

CrowdHuman: A Benchmark for Detecting Human in a Crowd

Shuai Shao* Zijian Zhao* Boxun Li Tete Xiao Gang Yu Xiangyu Zhang Jian Sun
Megvii Inc. (Face++)

{shaoshuai, zhaozijian, liboxun, xtt, yugang, zhangxiangyu, sunjian}@megvii.com

Abstract

Human detection has witnessed impressive progress in recent years. However, the occlusion issue of detecting human in highly crowded environments is far from solved. To make matters worse, crowd scenarios are still under-represented in current human detection benchmarks. In this paper, we introduce a new dataset, called CrowdHuman¹, to better evaluate detectors in crowd scenarios. The CrowdHuman dataset is large, rich-annotated and contains high diversity. There are a total of 470K human instances from the train and validation subsets, and 22.6 persons per image, with various kinds of occlusions in the dataset. Each human instance is annotated with a head bounding-box, human visible-region bounding-box and human full-body bounding-box. Baseline performance of state-of-the-art detection frameworks on CrowdHuman is presented. The cross-dataset generalization results of CrowdHuman dataset demonstrate state-of-the-art performance on previous dataset including Caltech-USA, CityPersons, and Brainwash without bells and whistles. We hope our dataset will serve as a solid baseline and help promote future research in human detection tasks.

1. Introduction

Detecting people in images is among the most important components of computer vision and has attracted increasing attention in recent years [29, 14, 32, 30, 10, 5, 4, 6, 18]. A system that is able to detect human accurately plays an essential role in applications such as autonomous cars, smart surveillance, robotics, and advanced human machine interactions. Besides, it is a fundamental component for research topics like multiple-object tracking [13], human pose estimation [28], and person search [24]. Coupled with the development and blooming of convolutional neural networks (CNNs) [12, 22, 8], modern human detectors [1, 29, 26] have achieved remarkable performance on several major hu-

man detection benchmarks.

However, as the algorithms improve, more challenging datasets are necessary to evaluate human detection systems in more complicated real world scenarios, where crowd scenes are relatively common. In crowd scenarios, different people occlude with each other with high overlaps and cause great difficulty of *crowd occlusion*. For example, when a target pedestrian T is largely overlapped with other pedestrians, the detector may fail to identify the boundaries of each person as they have similar appearances. Therefore, detector will treat the crowd as a whole, or shift the target bounding box of T to other pedestrians mistakenly. To make matters worse, even though the detectors are able to discriminate different pedestrians in the crowd, the highly overlapped bounding boxes will also be suppressed by the post process of **non-maximum suppression** (NMS). As a result, crowd occlusion makes the detector sensitive to the threshold of NMS. A lower threshold may lead to drastically drop on recall, while a higher threshold brings more false positives.

Current datasets and benchmarks for human detection, such as Caltech-USA [6], KITTI [25], CityPersons [31], and “person” subset of MSCOCO [17], have contributed to a rapid progress in the human detection. Nevertheless, crowd scenarios are still under-represented in these datasets. For example, the statistical number of persons per image is only 0.32 in Caltech-USA, 4.01 in COCOPersons, and 6.47 in CityPersons. And the average of pairwise overlap between two human instances (larger than 0.5 IoU) in these datasets is only 0.02, 0.02, and 0.32, respectively. Furthermore, the annotators for these datasets are more likely to annotate crowd human as a whole ignored region, which cannot be counted as valid samples in training and evaluation.

Our goal is to push the boundary of human detection by specifically targeting the challenging crowd scenarios. We collect and annotate a rich dataset, termed *CrowdHuman*, with considerable amount of crowded pedestrians. CrowdHuman contains 15,000, 4,370 and 5,000 images for training, validation, and testing respectively. The dataset is exhaustively annotated and contains diverse scenes. There are

*Equal contribution.

¹Our CrowdHuman dataset can be downloaded from <https://sshao0516.github.io/CrowdHuman/>



Figure 1. Illustrative examples from different human dataset benchmarks. The images inside the green, yellow, blue boxes are from the COCO [17], Caltech [6], and CityPersons [31] datasets, respectively. The images from the second row inside the red box are from our CrowdHuman benchmark with full body, visible body, and head bounding box annotations for each person.

totally 470k individual persons in the train and validation subsets, and the average number of pedestrians per image reaches 22.6. We also provide the visible region bounding-box annotation, and head region bounding-box annotation along with its full body annotation for each person. Fig. 1 shows examples in our dataset compared with those in other human detection datasets.

To summarize, we propose a new dataset called CrowdHuman with the following three contributions:

- To the best of our knowledge, this is the first dataset which specifically targets to address the crowd issue in human detection task. More specifically, the average number of persons in an image is 22.6 and the average of pairwise overlap between two human instances (larger than 0.5 IoU) is 2.4, both of which are much larger than the existing benchmarks like CityPersons, KITTI and Caltech.
- The proposed CrowdHuman dataset provides annotations with three categories of bounding boxes: head bounding-box, human visible-region bounding-box, and human full-body bounding-box. Furthermore, these three categories of bounding-boxes are bound for each human instance.
- Experiments of cross-dataset generalization ability

demonstrate our dataset can serve as a powerful pre-training dataset for many human detection tasks. A framework originally designed for general object detection without any specific modification provides state-of-the-art results on every previous benchmark including Caltech and CityPersons for pedestrian detection, COCOPerson for person detection, and Brainwash for head detection.

2. Related Work

2.1. Human detection datasets.

Pioneer works of pedestrian detection datasets involve INRIA [3], TudBrussels [27], and Daimler [7]. These datasets have contributed to spurring interest and progress of human detection. However, as algorithm performance improves, these datasets are replaced by larger-scale datasets like Caltech-USA [6] and KITTI [25]. More recently, Zhang *et al.* build a rich and diverse pedestrian detection dataset CityPersons [31] on top of CityScapes [2] dataset. It is recorded by a car traversing various cities, contains dense pedestrians, and is annotated with high-quality bounding boxes.

Despite the prevalence of these datasets, they all suffer a problem of from low density. Statistically, the Caltech-USA and KITTI datasets have less than one person per

image, while the CityPersons has ~ 6 persons per image. In these datasets, the crowd scenes are significantly under-represented. Even worse, protocols of these datasets allow annotators to ignore and discard the regions with a large number of persons as exhaustively annotating crowd regions is incredibly difficult and time consuming.

Human detection frameworks. Traditional human detectors, such as ACF [4], LDCF [19], and Checkerboard [32], exploit various filters based on Integral Channel Features (IDF) [5] with sliding window strategy.

Recently, the CNN-based detectors have become a predominating trend in the field of pedestrian detection. In [29], self-learned features are extracted from deep neural networks and a boosted decision forest is used to detect pedestrians. Cai *et al.* [1] propose an architecture which uses different levels of features to detect persons at various scales. Mao *et al.* [18] propose a multi-task network to further improve detection performance. Hosang *et al.* [9] propose a learning method to improve the robustness of NMS. Part-based models are utilized in [20, 33] to alleviate occlusion problem. Repulsion loss is proposed to detect persons in crowd scenes [26].

3. CrowdHuman Dataset

In this section, we describe our CrowdHuman dataset including the collection process, annotation protocols, and informative statistics.

3.1. Data Collection

We would like our dataset to be diverse for real world scenarios. Thus, we crawl images from Google image search engine with ~ 150 keywords for query. Exemplary keywords include “Pedestrians on the Fifth Avenue”, “people crossing the roads”, “students playing basketball” and “friends at a party”. These keywords cover more than 40 different cities around the world, various activities (*e.g.*, party, traveling, and sports), and numerous viewpoints (*e.g.*, surveillance viewpoint and horizontal viewpoint). The number of images crawled from a keyword is limited to 500 to make the distribution of images balanced. We crawl $\sim 60,000$ candidate images in total. The images with only a small number of persons, or with small overlaps between persons, are filtered. Finally, $\sim 25,000$ images are collected in the CrowdHuman dataset. We randomly select 15,000, 4,370 and 5,000 images for training, validation, and testing, respectively.

3.2. Image Annotation

We annotate individual persons in the following steps.

- We annotate a *full bounding box* of each individual exhaustively. If the individual is partly occluded, the annotator is required to complete the invisible part and

draw a full bounding box. Different from the existing datasets like CityPersons, where the bounding boxes annotated are generated via drawing a line from top of the head and the middle of feet with a fixed aspect ratio (0.41), our annotation protocol is more flexible in real world scenarios which have various human poses. We also provide bounding boxes for human-like objects, *e.g.*, statue, with a specific label. Following the metrics of [6], these bounding-boxes will be ignored during evaluation.

- We crop each annotated instance from the images, and send these cropped regions for annotators to draw a *visible bounding box*.
- We further send the cropped regions to annotate a *head bounding box*. All the annotations are double-checked by at least one different annotator to ensure the annotation quality.

Fig. 2 shows the three kinds of bounding boxes associated with an individual person as well as an example of annotated image.

We compare our CrowdHuman dataset with previous datasets in terms of annotation types in Table 1. Besides from the popular pedestrian detection datasets, we also include the COCO [17] dataset with only a “person” class. Compared with CrowdHuman, which provides various types of annotations, Caltech and CityPersons have only normalized full bounding boxes and visible boxes, KITTI has only full bounding boxes, and COCOPersons has only visible bounding boxes. More importantly, none of them has head bounding boxes associated with each individual person, which may serve as a possible means to address the crowd occlusion problem.

3.3. Dataset Statistics

Dataset Size. The volume of the CrowdHuman training subset is illustrated in the first three lines of Table 2. In a total of 15,000 images, there are $\sim 340k$ person and $\sim 99k$ ignore region annotations in the CrowdHuman training subset. The number is more than 10x boosted compared with previous challenging pedestrian detection dataset like CityPersons. The total number of persons is also noticeably larger than the others.

Density. In terms of density, on average there are ~ 22.6 persons per image in CrowdHuman dataset, as shown in the fourth line of Table 2. We also report the density from the existing datasets in Table 3. Obviously, CrowdHuman dataset is of much higher crowdness compared with all previous datasets. Caltech and KITTI suffer from extremely low-density, for that on average there is only ~ 1 person per

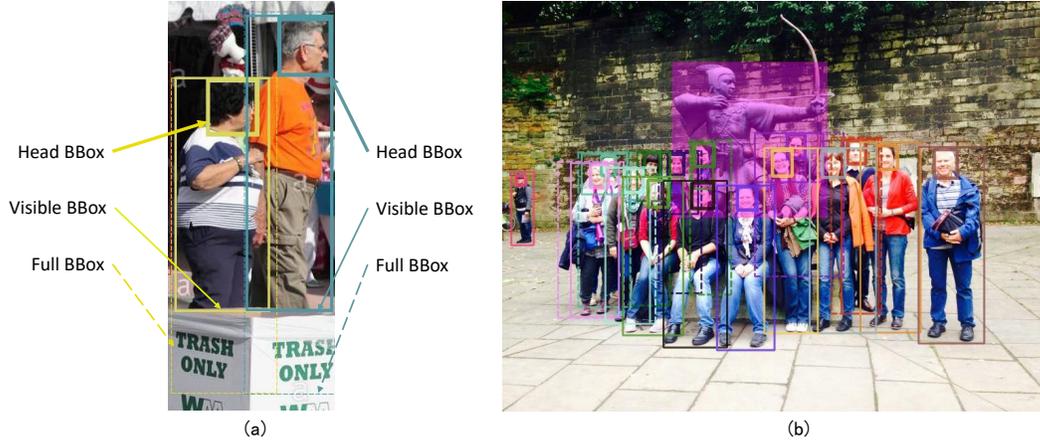


Figure 2. (a) provides an illustrative example of our three kinds of annotations: Head Bounding-Box, Visible Bounding-Box, and Full Bounding-Box. (b) is an example image with our human annotations where magenta mask illustrates the ignored region.

	Caltech	KITTI	CityPersons	COCOPersons	CrowdHuman
Full BBox	✓	✓	✓ [†]	×	✓
Visible BBox	✓	×	✓	✓	✓
Head BBox	×	×	×	×	✓

Table 1. Comparison of different annotation types for the popular human detection benchmarks.[†] : Aligned to a certain ratio.

	Caltech	KITTI	CityPersons	COCOPersons	CrowdHuman
# images	42, 782	3, 712	2, 975	64, 115	15, 000
# persons	13, 674	2, 322	19, 238	257, 252	339, 565
# ignore regions	50, 363	45	6, 768	5, 206	99, 227
# person/image	0.32	0.63	6.47	4.01	22.64
# unique persons	1, 273	< 2, 322	19, 238	257, 252	339, 565

Table 2. Volume, density and diversity of different human detection datasets. For fair comparison, we only show the statistics of training subset.

image. The number in CityPersons reaches ~ 7 , a significant boost while still not dense enough. As for COCOPersons, although its volume is relatively large, it is insufficient to serve as a ideal benchmark for the challenging crowd scenes. Thanks to the pre-filtering and annotation protocol of our dataset, CrowdHuman can reach a much better density.

Diversity. Diversity is an important factor of a dataset. COCOPersons and CrowdHuman contain people in unlimited poses in a wide range of domains, while Caltech, KITTI and CityPersons are all recorded by a car traversing on streets. The number of identical persons is also critical. As reported in the fifth line in Table 2, this number amounts to $\sim 33k$ in CrowdHuman while images in Caltech and KITTI are *not* sparsely sampled, resulting in less amount of identical persons.

Occlusion. To better analyze the distribution of occlusion levels, we divide the dataset into the “bare” sub-

set (occlusion $\leq 30\%$), the “partial” subset ($30\% < \text{occlusion} \leq 70\%$), and the “heavy” subset (occlusion $> 70\%$). In Fig. 3, we compare the distribution of persons at different occlusion levels for CityPersons². The bare subset and partial subset in CityPersons constitute 46.79% and 24.19% of entire dataset respectively, while the ratios for CrowdHuman are 29.89% and 32.13%. The occlusion levels are more balanced in CrowdHuman, in contrary to those in CityPersons, which have more persons with low occlusion.

We also provide statistics on pair-wise occlusion. For each image, We count the number of person pairs with different intersection over union (IoU) threshold. The results are shown in Table 4. In average, few person pairs with an IoU threshold of 0.3 are included in Caltech, KITTI or COCOPersons. For CityPersons dataset, the number is less than one pair per image. However, the number is 9 for CrowdHuman. Moreover, There are averagely 2.4 pairs whose IoU is greater than 0.5 in the CrowdHuman dataset.

²The statistics is computed without group people

We further count the occlusion levels for triples of persons. As shown in Table 5, such cases can be hardly found in previous datasets, while they are well-represented in CrowdHuman.

4. Experiments

In this section, we will first discuss the experiments on our CrowdHuman dataset, including full body detection, visible body detection and head detection. Meanwhile, the generalization ability of our CrowdHuman dataset will be evaluated on standard pedestrian benchmarks like Caltech and CityPersons, person detection benchmark on COCOPersons, and head detection benchmark on Brainwash dataset. We use FPN [15] and RetinaNet [16] as two baseline detectors to represent the two-stage algorithms and one-stage algorithms, respectively.

4.1. Baseline Detectors

Our baseline detectors are Faster R-CNN [21] and RetinaNet [16], both based on the Feature Pyramid Network (FPN) [15] with a ResNet-50 [8] back-bone network. Faster R-CNN and RetinaNet are both proposed for general object detection, and they have dominated the field of object detection in recent years.

4.2. Evaluation Metric

The training and validation subsets of CrowdHuman can be downloaded from our website. In the following experiments, our algorithms are trained based on CrowdHuman train subset and the results are evaluated in the validation subset. An online evaluation server will help to evaluate the performance of the testing subset and a leaderboard will be maintained. The annotations of testing subset will not be made publicly available.

We follow the evaluation metric used for Caltech [6], denoted as mMR, which is the average log miss rate over false positives per-image ranging in $[10^{-2}, 10^0]$. mMR is a good indicator for the algorithms applied in the real world applications. Results on ignored regions will not be considered in the evaluation. Besides, Average Precision (AP) and recall of the algorithms are included for reference.

4.3. Implementation Details

We use the same setting of anchor scales as [15] and [16]. For all the experiments related to full body detection, we modify the height v.s. width ratios of anchors as $\{1 : 1, 1.5 : 1, 2 : 1, 2.5 : 1, 3 : 1\}$ in consideration of the human body shape. While for visible body detection and human head detection, the ratios are set to $\{1 : 2, 1 : 1, 2 : 1\}$, in comparison with the original papers. The input image sizes of Caltech and CityPersons are set to $2\times$ and $1\times$ of the original images according to [31]. As the

images of CrowdHuman and MSCOCO are both collected from the Internet with various sizes, we resize the input so that their short edge is at 800 pixels while the long edge should be no more than 1400 pixels at the same time. The input sizes of Brainwash is set as 640×480 .

We train all datasets with $600k$ and $750k$ iterations for FPN and RetinaNet, respectively. The base learning rate is set to 0.02 and decreased by a factor of 10 after $150k$ and $450k$ for FPN, and $180k$ and $560k$ for RetinaNet. The Stochastic Gradient Descent (SGD) solver is adopted to optimize the networks on 8 GPUs. A mini-batch involves 2 images per GPU, except for CityPersons where a mini-batch involves only 1 image due to the physical limitation of GPU memory. Weight decay and momentum are set to 0.0001 and 0.9. We do not finetune the batch normalization [11] layers. Multi-scale training/testing are not applied to ensure fair comparisons.

4.4. Detection results on CrowdHuman

Visible Body Detection As the human have different poses and occlusion conditions, the visible regions may be quite different for each individual person, which brings many difficulties to human detection. Table 6 illustrates the results for the visible part detection based on FPN and RetinaNet. FPN outperforms RetinaNet in this case. According to Table 6, the proposed CrowdHuman dataset is a challenging benchmark, especially for the state-of-the-art human detection algorithms. The illustrative examples of visible body detection based on FPN are shown in Fig. 5.

Full Body Detection Detecting full body regions is more difficult than detecting the visible part as the detectors should predict the occluded boundaries of the full body. To make matters worse, the ground-truth annotation might be suffered from high variance caused by different decision-makings by different annotators.

Different from the visible part detection, the aspect ratios of the anchors for the full body detection are set as $[1.0, 1.5, 2.0, 2.5, 3.0]$ to make the detector tend to predict the slim and tall bounding boxes. Another important thing is that the RoIs are not clipped into the limitation of the image boundaries, as there are many full body bounding boxes extended out of images. The results are shown in Table 7 and the illustrative examples of FPN are shown in Fig. 4. Similar to the Visible body detection, FPN has a significant gain over RetinaNet.

In Table 7, we also report the FPN pedestrian detection results³ on Caltech, i.e., 10.08 mMR, and CityPersons, i.e., 14.81 mMR. It shows that our CrowdHuman dataset is much challenging than the standard pedestrian detection benchmarks based on the detection performance.

³The results are evaluated on the standard reasonable set

person/img \geq	Caltech		KITTI		CityPersons		COCOPersons		CrowdHuman	
1	7839	18.3%	969	26.1%	2482	83.4%	64115	100.0%	15000	100.0%
2	3257	7.6%	370	10.0%	2082	70.0%	39283	61.3%	15000	100.0%
3	1265	3.0%	273	7.4%	1741	58.5%	28553	44.5%	14996	100.0%
5	282	0.7%	164	4.4%	1225	41.2%	18775	29.3%	14220	94.8%
10	36	0.1%	19	0.5%	610	20.5%	9604	15.0%	10844	72.3%
20	0	0.0%	0	0.0%	227	7.6%	0	0.0%	5907	39.4%
30	0	0.0%	0	0.0%	94	3.2%	0	0.0%	3294	21.9%

Table 3. Comparison of the human density against the widely used human detection dataset. The first column refers to the number of human instances in the image.

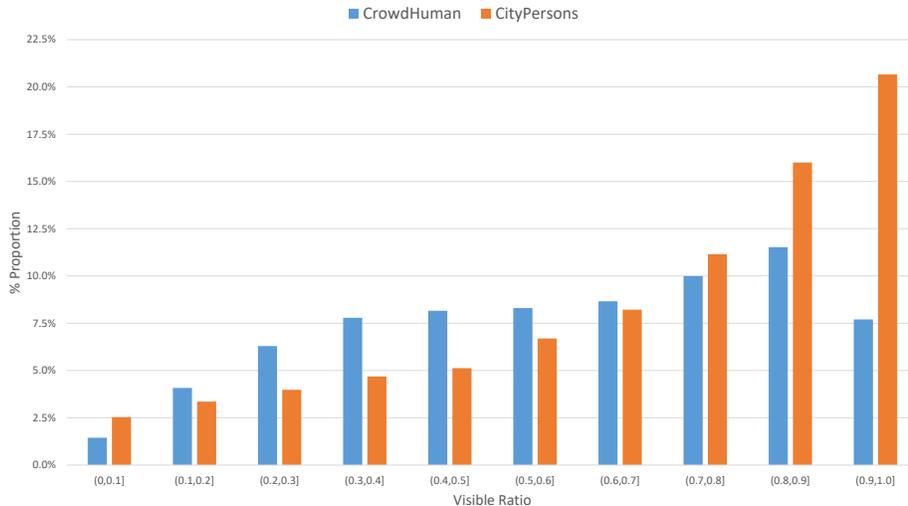


Figure 3. Comparison of the visible ratio between our CrowdHuman and CityPersons dataset. Visible Ratio is defined as the ratio of visible bounding box to the full bounding box.

pair/img	Cal	City	COCO	CrowdHuman
iou>0.3	0.06	0.96	0.13	9.02
iou>0.4	0.03	0.58	0.05	4.89
iou>0.5	0.02	0.32	0.02	2.40
iou>0.6	0.01	0.17	0.01	1.01
iou>0.7	0.00	0.08	0.00	0.33
iou>0.8	0.00	0.02	0.00	0.07
iou>0.9	0.00	0.00	0.00	0.01

Table 4. Comparison of pair-wise overlap between two human instances.

pair/img	Cal	City	COCO	CrowdHuman
iou>0.1	0.02	0.30	0.02	8.70
iou>0.2	0.00	0.11	0.00	2.09
iou>0.3	0.00	0.04	0.00	0.51
iou>0.4	0.00	0.01	0.00	0.12
iou>0.5	0.00	0.00	0.00	0.03

Table 5. Comparison of high-order overlaps among three human instances.

Table 6. Evaluation of visible body detection on CrowdHuman benchmark.

	Recall	AP	mMR
FPN [15]	91.51	85.60	55.94
RetinaNet [16]	90.96	77.19	65.47

Head Detection Head is one of the most obvious parts of a whole body. Head detection is widely used in the practical applications such as people number counting, face detection and tracking. We compare the results of FPN and RetinaNet as shown in Table 8. The illustrative examples of head detection on CrowdHuman by FPN detector are shown in Fig. 6.

4.5. Cross-dataset Evaluation

As shown in Section 3, the size of CrowdHuman dataset is obviously larger than the existing benchmarks, like Caltech and CityPersons. In this section, we evaluate that the



Figure 4. Qualitative results for the full body detection of FPN based on CrowdHuman dataset.



Figure 5. Qualitative results for the visible body detection of FPN based on CrowdHuman dataset.

Table 7. Evaluation of full body detection on CrowdHuman benchmark.

	Recall	AP	mMR
FPN [15]	90.24	84.95	50.42
RetinaNet [16]	93.80	80.83	63.33
FPN on Caltech	99.76	89.95	10.08
FPN on CityPersons	97.97	94.35	14.81

Table 8. Evaluation of Head detection on CrowdHuman benchmark.

	Recall	AP	mMR
FPN [15]	81.10	77.95	52.06
RetinaNet [16]	78.43	71.36	60.64

generalization ability of our CrowdHuman dataset. More specifically, we first train the model on our CrowdHuman dataset and then finetune it on the visible body detection benchmarks like COCOPersons [17], full body detection benchmarks like Caltech [6] and CityPersons [31], and head detection benchmarks like Brainwash [23]. As reported in Section 4.4, FPN is superior to RetinaNet in all three cases. Therefore, in the following experiments, we adopt FPN as our baseline detector.

COCOPersons COCOPersons is a subset of MSCOCO from the images with groundtruth bounding box of “person”. The other 79 classes are ignored in our evaluation. After the filtering process, there are 64115 images from the trainval minus minival for training, and the other 2639 images from minival for validation. All the persons in COCOPersons are annotated as the visible body with different type of human poses. The results are illustrated in Table 9. Based on the pretraining of our CrowdHuman dataset, our algorithm has superior performance on the COCOPersons

Table 9. Experimental results on COCOPersons.

Train-set	Recall	AP	mMR
COCOPersons	95.57	83.83	41.89
Crowd⇒COCO	95.87	85.02	39.79

benchmark against the one without CrowdHuman pretraining.

Caltech and CityPersons Caltech and CityPersons are widely used benchmarks for pedestrian detection, both of them are usually adopted to evaluate full body detection algorithms. We use the reasonable set for Caltech dataset where the object size is larger than 50 pixels. Table 10 and Table 11 show the results on Caltech and CityPersons, respectively. We compare the algorithms in the first part of the tables with:

- FPN trained on the Caltech
- FPN trained on CityPersons
- FPN trained on CrowdHuman
- FPN model pretrained on CrowdHuman and then finetuned on the corresponding target training set

Also, state-of-art algorithms on Caltech and CityPersons are reported in the second part of tables as well. To summarize, the results illustrated in Table 10 and Table 11 demonstrate that our CrowdHuman dataset can serve as an effective pre-training dataset for pedestrian detection task on Caltech and CityPersons⁴ for full body detection.

Brainwash Brainwash [23] is a head detection dataset whose images are extracted from the video footage at every 100 seconds. Following the step of [23], the training

⁴The evaluation is based on $1 \times$ scale.



Figure 6. Qualitative results for the head detection of FPN based on CrowdHuman dataset.

Table 10. Experimental results on Caltech dataset.

Train-set	Recall	AP	mMR
Caltech	99.76	89.95	10.08
CityPersons	99.05	85.81	14.69
CrowdHuman	99.88	90.58	8.81
Crowd⇒Calt	99.88	95.69	3.46
CityPersons⇒Calt [31]	-	-	5.1
Repulsion [26]	-	-	4.0
[18]	-	-	5.5

Table 12. Experimental results on Brainwash.

Train-set	Recall	AP	mMR
Brainwash	98.52	95.74	19.77
Crowd⇒Brain	98.66	96.15	17.24
[23]	-	78.0	-

set has 10,917 images with 82,906 instances and the validation set has 500 images with 3318 instances. Similar to visible body detection and full body detection, Brainwash dataset is evaluated to validate the generalization ability of our CrowdHuman dataset for head detection.

Table 12 shows the results of head detection task on Brainwash dataset. By using the FPN as the head detector, the performance is already much better than the state-of-art in [23]. On top of that, pretraining on the CrowdHuman dataset further boost the result by 2.5% of mMR, which validates the generalization ability of our CrowdHuman dataset for head detection.

5. Conclusion

In this paper, we present a new human detection benchmark designed to address the crowd problem. There are three contributions of our proposed CrowdHuman dataset. Firstly, compared with the existing human detection benchmark, the proposed dataset is larger-scale with much higher crowdness. Secondly, the full body bounding box, the visible bounding box, and the head bounding box are annotated for each human instance. The rich annotations enables a lot of potential visual algorithms and applications. Last but not least, our CrowdHuman dataset can serve as a powerful pretraining dataset. State-of-the-art results have been reported on benchmarks of pedestrian detection benchmarks

Table 11. Experimental results on CityPersons.

Train-set	Recall	AP	mMR
Caltech	87.21	65.87	45.52
CityPersons	97.97	94.35	14.81
CrowdHuman	98.73	98.10	21.18
Crowd⇒City	97.78	95.58	10.67
CityPersons [31]	-	-	14.8
Repulsion [26]	-	-	13.2

like Caltech and CityPersons, and Head detection benchmark like Brainwash. The dataset as well as the code and models discussed in the paper will be released⁵.

References

- [1] Zhaowei Cai, Quanfu Fan, Rogerio S Feris, and Nuno Vasconcelos. A unified multi-scale deep convolutional neural network for fast object detection. *arXiv preprint arXiv:1607.07155*, 2016. **1, 3**
- [2] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. **2**
- [3] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 886–893. IEEE, 2005. **2**
- [4] Piotr Dollár, Ron Appel, Serge Belongie, and Pietro Perona. Fast feature pyramids for object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(8):1532–1545, 2014. **1, 3**
- [5] Piotr Dollár, Zhuowen Tu, Pietro Perona, and Serge Belongie. Integral channel features. 2009. **1, 3**
- [6] Piotr Dollár, Christian Wojek, Bernt Schiele, and Pietro Perona. Pedestrian detection: A benchmark. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 304–311. IEEE, 2009. **1, 2, 3, 5, 7**
- [7] Markus Enzweiler and Dariu M Gavrila. Monocular pedestrian detection: Survey and experiments. *IEEE transactions on pattern analysis and machine intelligence*, 31(12):2179–2195, 2009. **2**

⁵<https://sshao0516.github.io/CrowdHuman/>

- [8] 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, 5
- [9] Jan Hosang, Rodrigo Benenson, and Bernt Schiele. Learning non-maximum suppression. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 3
- [10] Jan Hosang, Mohamed Omran, Rodrigo Benenson, and Bernt Schiele. Taking a deeper look at pedestrians. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4073–4082, 2015. 1
- [11] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. 5
- [12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 1
- [13] Ian Reid Stefan Roth Konrad Schindler Laura Leal-Taix, Anton Milan. Motchallenge 2015: Towards a benchmark for multi-target tracking. 2015. 1
- [14] Jianan Li, Xiaodan Liang, ShengMei Shen, Tingfa Xu, and Shuicheng Yan. Scale-aware fast r-cnn for pedestrian detection. *arXiv preprint arXiv:1510.08160*, 2015. 1
- [15] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, volume 1, page 4, 2017. 5, 6, 7
- [16] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *arXiv preprint arXiv:1708.02002*, 2017. 5, 6, 7
- [17] 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, 3, 6
- [18] Jiayuan Mao, Tete Xiao, Yuning Jiang, and Zhimin Cao. What can help pedestrian detection? In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 3, 8
- [19] Woonhyun Nam, Piotr Dollár, and Joon Hee Han. Local decorrelation for improved detection. *arXiv preprint arXiv:1406.1134*, 2014. 3
- [20] Wanli Ouyang and Xiaogang Wang. A discriminative deep model for pedestrian detection with occlusion handling. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 3258–3265. IEEE, 2012. 3
- [21] 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. 5
- [22] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1
- [23] Russell Stewart, Mykhaylo Andriluka, and Andrew Y Ng. End-to-end people detection in crowded scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2325–2333, 2016. 7, 8
- [24] Bochao Wang Liang Lin Xiaogang Wang Tong Xiao, Shuang Li. Joint detection and identification feature learning for person search. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1
- [25] Andreas Geigerand Philip Lenzand Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 1, 2
- [26] Xinlong Wang, Tete Xiao, Yuning Jiang, and Shen Chunhua Sun, Jian. Repulsion loss: Detecting pedestrians in a crowd. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 3, 8
- [27] Christian Wojek, Stefan Walk, and Bernt Schiele. Multi-cue onboard pedestrian detection. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 794–801. IEEE, 2009. 2
- [28] Yuxiang Peng Zhiqiang Zhang Gang Yu Jian Sun Yilun Chen, Zhicheng Wang. Cascaded pyramid network for multi-person pose estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1
- [29] Liliang Zhang, Liang Lin, Xiaodan Liang, and Kaiming He. Is faster r-cnn doing well for pedestrian detection? *arXiv preprint arXiv:1607.07032*, 2016. 1, 3
- [30] Shanshan Zhang, Rodrigo Benenson, Mohamed Omran, Jan Hosang, and Bernt Schiele. How far are we from solving pedestrian detection? In *IEEE Conference on Computer Vision and Pattern Recognition*. IEEE Computer Society, 2016. 1
- [31] Shanshan Zhang, Rodrigo Benenson, and Bernt Schiele. Citypersons: A diverse dataset for pedestrian detection. 1, 2, 5, 7, 8
- [32] Shanshan Zhang, Rodrigo Benenson, and Bernt Schiele. Filtered channel features for pedestrian detection. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1751–1760. IEEE, 2015. 1, 3
- [33] Chunlun Zhou and Junsong Yuan. Multi-label learning of part detectors for heavily occluded pedestrian detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3486–3495, 2017. 3