

PANDA: A Gigapixel-level Human-centric Video Dataset

Xueyang Wang^{1*}, Xiya Zhang^{1*}, Yinheng Zhu^{1*}, Yuchen Guo^{1*}, Xiaoyun Yuan¹, Liuyu Xiang¹,
Zerun Wang¹, Guiguang Ding¹, David Brady², Qionghai Dai¹, Lu Fang¹✉
¹Tsinghua University ²Duke University




Figure 1. A representative video *Marathon* of PANDA dataset. The characteristic of joint wide field-of-view and high spatial resolution empowers the large-scale, long-term, and multi-object visual analysis.

Abstract

We present PANDA, the first gigaPixel-level human-centric video dataset, for large-scale, long-term, and multi-object visual analysis. The videos in PANDA were captured by a gigapixel camera and cover real-world scenes with both wide field-of-view ($\sim 1 \text{ km}^2$ area) and high-resolution details (\sim gigapixel-level/frame). The scenes may contain 4k head counts with over $100\times$ scale variation. PANDA provides enriched and hierarchical ground-truth annotations, including 15,974.6k bounding boxes, 111.8k fine-grained attribute labels, 12.7k trajectories, 2.2k groups and 2.9k interactions. We benchmark the human detection and tracking tasks. Due to the vast variance of pedestrian pose, scale, occlusion and trajectory, existing approaches are challenged by both accuracy and efficiency. Given the uniqueness of PANDA with both wide FoV and high resolution, a new task of interaction-aware group detection is introduced. We design a ‘global-to-local zoom-in’ framework, where global trajectories and local interactions are simultaneously encoded, yielding promising results. We believe PANDA will contribute to the community of artificial intelligence and praxeology by understanding human behaviors and interactions in large-scale real-world scenes. PANDA Website: <http://www.panda-dataset.com>.

* These authors have contributed equally to this work.

✉ Corresponding author. Mail: fanglu@sz.tsinghua.edu.cn.

This work is supported in part by Natural Science Foundation of China (NSFC) under contract No. 61722209, 6181001011, 61971260 and U1936202, in part by Shenzhen Science and Technology Research and Development Funds (JCYJ20180507183706645).

1. Introduction

It has been widely recognized that the recent conspicuous success of computer vision techniques, especially the deep learning based ones, rely heavily on large-scale and well-annotated datasets. For example, ImageNet [58] and CIFAR-10/100 [64] are important catalyst for deep convolutional neural networks [33, 41], Pascal VOC [26] and MS COCO [47] for common object detection and segmentation, LFW [34] for face recognition, and Caltech Pedestrians [21] and MOT benchmark [51] for person detection and tracking. Among all these tasks, human-centric visual analysis is fundamentally critical yet challenging. It relates to many sub-tasks, e.g., pedestrian detection, tracking, action recognition, anomaly detection, attribute recognition etc., which attract considerable interests in the last decade [55, 9, 46, 68, 50, 63, 18, 71]. While significant progress has been made, **there is a lack of the long-term analysis of crowd activities at large-scale spatio-temporal range with clear local details.**

Analyzing the reasons behind, existing datasets [47, 21, 51, 28, 56, 6] suffer an inherent trade-off between the wide FoV and high resolution. Taking the football match as an example, a wide-angle camera may cover the panoramic scene, yet each player faces significant scale variation, suffering very low spatial resolution. Whereas one may use a telephoto lens camera to capture the local details of the particular player, the scope of the contents will be highly restricted to a small FoV. Even though the multiple surveillance camera setup may deliver more information, the requisite of re-identification on scattered video clips highly af-

fects the continuous analysis of the real-world crowd behaviour. All in all, existing human-centric datasets remain constrained by the limited spatial and temporal information provided. The problems of low spatial resolution [44, 54, 28], lack of video information [14, 72, 18, 4], unnatural human appearance and actions [1, 37, 36], and limited scope of activities with short-term annotations [6, 49, 15, 53] lead to inevitable influence for understanding the complicated behaviors and interactions of crowd.

To address aforementioned problems, we propose a new gigaPixel-level humAN-centric viDeo dAtaset (PANDA). The videos in PANDA are captured by a gigacamera [70, 7], which is capable of covering a large-scale area full of high resolution details. A representative video example of *Marathon* is presented in Fig. 1. Such rich information enables PANDA to be a competitive dataset with **multi-scale features**: (1) **globally wide field-of-view where visible area may beyond 1 km²**, (2) **locally high resolution details with gigapixel-level spatial resolution**, (3) **temporally long-term crowd activities with 43.7k frames in total**, (4) **real-world scenes with abundant diversities in human attributes, behavioral patterns, scale, density, occlusion, and interaction**. Meanwhile, PANDA is provided with rich and hierarchical ground-truth annotations, with 15,974.6k bounding boxes, 111.8k fine-grained labels, 12.7k trajectories, 2.2k groups and 2.9k interactions in total.

Benefiting from the comprehensive and multiscale information, PANDA facilitates a variety of fundamental yet challenging tasks for image/video based human-centric analysis. We start with the most fundamental detection task. Yet detection on PANDA has to address both the accuracy and efficiency issues. The former one is challenged by the significant scale variation and complex occlusion, while the latter one is highly affected by the gigapixel resolution. Whereafter, the task of tracking is benchmarked. Equipped with the simultaneous large-scale, long-term and multi-object properties, our tracking task is heavily challenged, due to the complex occlusion as well as large-scale and long-term activity existing in real-world scenarios. Moreover, PANDA enables a distinct task of identifying the group relationship of the crowd in the video, termed as interaction-aware group detection. In this task, we propose a novel global-to-local zoom-in framework to reveal the mutual effects between global trajectories and local interactions. Note that these three tasks are inherently correlated. Although detection may bias to local high-resolution detail and tracking may focus on global trajectories, the former promotes the latter significantly. Meanwhile, the spatial-temporal trajectories deduced from detection and tracking serve for group analysis.

In summary, PANDA aims to contribute a standardized dataset to the community, for investigating new algorithms

to understand the complicated crowd social behavior in large-scale real-world scenarios. The contributions are summarized as follows.

- We propose a new video dataset with gigapixel-level resolution for human-centric visual analysis. It is the first video dataset endowed with wide FoV and high spatial resolution simultaneously, which is capable of providing sufficient spatial and temporal information from both global scene and local details. Complete and accurate annotations of location, trajectory, attribute, group and intra-group interaction information of crowd are provided.
- We benchmark several state-of-the-art algorithms on PANDA for the fundamental detection and tracking tasks. The results demonstrate that existing methods are heavily challenged from both accuracy and efficiency perspectives, and indicate that it is quite difficult to accurately detect objects in a scene with significant scale variation and track objects that move continuously for a long distance under complex occlusion.
- We introduce a new visual analysis task, termed as interaction-aware group detection, based on the spatial and temporal multi-object interaction. A global-to-local zoom-in framework is proposed to utilize the multi-modal annotations in PANDA, including global trajectories, local face orientations and interactions. Promising results further validate the collaborative effectiveness of global scene and local details provided by PANDA.

By serving the visual tasks related to the long-term analysis of crowd activities at a large-scale spatial-temporal range, we believe PANDA will definitely contribute to the community for understanding the complicated behaviors and interactions of crowd in large-scale real-world scenes, and further boost the intelligence of unmanned systems.

2. Related Work

2.1. Image-based Datasets

The most representative human-centric task on image datasets is human (person or pedestrian) detection. The common object detection datasets, such as PASCAL VOC [26], ImageNet [58], MS COCO [47], Open Images [42] and Objects365 [60] datasets, are not initially designed for human-centric analysis, although they contain human object categories¹. However, restricted by the narrow FoV, each image only contains limited number of objects, far from enough to describe the crowd behaviour and interaction.

¹Different terms are used in these datasets, such as “person”, “people”, and “pedestrian”. We uniformly use “human” when there is no ambiguity.

Pedestrian Detection. Some pioneer representatives include INRIA [19], ETH [25], TudBrussels [67], and Daimler [23]. Later, Caltech [21], KITTI-D [31], CityPersons [75] and EuroCity Persons [8] datasets with higher quality, larger scale, and more challenging contents are proposed. Most of them are collected via a vehicle-mounted camera through the regular traffic scenario, with limited diversity of pedestrian appearances and occlusions. The latest WiderPerson [76] and CrowdHuman [61] datasets focus on crowd scenes with many pedestrians. Due to the trade-off between spatial resolution and field of view, existing datasets cannot provide sufficient high resolution local details if the scene becomes larger.

Group Detection. Starting with the free-standing conversational groups (FCGs) decades ago [24], the subsequent works try to study the interacting persons characterized by mutual scene locations and poses, known as F-formations [40]. Representatives ones include IDIAP Poster [35], Cocktail Party [72], Coffee Break [18] and GDet [4]. In [14], the problem of structure group along with a dataset is proposed, which defines the way people spatially interact with each other. Recently, pedestrian group Re-Identification (G-ReID) benchmarks like DukeMTMC Group [48] and Road Group [48] are proposed to match a group of persons across different camera views. However, these datasets only support position-aware group detection, lack of the important dynamic interactions.

2.2. Video-based Datasets

Pedestrian Tracking. It locates pedestrians in a series of frames and find the trajectories of them. MOT Challenge benchmarks [43, 51] were launched to establish a standardized evaluation of multiple objects tracking algorithms. The latest MOT19 benchmark [20] consists of 8 new sequences with very crowded challenging scenes. Besides, some datasets are designed for specific applications, e.g., Campus [57] and VisDrone2018 [79], which are drone-platform-based benchmarks. PoseTrack [2] contains joint position annotations for multiple persons in videos. To increase the FoV for long-term tracking, a network of cameras is adopted, leading to the multi-target multi-camera (MTMC) tracking problem. MARS [77], DukeMTMC [56] are representative ones.

On the other hand, to investigate pedestrians in surveillance perspectives, UCY Crowds-by-Example [44], ETH BIWI Walking Pedestrians [54], Town Center [5] and Train Station [78] are proposed for trajectory prediction, abnormal behaviour detection, and pedestrian motion analysis. PETS09 [28] was collected by eight cameras in a campus for person density estimation, people tracking, event recognition, etc. Recently, CUHK [59] and WorldExpo'10 [74] serve for evaluating the performance of crowd segmentation, crowd density, collectiveness, and cohesiveness esti-

mation. However, these datasets are insufficient of the richness and complexity of the scenes, and can hardly provide high resolution local details, which is critical to further analyze the human interactions in crowd.

Interaction Analysis. SALSA [1] contains uninterrupted multi-modal recordings of indoor social events with 18 participants for over 60 minutes. Panoptic Studio [37] uses 480 synchronized VGA cameras to capture social interactions, with 3D body poses annotated. BEHAVE [6], CAVIAR [49], Collective Activity [15] and Volleyball [36] are widely used datasets to evaluate human group activity recognition approaches. VIRAT [53] is a real-world surveillance video dataset containing diverse examples of multiple types of complex visual events. However, for the sake of local details, the group interactions are usually restricted in small scenes or unnatural human behaviors.

3. Data Collection and Annotation

3.1. Data Collection and Pre-processing

It is known that single camera based imaging suffers inevitable contradiction between wide FoV and high spatial resolution. The recently developed array-camera-based gigapixel videography techniques significantly boost the feasibility of high performance imaging [7, 70]. By designing the advanced computational algorithms, a number of micro-cameras work simultaneously to generate a seamless gigapixel-level video in realtime. As a result, the sacrifice in either field of view or spatial resolution can be eliminated. We adopt the latest gigacameras [3, 70] to collect the data for PANDA, where the FoV is around 70 degree horizontally, and the video resolution reaches $25k \times 14k$, working in 30Hz. A representative video *Marathon* in Fig. 1 fully reflects the uniqueness of PANDA with both globally wide FoV and locally high resolution details.

Currently, PANDA is composed by 21 real-world outdoor scenes², by taking scenario diversity, pedestrians density, trajectory distribution, group activity, etc. into account. In each scene, we collected approximately 2 hours of 30Hz video as the raw data pool. Afterwards, around 3600 frames (approximately two minutes long segments) are extracted. For the images to be annotated, around 30 representative frames per video, 600 in total are selected, covering different crowd distributions and activities.

3.2. Data Annotation

Annotating PANDA images and videos faces the difficulty of full image annotation due to the gigapixel-level resolution. Herein, following the idea of divide-and-merge, the

²We are continuously collecting more videos to enrich our dataset. Note that all the data was collected in public areas where photography is officially approved, and it will be published under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 License [17].

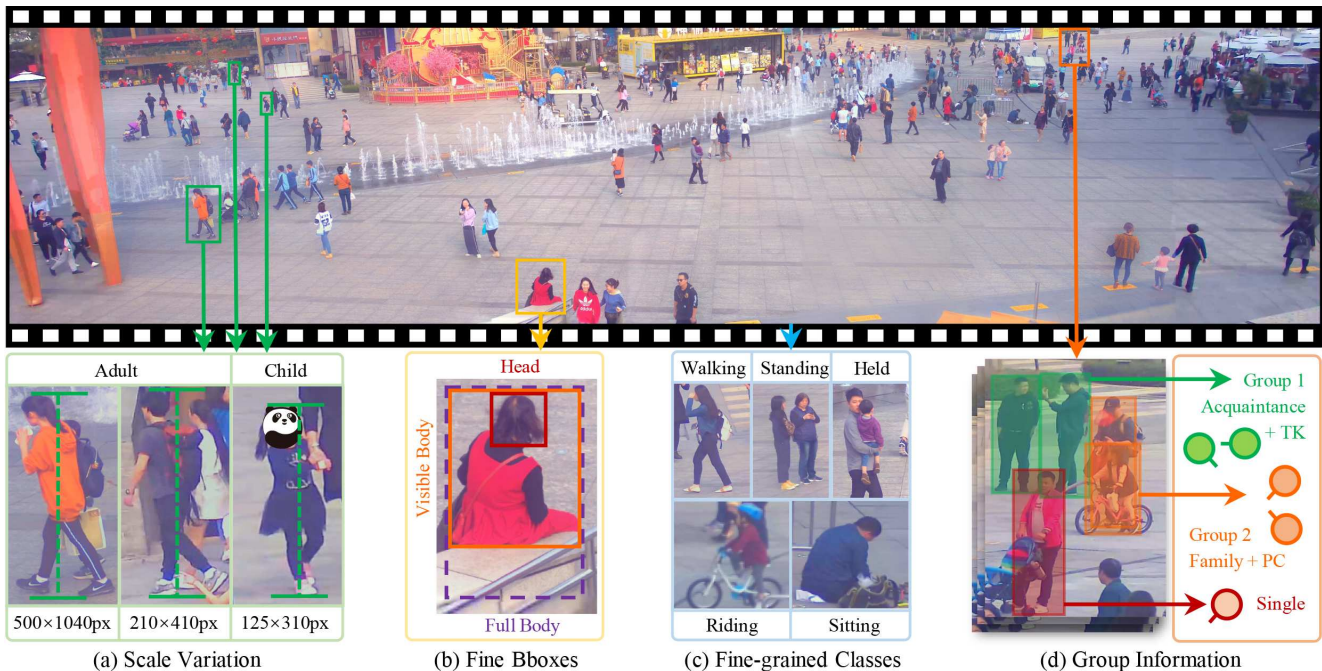


Figure 2. Visualization of annotations in PANDA dataset. (a) The scale variation of pedestrians in a large-scale scene. (b) Three fine-grained bounding boxes on human body. (c) Five categories for human body postures. (d) Group information along with the intra-group interactions (TK=Talking, PC=Physical contact), where the circle and short line denote pedestrian and their face orientation.

| | Caltech | CityPersons | PANDA | PANDA-C |
|-----|---------|-------------|----------|----------|
| Res | 480P | 2048×1024 | >25k×14k | >25k×14k |
| #Im | 249.9k | 5k | 555 | 45 |
| #Ps | 289.4k | 35.0k | 111.8k | 122.1k |
| Den | 1.16 | 7.0 | 201.4 | 2,713.8 |

Table 1. Pedestrian datasets comparison (statistics of CityPersons only contain public available training set). Res is the image resolution, #Im is the total number of images, #Ps is the total number of persons, Den denotes person density (average number of person per image) and PANDA-C is the PANDA-Crowd subset.

full image is partitioned into 4 to 16 subimages by considering the pedestrian density and size. After the labels are annotated on the subimages separately, the annotation results are mapped back to the full image. The objects cut by block borders are labeled with special status, which will be re-labeled after merging all blocks together. All labels are provided by a well trained professional annotation team.

3.2.1 Image Annotation

PANDA has 600 well annotated images captured from 21 diverse scenes for the multi-object detection task. Among them, PANDA-Crowd subset are composed by 45 images labeled with human head bounding boxes, which are selected from 3 extremely crowded scenes that full of head-counts. The remaining 555 images from 18 real-world daily scenes own 111k pedestrians in total, labeled with head point, head bounding box, visible body bounding box, and the estimated full body bounding box close to the border

of the pedestrian. For the crowd that are too far or too dense to be individually distinguished, the glass reflected persons, and the persons with more than 80% occluded area are marked as ‘ignore’ and disabled for benchmarking.

Fig. 2 presents a typical large-scale real-world scene *OCT Harbour* in PANDA, where the crowd shows a significant diversity in the scale, location, occlusion, activity, interaction, and so on. Beside the fine bounding boxes in (b), each pedestrian is further assigned a fine-grained label showing the detailed attributes in (c). Five categories are used, i.e., walking, standing, sitting, riding, and held in arms (for child), based on the daily postures. Pedestrians whose key parts occluded are marked as ‘unsure’. The ‘riding’ label is further subdivided into bicycle rider, tricycle rider and motorcycle rider. Another detailed attribute is termed as ‘child’ or ‘adult’, distinguished from the appearance and behavior, as shown in (a).

The comparisons with the representative Caltech [21] and CityPersons image datasets [75] are provided quantitatively (Tab. 1) and statistically (Fig. 3). From Tab. 1, each image of PANDA owns gigapixel-level resolution, which is around 100 times of existing datasets. Although the number of images is much smaller than other datasets, benefiting from the joint high resolution and wide FoV, PANDA has much higher pedestrian density per image than others especially in the extremely crowded PANDA-Crowd, and maintaining the total number of pedestrian in PANDA comparable to Caltech,

Some detailed statistics about image annotation are

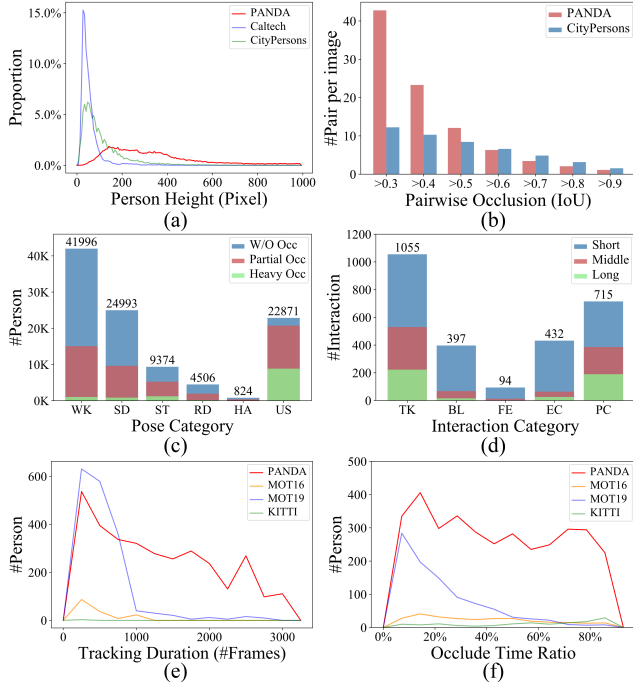


Figure 3. (a) Distribution of person scale (height in pixel). (b) Distribution of the number of person pairs with different occlusion (measured by IoU) threshold per image. (c) Distribution of persons’ pose labels in PANDA (WK=Walking, SD=Standing, ST=Sitting, RD=Ridding, HA=Held in arms, US=Unsure; The visible ratio is divided into W/O Occ (>0.9), Partial Occ ($0.5 - 0.9$), and Heavy Occ (<0.5)). (d) Distribution of categories and duration of inter-group interactions in PANDA (PC=Physical contact, BL=Body language, FE=Face expressions, EC=Eye contact, TK=Talking; The duration is divided into Short (< 10 s), Middle (10 s - 30 s), and Long (≥ 30 s)). (e) Distribution of person tracking duration. (f) Distribution of person occluded time ratio. The comparisons in (a), (b), (e) and (f) are limited to training sets.

shown in Fig. 3. In particular, Fig. 3(a) shows the distribution of person scale in pixel of PANDA, Caltech and CityPersons. As we can see, the height of persons in Caltech and CityPersons is mostly between 50px and 300px due to the limited spatial resolution, while PANDA has more balanced distribution from 100px to 600px. The larger scale variation in PANDA necessitates powerful multi-scale detection algorithms. In Fig. 3(b), the pairwise occlusion between persons measured by bounding box IoU of PANDA and CityPersons is given. The fine-grained label statistics for different poses and occlusion conditions are summarized in Fig. 3(c).

3.2.2 Video Annotation

Video annotation pays more attention on the labels revealing the activity/interaction. In addition to the bounding box of each person, we also label the face orientation (quantified into eight bins) and the occlusion ratio (without, partial and

| | KITTI-T | MOT16 | MOT19 | PANDA |
|-----|-----------|--------|----------|--------------------|
| Res | 1392×512 | 1080p | 1080p | >25k×14k |
| #V | 20 | 14 | 8 | 15 |
| #F | 19.1k | 11.2k | 13.4k | 43.7k |
| #T | 204 | 1.3k | 3.9k | 12.7k |
| #B | 13.4k | 292.7k | 2,259.2k | 15,480.5k |
| Den | 0.7 | 26.1 | 168.6 | 354.6 |

Table 2. Comparison of multi-object tracking datasets (statistics of KITTI only contain public available training set). Res means video resolution. #V, #F, #T and #B denote the number of video clips, video frames, tracks and bounding boxes respectively. Den means Density (average number of person per frame).

heavy). For pedestrians who are completely occluded for a short time, we label a virtual body bounding box and mark it as ‘disappearing’. MOT annotations are available for all the videos in PANDA except for PANDA-Crowd.

The comparisons with KITTI-T [31] and MOT [20] video datasets are provided quantitatively (Tab. 2) and statistically (Fig. 3). Apparently, PANDA is competitive with the largest number of frames, tracks and bounding boxes³. Moreover, in Fig. 3(e), we show the distribution of tracking duration of different datasets. It demonstrates that the tracking duration in PANDA is many times longer than that in KITTI-T and MOT because PANDA has wider FoV. This property makes PANDA an excellent dataset for large-scale and long-term tracking. Moreover, we also investigate the duration that each person is occluded and summarize the distribution in Fig. 3(f). It shows that more tracks in PANDA suffer from partial or heavy occlusions, in both absolute number and relative portion, making the tracking task more challenging.

For group annotation, the advance of PANDA with wide-FoV global information, high-resolution local details and temporal activities ensures more reliable annotations for group detection. Unlike existing group-based datasets that focus on either the similarity of global trajectories [54] or the stability of local spatial structure [14], we utilize the social signal processing [65] to label the group attributes at the interaction level.

More specifically, with the annotated bounding boxes, we firstly label the group members based on scene characteristics and social signals such as interpersonal distance [22] and interaction [65]. Afterwards, each group is assigned an category label denoting the relationship, such as acquaintance, family, business, etc., as shown in Fig. 2(d). To enrich the features for group identification, we further label the interactions between members within the group,

³Since the moving speed of pedestrians is relatively slow and stable, and the posture of pedestrians rarely changes rapidly and dramatically, we label them sparsely on every k frames ($k = 6$ to 15 based on the scene content) from the perspective of labeling cost. Here we compare the number of bounding boxes after linear interpolation to the original frame rate.

| | Sub | Visible Body | | Full Body | | Head | |
|---------|-----|-------------------|-------|-------------------|-------|-------------------|-------|
| | | AP _{.50} | AR | AP _{.50} | AR | AP _{.50} | AR |
| FR [55] | S | 0.201 | 0.137 | 0.190 | 0.128 | 0.031 | 0.023 |
| | M | 0.560 | 0.381 | 0.552 | 0.376 | 0.157 | 0.088 |
| | L | 0.755 | 0.523 | 0.744 | 0.512 | 0.202 | 0.105 |
| CR [9] | S | 0.204 | 0.140 | 0.227 | 0.160 | 0.028 | 0.018 |
| | M | 0.561 | 0.388 | 0.579 | 0.384 | 0.168 | 0.091 |
| | L | 0.747 | 0.532 | 0.765 | 0.518 | 0.241 | 0.116 |
| RN [46] | S | 0.171 | 0.121 | 0.221 | 0.150 | 0.023 | 0.018 |
| | M | 0.547 | 0.370 | 0.561 | 0.360 | 0.143 | 0.081 |
| | L | 0.725 | 0.482 | 0.740 | 0.479 | 0.259 | 0.149 |

Table 3. Performance of detection methods on PANDA. FR, CR, and RN denote Faster R-CNN, Cascade R-CNN and RetinaNet respectively. Sub means subset of different target sizes, where Small, Middle, and Large indicate object size being $< 32 \times 32$, $32 \times 32 - 96 \times 96$, and $> 96 \times 96$.

including the interaction category (including physical contact, body language, face expressions, eye contact and talking; multi-label annotation) and its begin/end time. The distribution and duration of interactions are shown in Fig. 3(d). The mean duration of interaction is 518 frames (17.3s). To avoid overly subjective or ambiguous cases, three rounds of cross-checking are performed.

4. Algorithm Analysis

We consider three human-centric visual analysis tasks on PANDA. The first is pedestrian detection, which biases local visual information. The second is multi-pedestrian tracking. In this task, global visual clues from different regions are taken into consideration. Based on these two well-defined tasks, we introduce the interaction-aware group detection task. In this task, both global trajectories and local interactions between persons are necessary.

4.1. Detection

Pedestrian detection is a fundamental task for human-centric visual analysis. The extremely high resolution of PANDA makes it possible to detect pedestrians from a long distance. However, the significant variance in scale, posture, and occlusion severely degrade the detection performance. In this paper, we benchmarked several state-of-the-art detection algorithms on PANDA⁴.

Evaluation metrics. For evaluation, we choose $AP_{.50}$ and AR as metrics: $AP_{.50}$ is the average precision at $IoU = 0.50$ and AR is the average recall with IoU ranging in $[0.5, 0.95]$ with a stride of 0.05.

Baseline detectors. We choose Faster R-CNN [55], Cascade R-CNN [9] and RetinaNet [46] as our baseline detectors with ResNet101 [33] backbone. The implementation is based on [11]. To train the gigapixel images on our network,

⁴For 18 ordinary scenes, 13 scenes are used for training and 5 scenes for testing. For 3 extremely crowded scenes, 2 scenes are for training and 1 scene for testing.




Figure 4. Left: False analysis for Faster R-CNN on Visible Body. C75, C50, Loc and BG denote PR-curve at $IoU=0.75$, $IoU=0.5$, ignoring localization errors and ignoring background false positives, respectively. Right: False negative instances (FN) v.s. All instances (ALL) in terms of person height (in pixel) distribution for Faster R-CNN on Visible Body.

we resize the original size image into multiple scales and partition the image into blocks with appropriate size as neural network input. For the objects cut by block borders, we retain them if the preserved area overs 50%. Similarly, for evaluation, we resize the original image into multiple scales and use the sliding window approach to generate proper size blocks for the detector. For a better analysis of detector performance and limitations, we split test results into subsets according to the object size.

Results. We train these 3 detectors from the COCO pre-trained weights and evaluate them on three tasks: visible body, full body, and head detection. As shown in Tab. 3, Faster R-CNN, Cascade R-CNN and RetinaNet show the difficulty in detecting small objects, resulting in very low precision and recall. We also apply false analysis on visible body using Faster R-CNN, as illustrated in Fig. 4 left. We can observe that the huge amount of false negatives is the most severe factor limiting the performance of the detector. We further analyze the height distribution of the false negative instances in Fig. 4 right. The results indicate false negative caused by missing detection of small objects is the main reason for poor recall. According to the results, it seems quite difficult to accurately detect objects in a scene with very large scale variation (most $100\times$ in PANDA) by the single detector based on existing architectures. More advanced optimization strategies and algorithms are highly demanded for the detection task on extra-large images with large object scale variation, such as scale self-adaptive detectors and efficient global-to-local multi-stage detectors.

4.2. Tracking

Pedestrian tracking aims to associate pedestrians at different spatial positions and temporal frames. The superior properties of PANDA make it naturally suitable for long-term tracking. Yet the complex scenes with crowded pedestrian impose various challenges as well.

Evaluation metrics. To evaluate the performance of multiple person tracking algorithms, we adopt the metrics of MOTChallenge [43, 51], including MOTA, MOTP,

| T | D | MOTA \uparrow | MOTP \uparrow | IDF1 \uparrow | FAR \downarrow | MT \uparrow |
|----------|----|-----------------|-----------------|-----------------|------------------|---------------|
| DS [68] | FR | 25.53 | 76.67 | 21.14 | 20.45 | 762 |
| | CR | 24.35 | 76.31 | 21.39 | 15.59 | 661 |
| | RN | 16.36 | 78.0 | 15.16 | 4.32 | 259 |
| DAN [63] | FR | 25.06 | 74.81 | 21.85 | 25.95 | 826 |
| | CR | 24.24 | 78.55 | 20.13 | 12.42 | 602 |
| | RN | 15.57 | 79.90 | 13.43 | 3.33 | 227 |
| MD [50] | FR | 13.51 | 78.82 | 14.92 | 6.52 | 257 |
| | CR | 13.54 | 80.25 | 14.89 | 4.41 | 255 |
| | RN | 10.77 | 80.62 | 11.86 | 1.90 | 162 |

Table 4. Performance of multiple object tracking methods on PANDA. T is tracker, D is detector, DS, DAN and MD denote the DeepSORT [68], DAN [63] and MOTDT [50] trackers, respectively. \uparrow denotes higher is better and vice versa.

IDF1, FAR, MT and Hz. Multiple Object Tracking Accuracy (MOTA) computes the accuracy considering three error sources: false positives, false negatives/missed targets and identity switches. Multiple Object Tracking Precision (MOTP) takes into account the misalignment between the groundtruth and the predicted bounding boxes. ID F1 score (IDF1) measures the ratio of correctly identified detections over the average number of ground-truth and computed detections. False alarm rate (FAR) measures the average number of false alarms per frame. The mostly tracked targets (MT) measures the ratio of ground-truth trajectories that are covered by a track hypothesis for at least 80% of their respective life span. The Hz indicates the processing speed of the algorithm.

Baseline trackers. Three representative algorithms DeepSORT [68], DAN [63] and MOTDT [50] are evaluated. All of them follow the tracking-by-detection strategy. In our experiments, the bounding boxes are generated from 3 detection algorithms [55, 9, 46] in the previous subsection. For the sake of fairness, We use the same pre-trained weights on the COCO dataset and detection threshold scores (0.7) for them. Default model parameters provided by the authors are used for evaluating three trackers.

Results. Tab. 4 shows the results of DeepSORT [68], MOTDT [50] and DAN [63] on PANDA. The time cost to process a single frame is 18.36s (0.054Hz), 19.13s (0.052Hz), 8.29s (0.121Hz) for DeepSORT, MOTDT and DAN, respectively. MOTDT shows better bounding box alignment according to MOTP and FAR. DAN leads on IDF1 and MT, implying its stronger capability to establish correspondence between the detected objects in different frames. The experimental results also demonstrate the challenge of PANDA dataset. The best MOTA for DeepSORT, DAN and MOTDT on MOT16 are 61.4, 52.42 and 47.6, while drop more than half on PANDA. With regards to object detectors, Faster R-CNN performs the best and Cascade R-CNN shows similar performance. Whereas the performance of RetinaNet is relatively poor except MOTP and FAR, the reason is that RetinaNet has low recall under

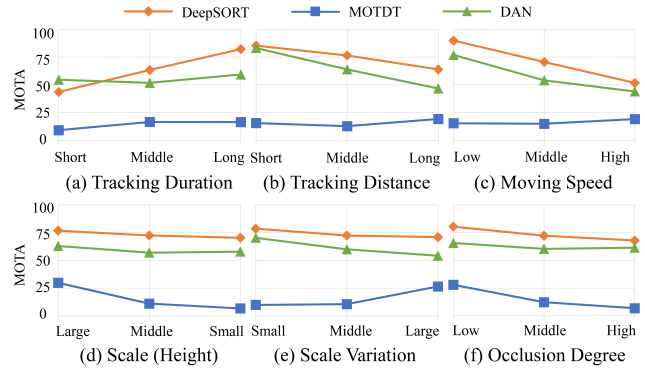


Figure 5. Influence of target properties on tracker’s MOTA. We divided the pedestrian targets into 3 subsets from easy to hard for each property.

confidence threshold 0.7 for detection results.

We further analyze the influence of different pedestrian properties, including: (a) tracking duration; (b) tracking distance; (c) moving speed; (d) scale (height); (e) scale variation (the standard deviation of height); (f) occlusion. For each property, we divided the pedestrian targets into 3 subsets from easy to hard. Besides, in order to eliminate the influence of detectors, we used the ground-truth bounding boxes as input here. Fig. 5(b)(c) show that the tracking distance and moving speed are the most influential factors to trackers’ performance. In Fig. 5(a), the impact of tracking duration on tracker performance is not obvious because there are many stationary or slow moving people in the scene.

4.3. Group Detection

Group detection aims at identifying groups of people from crowds. Unlike existing datasets that focus on either the similarity of global trajectories [54] or the stability of local spatial structure [14], the advance of PANDA with joint wide-FoV global information, high-resolution local details and temporal activities imposes rich information for group detection.

Furthermore, as indicated by the recent advances on trajectory embedding [30, 16], trajectory prediction [10, 13] and interaction modeling in video recognition [36, 62, 27], these tasks are strongly correlated to the group detection task. For example, modeling group interaction can help improve the trajectory prediction performance [69, 10, 13], while learning a good trajectory embedding is also beneficial for video action recognition [66, 30, 16]. However, none of previous research has investigated how those multi-modal information can be incorporated into the group detection task. Hence, we propose the *interaction-aware group detection task*, where video data and multi-modal annotations (spatial-temporal human trajectory, face orientation, and human interaction) are provided as input for group detection.

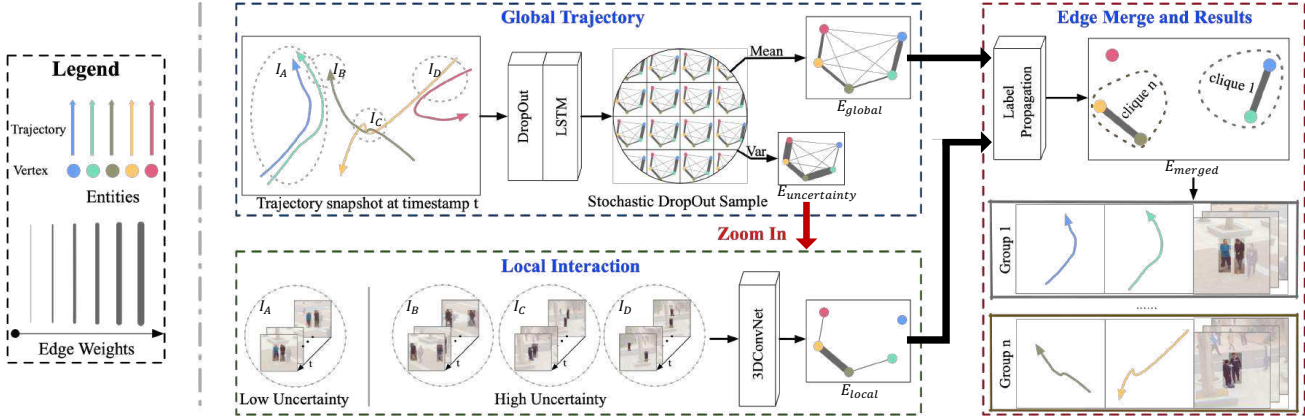


Figure 6. Global-to-local zoom-in framework for interaction-aware group detection. The Global Trajectory, Local Interaction, Zoom In, and Edge Merging modules are associated. Different color vertices and trajectories stand for different human entities. Line thickness represents the edge weights in the graph. **(1) Global Trajectory:** Trajectories are firstly fed into LSTM encoder with dropout layer to obtain embedding vectors and then construct a graph where the edge weight is L2 distance between embedding vectors. **(2) Zoom In:** By repeating inference with dropout activated as Stochastic Sampling [38], E_{global} and $E_{uncertainty}$ are obtained from sample mean and variance respectively. **(3) Local Interaction:** The local interaction videos corresponding to high uncertainty edges ($I_B \sim I_D$) are further checked using video interaction classifier (3DConvNet [32]). **(4) Edge Merge and Results:** Edges are merged using label propagation [73], and cliques remaining in the graph are the group detection results.

Framework. We further design a global-to-local zoom-in framework as shown in Fig. 6 to validate the incremental effectiveness of local visual clues to global trajectories. More specifically, human entities and their relationships are represented as vertices and edges respectively in graph $G = (V, E)$. And features from multiple scales and modalities such as the global trajectory, face orientation vector, and local interaction video are used to generate edge set E_{global} and E_{local} . Following a global-to-local strategy [52, 29, 45, 12], E_{global} is firstly obtained by calculating L2 distance in feature space for each trajectory embedding vector, which comes from LSTM encoder like common practice [30]. After that, uncertainty-based [38, 39] and random selection policies are adopted to determine the sub-set of edges that need to be further checked using visual clues. Then, video interaction scores among entities are estimated by spatial-temporal ConvNet [32]. The combinations of obtained edge sets, e.g., $E_{global} \cup E_{local}$ or E_{global} , are merged using label propagation [73], and the cliques remaining in the graph are the group detection results. Finally, we can estimate the incremental effectiveness with the performance metrics specified in [14] under different combinations.

Results. The details about network structure, zoom-in policy, and edge merging strategy are placed in supplemental material S.4.3.1, S.4.3.2, and S.4.3.3 respectively, while experimental results are shown in Tab. 5. Half metrics [14] including precision, recall, and F1 where group member $IoU = 0.5$ are used for evaluation. The performance is improved significantly by leveraging E_{local} as well as uncertainty estimation, which further validates the effectiveness of local visual clues provided by PANDA.

| Edge Sets | Zoom In | Precision | Recall | F1 |
|-----------------------------|-------------|-----------|--------|-------|
| E_{global} | / | 0.237 | 0.120 | 0.160 |
| $E_{global} \cup E_{local}$ | Random | 0.244 | 0.133 | 0.172 |
| $E_{global} \cup E_{local}$ | Uncertainty | 0.293 | 0.160 | 0.207 |

Table 5. Incremental Effectiveness (half metric [14]). The random zoom-in policy randomly selects several local videos to estimate interaction score while the uncertainty-based one selects local videos depending on the uncertainty estimation from Stochastic Dropout Sample [38].

5. Conclusion

In this paper, we introduced a gigapixel-level video dataset (PANDA) for large-scale, long-term, and multi-object human-centric visual analysis. The videos in PANDA are equipped with both wide FoV and high spatial resolution. Rich and hierarchical annotations are provided. We benchmarked several state-of-the-art algorithms for the fundamental human-centric tasks, pedestrian detection and tracking. The results demonstrate that they are heavily challenged for accuracy due to the significant variance of pedestrian pose, scale, occlusion and trajectory, etc., and efficiency due to the large image size and the huge amount of objects in single frame. Besides, we introduced a new task, termed as interaction-aware group detection based on the characteristics of PANDA. We proposed a global-to-local zoom-in framework which combines both global trajectories and local interactions, yielding promising group detection performance. Based on PANDA, we believe the community will develop new effective and efficient algorithms for understanding complicated behaviors and interactions of crowd in large-scale real-world scenes.

References

- [1] Xavier Alameda-Pineda, Jacopo Staiano, Ramanathan Subramanian, Ligia Batrinca, Elisa Ricci, Bruno Lepri, Oswald Lanz, and Nicu Sebe. Salsa: A novel dataset for multimodal group behavior analysis. *IEEE transactions on pattern analysis and machine intelligence*, 38(8):1707–1720, 2015. 2, 3
- [2] Mykhaylo Andriluka, Umar Iqbal, Eldar Insafutdinov, Leonid Pishchulin, Anton Milan, Juergen Gall, and Bernt Schiele. Posetrack: A benchmark for human pose estimation and tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5167–5176, 2018. 3
- [3] Inc. Aqueti. Aqueti mantis 70 array cameras webpage. <https://www.aqueti.com/>. Accessed 2019. 3
- [4] Loris Bazzani, Marco Cristani, Diego Tosato, Michela Farenzena, Giulia Paggetti, Gloria Menegaz, and Vittorio Murino. Social interactions by visual focus of attention in a three-dimensional environment. *Expert Systems*, 30(2):115–127, 2013. 2, 3
- [5] Ben Benfold and Ian Reid. Stable multi-target tracking in real-time surveillance video. In *CVPR 2011*, pages 3457–3464. IEEE, 2011. 3
- [6] Scott Blunsden and RB Fisher. The behave video dataset: ground truthed video for multi-person behavior classification. *Annals of the BMVA*, 4(1-12):4, 2010. 1, 2, 3
- [7] David J. Brady, Michael E. Gehm, Ronald A. Stack, Daniel L. Marks, David S. Kittle, Dathon R. Golish, Esteban Vera, and Steven D. Feller. Multiscale gigapixel photography. *Nature*, 486:386–389, 2012. 2, 3
- [8] Markus Braun, Sebastian Krebs, Fabian Flohr, and Dariu Gavrilă. Eurocity persons: a novel benchmark for person detection in traffic scenes. *IEEE transactions on pattern analysis and machine intelligence*, 2019. 3
- [9] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6154–6162, 2018. 1, 6, 7
- [10] Rohan Chandra, Uttaran Bhattacharya, Aniket Bera, and Dinesh Manocha. Taphic: Trajectory prediction in dense and heterogeneous traffic using weighted interactions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8483–8492, 2019. 7
- [11] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 6
- [12] Wuyang Chen, Ziyu Jiang, Zhangyang Wang, Kexin Cui, and Xiaoning Qian. Collaborative global-local networks for memory-efficient segmentation of ultra-high resolution images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8924–8933, 2019. 8
- [13] Chiho Choi and Behzad Dariush. Learning to infer relations for future trajectory forecast. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 7
- [14] Wongun Choi, Yu-Wei Chao, Caroline Pantofaru, and Silvio Savarese. Discovering groups of people in images. In *European conference on computer vision*, pages 417–433. Springer, 2014. 2, 3, 5, 7, 8
- [15] Wongun Choi, Khuram Shahid, and Silvio Savarese. What are they doing?: Collective activity classification using spatio-temporal relationship among people. In *2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops*, pages 1282–1289. IEEE, 2009. 2, 3
- [16] John D Co-Reyes, YuXuan Liu, Abhishek Gupta, Benjamin Eysenbach, Pieter Abbeel, and Sergey Levine. Self-consistent trajectory autoencoder: Hierarchical reinforcement learning with trajectory embeddings. *arXiv preprint arXiv:1806.02813*, 2018. 7
- [17] Creative Commons. Commons attribution-noncommercial-sharealike 4.0 license. <https://creativecommons.org/licenses/by-nc-sa/4.0/>. 3
- [18] Marco Cristani, Loris Bazzani, Giulia Paggetti, Andrea Fos-sati, Diego Tosato, Alessio Del Bue, Gloria Menegaz, and Vittorio Murino. Social interaction discovery by statistical analysis of f-formations. In *BMVC*, volume 2, page 4, 2011. 1, 2, 3
- [19] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. 2005. 3
- [20] Patrick Dendorfer, Hamid Reza Tofighi, Anton Milan, Javen Shi, Daniel Cremers, Ian Reid, Stefan Roth, Konrad Schindler, and Laura Leal-Taixe. Cvpr19 tracking and detection challenge: How crowded can it get? *arXiv preprint arXiv:1906.04567*, 2019. 3, 5
- [21] Piotr Dollár, Christian Wojek, Bernt Schiele, and Pietro Perona. Pedestrian detection: An evaluation of the state of the art. *IEEE transactions on pattern analysis and machine intelligence*, 34(4):743–761, 2011. 1, 3, 4
- [22] Marshall P Duke and Stephen Nowicki. A new measure and social-learning model for interpersonal distance. *Journal of Experimental Research in Personality*, 1972. 5
- [23] Markus Enzweiler and Dariu M Gavrilă. Monocular pedestrian detection: Survey and experiments. *IEEE transactions on pattern analysis and machine intelligence*, 31(12):2179–2195, 2008. 3
- [24] Goffman Erving. Behavior in public places: notes on the social organization of gatherings. *New York*, 1963. 3
- [25] Andreas Ess, Bastian Leibe, Konrad Schindler, and Luc Van Gool. A mobile vision system for robust multi-person tracking. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. 3
- [26] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010. 1, 2
- [27] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. *arXiv preprint arXiv:1812.03982*, 2018. 7

- [28] James Ferryman and Ali Shahrokni. Pets2009: Dataset and challenge. In *2009 Twelfth IEEE international workshop on performance evaluation of tracking and surveillance*, pages 1–6. IEEE, 2009. 1, 2, 3
- [29] Mingfei Gao, Ruichi Yu, Ang Li, Vlad I. Morariu, and Larry S. Davis. Dynamic zoom-in network for fast object detection in large images. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6926–6935, 2018. 8
- [30] Qiang Gao, Fan Zhou, Kunpeng Zhang, Goce Trajcevski, Xucheng Luo, and Fengli Zhang. Identifying human mobility via trajectory embeddings. In *IJCAI*, volume 17, pages 1689–1695, 2017. 7, 8
- [31] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361. IEEE, 2012. 3, 5
- [32] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6546–6555, 2018. 8
- [33] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 6
- [34] Gary B Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. 2008. 1
- [35] Hayley Hung and Ben Kröse. Detecting f-formations as dominant sets. In *Proceedings of the 13th international conference on multimodal interfaces*, pages 231–238. ACM, 2011. 3
- [36] Mostafa S Ibrahim, Srikanth Muralidharan, Zhiwei Deng, Arash Vahdat, and Greg Mori. A hierarchical deep temporal model for group activity recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1971–1980, 2016. 2, 3, 7
- [37] Hanbyul Joo, Hao Liu, Lei Tan, Lin Gui, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social motion capture. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3334–3342, 2015. 2, 3
- [38] Alex Kendall, Vijay Badrinarayanan, and Roberto Cipolla. Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. *arXiv preprint arXiv:1511.02680*, 2015. 8
- [39] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in neural information processing systems*, pages 5574–5584, 2017. 8
- [40] Adam Kendon. *Conducting interaction: Patterns of behavior in focused encounters*, volume 7. CUP Archive, 1990. 3
- [41] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012. Proceedings of a meeting held December 3-6, 2012, Lake Tahoe, Nevada, United States.*, pages 1106–1114, 2012. 1
- [42] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Tom Duerig, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *arXiv preprint arXiv:1811.00982*, 2018. 2
- [43] Laura Leal-Taixé, Anton Milan, Ian Reid, Stefan Roth, and Konrad Schindler. Motchallenge 2015: Towards a benchmark for multi-target tracking. *arXiv preprint arXiv:1504.01942*, 2015. 3, 6
- [44] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. In *Computer graphics forum*, volume 26, pages 655–664. Wiley Online Library, 2007. 2, 3
- [45] Hongyang Li, Yu Liu, Wanli Ouyang, and Xiaogang Wang. Zoom out-and-in network with map attention decision for region proposal and object detection. *International Journal of Computer Vision*, 127(3):225–238, 2019. 8
- [46] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 1, 6, 7
- [47] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 1, 2
- [48] Weiyao Lin, Yuxi Li, Hao Xiao, John See, Junni Zou, Hongkai Xiong, Jingdong Wang, and Tao Mei. Group re-identification with multi-grained matching and integration. *arXiv preprint arXiv:1905.07108*, 2019. 3
- [49] Thor List, Jos Bins, Jose Vazquez, and Robert B Fisher. Performance evaluating the evaluator. In *2005 IEEE International Workshop on Visual Surveillance and Performance Evaluation of Tracking and Surveillance*, pages 129–136. IEEE, 2005. 2, 3
- [50] Chen Long, Ai Haizhou, Zhuang Zijie, and Shang Chong. Real-time multiple people tracking with deeply learned candidate selection and person re-identification. In *ICME*, 2018. 1, 7
- [51] Anton Milan, Laura Leal-Taixé, Ian Reid, Stefan Roth, and Konrad Schindler. Mot16: A benchmark for multi-object tracking. *arXiv preprint arXiv:1603.00831*, 2016. 1, 3, 6
- [52] Mahyar Najibi, Bharat Singh, and Larry S. Davis. Autofocus: Efficient multi-scale inference. *arXiv preprint arXiv:1812.01600*, 2018. 8
- [53] Sangmin Oh, Anthony Hoogs, Amitha Perera, Naresh Cuntoor, Chia-Chih Chen, Jong Taek Lee, Saurajit Mukherjee, JK Aggarwal, Hyungtae Lee, Larry Davis, et al. A large-scale benchmark dataset for event recognition in surveillance video. In *CVPR 2011*, pages 3153–3160. IEEE, 2011. 2, 3
- [54] Stefano Pellegrini, Andreas Ess, Konrad Schindler, and Luc Van Gool. You’ll never walk alone: Modeling social be-

- havior for multi-target tracking. In *2009 IEEE 12th International Conference on Computer Vision*, pages 261–268. IEEE, 2009. 2, 3, 5, 7
- [55] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 1, 6, 7
- [56] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *European Conference on Computer Vision*, pages 17–35. Springer, 2016. 1, 3
- [57] Alexandre Robicquet, Amir Sadeghian, Alexandre Alahi, and Silvio Savarese. Learning social etiquette: Human trajectory understanding in crowded scenes. In *European conference on computer vision*, pages 549–565. Springer, 2016. 3
- [58] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 1, 2
- [59] Jing Shao, Chen Change Loy, and Xiaogang Wang. Scene-independent group profiling in crowd. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2219–2226, 2014. 3
- [60] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 2
- [61] Shuai Shao, Zijian Zhao, Boxun Li, Tete Xiao, Gang Yu, Xiangyu Zhang, and Jian Sun. Crowdhuman: A benchmark for detecting human in a crowd. *arXiv preprint arXiv:1805.00123*, 2018. 3
- [62] Tianmin Shu, Sinisa Todorovic, and Song-Chun Zhu. Cern: confidence-energy recurrent network for group activity recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5523–5531, 2017. 7
- [63] ShiJie Sun, Naveed Akhtar, HuanSheng Song, Ajmal S Mian, and Mubarak Shah. Deep affinity network for multiple object tracking. *IEEE transactions on pattern analysis and machine intelligence*, 2019. 1, 7
- [64] Antonio Torralba, Robert Fergus, and William T. Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(11):1958–1970, 2008. 1
- [65] Alessandro Vinciarelli, Maja Pantic, and Hervé Bourlard. Social signal processing: Survey of an emerging domain. *Image Vision Comput.*, 27:1743–1759, 2009. 5
- [66] Limin Wang, Yu Qiao, and Xiaoou Tang. Action recognition with trajectory-pooled deep-convolutional descriptors. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4305–4314, 2015. 7
- [67] Christian Wojek, Stefan Walk, and Bernt Schiele. Multi-cue onboard pedestrian detection. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 794–801. IEEE, 2009. 3
- [68] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple online and realtime tracking with a deep association metric. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 3645–3649. IEEE, 2017. 1, 7
- [69] Yanyu Xu, Zhixin Piao, and Shenghua Gao. Encoding crowd interaction with deep neural network for pedestrian trajectory prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5275–5284, 2018. 7
- [70] Xiaoyun Yuan, Lu Fang, Qionghai Dai, David J Brady, and Yebin Liu. Multiscale gigapixel video: A cross resolution image matching and warping approach. In *Computational Photography (ICCP), 2017 IEEE International Conference on*, pages 1–9. IEEE, 2017. 2, 3
- [71] Francesco Zanlungo, Dražen Brščić, and Takayuki Kanda. Pedestrian group behaviour analysis under different density conditions. *Transportation Research Procedia*, 2:149–158, 2014. 1
- [72] Gloria Zen, Bruno Lepri, Elisa Ricci, and Oswald Lanz. Space speaks: towards socially and personality aware visual surveillance. In *Proceedings of the 1st ACM international workshop on Multimodal pervasive video analysis*, pages 37–42. ACM, 2010. 2, 3
- [73] Xiaohang Zhan, Ziwei Liu, Junjie Yan, Dahua Lin, and Chen Change Loy. Consensus-driven propagation in massive unlabeled data for face recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 568–583, 2018. 8
- [74] Cong Zhang, Kai Kang, Hongsheng Li, Xiaogang Wang, Rong Xie, and Xiaokang Yang. Data-driven crowd understanding: A baseline for a large-scale crowd dataset. *IEEE Transactions on Multimedia*, 18(6):1048–1061, 2016. 3
- [75] Shanshan Zhang, Rodrigo Benenson, and Bernt Schiele. Citypersons: A diverse dataset for pedestrian detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3213–3221, 2017. 3, 4
- [76] Shifeng Zhang, Yiliang Xie, Jun Wan, Hansheng Xia, Stan Z Li, and Guodong Guo. Widerperson: A diverse dataset for dense pedestrian detection in the wild. *IEEE Transactions on Multimedia*, 2019. 3
- [77] Liang Zheng, Zhi Bie, Yifan Sun, Jingdong Wang, Chi Su, Shengjin Wang, and Qi Tian. Mars: A video benchmark for large-scale person re-identification. In *European Conference on Computer Vision*, pages 868–884. Springer, 2016. 3
- [78] Bolei Zhou, Xiaogang Wang, and Xiaoou Tang. Understanding collective crowd behaviors: Learning a mixture model of dynamic pedestrian-agents. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2871–2878. IEEE, 2012. 3
- [79] Pengfei Zhu, Longyin Wen, Xiao Bian, Haibin Ling, and Qinghua Hu. Vision meets drones: A challenge. *arXiv preprint arXiv:1804.07437*, 2018. 3