming problem, which costs $O(n^3 \log(n))$, $n = \max(n_i, n_j)$ [23]. In the training phase, we randomly sample a small number of images for each person and avoid solving a large-scale problem. In the testing phase, however, we need to process all the images of each person in order to achieve the best performance. Practically, the number of images of each person, especially in the gallery set, could be very large. Hence, to make the inference process efficient, we adopt two types of approximation.

The first approach is to smooth the original Wasserstein distance with an entropy regularization term [3]:

$$D_S(S_i, S_j) = \min_{P \in \Pi(\hat{v}_i, \hat{v}_j)} \langle P, M \rangle_F - \frac{1}{\lambda} h(P) \quad (6)$$

where $h(P) = \sum_{k,l} P_{k,l} \log(P_{k,l})$ is the entropy of joint probability. It can be solved by iterating Sinkhorn's update [3], which costs only $O(n^2)$ in terms of computation.

Entropy regularized Wasserstein distance is also called Sinkhorn distance. It still takes the distances between all the image pairs into account, but gives more weight to nearest neighbors. When $\lambda \rightarrow \infty$, it is equivalent to the original Wasserstein distance.

The second approach is to make a parametric assumption: the extracted appearance features of an image set follow a Gaussian distribution. In the case of comparing two Gaussian distributions $\hat{\mu}_i \sim N(m_i, \Sigma_i)$ and $\hat{\mu}_j \sim N(m_j, \Sigma_j)$, Wasserstein distance can be calculated by the following [6]:

$$D_G(S_i, S_j) = W^2(\hat{\mu}_i, \hat{\mu}_j) = ||m_i - m_j||^2 + \text{Tr}||\Sigma_i + \Sigma_j - 2(\Sigma_i^{1/2} \Sigma_j \Sigma_i^{1/2})^{1/2}|| \quad (7)$$

m_i, m_j, Σ_i and Σ_j can be estimated directly by the mean and covariance of appearance features of S_i and S_j . Comparing to entropy regularization, this approach does not require accessing the Euclidean distance matrix M , so it is much more efficient. Obviously, to use this approximation in our framework, we also need to choose the probability estimator with Gaussian assumption (GE).

Figure 3: The architecture of our model in the training and testing phases

3.4 Model Details

Figure 3 illustrates the architecture of our model on the strength of visual distributional representation learning. In the training phase, our model extracts appearance features, estimates the visual distributional representation by KDE with Dirac delta, and optimize the parameters in appearance feature extractor based on training objective function. The training objective is defined with exact 2-Wasserstein distance (Equation (5)). In the testing phase, given a person tracklet as query, our model calculates the dissimilarity between the input tracklet and all tracklets in gallery set. The dissimilarity can be determined by two types of Wasserstein distance approximation, Equation (6) and Equation (7), which require the distributions estimated by KDE with Dirac delta and GE, respectively. To eliminate the influence of outlier images in person tracklet, we apply a moving average filter to appearance feature sequences in the testing phase:

$$f_i^{r(k)} = \frac{1}{K} \sum_{l=1}^K f(I_{k+l}^i) \quad (8)$$

and then estimate the distribution of this tracklet on smoothed appearance features $f_i^{r(k)}$.

3.4.1 Training Objective Function. The overall objective function in the training phase consists of two parts: triplet loss and identification loss.

Let (S_a, S_p, S_n) be a triplet input, where S_a , S_p and S_n are the anchor, positive, and negative examples, respectively. Anchor and positive examples belong to the same identity, while the negative example is from a different one. Triplet loss [31] forces the distance between the positive pair $D_E(S_a, S_p)$ to be smaller than the distance between the negative pair $D_E(S_a, S_n)$ with a margin Δ :

$$L_{triplet} = \sum_{a=1}^B \max(0, D_E(S_a, S_p) - D_E(S_a, S_n) + \Delta) \quad (9)$$

where B is the batch size. In practice, we apply batch-wise hard-negative mining [12], selecting the most distant positive example and most similar negative example within a batch for each anchor.

The identification loss aims to categorize each set of images to the correct identity:

$$L_{ID} = \sum_{a=1}^B \sum_{k=1}^{N_{id}} q_{ak} \log(p(k|F(v_a))) \quad (10)$$

where N_{id} is the total number of identities, and v_a is the visual distributional representation of input image set S_a . $q_{ak} = 1, 0$ indicates if S_a belongs to k -th identity. $\log(p(k|F(v_a)))$ is obtained from a logistic regression layer taking $F(v_a)$, the mean estimation of v_a , as input, which is equivalent to the output of mean pooling strategy.

The total loss is the combination of the two: $L_{total} = L_{triplet} + L_{ID}$. By observing the log during the training phase, we find that the network can usually decrease triplet loss to zero, so we don't need to introduce a hyper-parameter to adjust the balance between two loss functions. Based on this total loss, the whole network is trained in an end-to-end manner using back-propagation.

3.5 Discussion

Attention mechanism [8, 18, 20, 28, 37] is another common approach trying to improve naive pooling methods like mean and max pooling. An additional branch of network is applied to calculate the weight of each image in the set, and assess the final representation by leveraging the weighted summation of the extracted features of the whole set. This type of methods can effectively deal with the noisy elements in an image set. However, the aggregation process of final representation is still independent of other identities in the training data. Thus, it is not likely to capture the sparse evidence for set-to-set comparison. There were also several recent approaches dealing with set-to-set comparison [27, 42] directly. They both introduced dual attention frameworks, performing image pair alignment. However, these frameworks use additional parameters to model the attention mechanism, and increase the risk of over-fitting. In our model, Wasserstein distance naturally incorporates feature alignment, and no additional parameters are required.

In many applications, Wasserstein distance estimation is done by Kantorovich-Rubinstein dual form [14]. It requires finding the solution of an optimization problem in 1-Lipschitz function space, which is often implemented by another deep neural network [2]. This dual form solution is not applicable to our case. We need to estimate the distance between each pair of persons. For each person tracklet, only several hundreds of examples are available at most, which is not enough to train a network.

While we choose Wasserstein distance in this work, our learning framework accepts other types of distance measure between distributions, such as total variance and symmetric KL divergence. One can investigate the advantages of these choices, although it is not the main focus of this work.

4 EXPERIMENTS

In this section, we evaluate our proposed model for multi-shot person re-id on three public datasets: iLIDS-VID [32], MARS [41] and DukeMTMC-VideoReID [34].

4.1 Datasets and Evaluation Protocol

The iLIDS-VID dataset contains 300 identities. For each identity, 2 tracklets are captured from two non-overlapping cameras, respectively. The number of frames of each tracklet ranges from 23 to 192, with an average number of 73. The bounding boxes of tracklets are annotated by humans. The MARS dataset contains 1,261 identities and 20,715 tracklets. The bounding boxes of tracklets are detected and tracked by DPM detector and GMMCP tracker, respectively. There are 6 camera views, and each identity is captured by at least two views. 3,248 distractor tracklets appeared due to false detection or tracking. The DukeMTMC-VideoReID is a subset of DukeMTMC tracking dataset [24]. It contains 4,832 tracklets from 1,812 identities. Each tracklet has 168 frames on average. The bounding boxes are annotated manually.

Following the standard evaluation protocol of iLIDS-VID, we randomly split 50% of identities for training and 50% for testing, and repeat the same experiment for 10 times. For MARS dataset, we apply the same experiment setup in [41], which selects 625 identities for the training set and the remaining for testing. For DukeMTMC-VideoReID, we follow the protocol in [34], using 702 identities for training, 702 identities for testing and the rest as distractors.

The performance of all the methods are reported in Cumulated Matching Characteristics (CMC), which measures the probability that an image set in the first rank k gallery set matches the query image set. For MARS and DukeMTMC-VideoReID datasets, we also report mean Average Precision (mAP) since multiple ground truth matches are available.

4.2 Experiment Setup

We apply two backbone networks: ResNet-50 [11] and DenseNet-121 [13] as appearance feature extractors of our model. Both networks are pretrained on ImageNet. We extract the spatial average pooling of their last convolutional layers, which generate appearance features of 2048 and 1024 dimensions, respectively. During data preprocessing, input bounding box images are resized to 256×128 . We adopt standard data augmentation steps, including horizontal flip and random crop. In the training phase, the number of tracklets in each minibatch is 28, and 4 images are randomly selected for each tracklet. The margin of triplet loss function is set to 0.3. The parameters in appearance feature extractor and identity classifier are optimized using Adaptive Moment Estimation (ADAM) algorithm [15]. The learning rate starts at 0.0003, and decreases by a factor 0.1 for each 100 epochs, until the model finishes training

Methods	Top-1	Top-5	Top-20	mAP
R + mean + Eu	84.3	94.6	97.4	79.7
R + RNN + Eu	83.0	93.7	96.6	77.3
R + Atten + Eu	84.9	94.8	97.9	80.1
R + GE + D_G	87.1	95.7	98.0	81.2
R + KDE + D_S	86.6	95.6	98.0	80.4
D + mean + Eu	84.3	95.2	97.8	77.9
D + RNN + Eu	83.2	94.3	95.6	76.4
D + Atten + Eu	85.2	95.7	97.7	79.6
D + GE + D_G	87.3	96.0	98.1	81.8
D + KDE + D_S	86.7	96.0	98.0	80.5

Table 1: Comparison with common pooling strategies on MARS. R: ResNet-50, D: DenseNet121, Eu: Euclidean Distance, GE: Gaussian estimation, KDE: kernel density estimation, D_G : Equation (7), D_S : Equation (6)

at 400 epochs. In the testing phase, all the images of a tracklet are used to extract the appearance feature set.

4.3 Comparison with Other Pooling Strategies

To understand the efficacy of our proposed visual distributional representation, we compare it with other common pooling methods, including mean pooling, RNN, and temporal attention. For a fair comparison, we adopt the same appearance feature extractors (ResNet-50 and DenseNet-121), and follow the same training scheme as our method. Euclidean distance serves as the dissimilarity measure.

For RNN pooling, we utilize vanilla RNN architecture, and set the number of hidden units to 512. For temporal attention pooling, we add a fully connected layer L_a on top of the appearance feature extractor to calculate the attention scores. The final representation z_i of the i -th image set can be obtained as followed:

$$z_i = \sum_{t=1}^{n_i} a_t f_t$$

$$a_t = \frac{\exp(L_a(f_t))}{\sum_{t=1}^{n_i} \exp(L_a(f_t))} \quad (11)$$

where a_t is the attention score of the t -th image.

Table 1 shows the experiment results of pooling method comparison. We use a 3-dim tuple to note each possible combination of appearance feature, pooling method, and distance function. For example, R+mean+Eu represents the model with Resnet-50 as appearance feature extractor, mean pooling strategy, and Euclidean distance as dissimilarity measure. "KDE" and "GE" refer to two probability estimators in visual distributional representation. D_S and D_G refer to two types of Wasserstein distance approximation: Sinkhorns iteration (Equation (6)) and Gaussian approximation (Equation (7)), respectively.

From the experiment results, we can make the following observations: (1) RNN does not outperform mean pooling baseline. (2) Temporal attention improves performance in terms of both top-k matching accuracy and mAP. The trends are similar when we use ResNet-50 and DenseNet-121 for feature extraction. (3) Our visual distributional representation methods outperform all the other

(a) Experiment results on top-1 matching accuracy

(b) Experiment results on mAP

Figure 4: Experiment results on MARS with different window size K in Equation (8). ResNet-50 is utilized as appearance feature extractor.

pooling strategies. Using Gaussian approximation for Wasserstein distance estimation achieve the best performance. (4) In most of the methods, appearance feature extracted from Resnet-50 provides better performance on mAP, while that from DenseNet-121 performs better on top-k matching accuracy. This set of experiments demonstrates that visual distributional representation is able to preserve more information than conventional pooling strategies in multi-shot person re-id scenario.

4.4 Sensitivity Analysis

There are two hyper-parameters in our proposed model. The first is λ in Equation (6), which regularizes the 2-Wasserstein distance. The second is the window size K of the moving average filter in the testing phase (Equation (8)). To understand the sensitivity of these two hyper-parameters, we apply different values to our model and observe the performances. In this experiment, we choose ResNet-50 as appearance feature extractor.

From Table 2, we can tell that the re-id model receives the best performance with proper level of regularization ($\lambda = 10$ or $\lambda = 20$). When $\lambda = 0$, the entropy regularization term in Equation (6)

	MARS	iLIDS-VID
R+KDE+ D_S	Top-1	Top-1
* $\lambda=0$	86.1	77.9
$\lambda=5$	86.2	78.3
$\lambda=10$	86.7	79.0
$\lambda=20$	86.7	79.4
$\lambda=30$	86.1	78.8
$\lambda=50$	85.8	78.1

Table 2: Top-1 accuracy on MARS and iLIDS-VID with different values of λ in Equation (6). *The results with $\lambda = 0$ is obtained by calculating the average of distance matrix M , instead of executing Sinkhorn iteration.

dominates, and the value of D_S becomes the average of distance matrix M .

Figure 4 illustrates the experiment results on MARS with different window size K . We find that the moving average filter does not affect models with Gaussian assumption (R+GE+ D_G), since the results keep almost the same when $K = 1$. The reason may be that the Gaussian assumption has already caused some smoothing effect. However, the moving average filter is favorable for models with KDE and D_S . We can see that the best performance happens when $K = 4$. The case of over-smoothing ($K = 8$) deteriorates performances for both models.

4.5 Comparison with State-of-the-art Methods

We compare our proposed method to the following previous state-of-the-art methods: TDL [39], CNN+XQDA [41], RNN [21], ASTPN [37], SeeForest [42], LDCAF [17], QAN [20], RQEN [28], TriNet [12], DRA [18], ETAP-Net [34], DuATM [27], TM [9] and STA [8]. As suggested in previous works [18, 27], we use DenseNet-121 for MARS dataset, and ResNet-50 for iLIDS-VID and Duke-VideoReID datasets, respectively.

Table 4 summarizes the performances of all methods on iLIDS-VID dataset. We achieve competitive result comparing to state-of-the-art methods. The reason that DRA [18] performs better may be that they incorporated additional image-based re-id datasets to help with training.

The overall performances of all methods on MARS dataset are reported in Table 3. Our implementation of conventional mean pooling is already competitive. The proposed visual distributional representation can still further improve the performance without introducing additional parameters. Our method achieves the best scores in terms of both top-k matching accuracy and mAP.

In Table 5, we show the performance of all methods on DukeMTMC-VideoReID dataset. Only a few results were published on this dataset. The recent work STA [8] has reported high scores in both CMC and mAP. This set of experiment results indicates that our approach is comparable to this strong baseline.

4.6 Visualization

We show two examples of temporal alignment between person tracklets in Figure 5. Two sets of three person tracklets are picked from iLID-VIDS and MARS datasets, respectively. Among these

Methods	Top-1	Top-5	Top-20	mAP
CNN+XQDA (ECCV2016)	68.3	82.6	89.4	49.3
SeeForest (CVPR2017)	70.6	90.0	97.6	50.7
LDCAF (CVPR2017)	71.8	86.6	93.1	56.5
RQEN (AAAI2018)	73.7	84.9	91.6	51.7
TriNet (Axiv2017)	79.8	91.4	-	67.7
*DRA (CVPR2018)	82.3	-	-	65.8
DuATM (Axiv2018)	81.2	92.5	-	67.7
TM (Axiv2018)	83.3	93.8	97.4	76.6
STA (AAAI2019)	86.3	95.7	97.1	80.8
Ours (D+mean+Eu)	84.3	95.2	97.8	77.9
Ours (D+KDE+ D_S)	86.5	96.0	98.0	80.2
Ours (D+GE+ D_G)	87.3	96.0	98.1	81.8

Table 3: Comparison with state-of-the-art methods on MARS. *: additional training datasets used.

Methods	Top-1	Top-5	Top-20
TDL (CVPR2016)	56.3	87.6	98.3
RNN (CVPR2016)	58	84	96
ASTPN (ICCV2017)	62	86	98
CNN+XQDA (ECCV2016)	53.0	81.4	95.1
SeeForest (CVPR2017)	55.2	86.5	97.0
QAN (CVPR2017)	68.0	86.8	97.4
RQEN (AAAI2018)	76.1	92.9	99.3
*DRA (CVPR2018)	80.2	-	-
Ours (R+mean+Eu)	77.0	93.3	98.4
Ours (R+GE+ D_G)	77.2	93.3	98.4
Ours (R+KDE+ D_S)	79.4	93.8	98.7

Table 4: Comparison to state-of-the-art methods on iLIDS-VID. *: additional training datasets used.

Methods	Top-1	Top-5	Top-20	mAP
*ETAP-Net (CVPR2018)	83.6	94.6	97.6	78.3
STA (AAAI2019)	96.2	99.3	99.6	94.9
Ours (R+mean+Eu)	93.6	98.4	99.1	92.9
Ours (R+KDE+ D_S)	94.2	99.1	99.9	94.2
Ours (R+GE+ D_G)	95.6	99.3	99.9	94.9

Table 5: Comparison to state-of-the-art methods on DukeMTMC-VideoReID. *: The supervised baseline in [34]

three tracklets, two of them belong to the same identity, and the rest one is different. We use color bounding boxes to mark the image pairs with highest joint probability, which is automatically assigned by the optimal P in Equation (6). From Figure 5, one can see that our proposed method attends to different images of a tracklet while comparing to different candidates. Please note that the images without color bounding boxes are manually selected to indicate the appearance of the whole tracklet.

5 CONCLUSION

In this paper, we propose to learn visual distributional representation with Wasserstein distance, and conduct set-to-set comparison

Figure 5: Visualization of the alignment automatically discovered by our proposed method. Image pairs assigned with highest joint probability are marked by bounding boxes in the same color.

for multi-shot person re-identification task. Our approach can effectively discover a proper alignment between input query and gallery examples without additional parameters. Experiment results show that our proposed method outperforms several common feature aggregation strategies such as mean pooling and temporal attention, and achieve competitive/superior performance comparing to previous state-of-the-art approaches.

The proposed visual distributional representation provides a general strategy to summarize a set of images, and can be easily plugged into an end-to-end learning architecture. Our future work will focus on the exploration of downstream tasks other than multi-shot person re-id.

ACKNOWLEDGMENTS

This work was partially supported by the financial assistance award 60NANB17D156 from U.S. Department of Commerce, National Institute of Standards and Technology (NIST). This work was also supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/ Interior Business Center (DOI/IBC) contract number D17PC00340. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation/herein. Disclaimer: The views and conclusions contained herein are those

of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, NIST, DOI/IBC, or the U.S. Government.

REFERENCES

- [1] Ejaz Ahmed, Michael Jones, and Tim K Marks. 2015. An improved deep learning architecture for person re-identification. In *CVPR*. 3908–3916.
- [2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein Generative Adversarial Networks. In *Proceedings of the 34th International Conference on Machine Learning*. 214–223.
- [3] Marco Cuturi. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in neural information processing systems*. 2292–2300.
- [4] Mert Dikmen, Emre Akbas, Thomas S Huang, and Narendra Ahuja. 2010. Pedestrian recognition with a learned metric. In *Asian conference on Computer vision*. Springer, 501–512.
- [5] Shengyong Ding, Liang Lin, Guangrun Wang, and Hongyang Chao. 2015. Deep feature learning with relative distance comparison for person re-identification. *Pattern Recognition* 48, 10 (2015), 2993–3003.
- [6] DC Dowson and BV Landau. 1982. The Fréchet distance between multivariate normal distributions. *Journal of multivariate analysis* 12, 3 (1982), 450–455.
- [7] Charlie Frogner, Chiyuan Zhang, Hossein Mobahi, Mauricio Araya, and Tomaso A Poggio. 2015. Learning with a Wasserstein loss. In *Advances in Neural Information Processing Systems*. 2053–2061.
- [8] Yang Fu, Xiaoyang Wang, Yunchao Wei, and Thomas Huang. 2019. STA: Spatial-Temporal Attention for Large-Scale Video-based Person Re-Identification. In *AAAI*.
- [9] Jiyang Gao and Ram Nevatia. 2018. Revisiting Temporal Modeling for Video-based Person ReID. *arXiv preprint arXiv:1805.02104* (2018).
- [10] Mehrtash Harandi, Mathieu Salzmann, and Mahsa Baktashmotlagh. 2015. Beyond gauss: Image-set matching on the riemannian manifold of pdfs. In *Proceedings of the IEEE International Conference on Computer Vision*. 4112–4120.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*. 770–778.
- [12] Alexander Hermans, Lucas Beyer, and Bastian Leibe. 2017. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737* (2017).
- [13] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. 2017. Densely Connected Convolutional Networks. In *CVPR*, Vol. 1. 3.
- [14] Leonid Vasilievich Kantorovich and Gennady S Rubinstein. 1958. On a space of completely additive functions. *Vestnik Leningrad. Univ* 13, 7 (1958), 52–59.
- [15] D Kinga and J Ba Adam. 2015. A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*.
- [16] Martin Koestinger, Martin Hirzer, Paul Wohlhart, Peter M Roth, and Horst Bischof. 2012. Large scale metric learning from equivalence constraints. In *CVPR. IEEE*, 2288–2295.
- [17] Dangwei Li, Xiaotang Chen, Zhang Zhang, and Kaiqi Huang. 2017. Learning deep context-aware features over body and latent parts for person re-identification. In *CVPR*. 384–393.
- [18] Shuang Li, Slawomir Bak, Peter Carr, and Xiaogang Wang. 2018. Diversity Regularized Spatiotemporal Attention for Video-based Person Re-identification. In *CVPR*. 369–378.
- [19] Shengcai Liao, Yang Hu, Xiangyu Zhu, and Stan Z Li. 2015. Person re-identification by local maximal occurrence representation and metric learning. In *CVPR*. 2197–2206.
- [20] Yu Liu, Junjie Yan, and Wanli Ouyang. 2017. Quality aware network for set to set recognition. In *CVPR*, Vol. 2. 8.
- [21] Niall McLaughlin, Jesus Martinez del Rincon, and Paul Miller. 2016. Recurrent convolutional network for video-based person re-identification. In *CVPR*. 1325–1334.
- [22] Sateesh Pedagadi, James Orwell, Sergio Velastin, and Boghos Boghossian. 2013. Local fisher discriminant analysis for pedestrian re-identification. In *CVPR*. 3318–3325.
- [23] Ofir Pele and Michael Werman. 2009. Fast and robust Earth Mover’s Distances. In *ICCV*, Vol. 9. 460–467.
- [24] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. 2016. Performance Measures and a Data Set for Multi-Target, Multi-Camera Tracking. In *European Conference on Computer Vision workshop on Benchmarking Multi-Target Tracking*.
- [25] Ergys Ristani and Carlo Tomasi. 2018. Features for Multi-Target Multi-Camera Tracking and Re-Identification. In *CVPR*.
- [26] Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. 2018. Wasserstein Distance Guided Representation Learning for Domain Adaptation. In *AAAI*.
- [27] Jianlou Si, Honggang Zhang, Chun-Guang Li, Jason Kuen, Xiangfei Kong, Alex C Kot, and Gang Wang. 2018. Dual attention matching network for context-aware feature sequence based person re-identification. *arXiv preprint arXiv:1803.09937* (2018).

- [28] Guanglu Song, Biao Leng, Yu Liu, Congrui Hetang, and Shaofan Cai. 2018. Region-based Quality Estimation Network for Large-scale Person Re-identification. *AAAI* (2018).
- [29] Yumin Suh, Jingdong Wang, Siyu Tang, Tao Mei, and Kyoung Mu Lee. 2018. Part-Aligned Bilinear Representations for Person Re-identification. *arXiv preprint arXiv:1804.07094* (2018).
- [30] Cédric Villani. 2008. *Optimal transport: old and new*. Vol. 338. Springer Science & Business Media.
- [31] Jiang Wang, Yang Song, Thomas Leung, Chuck Rosenberg, Jingbin Wang, James Philbin, Bo Chen, and Ying Wu. 2014. Learning fine-grained image similarity with deep ranking. In *CVPR*. 1386–1393.
- [32] Taiqing Wang, Shaogang Gong, Xiatian Zhu, and Shengjin Wang. 2014. Person re-identification by video ranking. In *ECCV*. Springer, 688–703.
- [33] Wen Wang, Ruiping Wang, Zhiwu Huang, Shiguang Shan, and Xilin Chen. 2015. Discriminant analysis on Riemannian manifold of Gaussian distributions for face recognition with image sets. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2048–2057.
- [34] Yu Wu, Yutian Lin, Xuanyi Dong, Yan Yan, Wanli Ouyang, and Yi Yang. 2018. Exploit the unknown gradually: One-shot video-based person re-identification by stepwise learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5177–5186.
- [35] Tong Xiao, Shuang Li, Bochao Wang, Liang Lin, and Xiaogang Wang. 2017. Joint detection and identification feature learning for person search. In *CVPR*. IEEE, 3376–3385.
- [36] Fei Xiong, Mengran Gou, Octavia Camps, and Mario Sznaier. 2014. Person re-identification using kernel-based metric learning methods. In *ECCV*. Springer, 1–16.
- [37] S. Xu, Y. Cheng, K. Gu, Y. Yang, S. Chang, and P. Zhou. 2017. Jointly Attentive Spatial-Temporal Pooling Networks for Video-Based Person Re-identification. In *ICCV*.
- [38] Jianbo Ye, James Z Wang, and Jia Li. 2016. A simulated annealing based inexact oracle for Wasserstein loss minimization. *arXiv preprint arXiv:1608.03859* (2016).
- [39] Jinjie You, Ancong Wu, Xiang Li, and Wei-Shi Zheng. 2016. Top-push video-based person re-identification. In *CVPR*. 1345–1353.
- [40] Rui Zhao, Wanli Ouyang, and Xiaogang Wang. 2013. Person re-identification by salience matching. In *ICCV*. 2528–2535.
- [41] Liang Zheng, Zhi Bie, Yifan Sun, Jingdong Wang, Chi Su, Shengjin Wang, and Qi Tian. 2016. MARS: A Video Benchmark for Large-Scale Person Re-identification. In *ECCV*. Springer.
- [42] Zhen Zhou, Yan Huang, Wei Wang, Liang Wang, and Tieniu Tan. 2017. See the forest for the trees: Joint spatial and temporal recurrent neural networks for video-based person re-identification. In *CVPR*. IEEE, 6776–6785.