

MessyTable: Instance Association in Multiple Camera Views

Zhongang Cai^{1*}, Junzhe Zhang^{1,2*}, Daxuan Ren^{1,2}, Cunjun Yu¹, Haiyu Zhao¹, Shuai Yi¹, Chai Kiat Yeo², and Chen Change Loy²

¹ SenseTime Research

{caizhongang,yucunjun,zhaohaiyu,yishuai}@sensetime.com

² Nanyang Technological University

{junzhe001,daxuan001}@e.ntu.edu.sg, {asckyeo,ccloy}@ntu.edu.sg

Abstract. We present an interesting and challenging dataset that features a large number of scenes with messy tables captured from multiple camera views. Each scene in this dataset is highly complex, containing multiple object instances that could be identical, stacked and occluded by other instances. The key challenge is to associate all instances given the RGB image of all views. The seemingly simple task surprisingly fails many popular methods or heuristics that we assume good performance in object association. The dataset challenges existing methods in mining subtle appearance differences, reasoning based on contexts, and fusing appearance with geometric cues for establishing an association. We report interesting findings with some popular baselines, and discuss how this dataset could help inspire new problems and catalyse more robust formulations to tackle real-world instance association problems.³

1 Introduction

We introduce a new and interesting dataset, MessyTable. It contains over 5,000 scenes, each of which captured by nine cameras in one of the 600 configurations of camera poses. Each scene is shot with a random cluttered background and different lighting conditions with about 30 general objects on average. The objects are chosen arbitrarily from 120 classes of possible instances. Figure 1 depicts some scene examples. The goal is to associate the different objects in a scene, i.e., finding the right match of the same instance across views.

The seemingly easy task on this dataset is surprisingly challenging. The relative pose between two cameras can be large, and therefore, an object may appear to be very different when viewed from different angles. Objects are heavily occluded while some of them can be elevated by other objects in a cluttered scene. Hence, appearance viewed from different cameras is always partial. The problem is further complicated with similar-looking or even identical instances.

* indicates equal contribution.

³ project page: <https://caizhongang.github.io/projects/MessyTable/>.




Fig. 1: MessyTable is a large-scale multi-camera general object dataset with cross-camera association labels. One of the 5,579 scenes is shown in this figure (with only 4 out of 9 cameras depicted). MessyTable is challenging due to severe partial occlusion (instance 2 in camera 9 and instance 31 in camera 2), large camera angle differences, similar or identical-looking instances (instances 31 and 27) and more. Objects are arbitrarily chosen from 120 classes that are commonly found on table. The brand names are pixelated in all illustrations

Solving the aforementioned problems is non-trivial. The geometric constraint is hard to use right away. Multi-view epipolar geometry is ambiguous when a pixel can correspond to all points on the epipolar line in the other view. Homographic projection assumes a reference plane, which is not always available. To associate an object across camera views, a method needs to distinguish subtle differences between similar-looking objects. Fine-grained recognition is still a challenging problem in computer vision. When a scene contains identical-looking objects, the method is required to search for neighbouring cues in the vicinity of the objects to differentiate them. The neighbouring configuration, however, can be occluded and highly non-rigid with changing relative position due to different camera views.

While the method developed from MessyTable can be applied to some immediate applications such as automatic check-out [29] in supermarkets, e.g., leveraging multiple views to prevent counting error on the merchandise, the instance association problem found in this dataset is reminiscent of many real-world problems such as person re-identification or object tracking across views. Both examples of real-world problems require some sort of association between objects, either through appearances, group configurations, or temporal cues. Solving these real-world problems requires one to train a model using domain-specific data. Nonetheless, they still share similar challenges and concerns as to the setting in MessyTable.

MessyTable is not expected to replace the functionalities of domain-specific datasets. It aims to be a general dataset offering fundamental challenges to existing vision algorithms, with the hope of inspiring new problems and encouraging novel solutions. In this paper, apart from describing the details of MessyTable, we

also present the results of applying some baselines and heuristics to address the instance association problem. We also show how a deep learning-based method developed from MessyTable can be migrated to other real-world multi-camera domains and achieve good results.

2 Related Work

Related Problems. Various computer vision tasks, such as re-identification and tracking, can be viewed as some forms of instance association. Despite differences in the problem settings, they share common challenges as featured in MessyTable, including subtle appearance difference, heavy occlusion and viewpoint variation. We take inspirations from methods in these fields for developing a potential solution for multi-camera instance association.

Re-identification (ReID) aims to associate the query instance (e.g., a person) and the instances in the gallery [9]. Association performance suffers from drastic appearance differences caused by viewpoint variation, and heavy occlusion in crowded scenes. Therefore, the appearance feature alone can be insufficient for satisfactory results. In this regard, instead of distinguishing individual persons in isolation (e.g., [26, 38, 40]), an alternative solution proposed by [39] exploits contextual cues: as people often walk in groups in crowded scenes, it associates the same group of people over space and time.

Multi-object Tracking (MOT) is the task to associate instances across sequential frames, leveraging the availability of both appearance features and temporal cues [19, 22, 24]. It suffers from ID switches and fragmentation primarily caused by occlusion [14, 17]. MOT in a multi-camera setting is formally referred to as Multi-Target Multi-Camera Tracking (MTMCT) [20], which also suffers from viewpoint variation in cross-camera tracklet association [8, 12, 37]. In addition, MTMCT with overlapping field of view [4, 32] is similar to MessyTable’s multi-camera setting. Thus, studies conducted on MessyTable might be inspiring for a better cross-camera association performance in MTMCT.

Related Datasets. Many datasets for ReID and MOT offer prominent challenges [10] that are common in real life. For instance, CUHK03 [13], MSMT17 [28], and MOT16 [18] feature occlusion and viewpoint variation, and many other datasets [31, 34] also feature illumination variations.

There are several multi-camera datasets. Though originally designed for different purposes, they can be used for evaluating instance association. MPII Multi-Kinect (MPII MK) [27] is designed for object instance detection and collected on a flat kitchen countertop with nine classes of kitchenwares captured in four fixed views. The dataset features some level of occlusion, but the scenes are relatively simple for evaluating general object association. EPFL Multi-View Multi-Class (EPFL MVMC) [21] contains only people, cars, and buses and is built from video sequences of six static cameras taken at the university campus (with the road, bus stop, and parking slots). WILDTRACK [4], captured with seven static cameras in a public open area, is the latest challenging dataset for multi-view people detection.




Fig. 2: Number of instances per class in MessyTable. MessyTable contains 120 classes: 60 supermarket merchandise, 23 agricultural products, 13 bakery products, and 24 dinning wares. A full list of classes is included in the Supplementary Materials

Compared to existing datasets, MessyTable aims to offer fundamental challenges that are not limited to specific classes. MessyTable also contains a large number of camera setup configurations for a larger variety of camera poses and an abundance of identical instances that are absent in other datasets.

3 MessyTable Dataset

MessyTable is a large-scale multi-camera general object dataset designed for instance association tasks. It comprises 120 classes of common on-table objects (Figure 2), encompassing a wide variety of sizes, colors, textures, and materials. Nine cameras are arranged in 567 different setup configurations, giving rise to 20,412 pairs of relative camera poses (Section 3.1). A total of 5,579 scenes, each containing 6 to 73 randomly selected instances, are divided into three levels of difficulties. Harder scenes have more occlusions, more similar- or identical-looking instances, and proportionally fewer instances in the overlapping areas (Section 3.2). The total 50,211 images in MessyTable are densely annotated with 1,219,240 bounding boxes. Each bounding box has an instance ID for instance association across cameras (Section 3.3). To make MessyTable more representative, varying light conditions and backgrounds are added. Details of the data collection can be found in the Supplementary Materials.

3.1 Variations in View Angles

For a camera pair with a large angle difference, the same instance may appear to be very different (e.g., instance ID 5 in Figure 1) in the two views. Existing multi-camera datasets typically have their cameras installed on static structures [21, 27], even at similar heights [4]. This significantly limits the variation as the data essentially collapses to a very limited set of modes. In contrast, MessyTable has not only a high number of cameras but also a large variation in cameras' poses. The camera stands are arbitrarily adjusted between scenes, resulting in

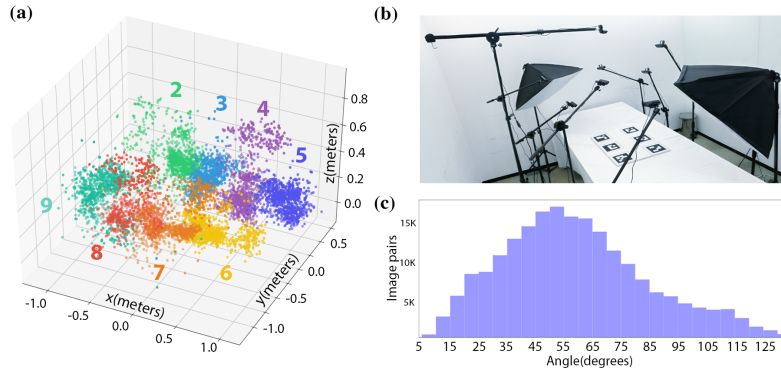


Fig. 3: MessyTable has a rich variety of camera poses. (a) The diverse distribution of camera positions covers the entire space. The positions of cameras #2-9 are projected to the coordinate system of camera #1, visualized by eight clusters of different colors. Camera #1 is fixed providing the bird’s eye view. (b) The studio for data collection. The checkerboard is used for camera extrinsic calibration (see the Supplementary Materials for details); (c) Distribution of the angle differences of image pairs. The angle is computed between line-of-sights of cameras in each image pair. A significant portion of image pairs have larger angles than that in wide-baseline stereo matching, which is typically capped at 45° [3]

an extremely diverse distribution of camera poses and large variance of angle difference between cameras, as shown in Figure 3.

3.2 Variations in Scenes

Partial and Full Occlusions. As shown in Figure 4(a) and (b), partial occlusion results in loss of appearance features [1], making matching more difficult across cameras; full occlusion completely removes the object from one’s view despite its existence in the scene. To effectively benchmark algorithms, in addition to dense clutter, artificial obstacles (such as cardboard boxes) are inserted into the scene.

Similar- or Identical-looking Instances. It is common to have similar and identical objects placed in the vicinity as illustrated in Figure 4(c) and (d). Similar or identical instances are challenging for association. MessyTable has multiple duplicates of the same appearance included in each scene, a unique feature that is not present in other datasets such as [27].

Variations in Elevation. Many works simplify the matching problem by assuming all objects are in contact with the same plane [1, 5, 7, 15]. However, this assumption often does not hold as the scene gets more complicated. To mimic the most general and realistic scenarios in real life, object instances in MessyTable are allowed to be stacked or placed on elevated surfaces as shown in Figure 4(e) and (f).

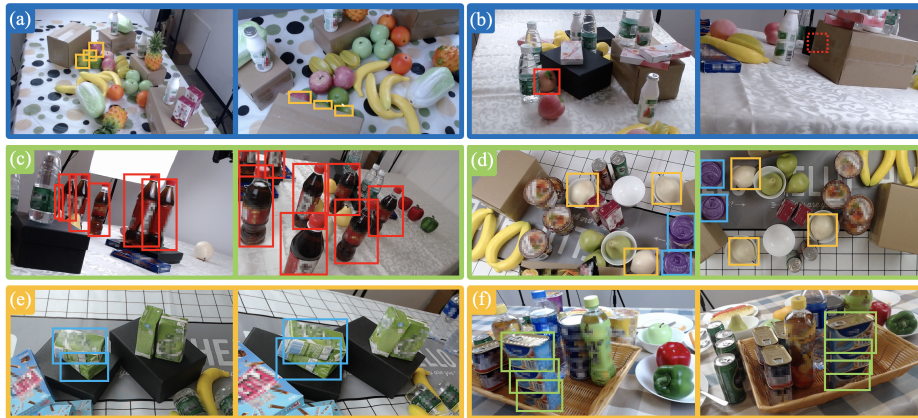


Fig. 4: Visualization of some of the design features that impose challenges for association: (a) Partial occlusion. (b) Full occlusion. (c) Similar-looking instances. (d) Identical-looking instances. (e) and (f): Object stacking and placing objects on elevated surfaces

Data Splits. The dataset is collected with three different complexity levels: Easy, Medium, and Hard, with each subset accounting for 30%, 50%, and 20% of the total scenes. For each complexity level, we randomly partition data equally (1:1:1) into the training, validation, and test sets.

The camera angle differences are similar among the three levels of complexity. But the number of instances, the fraction of overlapped instances and the fraction of identical instances are significantly different as shown in Figure 5. Furthermore, as shown in the example scenes, the larger number of instances in a harder scene significantly increases the chance of heavy occlusion. We empirically show that these challenges undermine the association performance of various methods.

3.3 Data Annotation

We use OpenCV [2] for calibrating intrinsic and extrinsic parameters. As for the instance association annotation, we gather a team of 40 professional annotators and design a three-stage annotation scheme to obtain reliable annotations in a timely manner. The annotators first annotate bounding boxes to enclose all foreground objects (localization stage), followed by assigning class labels to the bounding boxes (classification stage). In the last stage, we develop an interactive tool for the annotators to group bounding boxes of the same object in all nine views and assign them the same instance ID (association stage). Bounding boxes in different views with the same instance ID are associated (corresponding to the same object instance in the scene) and the ID is unique in each view. For each stage, the annotators are split into two groups in the ratio of 4:1 for annotation and quality inspection.

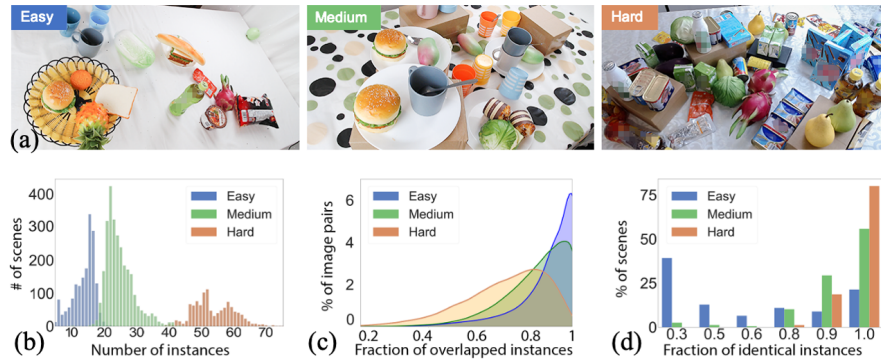


Fig. 5: The comparison amongst Easy, Medium and Hard subsets. (a) Example scenes by difficulty: harder scenes have more severe dense clutter and heavy occlusion. (b) Distributions of scenes by the number of instances per scene: harder scenes have more instances. (c) Distribution of image pairs by the extent of instance overlap: the graph peak of harder scenes is shifted leftward, indicating they have less fraction of instances visible in both cameras. (d) Distribution of scenes by the fraction of identical instances in a scene: harder scenes have more identical-looking instances

It is worth mentioning that the interactive tool has two desirable features to boost efficiency and minimize errors: first, the class labels are used to filter out irrelevant bounding boxes during the association stage; second, the association results correct errors in the classification stage as the disagreement of classification labels from different views triggers reannotation. The details of the data annotation can be found in the Supplementary Materials. In short, MessyTable provides the following annotations:

- intrinsic and extrinsic parameters for the cameras;
- regular bounding boxes (with class labels) for all 120 foreground objects;
- instance ID for each bounding box

4 Baselines

In this section, we describe a few popular methods and heuristics that leverage appearance and geometric cues, and a new baseline that additionally exploits contextual information. We adopt a multiple instance association framework, in which pairwise distances between two sets of instances are computed first. After that, the association is formulated as a maximum bipartite matching problem and solved by the Kuhn-Munkres (Hungarian) algorithm.

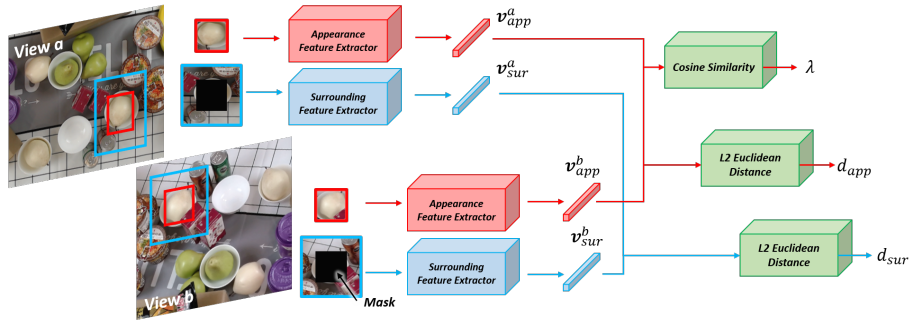


Fig. 6: Appearance-Surrounding Network (ASNet) has an appearance branch (red) and a surrounding branch (blue). Each instance and its immediate surrounding are cropped as the input to their respective branches

4.1 Appearance Information

The appearance feature of the instance itself is the most fundamental information for instance association. As the instances are defined by bounding boxes, which are essentially image patches, we find instance association largely resembles the patch matching (local feature matching) problem.

Local feature matching is one of the key steps for low-level multiple camera tasks [3, 30, 36]. Various hand-crafted feature descriptors, e.g., SIFT [16], have been widely used in this task. We implement a classical matching pipeline including SIFT keypoint description of the instances and K-means clustering for the formation of a visual bag of words (VBoW). The distance between two VBoW representations is computed via chi-square (χ^2).

The application of deep learning has led to significant progress in patch matching [6, 11, 35]. The recent works highlight the use of CNN-based discriminative feature extractors such as DeepDesc [25] to directly output feature vectors, and the distance between two vectors can be computed using L2 distance; MatchNet [11] uses metric networks instead of L2 distance for better performance; DeepCompare [35] proposes to use both multi-resolution patches and metric networks. We use these three works as baselines.

We also implement a standard triplet network architecture with a feature extractor supervised by the triplet loss during training, which has been proven effective to capture subtle appearance difference in face recognition [23]. It is referred to as TripletNet in the experiments and L2 distance is used to measure the feature dissimilarity.

4.2 Surrounding Information

Inspired by Zheng et al. [39] in ReID, which addresses occlusion and view variation by associating a group of people instead of the individuals, we propose

to look outside of the tight bounding box and involve the neighboring information, hoping to tackle not only occlusion and viewpoint variation, but also the existence of similar-looking instances.

The most intuitive idea is to expand the receptive field by a linear scaling ratio, i.e., cropping an area larger than the actual bounding box. This modification on the network input is referred to in the experiments as **zoom-out**, and the ratio as **zoom-out ratio**.

We take further inspiration from the human behavior: one may look at the surroundings for more visual cues *only if* the appearance of the instances themselves are not informative. Hence, we design a simple network (named Appearance-Surrounding Network, Figure 6), which has two branches for appearance and surrounding feature extraction, fused as follows:

$$d^{ab} = (1 - \lambda) \times D_{l_2}(\mathbf{v}_{app}^a, \mathbf{v}_{app}^b) + \lambda \times D_{l_2}(\mathbf{v}_{sur}^a, \mathbf{v}_{sur}^b) \quad (1)$$

$$\lambda = S_c(\mathbf{v}_{app}^a, \mathbf{v}_{app}^b) \quad (2)$$

where a and b are superscripts for two patches, \mathbf{v}_{app} and \mathbf{v}_{sur} are appearance and surrounding feature vectors, respectively. D_{l_2} is L2 distance, and S_c is cosine similarity. λ is the weighting factor to fuse the appearance and surrounding branches. The fusion is designed such that, if the appearance features are similar, the network will place more weight on the surrounding than the appearance. Note that λ is jointly optimized in an end-to-end network; it is not a hyperparameter to be set manually.

4.3 Geometric Methods

Homographic projection-based methods are very popular and used extensively in past works on Multi-Target Multi-Camera Tracking (MTMCT) [32, 33] and Multi-View People Detection [1, 5, 7, 15]. The mid-points of the bottom edges of the bounding boxes [15] are typically projected to a common coordinate system. The instances can thus be associated based on L2 distance between two sets of projected points. It assumes that all instances are placed on one reference 2D plane (e.g., the ground) and this simplification allows for an unambiguous pixel-to-pixel projection across cameras.

We also make use of epipolar geometry, which does not assume a reference plane. For a pair of bounding boxes in two views, the bounding box center in the first view is used to compute an epipolar line in the second view using the calibrated camera parameters. The distance between the bounding box center in second view and the epipolar line is added to the overall distance between the two bounding boxes. It is a soft constraint, since it does not accept or reject the matches, but penalizes unlikely matches by a large distance.

5 Experiments

Unless specified otherwise, we choose ResNet-18 as a light-weight backbone for all models, zoom-out ratio of 2 for models with zoom-out, a mixture of Easy, Medium, and Hard sets are used for training and evaluation.

Table 1: Baseline performances on MessyTable shows a combination of appearance, surrounding and geometric cues is most effective for instance association. †: with metrics learning; ★: upgraded backbone for a fair comparison; ZO: zoom-out; ESC: epipolar soft constraint; ‡: triplet architecture. ↑: the higher the value, the better; ↓: the lower the value, the better; the same applies for the rest of the tables and figures

Model/Method	AP↑	FPR-95↓	IPAA-100↑	IPAA-90↑	IPAA-80↑
Homography	0.049	0.944	0	0	0
SIFT	0.063	0.866	0	0	0
MatchNet†	0.193	0.458	0.010	0.012	0.033
MatchNet†★	0.138	0.410	0.002	0.003	0.010
DeepCompare†	0.202	0.412	0.023	0.025	0.063
DeepCompare†★	0.129	0.402	0.005	0.005	0.01
DeepDesc	0.090	0.906	0.011	0.011	0.018
DeepDesc★	0.171	0.804	0.027	0.032	0.058
TripletNet‡	0.467	0.206	0.168	0.220	0.376
TripletNet‡+ZO	0.430	0.269	0.047	0.062	0.161
ASNet‡	0.524	0.209	0.170	0.241	0.418
ASNet‡+ESC	0.577	0.157	0.219	0.306	0.499

5.1 Evaluation Metrics

AP: Class-agnostic Average Precision is used to evaluate the algorithm’s ability to differentiate positive and negative matches, independent of the choice of the threshold value. All distances are scaled into a range of 0 and 1, and the confidence score is obtained by $1 - x$, where x is the scaled distance.

FPR-95: False Positive Rate at 95% recall [11] is commonly used in patch-based matching tasks and is adopted as a supplement to AP. However, it is worth noting that in the patch matching problem, the positive and negative examples are balanced in the evaluation, which is not the case in our task where the negative examples largely outnumber the positive ones.

IPAA: We introduce a new metric, Image Pair Association Accuracy (IPAA), that evaluates the image-pair level association results instead of the instance-pair level confidence scores. IPAA is computed as the fraction of image pairs with no less than X% of the objects associated correctly (written as IPAA-X). In our experiments, we observed that IPAA is more stringent than AP, making it ideal for showing differences between models with reasonably high AP values. Details can be found in the Supplementary Materials.

5.2 Benchmarking Baselines on MessyTable

In this section, we analyze and provide explanations for the performances of baselines on MessyTable, collated in Table 1.

Homographic projection performs poorly on MessyTable. The result is not surprising as objects in MessyTable can be placed on different elevated surfaces, violating the 2D reference plane assumption that is critical to accurate projection.

The SIFT-based classical method gives a poor performance as the hand-crafted key points tend to cluster around edges and texture-rich areas, leading to an unbalanced distribution. Hence, texture-less instances have very scarce key points, resulting in ineffective feature representation.

Deep learning-based patch matching SOTAs such as MatchNet [11], DeepCompare [35], and DeepDesc [25] give sub-optimal results as they struggle in distinguishing identical objects, which are abundant in MessyTable. Interestingly, our experiments show that a deeper backbone does not improve performance for MatchNet and DeepCompare, as their performances may be bottlenecked by their simple metric network designs. TripletNet with a triplet architecture outperforms these three models with a Siamese architecture by a clear margin (around a 0.25 increment in AP).

We compare TripletNet and ASNet on surrounding information extraction. Naive inclusion of surrounding information (TripletNet+ZO) worsens the association results, as a larger receptive field may introduce noises. In contrast, ASNet trains a specialized branch for the surrounding information to extract meaningful features. Figure 7 visualizes the feature map activations, showing that ASNet effectively learns to use surrounding information whereas TripletNet+ZO tends to focus on the instance itself. However, we highlight that despite a considerable improvement, the ASNet only achieves a moderate AP of 0.524. This leaves a great potential for improvements.

We also show that adding soft geometric constraints to ASNet gives further improvement (around 0.05 improvement in AP), indicating that the geometric information is complementary to appearance and surrounding information. However, the performance, especially in terms of the stringent metric IPAA-100, is still unsatisfactory.

5.3 Effects of View Variation and Scene Complexity

We ablate the challenges featured in MessyTable and their effects on instance association.

We compare the performances of several relatively strong baseline methods at various angle differences in Figure 8. It is observed that the performance by all three metrics deteriorate rapidly with an increase in the angle differences. As shown in Figure 1, large angle difference leads to differences in not only the appearance of an instance itself, but also its relative position within its context.

In addition, we test the same trained model on Easy, Medium, and Hard test sets. The three test sets have the same distribution of angle differences, but different scene complexity in terms of the number of instances, percentage of identical objects, and the extent of overlapping (Figure 5). The performance drops significantly in harder scenes, as shown in Table 2. We offer explanations: first, with more instances on the table, harder scenes contain more occlusion, as




Fig. 7: Visualization of feature map activations after conv_5 of ResNet-18 for TripletNet+ZO (zoom-out) and ASNet. We normalize feature maps to the same scale, then map the values to colors. The higher activation in the surroundings for ASNet indicates that it effectively learns to use surrounding features compared to TripletNet with naive zoom-out




Fig. 8: Effects of angle differences. As the angle difference increases, cross-camera association performance deteriorates for all metrics. (a) AP \uparrow . (b) FPR-95 \downarrow . (c) IPAA-80 \uparrow

shown in Figure 5(a). Second, it is more common to have identical objects closely placed or stacked together, leading to similar surrounding features and geometric distances (Figure 9), making such instances indistinguishable. Third, harder scenes have a smaller fraction of instances in the overlapping area, this may lead to more false positive matches between non-overlapped similar or identical objects, which contributes to higher FPR-95 values.

The above challenges demand a powerful feature extractor that is invariant to viewpoint changes, robust under occlusion, and able to learn the surrounding feature effectively, yet, all baselines have limited capabilities.

5.4 MessyTable as a Benchmark and a Training Source

We further validate the usefulness of MessyTable by conducting experiments on three public multi-camera datasets (Table 3), which gives the following insights:

First, methods that saturate MPII MK and EPFL MVMC are far from saturating MessyTable (Table 1). Note that both datasets have a limited number of classes and instances. Hence, this result highlights the need for MessyTable, a more realistic and challenging dataset for research of instance association.

Table 2: Results on Easy, Medium, and Hard test sets separately. As the scene get more complicated, the performances of all models worsen. Models are trained on a mixture of Easy, Medium and Hard train sets

Subsets	Model	AP \uparrow	FPR-95 \downarrow	IPAA-100 \uparrow	IPAA-90 \uparrow	IPAA-80 \uparrow
Easy	TripletNet	0.618	0.156	0.266	0.342	0.561
	ASNet	0.667	0.151	0.408	0.427	0.660
	ASNet+ESC	0.709	0.122	0.497	0.500	0.734
Medium	TripletNet	0.494	0.211	0.063	0.150	0.290
	ASNet	0.547	0.207	0.101	0.226	0.391
	ASNet+ESC	0.594	0.163	0.151	0.296	0.476
Hard	TripletNet	0.341	0.259	0.003	0.023	0.078
	ASNet	0.396	0.255	0.003	0.026	0.100
	ASNet+ESC	0.457	0.185	0.007	0.048	0.155




Fig. 9: Failure cases: wrong associations are visualized with bounding boxes and instance IDs. Most failure cases occur when identical objects are close to one another such as being (a) placed in a heavy clutter or (b) stacked, leading to indistinguishable surrounding features and similar penalty distances by the soft bipolar constraint. This remains as a challenge worth future research

Second, it is observed that algorithms show consistent trends on MessyTable and other datasets, that is, an algorithm that performs better on MessyTable also performs better on all other datasets. This shows MessyTable can serve as a highly indicative benchmark for multi-camera instance association.

Third, models pretrained on MessyTable consistently perform better than those pretrained on ImageNet, showing MessyTable is a better training source for instance association tasks. Note that EPFL MVMC has three classes (people, cars, and buses) and WILDTRACK is a challenging people dataset. It shows that a model trained on the general objects in MessyTable, learns feature extraction that is readily transferable across domains without feature engineering.

6 Discussion

In this work, we have presented MessyTable, a large-scale multi-camera general object dataset for instance association. MessyTable features the prominent challenges for instance association such as appearance inconsistency due to view angle differences, partial and full occlusion, similar and identical-looking objects,

Table 3: Experiments on three other multi-camera datasets. MessyTable is highly indicative as algorithms benchmarked on MessyTable show consistent trends on all other datasets. Moreover, models pretrained on MessyTable consistently give better performances, showing MessyTable’s potential as a training source

Dataset	Model	Pretraining	AP \uparrow	FPR-95 \downarrow	IPAA-100 \uparrow
MPII MK	TripletNet	ImageNet	0.847	0.196	0.333
	ASNet	ImageNet	0.881	0.167	0.696
	ASNet	MessyTable	0.905	0.119	0.765
EPFL MVMC	TripletNet	ImageNet	0.921	0.056	0.529
	ASNet	ImageNet	0.950	0.038	0.559
	ASNet	MessyTable	0.969	0.031	0.575
WILDTRACK	TripletNet	ImageNet	0.616	0.207	0.095
	ASNet	ImageNet	0.718	0.094	0.304
	ASNet	MessyTable	0.734	0.083	0.321

difference in elevation, and limited usefulness of geometric constraints. We show in the experiments that it is useful in two more ways. First, MessyTable is a highly indicative benchmark for instance association algorithms. Second, it can be used as a training source for domain-specific instance association tasks.

By benchmarking baselines on MessyTable, we obtain important insights for instance association: appearance feature is insufficient especially in the presence of identical objects; our proposed simple baseline, ASNet, incorporates the surrounding information into association and effectively improves the association performance. In addition, we show that epipolar geometry as a soft constraint is complementary to ASNet.

Although the combined use of appearance features, context information, and geometric cues achieves reasonably good performance, ASNet is still inadequate to tackle all challenges. Therefore, we ask three important questions: (1) how to extract stronger appearance, neighbouring and geometric cues? (2) is there a smarter way to fuse these cues? (3) is there more information that we can leverage to tackle instance association?

The experiment results on MessyTable set many directions worth exploring. First, increasing view angle difference leads to a sharp deterioration of instance association performance of all baselines, highlighting the need for research on methods that capture view-invariant features and non-rigid contextual information. Second, methods give poorer performances as the scenes get more complicated; failure cases show that identical instances placed close to each other are extremely difficult to address despite that the strongest baseline already leverages appearance, surrounding and geometric cues. Hence, more in-depth object relationship reasoning may be helpful to distinguish such instances.

Acknowledgements. This research was supported by SenseTime-NTU Collaboration Project, Singapore MOE AcRF Tier 1 (2018-T1-002-056), NTU SUG, and NTU NAP.

References

1. Baqué, P., Fleuret, F., Fua, P.: Deep occlusion reasoning for multi-camera multi-target detection. In: ICCV (2017)
2. Bradski, G.: The OpenCV library. Dr. Dobb's Journal of Software Tools (2000)
3. Caliskan, A., Mustafa, A., Imre, E., Hilton, A.: Learning dense wide baseline stereo matching for people. In: ICCVW (2019)
4. Chavdarova, T., Baqué, P., Bouquet, S., Maksai, A., Jose, C., Bagautdinov, T., Lettry, L., Fua, P., Van Gool, L., Fleuret, F.: WILDTRACK: A multi-camera HD dataset for dense unscripted pedestrian detection. In: CVPR (2018)
5. Chavdarova, T., et al.: Deep multi-camera people detection. In: ICMLA (2017)
6. Csurka, G., Humenberger, M.: From handcrafted to deep local features for computer vision applications. CoRR **abs/1807.10254** (2018)
7. Fleuret, F., Berclaz, J., Lengagne, R., Fua, P.: Multicamera people tracking with a probabilistic occupancy map. PAMI (2007)
8. Gao, J., Nevatia, R.: Revisiting temporal modeling for video-based person ReID. CoRR **abs/1805.02104** (2018)
9. Gong, S., Cristani, M., Yan, S., Loy, C.C.: Person Re-Identification. Springer (2014)
10. Gou, M., Wu, Z., Rates-Borras, A., Camps, O., Radke, R.J., et al.: A systematic evaluation and benchmark for person re-identification: Features, metrics, and datasets. PAMI (2018)
11. Han, X., Leung, T., Jia, Y., Sukthankar, R., Berg, A.C.: Matchnet: Unifying feature and metric learning for patch-based matching. In: CVPR (2015)
12. Hsu, H.M., Huang, T.W., Wang, G., Cai, J., Lei, Z., Hwang, J.N.: Multi-camera tracking of vehicles based on deep features Re-ID and trajectory-based camera link models. In: CVPRW (2019)
13. Li, W., Zhao, R., Xiao, T., Wang, X.: DeepReID: Deep filter pairing neural network for person re-identification. In: CVPR (2014)
14. Li, W., Mu, J., Liu, G.: Multiple object tracking with motion and appearance cues. In: ICCVW (2019)
15. López-Cifuentes, A., Escudero-Viñolo, M., Bescós, J., Carballeira, P.: Semantic driven multi-camera pedestrian detection. CoRR **abs/1812.10779** (2018)
16. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. IJCV (2004)
17. Luo, W., Xing, J., Milan, A., Zhang, X., Liu, W., Zhao, X., Kim, T.K.: Multiple object tracking: A literature review. CoRR **abs/1409.7618** (2014)
18. Milan, A., Leal-Taixé, L., Reid, I., Roth, S., Schindler, K.: MOT16: A benchmark for multi-object tracking. CoRR **abs/1603.00831** (2016)
19. Milan, A., Rezatofghi, S.H., Dick, A., Reid, I., Schindler, K.: Online multi-target tracking using recurrent neural networks. In: AAAI (2017)
20. Ristani, E., Solera, F., Zou, R., Cucchiara, R., Tomasi, C.: Performance measures and a data set for multi-target, multi-camera tracking. In: ECCV (2016)
21. Roig, G., Boix, X., Shitrit, H.B., Fua, P.: Conditional random fields for multi-camera object detection. In: ICCV (2011)
22. Sadeghian, A., Alahi, A., Savarese, S.: Tracking the untrackable: Learning to track multiple cues with long-term dependencies. In: ICCV (2017)
23. Schroff, F., Kalenichenko, D., Philbin, J.: FaceNet: A unified embedding for face recognition and clustering. In: CVPR (2015)
24. Schuster, S., Vernaza, P., Choi, W., Chandraker, M.: Deep network flow for multi-object tracking. In: CVPR (2017)

25. Simo-Serra, E., Trulls, E., Ferraz, L., Kokkinos, I., Fua, P., Moreno-Noguer, F.: Discriminative learning of deep convolutional feature point descriptors. In: ICCV (2015)
26. Sun, Y., Zheng, L., Yang, Y., Tian, Q., Wang, S.: Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In: ECCV (2018)
27. Susanto, W., Rohrbach, M., Schiele, B.: 3D object detection with multiple kinects. In: ECCV (2012)
28. Wei, L., Zhang, S., Gao, W., Tian, Q.: Person transfer GAN to bridge domain gap for person re-identification. In: CVPR (2018)
29. Wei, X.S., Cui, Q., Yang, L., Wang, P., Liu, L.: RPC: A large-scale retail product checkout dataset. CoRR **abs/1901.07249** (2019)
30. Winder, S., Hua, G., Brown, M.: Picking the best DAISY. In: CVPR (2009)
31. Xu, Y., Zhou, X., Chen, S., Li, F.: Deep learning for multiple object tracking: a survey. IET Computer Vision (2019)
32. Xu, Y., Liu, X., Liu, Y., Zhu, S.C.: Multi-view people tracking via hierarchical trajectory composition. In: CVPR (2016)
33. Xu, Y., Liu, X., Qin, L., Zhu, S.C.: Cross-view people tracking by scene-centered spatio-temporal parsing. In: AAAI (2017)
34. Ye, M., Shen, J., Lin, G., Xiang, T., Shao, L., Hoi, S.C.: Deep learning for person re-identification: A survey and outlook. CoRR **abs/2001.04193** (2020)
35. Zagoruyko, S., Komodakis, N.: Learning to compare image patches via convolutional neural networks. In: CVPR (2015)
36. Zbontar, J., LeCun, Y.: Computing the stereo matching cost with a convolutional neural network. In: CVPR (2015)
37. Zhang, Z., Wu, J., Zhang, X., Zhang, C.: Multi-target, multi-camera tracking by hierarchical clustering: recent progress on DukeMTMC project. CoRR **abs/1712.09531** (2017)
38. Zhao, H., Tian, M., Sun, S., Shao, J., Yan, J., Yi, S., Wang, X., Tang, X.: Spindle net: Person re-identification with human body region guided feature decomposition and fusion. In: CVPR (2017)
39. Zheng, W.S., Gong, S., Xiang, T.: Associating groups of people. In: BMVC (2009)
40. Zhou, Y., Shao, L.: Aware attentive multi-view inference for vehicle re-identification. In: CVPR (2018)