# Unified Framework for Automated Person Re-identification and Camera Network Topology Inference in Camera Networks

Yeong-Jun Cho, Jae-Han Park\*, Su-A Kim\*, Kyuewang Lee and Kuk-Jin Yoon Computer Vision Laboratory, GIST, South Korea

{yjcho, qkrwogks, suakim, kyuewang, kjyoon}@gist.ac.kr

#### Abstract

Person re-identification in large-scale multi-camera networks is a challenging task because of the spatio-temporal uncertainty and high complexity due to large numbers of cameras and people. To handle these difficulties, additional information such as camera network topology should be provided, which is also difficult to automatically estimate. In this paper, we propose a unified framework which jointly solves both person re-id and camera network topology inference problems. The proposed framework takes general multi-camera network environments into account. To effectively show the superiority of the proposed framework, we also provide a new person re-id dataset with full annotations, named SLP, captured in the synchronized multicamera network. Experimental results show that the proposed methods are promising for both person re-id and camera topology inference tasks.

# 1. Introduction

Person re-identification (*i.e.* re-id) is the task of *automatically* recognizing and identifying a person across multiple views in multi-camera networks, and has been studied for last decades. Nevertheless, the re-id in large-scale multi-camera networks still remains a challenging task because of the large spatio-temporal ambiguity and high complexity due to large numbers of cameras and people. Especially, it becomes more challenging when camera views are not overlapped each other. As shown in Fig. 1, spatio-temporal uncertainty due to the unknown geometrical relationship between cameras in a multi-camera network makes the re-id difficult. Unless some prior knowledge about the camera



Figure 1. Challenges of large-scale person re-identification : spatio-temporal uncertainties between cameras.

network is given, re-id should be done by thoroughly matching the person of interest with all other people appeared in the other camera within some time interval. This exhaustive search method is slow in general and shows unsatisfactory results in many cases because it is hard to find a correct match among a large number of candidates — among the large number of candidates, there might be many people having similar appearances with the person of interest.

However, most of the previous works [11, 13] conduct the exhaustive search to re-identify people across multiple cameras, relying solely on the appearance information of people. These methods work quite well when the numbers of cameras and people are small, but cannot effectively handle the aforementioned challenges in the large-scale problem. To resolve the complexity problem and to improve the accuracy of the large-scale person re-id, the number of matching candidates of a person of interest should be constrained and reduced by inferring and exploiting the spatiotemporal relation between cameras, referred to as the camera network topology. In recent years, several camera network topology inference methods [20, 22] have been proposed. Those methods infer the topology of a camera network based on the simple occurrence correlation between entering and exiting events of people. However, since they do not perform any appearance-based validation for topol-

<sup>\*</sup> These authors contributed equally to this work. This work has been done while Su-A Kim and Kyuewang Lee were with Computer Vision Lab in GIST. Currently Su-A Kim is with Intel VCI (Saarland Informatics Campus) and supported by the European Commission through the H2020-MSCA Distributed 3D Object Design (Grant No. 642841) since April 2017. Currently Kyuewang Lee is with ASRI in seoul national university. Kuk-Jin Yoon is the corresponding author.

ogy inference, the inferred topology can be inaccurate in crowded scenes.

The main idea of this paper is that the camera network topology inference and person re-id can be solved jointly while complementing each other. Based on this idea, we propose a unified framework which automatically solves both person re-id and camera network topology inference problems together. To the best of our knowledge, this is the first attempt to solve both problems jointly. In the proposed framework, we first infer the initial camera network topology using only highly reliable re-id results obtained by the proposed multi-shot matching method. This initial topology is used to improve the person re-id results, and the improved re-id results are then used to refine the camera network topology. This procedure is repeated until the estimated camera network topology converges. Once we estimate the reliable camera network topology in the training stage, we can utilize it for the online person re-id and update the camera network topology with time.

To sum up, we propose a multi-shot person re-id method which exploits time-efficient random forest (Sec. 3.1). We also propose fast and accurate camera network topology inference method in Sec. 3.2. It is worthy to note that our proposed framework runs fully automatic with minimal prior knowledge about the environments. Besides the proposed methods, we also provide a new synchronized large-scale person re-id dataset named SLP (Sec. 4). To validate our unified framework, we extensively evaluate the performance of the proposed method and compare with other state-of-the-art methods.

# 2. Previous Works

Person re-id methods can be categorized into noncontextual and contextual methods as summarized in [3]. In general, non-contextual methods rely only on appearances of people and measure visual similarities between people to establish correspondences, while contextual methods exploit additional contexts such as human pose prior, camera parameters, geometry, and camera topology.

**Non-contextual Methods** In order to identify people across non-overlapping views, most of works generally rely on appearances of people by utilizing appearance-based matching methods with feature learning or metric learning. For the feature learning, many works [13, 19, 30] have tried to design visual descriptors to well describe the appearance of people. Regarding the metric learning, several methods such as KISSME [16] and LMNN-R [12] have been proposed and applied to the re-id problem [25, 28].

Although many non-contextual methods have improved the performance of person re-id, the challenges such as spatio-temporal uncertainty between non-overlapping cameras and high computational complexity still remain.

**Contextual Methods** Several works [1, 10, 29] using hu-

man pose priors have been proposed recently. Exploiting human poses mitigates ambiguities by seeking pose variations of people but spatio-temporal ambiguities between cameras still remain.

To resolve the spatio-temporal ambiguities, many works have tried to employ camera network topology and camera geometry. Several works [6, 15, 24] assume that the camera network topology is given, and show the effectiveness of the topological information. However, the topology is not given in the real-world scenario; thus, many works have tried to infer the camera network topology in an unsupervised manner. Makris *et al.* [22] proposed a topology inference method that simply observes entering and exiting events of targets and measures correlations between the events to establish the camera network topology. This method was extended in [8,23,26]. Similarly, Loy *et al.* [20] also proposed topology inference methods to understand multi-camera activity by measuring correlation or mutual information between simple activity patterns.

The aforementioned topology inference methods [8, 20, 22, 26], so-called, event-based approaches, are practical since they do not require any appearance matching steps such as re-id or inter-camera tracking for topology inference. However, the topology inferred by the event-based approach can be inaccurate since the topology may be inferred from false event correlations.

## 3. Proposed Unified Framework

Figure 2 illustrates the proposed unified framework for person re-id and camera network topology inference. In the proposed framework, we first train random forest-based person classifiers (Sec. 3.1) for efficient person re-id. Subsequently, we jointly estimate and refine the camera network topology and person re-id results (Sec. 3.2) using the trained random forests.

## 3.1. Random Forest-based Person Re-identification

Most of previous works mainly focused on enhancing reid performance. However, when handling a large number of people, time complexity is also very important for building a practical re-id framework. To this end, we utilize a random forest algorithm [5] for the multi-shot person re-id and incorporate it into our framework. We denote the k-th appearance of person i in camera  $c_A$  as  $\mathbf{v}_{i,k}^{CA}$ . A set of the appearances of people in the camera is expressed as,

$$\mathcal{D}^{c_A} = \left\{ \left( \mathbf{v}_{i,k}^{c_A}, y_i \right) | 1 \le i \le N^{c_A}, 1 \le k \le K_i^{c_A} \right\}, \quad (1)$$

where  $y_i$  is the label of person *i* is the number of people in camera  $c_A$  and  $K_i^{c_A}$  is the number of appearances of person *i* in camera  $c_A$ . We then train a random forest classifier using the appearance set  $\mathcal{D}^{c_A}$ . After the random forest



Figure 2. The proposed unified framework for person reidentification and camera network topology inference.

classifier is trained, we have the probability distribution of classification as  $p^{c_A}(y|\mathbf{v})$ .

To obtain a multi-shot person re-id result, we test multiple appearances of each person and average the multiple results as  $p^{c_A}(y|\mathbf{v}_j^{c_B}) = \frac{1}{K_j^{c_B}} \sum_{l=1}^{K_j^{c_B}} p^{c_A}(y|\mathbf{v}_{j,l}^{c_B})$ , where  $K_j^{c_B}$  is the number of appearances of the probe. Among the probability distribution  $p^{c_A}(y|\mathbf{v}_j^{c_B})$ , we choose a final matched label  $y_i^*$  which maximizes  $p^{c_A}(y_i|\mathbf{v}_j^{c_B})$ . As the result of the multiple appearance matching test, we have a corresponding pair  $(\mathbf{v}_{y_i^*}^{c_A}, \mathbf{v}_j^{c_B})$  between camera  $c_A$  and  $c_B$ . Finally, we calculate a similarity score of the corresponding pair by selecting the smallest matching score as in [13]. We denote the similarity score as  $S\left(\mathbf{v}_{y_i^*}^{c_A}, \mathbf{v}_j^{c_B}\right)$ . The score lies on [0, 1] and it is used in Sec. 3.2 for inferring the topology. The tree structure of the random forest method makes the multi-shot test very fast. Besides the superiority of the computational cost, our method gives high person re-id accuracy as shown in Sec. 5.1.

#### 3.2. Camera Network Topology Initialization

Camera network topology represents spatio-temporal relations and connections between cameras in the network. It involves with the inter-camera transition distributions between two cameras, which denote transition distributions of objects across cameras according to time, and represents the strength of connectivity between cameras. In general, the topology is represented as a graph G = (V, E), where vertices V denote cameras and edges E denote inter-camera transition distributions as shown in Fig. 5 (b).

#### 3.2.1 CAM-to-CAM topology inference

First, we estimate transition distributions between cameras to build the CAM-to-CAM topology. In this work, we estimate the transition distributions based on the person re-id results. We first split a whole group of people into several sub-groups according to their time-stamps and train a series of random forest classifiers with time window T. Next, we search correspondences of people disappeared in a camera using the trained random forest classifiers of other cameras. Initially, we have no transition distributions between cameras to utilize. Hence, we consider every pair of cameras in the camera network and search correspondences within wide time interval. When a person disappears at time t in a certain camera, we search the correspondence of the person from the other cameras within time range [t - T, t + T].

When these initial correspondences are given, we infer transition distributions between cameras using highly reliable correspondences only. We regard a correspondence as a reliable one with a high similarity score when  $S(\mathbf{v}_{y_i^*}^{CA}, \mathbf{v}_j^{CB}) > \theta_{sim}$ . Transition distribution inference procedure is as follows: (1) calculating time difference between correspondences and making a histogram of the time difference; (2) normalizing the histogram by the total number of reliable correspondences. We denote the transition distribution as  $p(\Delta t)$ . Figure 3 shows two distributions: Fig. 3 (a) comes from a pair of cameras having strong connection, and Fig. 3 (b) from a pair of cameras having a weak or no connection.

**Connectivity Check** Based on the estimated transition distributions, we automatically identify whether a pair of cameras is connected or not. We assume that the transition distribution follows a normal distribution if there is a topological connection. Based on this assumption, we fit a Gaussian model  $N(\mu, \sigma^2)$  to the distribution  $p(\Delta t)$  and measure the connectivity of a pair of cameras based on the following observations:

• Variance of  $p(\Delta t)$ : In general, most of people re-appear around the certain transition time  $\mu$ ; therefore, the variance of the transition distribution ( $\sigma^2$ ) is not extremely large and the distribution shows a clear peak.

• Fitting error: Although the distribution comes from a pair cameras with a weak connection, the variance of the distribution can be small and the distribution can have a clear peak due to noise. In order to measure the connectivity robust to noise, we consider the model fitting error  $\mathcal{E}(p(\Delta t)) \in [0, 1]$  calculated by R-squared statistics.

Based on the above observations, we newly define a connectivity confidence between a pair of cameras as

$$conf\left(p\left(\Delta t\right)\right) = e^{-\sigma} \cdot \mathcal{E}\left(p\left(\Delta t\right)\right).$$
(2)

The connectivity confidence lies on [0, 1]. We regard a pair of cameras as a valid link when  $conf(p(\Delta t)) > \theta_{conf}$ . Compared with a previous method [22] which only considers the variance of a distribution, our method is more robust to noise of distributions. Using the defined confidences, we check every pair of cameras and reject invalid links as shown in Fig. 5. We can see that many camera pairs in the



Figure 3. Examples of estimated transition distributions with connectivity confidences.

camera network have weak connection; therefore, we can greatly reduce computation time and save resources. Only the valid pairs of cameras are proceeding to the next step.

#### 3.2.2 Zone-to-Zone topology inference

In this step, we estimate transition distributions between zones in cameras and build a Zone-to-Zone topology. For each camera, a set of entry and exit zones is automatically learned by [21]. Note that, we only consider exit-to-entry zone pairs when two zones belong to different cameras. Other pairs of zones such as exit-to-exit, entry-to-entry, and entry-to-exit are not considered.

A person disappeared at an exit zone at time t is likely to appear at entry zone in a different camera within the certain time interval T. Therefore, we search the correspondence of the disappeared person from entry zones in different cameras within time range [t, t + T]. Similarly as in Sec. 3.2.1, we train a series of random forest classifiers for each entry zone and measure connectivity confidences of all possible pairs of zones using only reliable correspondences. Through this step, many invalid pairs of zones between cameras are ignored. In the next section, we iteratively update the valid links between zones and build a camera topology map of the camera network.

# **3.2.3** Iterative update of person re-identification and camera network topology

After the Zone-to-Zone topology inference, we have an initial topology map between every pair of zones in the camera network. However, the initial topology map can be inaccurate, since it is inferred by noisy initial re-id results. As mentioned before, camera network topology information and re-id results can be used for each other: the inferred topology of the camera network can enhance the person reid performance, and person re-id can assist the topology inference. Therefore, we update person re-id results and the camera network topology in an iterative manner as follows:

• Step 1. Update the time window T. The initial time window T was set quite wide, but now we can narrow it down based on the inferred topology. To this end, we find lower

and upper time bounds  $(T_L, T_U)$  of the transition distribution  $p(\Delta t)$  with a constant R as,

$$p\left(T_L \le \Delta t \le T_U\right) = \frac{R}{100}.$$
(3)

We set R as 95, following 3-sigma rule, in order to cover the most of the distribution (95%) and ignore some outliers (5%). Then, using the obtained time bounds, the time window T is updated as,

$$T = \frac{1}{1 - \mathcal{E}\left(p\left(\Delta t\right)\right)} \left(T_U - T_L\right),\tag{4}$$

where  $\mathcal{E}(p(\Delta t))$  is a topology fitting error rate. When the fitting error is large, the time window T becomes large. Thanks to our update strategy, we can avoid the overfitting of the topology during the update steps.

- Step 2. Re-train a series of random forests of an entry zone with the updated time window T.
- Step 3. Find correspondences of disappeared people at an exit zone. Based on the topology, a person disappeared at time t at the exit zone is expected to appear around the time (t + μ) at the entry zone of the other camera. Using the topological information, we search the correspondence of the person from a trained random forest having the center of time slot close to (t + μ).
- Step 4. Update a topology using reliable correspondences with a high similarity score S(**v**<sup>cA</sup><sub>y<sup>cA</sup><sub>i</sub></sub>, **v**<sup>cB</sup><sub>j</sub>) > θ<sub>sim</sub>.

This procedure (Step 1 – Step 4) is repeated until the topology converges. The above procedure improves the performance of re-id as well as the accuracy of the topology inference. We empirically set parameters  $\theta_{sim}$ ,  $\theta_{conf}$ , T as 0.7, 0.4, 600 based on an extensive evaluation.

# 4. A New Person Re-id Dataset : SLP

To validate the performance of person re-id methods, numerous datasets have been published. For example, [14] was constructed by two cameras and contains 632 people. Each camera provides one single image for one person. On the contrary, [4] includes 150 people captured from eight cameras. In this dataset, each camera provides multiple images for a person. However, despite the outburst of published datasets, none of them reflect practical large-scale surveillance scenarios, in which (1) video frames captured from multiple synchronized cameras are available, and (2) both the numbers of people and cameras are large.

Most of the public datasets include a small number of people (# IDs < 200) [2,4,9,32] or cameras (# cam < 5) [7, 9, 14, 17, 27, 32]. Moreover, some datasets provide single-shot of each person [14, 20] or do not provide annotation information of people (*track gt*) throughout the entire video sequences [2,9,17,31,32]. Furthermore, there are only a few

Table 1. Details of our new dataset: SLP.

Index	CAM 1	CAM 2	CAM 3	CAM 4	CAM 5	CAM 6	CAM 7	CAM 8	CAM 9	Total
# ID	256	661	1,175	243	817	324	516	711	641	2,632
# frames	19,545	65,518	104,639	41,824	78,917	79,974	93,978	53,621	42,347	580,363
# annotated box	47,870	205,003	310,262	65,732	307,156	160,367	78,259	176,406	117,087	1,468,142
Duration	2h 13m	2h 12m	2h 22m	2h	2h 21m	2h	2h 38m	2h 29m	2h 28m	-



(a) Layout of a camera network



(b) Example frames of nine cameras

Figure 4. A new synchronized large-scale person reidentification dataset: SLP.

datasets which provide camera synchronization information or time stamps of all frames (*sync*) [7].

In this paper, we provide a new synchronized large-scale person re-id dataset called SLP constructed for practical large-scale surveillance scenarios. The main characteristics of our dataset are as follows: The total number of people in the dataset is 2,632. The layout of the camera network and example frames are shown in Fig. 4. It provides extracted feature descriptor of each person<sup>1</sup>. The ground truth of every person is available. Table 1 shows the details of SLP. It is available on online. https://sites.google.com/view/yjcho/project-pages/re-id\_topology.

# 5. Experimental Results

#### **Experimental settings**

Since we mainly focus on person re-identification and camera topology inference problems, we assume that person detection and tracking results are given. We divide our dataset into two subsets according to time: The first subset contains 1-hour data starting from the global start time (AM 11:20). It is used in an camera network topology training stage. The latter subset including the remaining data is utilized in a person re-id test stage. We used LOMO feature [18] to describe the appearances of people. Note that our method can adopt any kind of feature extraction methods.

#### **Evaluation methodology**

To evaluate the performance of person re-id, we measure the re-id accuracy (Re-id acc) defined as  $\frac{TP}{T_{gt}}$ , where TP is the number of true matching results and  $T_{gt}$  is the total number of ground truth re-id pairs in the camera network. To evaluate the accuracy of the camera network topology, we measure topology distance (Top dist). When an inferred transition distribution and a ground truth are given as  $(p(\Delta t) \sim N(\mu, \sigma^2), p_{gt}(\Delta t) \sim N(\mu_{gt}, \sigma_{gt}^2))$ , we defined the topology distance based on *Bhattacharyya* distance which measures the difference between two probability distributions as  $d_B(p, p_{gt}) = -\ln\left(\int \sqrt{p(\Delta t)}p_{gt}(\Delta t) \, \mathrm{d}\Delta t\right)$ . If there are multiple links in the camera network, we measure each evaluation metric for all links and average them to get the final topology distance.

#### 5.1. Camera Network Topology Training Results

#### CAM-to-CAM topology inference result

For all of camera pairs, we illustrate a color map of estimated CAM-to-CAM connectivity confidence in Fig. 5 (a). Each row and column indicate the index of the camera. When the confidence value is greater than  $\theta_{conf}$ , we regard the corresponding camera pair as a valid link. As a result, the valid camera links are drawn as Fig. 5 (b). Each vertex indicates the index of the camera and valid links are represented by edges. Unfortunately, CAM6 failed to be linked to CAM5. That is because the size of the person image patches is very small due to far distance from the camera; therefore it is hard to distinguish the appearances of people. In addition, CAM6 is quite isolated from other cameras.

# Camera topology training results

Figure 6 (a) represents the accuracy of person re-id in each of proposed training steps such as CAM-to-CAM, Zone-to-Zone, and iterative update steps (Sec. 3.2.1–3.2.3).

<sup>&</sup>lt;sup>1</sup>In this version, we do not provide the entire video frames but provide extracted feature descriptors due to legal problem. However we will provide entire video frames in the near future.



Figure 7. Comparison of inferred transition distributions and ground truth. First row: Valid link (Exit: CAM8, ZONE1) – (Entry: CAM9, ZONE2). Second row: Invalid link (Exit: CAM3, ZONE1) – (Entry: CAM7, ZONE2).



Figure 5. Results of CAM-to-CAM connectivity check. (a) Connectivity confidence map. (b) Valid camera links of two methods (*solid line*: true link, *dotted line*: missing).



Figure 6. Results of the person re-identification through the camera topology training (with iteration).

The accuracy of person re-id is 28.54% at the beginning, but it is consistently improved by using inferred and refined camera topology information. As a result, our method reaches 62.55% accuracy at the last step of the topology training. In addition, it took only 112.46 seconds to conduct both person re-id and camera topology inference tasks with a large number of people in the nine cameras (using Intel i7 CPU in MATLAB). Figure 6 (b) shows the comparison with a conventional approach which fully compares multiple appearances between people and exhaustively searches the correspondences of people between the entry/exit zones without using camera topology information. It shows lower performance (52.06% person re-id accuracy) compared to the proposed method, and moreover it takes much more

Table 2. Valid Zone-to-Zone links and ground-truths.

							U				
Exit	Entry	$\mu$	$\mu_{gt}$	$\sigma$	$\sigma_{gt}$	Exit	Entry	$\mu$	$\mu_{gt}$	$\sigma$	$\sigma_{gt}$
C1,Z1	C2,Z5	34.4	34.7	6.25	6.04	C2,Z5	C1,Z1	40.4	40.4	7.62	5.93
C2,Z2	C3,Z1	36.7	36.3	8.03	5.79	C3,Z1	C2,Z2	37.6	37.0	10.3	8.90
C3,Z2	C5,Z6	-0.42	-0.57	3.49	3.23	C5,Z6	C3,Z2	0.70	1.59	3.43	2.32
C3,Z3	C7,Z3	4.8	4.3	4.8	3.5	C7,Z3	C3,Z3	3.75	4.68	2.16	3.04
C4,Z4	C5,Z2	30.2	30.1	13.4	12.5	C5,Z2	C4,Z4	39.5	28.6	3.82	14.8
C7,Z1	C8,Z2	28.2	28.4	21.3	6.36	C8,Z2	C7,Z1	31.9	30.0	2.41	4.02
C8,Z1	C9,Z2	11.6	11.7	4.82	4.24	C9,Z2	C8,Z1	10.5	10.5	4.03	4.08

time (337.27 seconds) than ours.

A list of valid Zone-to-Zone links inferred by the proposed methods is summarized in Table. 2, and the overall results of our method are close to ground-truth  $N(\mu_{gt}, \sigma_{gt}^2)$ . The previous methods [8, 22, 23] showed unclear and noisy distributions for both valid and invalid links as shown in Fig. 7 (a-c). On the other hand, our results are very similar with the ground truth (Fig. 7 (d-e)).

## 5.2. Person Re-identification Test Results

Based on the inferred camera topology in the training stage, we conducted person re-id for the remaining sequence and compared with two different approaches. The first approach estimates the camera topology based on Ex*haustive search* in the training stage (11:20 - 12:20) and uses the inferred topology for the re-id test (12:21 - 13:20). Note that this approach exploits inferred topology but still fully compares multiple appearances of people to find correspondences in the test stage. The second approach based on True matching estimates the camera topology using ground truth re-id pairs in the training stage and perform the re-id test in the same way with our method. As shown in Table. 3, the performance of our re-id test is comparable to that of True matching. Our method also outperforms Exhaustive search in terms of both re-id and topology accuracies. It supports that our iterative topology update and re-id methods are effective and complement each other.

-		•	-
Training	Test stage		
Re-id acc	Top dist	Re-id acc	_
52.1 %	5.620	65.6 %	-
62.5 %	0.076	72.3 %	_
100 %	0	75.6 %	_
	Training Re-id acc 52.1 % 62.5 % 100 %	Training stage           Re-id acc         Top dist           52.1 %         5.620           62.5 %         0.076           100 %         0	Test stage           Test stage         Test stage           Re-id acc         Top dist         Re-id acc           52.1 %         5.620         65.6 %           62.5 %         0.076         72.3 %           100 %         0         75.6 %

Table 3. Performance comparison in training & test stages.

# 6. Conclusions

In this paper, we proposed a unified framework to automatically solve both person re-id and camera network topology inference problems. Besides, in order to validate the performance of person re-id in the practical largescale surveillance scenarios, we provided a new person reid dataset called SLP. We qualitatively and quantitatively evaluated and compared the performance of the proposed framework with state-of-the-art methods. The results show that the proposed framework is promising for both person re-id and camera topology inference and superior to other frameworks in terms of both speed and accuracy.

# Acknowledgement

This work was supported by Institute for Information & communications Technology Promotion(IITP) grant funded by the Korea government(MSIP)(2014-0-00059, Development of Predictive Visual Intelligence Technology).

#### References

- S. Bak, F. Martins, and F. Bremond. Person re-identification by pose priors. In *IS&T/SPIE Electronic Imaging*, 2015. 2
- [2] D. Baltieri, R. Vezzani, and R. Cucchiara. 3dpes: 3d people dataset for surveillance and forensics. In *MA3HO*, 2011. 4
- [3] A. Bedagkar-Gala and S. K. Shah. A survey of approaches and trends in person re-identification. *Image and Vision Computing*, 2014. 2
- [4] A. Bialkowski, S. Denman, S. Sridharan, C. Fookes, and P. Lucey. A database for person re-identification in multicamera surveillance networks. In *DICTA*, 2012. 4
- [5] L. Breiman. Random forests. Machine learning, 2001. 2
- [6] Y. Cai and G. Medioni. Exploring context information for inter-camera multiple target tracking. In WACV, 2014. 2
- [7] W. Chen, L. Cao, X. Chen, and K. Huang. An equalised global graphical model-based approach for multi-camera object tracking. arXiv:1502.03532, 2015. 4, 5
- [8] X. Chen, K. Huang, and T. Tan. Object tracking across nonoverlapping views by learning inter-camera transfer models. *Pattern Recognition*, 2014. 2, 6
- [9] D. S. Cheng, M. Cristani, M. Stoppa, L. Bazzani, and V. Murino. Custom pictorial structures for re-identification. In *BMVC*, 2011. 4
- [10] Y.-J. Cho and K.-J. Yoon. Improving person re-identification via pose-aware multi-shot matching. In CVPR, 2016. 2

- [11] J. V. Davis, B. Kulis, P. Jain, S. Sra, and I. S. Dhillon. Information-theoretic metric learning. In *ICML*, 2007. 1
- [12] M. Dikmen, E. Akbas, T. S. Huang, and N. Ahuja. Pedestrian recognition with a learned metric. In ACCV. 2011. 2
- [13] M. Farenzena, L. Bazzani, A. Perina, V. Murino, and M. Cristani. Person re-identification by symmetry-driven accumulation of local features. In *CVPR*, 2010. 1, 2, 3
- [14] D. Gray, S. Brennan, and H. Tao. Evaluating appearance models for recognition, reacquisition, and tracking. In *PETS*, 2007. 4
- [15] O. Javed, Z. Rasheed, K. Shafique, and M. Shah. Tracking across multiple cameras with disjoint views. *ICCV*, 2003. 2
- [16] M. Koestinger, M. Hirzer, P. Wohlhart, P. M. Roth, and H. Bischof. Large scale metric learning from equivalence constraints. In *CVPR*, 2012. 2
- [17] W. Li, R. Zhao, T. Xiao, and X. Wang. Deepreid: Deep filter pairing neural network for person re-identification. In *CVPR*, 2014. 4
- [18] S. Liao, Y. Hu, X. Zhu, and S. Z. Li. Person re-identification by local maximal occurrence representation and metric learning. In *CVPR*, 2015. 5
- [19] C. Liu, S. Gong, C. C. Loy, and X. Lin. Person re-identification: What features are important? In ECCV, 2012. 2
- [20] C. C. Loy, T. Xiang, and S. Gong. Time-delayed correlation analysis for multi-camera activity understanding. *IJCV*, 2010. 1, 2, 4
- [21] D. Makris and T. Ellis. Automatic learning of an activitybased semantic scene model. In AVSS, 2003. 4
- [22] D. Makris, T. Ellis, and J. Black. Bridging the gaps between cameras. In CVPR, 2004. 1, 2, 3, 6
- [23] C. Niu and E. Grimson. Recovering non-overlapping network topology using far-field vehicle tracking data. In *ICPR*, 2006. 2, 6
- [24] A. Rahimi, B. Dunagan, and T. Darrell. Simultaneous calibration and tracking with a network of non-overlapping sensors. In *CVPR*, 2004. 2
- [25] P. M. Roth, M. Hirzer, M. Köstinger, C. Beleznai, and H. Bischof. Mahalanobis distance learning for person reidentification. In *Person Re-Identification*. 2014. 2
- [26] C. Stauffer. Learning to track objects through unobserved regions. In WACV/MOTIONS, 2005. 2
- [27] T. Wang, S. Gong, X. Zhu, and S. Wang. Person reidentification by video ranking. In ECCV, 2014. 4
- [28] K. Q. Weinberger, J. Blitzer, and L. K. Saul. Distance metric learning for large margin nearest neighbor classification. In *NIPS*, 2005. 2
- [29] Z. Wu, Y. Li, and R. J. Radke. Viewpoint invariant human re-identification in camera networks using pose priors and subject-discriminative features. *TPAMI*, 2015. 2
- [30] R. Zhao, W. Ouyang, and X. Wang. Learning mid-level filters for person re-identification. In CVPR, 2014. 2
- [31] L. Zheng, H. Zhang, S. Sun, M. Chandraker, and Q. Tian. Person re-identification in the wild. arXiv preprint arXiv:1604.02531, 2016. 4
- [32] W.-S. Zheng, S. Gong, and T. Xiang. Associating groups of people. In *BMVC*, 2009. 4