Multi-Target, Multi-Camera Tracking

Louis Israel and Amanda Bolton

Abstract

Albeit numerous techniques perform well in single camera following, multi-camera following stays a difficult issue with less consideration. DukeMTMC is an enormous scope, all around explained multi-camera following benchmark which gains extraordinary ground in this field. This report is devoted to quickly present our technique on DukeMTMC and show that straightforward progressive grouping with very much prepared individual re-distinguishing proof highlights can get great outcomes on this dataset.

1 Introduction

Multi-target, multi-camera tracking is an useful but challenging problem in computer vision. By combining the trajectory information from multiple cameras, it overcomes the limitation of the field of view in single camera and allows people to conduct a more detailed analysis of the target. This technology can be used in video surveillance, sports analysis and many other fields. However, compared to single camera tracking, multi-camera tracking has less location information and less strict time constraints, which makes it even more difficult. Some algorithms have been proposed to solve this knotty problem. E Ristani et al. [11] formulated the tracking problem as correlation clustering problem. YT Tesfaye et al. [13] formulated the tracking problem as finding constrained dominant sets from a graph. A Maksai et al. [10] put the tracking problem under global, non-Markovian behavioral constraints. In this report, we study a strategy based on the hierarchical clustering and show that it works well with person re-identification features.

2 DukeMTMC Dataset

DukeMTMC [11] dataset is a multi-target, multi-camera tracking dataset. It consists of 85 minutes of 1080p, 60 fps videos recorded by 8 calibrated and synchronized static outdoor cameras, together with manual annotations as ground truth. This dataset is split into one training/validation set and two test sets, the easy test set and the hard test set. The training/validation set is 50 minutes long and the test sets totalled 35 minutes long. More than 7,000 single camera trajectories and over 2,000 identities appeared in the dataset in total, which is a very large number compared to other existing

MTMC datasets, allowing to better evaluate the performance of tracking methods.

3 Proposed Method

In this section we show the details of our method. It includes four main steps: detection, feature extraction, single camera tracking and multi-camera tracking.

- Detection: we studied both public detections and private detections. We use a Faster R-CNN [2] detector to generate private detections.
- Feature Extraction: we use a person re-identification model to extract the appearance features.
- Single camera tracking: we merge neighbor bounding boxes into tracklets (short but reliable trajectories) first, then generate trajectories by hierarchically clustering tracklets.
- Multi-camera tracking: we merge the trajectories by a hierarchical clustering without distance matrix updating.

3.1 Detection

Private person detections are generated using the Faster R-CNN object detector, which works in two stages: First, it extracts candidate regions of interest (RoI) using a Region Proposal Network (RPN); Then it performs binary classification into person and background after applying ROI pooling on the candidate regions. We trained our detection model on COCO [8] dataset and some other private datasets. Note that our detection model have not fine-tuned on DukeMTMC dataset though it can definitely improve the overall performance. We set a threshold value of 0.9 and perform non-maximum suppression (NMS) for each frame with iou threshold 0.3. Our detections can be download at <https://pan.baidu.com/s/1hrPxYNq>. 8 text files are provided. Each one contains detections in each camera with the following line format (each line represents a bounding box): [camera, frame, left, top, right, bottom, confidence].

For public detections, we set a threshold value of 0 and the nms threshold as 0.3.

3.2 Person Re-Identification

In order to combine the bounding boxes together with their appearance, we trained a person re-identification network

Benchmark	big model	small model
CUHK03	92.5	88.1
CUHK-SYSU	96.0	92.1
DukeMTMC-reID	81.9	78.9
Market1501	93.9	91.8
MARS	81.9	83.5

Table 1: Top1 accuracy of our reid models without re-ranking

to give an appearance feature for each bounding box. We trained our network following the AlignedReID [22] which gives an impressive result on reid task recently. Our model is trained on 7 public reid datasets, including CUHK03 [7], CUHK-SYSU [21], Market1501 [23], MARS [24], VIPeR [3], PRID [4], as well as DukeMTMC-reID [25] which is an extension of DukeMTMC.

In practice, we found that the reid model with higher top1 accuracy does not always lead to a better multi-camera tracking result. Hence, we trained a smaller and faster reid model, whose top1 accuracy is slightly lower, but performs as well as the big one for multi-camera tracking.

Table 1 shows the top 1 accuracy of our two reid models on the datasets, without re-ranking.

3.3 Single Camera Tracking

Single camera tracking is a popular research topic in computer vision and a lot of algorithms have been proposed [6]. In this work, we proposed a simple method to aggregate bounding boxes into some trajectories in single camera and these trajectories will be further merged by our multi-camera tracking algorithm. Our single camera tracking strategy is a near online [1] algorithm. We first merge bounding boxes of adjacent frames into tracklets by the Kuhn-Munkras algorithm [5] in a small window. Then we use a hierarchical clustering to merge the tracklets into trajectories.

Merging Neighbors Let F_t and F_{t+1} be two adjacent frames, where $t_{start} \leq t \leq t_{end}$ and t_{start}, t_{end} are the range of the tracking window. Let $\{B_i^t\}_{i=1\dots N_t}$ be the bounding boxes in F_t , where N_t is the number of bounding boxes in F_t . For each possible assignment between $\{B_i^t\}_{i=1\dots N_t}$ and $\{B_j^{t+1}\}_{j=1\dots N_{t+1}}$, we define a weight W_{ij}^t as the likelihood that the two bounding boxes are the same person. W_{ij}^t is decided by two factors: the location closeness and the appearance similarity. For location closeness, we force the IoU of the bounding boxes greater than a threshold, which is set to 0.5 in this work. For appearance similarity, we apply the L2 distance of their reid features. Hence, W_{ij}^t is computed by the following equation:

$$W_{ij}^t = s - \|f_i^t - f_j^{t+1}\|_2 - g(IoU(B_i^t, B_j^{t+1}), 0.5), \quad (1)$$

where s is a given threshold, f_i^t and f_j^{t+1} are the reid features of B_i^t and B_j^{t+1} , IoU is their intersection over union, and g is defined as

$$g(x, k) = \begin{cases} \infty, & x < k, \\ 0, & x \geq k, \end{cases} \quad (2)$$

Then it can be solved as a maximum weighted bipartite matching problem by the Kuhn-Munkras algorithm. Bounding boxes that are not matched are considered as false alarms and discarded.

Hierarchical Clustering Hierarchical clustering is used to merge trajectories and tracklets together. We first calculate the distance matrix between them. The distance between a trajectory and a tracklet is comprised of three parts: The distance for appearance similarity, the distance of the separation part and the distance of the overlapping part.

For single camera tracking, the reid features of the same person are quite close to each other in nearby frames. Hence the distance of appearance similarity is decided by the neighboring frames around the endpoints of trajectories. In more details, given two trajectories T_i and T_j , assuming the last frame of T_i , which is noted as t_i , is before the last frame of T_j , then find the first frame after t_i in the trajectory T_j , denoted as t_u . The reid features of T_i and T_j are calculated as the average of reid features in neighboring frames around t_i and t_u separately. The distance of appearance similarity is the L2 distance of the reid features of T_i and T_j . For separation part, the speed is calculated by the bounding boxes in frame t_i of trajectory T_i and frame t_u of trajectory T_j , where t_i and t_u have the same definition as before. If the speed is less than a threshold, the distance is set to zero; otherwise it is set to ∞ . For overlapping part, the average IoU is calculated. The distance is set to zero or ∞ , depending on whether the average IoU is greater than a threshold.

Then, as clustering progresses, trajectories are merged and the distance matrix updated until the minimum distance is greater than a threshold. In general, this threshold can be the same as s in equation 1. After merged, trajectory interpolation and smoothing are performed but new reid feature of interpolated bounding boxes will not be extracted for the consideration of speed.

3.4 Multi-Camera Tracking

Our multi-camera tracking is a greedy procedure or a hierarchical clustering without distance matrix updating when clustering. We put all the trajectories in all the cameras together, for each trajectory, we average the reid features to get a more robust one. The distance matrix is defined as the Euclidean distance between the averaged features. Then we perform re-ranking [26] on the distance matrix. Re-ranking is a critical step to improve the accuracy of person re-identification. In our experiments, re-ranking can improve IDF1 of multi-camera tracking on training/validation set by about 2%-3%. After re-ranking, we merge the trajectories according to the distance from small to large until minimum distance reaches a threshold. In addition, in order to make tracking more accuracy, some restrictions are added to the merging procedure:

- Two trajectories should be separated less than one minute. One minute is a statistic on the training/validation set.
- Person can not appear in different cameras at the same time, except for 3 pairs of cameras with overlapping areas, which are camera 2 and camera 8, camera 3 and camera 5 and camera 5 and camera 7.

Note in multi-camera tracking, we also do in-camera trajectories merging. We found it is very helpful for long-term tracking since the averaged reid features are more robust.

4 Results

We show benchmark results for our method and compare them with prior ones in this section. The measures and results are given by the benchmark available on motchallenge.net.

Table 2, Table 3, Table 4 and Table 5 compared the trackers performance on easy test set and hard test set, where MTMC_ReIDp is our method on public detections and MTMC_ReID is our method on private detections. We also compared to some anonymous methods. It is shown that, our method has achieved the state of the art in all the test settings.

Tracker	IDF1↑	IDP↑	IDR↑
MTMC_ReIDp	74.4	84.4	66.4
MYTRACKER	65.4	71.1	60.6
MTMC_CDSC [13]	60.0	68.3	53.5
W. Liu et al. [9]	55.5	78.89	44.6
lx_b	58.0	72.6	48.2
BIPCC [12]	56.2	67.0	48.4
dirBIPCC	52.1	62.0	45.0
PT_BIPCC [10]	34.9	41.6	30.1

Table 2: Multi-Camera Tracking Results of Easy Test Set on Public Detections

Tracker	IDF1↑	IDP↑	IDR↑
MTMC_ReID	83.2	85.2	81.2
DeepCC	82.0	84.4	79.8

Table 3: Multi-Camera Tracking Results of Easy Test Set on Private Detections

Tracker	IDF1↑	IDP↑	IDR↑
MTMC_ReIDp	65.6	78.1	56.5
MTMC_CDSC [13]	50.9	63.2	42.6
MYTRACKER	50.1	58.3	43.9
lx_b	48.3	60.6	40.2
BIPCC [12]	47.3	59.6	39.2
dirBIPCC	45.0	56.3	37.5
PT_BIPCC [10]	32.9	41.3	27.3

Table 4: Multi-Camera Tracking Results of Hard Test Set on Public Detections

The hard test set and the easy test set are split according to the scale and the crowd, hence in the hard test set, location information will be more unreliable. In such a condition, more reliable reid feature can help a lot on tracking, which is the reason why we have a larger advantage on the hard test set.

Tracker	IDF1↑	IDP↑	IDR↑
MTMC_ReID	74.0	81.4	67.8
DeepCC	68.5	75.9	62.4

Table 5: Multi-Camera Tracking Results of Hard Test Set on Private Detections

Figure 1 shows some examples of out results. According to our observation, most of the failures in tracking came from different people with similar appearances. Some cases may be even difficult for human to distinct.

Note that our method is not an online algorithm, especially the multi-camera tracking part, but we can easily transform it to a near online one like single-camera tracking part does. We believe this will not lose too much accuracy.

5 Conclusion

In this report, we have presented our method for multi-target, multi-camera tracking on DukeMTMC dataset. We use hierarchical clustering algorithm to merge trajectories. The distance matrix is mainly based on high quality reid features. Benchmark results on easy test set and hard set set show that our method achieves new state of the art results. We hope this report can help the development of multi-camera tracking.

References

- [1] Choi, W. 2015. Near-online multi-target tracking with aggregated local flow descriptor. *2015 IEEE International Conference on Computer Vision (ICCV)* 3029–3037.
- [2] Girshick, R. B. 2015. Fast r-cnn. *2015 IEEE International Conference on Computer Vision (ICCV)* 1440–1448.
- [3] Gray, D., and Tao, H. 2008. Viewpoint invariant pedestrian recognition with an ensemble of localized features. In *ECCV*.
- [4] Hirzer, M.; Belezna, C.; Roth, P. M.; and Bischof, H. 2011. Person re-identification by descriptive and discriminative classification. In *SCIA*.
- [5] Kuhn, H. W. 2010. The hungarian method for the assignment problem. In *50 Years of Integer Programming*.
- [6] Leal-Taixé, L.; Milan, A.; Schindler, K.; Cremers, D.; Reid, I.; and Roth, S. 2017. Tracking the trackers: An analysis of the state of the art in multiple object tracking. *arXiv:1704.02781 [cs]*. arXiv: 1704.02781.
- [7] Li, W.; Zhao, R.; Xiao, T.; and Wang, X. 2014. Deep-reid: Deep filter pairing neural network for person re-identification. *2014 IEEE Conference on Computer Vision and Pattern Recognition* 152–159.
- [8] Lin, T.-Y.; Maire, M.; Belongie, S. J.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *ECCV*.

Figure 1: Examples of our results. (a), (b) are successful examples and (c), (d) are failure examples. The curves in the right are the trajectories and the dotted parts represent interpolated trajectories between the cameras. For better effect, in all the examples, cameras 1-8 have been ren numbered to cameras 8, 7, 6, 5, 3, 1, 2, 4.

- [9] Liu, W.; Camps, O.; and Sznai, M. 2017. Multi-camera multi-object tracking. *arXiv preprint arXiv:1709.07065*.
- [10] Maksai, A.; Wang, X.; Fleuret, F.; and Fua, P. 2017. Non-markovian globally consistent multi-object tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2544–2554.
- [11] Ristani, E.; Solera, F.; Zou, R.; Cucchiara, R.; and Tomasi, C. 2016a. Performance measures and a data set for multi-target, multi-camera tracking. In *European Conference on Computer Vision workshop on Benchmarking Multi-Target Tracking*.
- [12] Ristani, E.; Solera, F.; Zou, R. S.; Cucchiara, R.; and Tomasi, C. 2016b. Performance measures and a data set for multi-target, multi-camera tracking. In *ECCV Workshops*.
- [13] Tesfaye, Y. T.; Zemene, E.; Prati, A.; Pelillo, M.; and Shah, M. 2017. Multi-target tracking in multiple non-overlapping cameras using constrained dominant sets. *arXiv preprint arXiv:1706.06196*.
- [14] Y. Liu, W. Wang, J. Chambers, V. Kilic, and A. Hilton, “Particle flow smc-phd filter for audio-visual multi-speaker tracking,” in *International Conference on Latent Variable Analysis and Signal Separation*. Springer, 2017, pp. 344–353.
- [15] Y. Liu, W. Wang, and Y. Zhao, “Particle flow for sequential monte carlo implementation of probability hypothesis density,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 4371–4375.
- [16] Y. Liu, A. Hilton, J. Chambers, Y. Zhao, and W. Wang, “Non-zero diffusion particle flow SMC-PHD filter for audio-visual multi-speaker tracking,” *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- [17] Y. Liu, W. Wang, and V. Kilic, “Intensity particle flow smc-phd filter for audio speaker tracking,” *arXiv preprint arXiv:1812.01570*, 2018.
- [18] Y. Liu, V. Kılıç, J. Guan, and W. Wang, “Audio-visual particle flow smc-phd filtering for multi-speaker tracking,” *IEEE Transactions on Multimedia*, vol. 22, no. 4, pp. 934–948, 2019.

- [19] Y. Liu, Q. Hu, Y. Zou, and W. Wang, "Labelled non-zero particle flow for smc-phd filtering," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5197–5201.
- [20] Y. Liu, W. Wang, J. Chambers, V. Kilic, and A. Hilton, "Particle flow SMC-PHD filter for audio-visual multi-speaker tracking," in *Proc. IEEE Intl Conf. Latent Variable Analysis and Signal Separation*, Mar. 2017, pp. 344–353.
- [21] Xiao, T.; Li, S.; Wang, B.; Lin, L.; and Wang, X. 2017. Joint detection and identification feature learning for person search. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 3376–3385.
- [22] Zhang, X.; Luo, H.; Fan, X.; Xiang, W.; Sun, Y.; Xiao, Q.; Jiang, W.; Zhang, C.; and Sun, J. 2017. Aligned-dreid: Surpassing human-level performance in person re-identification. *arXiv preprint arXiv:1711.08184*.
- [23] Zheng, L.; Shen, L.; Tian, L.; Wang, S.; Wang, J.; and Tian, Q. 2015. Scalable person re-identification: A benchmark. *2015 IEEE International Conference on Computer Vision (ICCV)* 1116–1124.
- [24] Zheng, L.; Bie, Z.; Sun, Y.; Wang, J.; Su, C.; Wang, S.; and Tian, Q. 2016. Mars: A video benchmark for large-scale person re-identification. In *ECCV*.
- [25] Zheng, Z.; Zheng, L.; and Yang, Y. 2017. Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In *Proceedings of the IEEE International Conference on Computer Vision*.
- [26] Zhong, Z.; Zheng, L.; Cao, D.; and Li, S. 2017. Re-ranking person re-identification with k-reciprocal encoding. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 3652–3661.