



Age Should Not Matter: Towards More Accurate Pedestrian Detection via Self-Training [†]

Shunsuke Kogure ^{1,†}, Kai Watabe ^{1,†}, Ryosuke Yamada ¹, Yoshimitsu Aoki ², Akio Nakamura ³
and Hirokatsu Kataoka ^{2,*} 

¹ National Institute of Advanced Industrial Science and Research Technology (AIST), Tokyo 135-0064, Japan; skogure@aoki-medialab.jp (S.K.); kai.watabe@aist.go.jp (K.W.); ryosuke.yamada@aist.go.jp (R.Y.)

² Department of Electronics and Electrical Engineering, Keio University, Yokohama 223-8522, Japan; aoki@elec.keio.ac.jp

³ Department of Robotics and Mechatronics, Tokyo Denki University, Tokyo 120-0026, Japan; nkmr-a@cck.dendai.ac.jp

* Correspondence: hirokatsu.kataoka@aist.go.jp

[†] Presented at the AAAI Workshop on Artificial Intelligence with Biased or Scarce Data (AIBSD), Online, 28 February 2022.

[‡] These authors contributed equally to this work.

Abstract: Why is there disparity in the miss rates of pedestrian detection between different age attributes? In this study, we propose to (i) improve the accuracy of pedestrian detection using our pre-trained model; and (ii) explore the causes of this disparity. In order to improve detection accuracy, we extend a pedestrian detection pre-training dataset, the Weakly Supervised Pedestrian Dataset (WSPD), by means of self-training, to construct our Self-Trained Person Dataset (STPD). Moreover, we hypothesize that the cause of the miss rate is due to three biases: (1) the apparent bias towards “adults” versus “children”; (2) the quantity of training data bias against “children”; and (3) the scale bias of the bounding box. In addition, we constructed an evaluation dataset by manually annotating “adult” and “child” bounding boxes to the INRIA Person Dataset. As a result, we confirm that the miss rate was reduced by up to 0.4% for adults and up to 3.9% for children. In addition, we discuss the impact of the size and appearance of the bounding boxes on the disparity in miss rates and provide an outlook for future research.

Keywords: computer vision; pedestrian detection; fairness



Citation: Kogure, S.; Watabe, K.; Yamada, R.; Aoki, Y.; Nakamura, A.; Kataoka, H. Age Should Not Matter: Towards More Accurate Pedestrian Detection via Self-Training. *Comput. Sci. Math. Forum* **2022**, *3*, 11. <https://doi.org/10.3390/cmsf2022003011>

Academic Editors: Kuan-Chuan Peng and Ziyang Wu

Published: 24 May 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Recently, research has frequently explored approaches to pedestrian detection which is expected to be applied in various fields. The remarkable progress that has been made in this area is partly due to the large-scale collection of human images from the Web.

However, there are still concerns about the safety of utilizing pedestrian detection in areas such as automated driving. One of these concerns is the disparity in detection rates based on human age and race; specifically, a disparity in detection rates between “adults” and “children” has been reported when using classical human detection methods. Brando [1] affirmed that the difference in the quantity of adult versus child data in the person detection dataset is a problem that naturally arises from demographics. There are a small number of “children” in the existing pedestrian dataset, which we assume is responsible for a sample bias and a detection rate disparity between “adults” and “children.”

In this paper, we constructed our Self-Trained Person Dataset (STPD) by extending the Weakly Supervised Pedestrian Dataset (WSPD) [2] to improve the accuracy of person detection. We studied the effect of each age attribute on detection performance using each pre-trained model generated by the WSPD and STPD. The INRIA Person Dataset [3] is used to evaluate the detection performance. We re-annotated both the training and test data

of the INRIA Person Dataset to rigorously investigate the effect of age on the accuracy of pedestrian detection. For this re-annotation, we added the age attribute and the bounding box (bbox). In this way, we constructed a dataset for pedestrian detection validation with the age attribute. In addition, we studied the reason for the disparity in detection rate by age. Specifically, we examined the age gap in the detection rate using three experiments: (i) we clarify whether there is a difference in appearance between “adults” and “children”; (ii) we study the impact of the data augmentation of children’s learning data alone on the missed rate; and (iii) finally, we compare the miss rate for each age attribute when the scale of the input image is changed. Our contributions are as follows:

- The STPD was constructed by extending the pedestrian dataset, WSPD, using self-training.
- In order to rigorously evaluate the detection performance for “adults” versus “children,” we constructed a new evaluation dataset.
- The person detector with STPD pre-training reduced the miss rate of “adults” and “children” compared to the detector with WSPD pre-training. Furthermore, we observed a mitigating effect of self-training on the detection rate gap.
- We studied three aspects to investigate the cause of the gap in detection rates by age: (i) the appearance of “adults” and “children”; (ii) the quantity of data for “children”; and (iii) the scale of the input images.

2. Related Work

2.1. Detector

In recent years, approaches to detection have been dramatically improved with the rise of deep neural networks (DNNs). In the literature, a two-step region identifier and DNN-based classification have been proposed [4]. The basic approach, called R-CNN, follows three steps when generating bounding boxes: (i) detecting areas in the image that may contain objects (region proposal); (ii) extracting CNN features from region candidates; and (iii) classifying objects based on the extracted features. Fast R-CNN [5] also generates region proposals, but it is more efficient than R-CNN because Fast R-CNN pools the CNN features corresponding to each region proposal. Faster R-CNN [6] adds a region proposal network (RPN) to generate a region proposal in the network. Current research focuses on widely divided one-shot detectors such as single-shot multi-box detector (SSD) [7] and you look only once (YOLO) [8].

Recent works have also focused on high-performance detectors, such as M2Det [9], RetinaNet [10] and instance segmentation with Mask R-CNN [11]. In this paper, we applied SSD as a method of detecting people in a dataset. Here, we used a WSPD pre-trained model for self-training.

2.2. Pedestrian Detection

In the past decade, approaches to person detection have dramatically improved. Recent work has proposed configurations to improve recognition and localization, including DNNs, semantic segmentation, combined methods and small image and cloud analysis. However, in order to train these models, it is necessary to prepare a large dataset and fine-tune its architecture (e.g., SSD or M2Det). Wilson et al. tested whether an object detector can correctly detect pedestrians with different skin colors [12]. In addition, they found that it is problematic to accurately detect children because their miss rate is higher than that of adults [1]. In this study, we were able to detect pedestrians more reliably than in previous studies.

3. Self-Training

3.1. Problem

A number of datasets for pedestrian detection have been proposed to date. However, as shown in Table 1, their scale is small compared to those used for object detection. Minoguchi et al. proposed a weakly supervised learning method that eliminates false positives using existing pre-trained models by referring to bounding boxes and SVM and

by constructing a labeled dataset called the Weakly Supervised Person Dataset (WSPD) [2], which far exceeds the scale of previous pedestrian detection datasets. To the best of our knowledge, the WSPD is the largest existing pedestrian dataset. Minoguchi et al. revealed the detection performance of the pre-trained model on that dataset but did not mention the disparity in the miss rate for each age attribute. Table 2 shows the attribute distribution of some bounding boxes in the WSPD. This distribution is based on our random selection of 5000 bounding boxes from the WSPD and their classification by attribute. The “Noise” label indicates that there is no person in the bounding box, while the “Multiple” label indicates that the bounding box contains multiple people. As such, we can see that the existing pedestrian dataset has a large bias in the distribution of the quantity of data; in particular, the data for children are excessively limited. Therefore, it is necessary to check whether this bias in the quantity of data contributes to the disparity in detection performance.

Table 1. Comparison of object detection and person detection datasets.

Dataset	Image	Bounding Box	Class
Pascal VOC	11,530	27,450	20
MS COCO	123,287	896,782	80
OpenImages v5	1,743,042	14,610,229	600
CityPersons	5000	35,016	2
EuroCity Persons	47,300	238,200	17
Caltech Pedestrian	250,000	350,000	2
WSPD	2,822,421	8,716,461	2
STPD (Ours)	3,461,024	9,739,996	1
FA-INRIA (Ours)	902	2993	2

Table 2. The age attribute statistics for people in bounding boxes in 5000 randomly sampled images from the WSPD dataset. The “Noise” label indicates that there is no person in the bounding box, whereas “Multiple” label means that one bounding box contains multiple people. In this paper, images labeled “Multiple” are not considered.

Annotation Type	Images	%
(i) Adult	2687	53.7
(ii) Children	169	3.4
(iii) Noise	536	10.7
(iv) Multiple	1608	32.2

3.2. Solution

As previously mentioned, we can see that the WSPD contains the largest number of images and bounding boxes among the available person detection datasets. Furthermore, the WSPD contains a wide variety of person images collected from various locations around the world. The semi-automatically collected dataset has millions of bounding boxes which may be useful for pre-training. We used a WSPD pre-trained model to apply self-training to another dataset to collect high-quality bounding boxes and to investigate the impact of each age attribute on the miss rate. Our self-training pipeline is shown in Figure 1. First, we input images from the Places365 dataset [13] to the SSD, a detector pre-trained with the WSPD, to estimate the location of the bounding box. We assigned a pseudo-label of “person” to the predicted bounding box. The determination of the location of the bounding box when generating the pseudo-label is expressed by the following equation:

$$(y', b'_{box}) = D(x; \theta), \quad (1)$$

where y' and b' represent the predicted values of the object category and bounding box, respectively, and θ represents the learned parameters of the detector. Our self-training approach allows us to automatically extend the dataset. We refer to the WSPD and the generated pseudo-labeled Places365 dataset together as the Self-Trained Person Dataset (STPD).

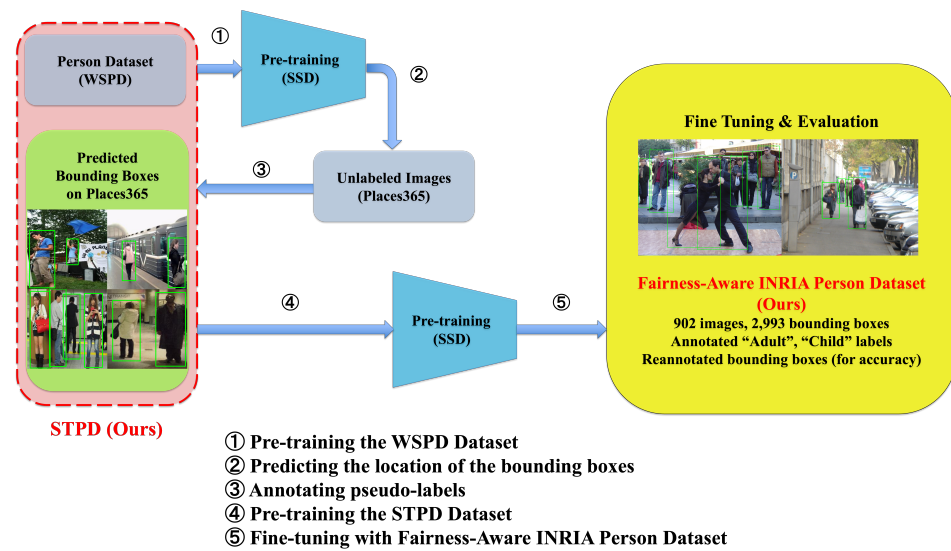


Figure 1. The self-training approach. We used a model pre-trained using the WSPD dataset with the SSD to infer the location of bounding boxes for unlabeled images from the Places365 dataset. We then gave each predicted bounding box a “person” attribute label. By combining these pseudo-labels with the WSPD labels and pre-training them with the SSD, we were able to build a larger model to verify miss rates.

Furthermore, we pre-trained the constructed STPD and compared its detection performance with the model pre-trained using the WSPD. In order to examine the disparity in the miss rate among age attributes, it is essential to add an age attribute to the bounding box. Then, in order to evaluate the miss rate for each age attribute, we assigned “adult” and “children” labels to the INRIA Person Dataset, which is commonly used for person detection, using the models pre-trained with the WSPD and STPD, respectively. We also re-annotated the location of the bounding box. These two age attributes follow the age categories defined by the Statistics Bureau of the Ministry of Internal Affairs and Communications in Japan for (i) children (0–14 years) and (ii) adults (15 years and older). As a result, we constructed a pedestrian detection dataset consisting of 902 images and 2993 bounding boxes for training and evaluation. We named this dataset the Fairness-Aware INRIA Person Dataset (FA-INRIA). An example of the annotations and the breakdown of the dataset attributes are shown in Figure 2 and Table 3, respectively.



Figure 2. Examples of age attribute annotation in Fairness-Aware INRIA Person Dataset (FA-INRIA).

Table 3. The age attributes in the Fairness-Aware INRIA Person Dataset (FA-INRIA).

Age	Images	Bounding Boxes
Adult	870	2672
Children	151	321
All	902	2993

3.3. Experimental Settings

In this paper, we compared the results under the same pre-training conditions. The batch sizes for pre-training the SSD were set to 64, 128 and 256, the number of epochs was set to 10, and the learning rate was set to 0.0005. When we conducted fine-tuning with the FA-INRIA using the pre-trained models on each dataset, the batch size was set to 4, the number of iterations was set to 12,000, and the learning rate was set to 0.0005. Furthermore, the training and test datasets were used with the same configuration as the original INRIA Person Dataset. The experimental settings described below also conform to these conditions.

3.4. Evaluation Metric

We only used the miss rate as an evaluation metric to assess the detection performance for adults and children. In person detection, the relationship between the miss rate and false positives is often evaluated for each image. However, our goal is to detect all ground truth bounding boxes. Therefore, we calculated the miss rate by examining the breakdown of the age attributes of the bounding boxes that could not be detected. The miss rate M is derived by the following equation:

$$MR = 1 - Recall \quad (2)$$

In this paper, we calculated the standard deviation to represent the miss rate disparity among age attributes:

$$MR_{std} = \frac{1}{n} \sum_{i=1}^n (MR_i - \bar{MR})^2, \quad (3)$$

where n refers to the number of classes of attributes, which in this study was two (“Adult” and “Children”).

3.5. Results

Table 4 shows the miss rate in the FA-INRIA Person Dataset using each of the pre-trained models. Compared to the model pre-trained with the WSPD, the model pre-trained with STPD reduced the miss rate by up to 0.4% for adults and up to 3.9% for children. In the WSPD pre-trained model, the disparity between the miss rates of adults and children was a maximum standard deviation of 4.6% and a minimum of 3.1%. In contrast, the STPD pre-trained model had a maximum standard deviation of 2.9% and a minimum standard deviation of 2.1%.

Table 4. Detection performance comparisons for our FA-INRIA. We use standard deviation to describe the disparity in detection rates between attributes. It is clear that our approach reduces the miss rate for all attributes.

Dataset	Batch Size, Epochs	MR_{adult}	MR_{child}	MR_{std}
WSPD	64, 10	13.9	23.1	4.6
	128, 10	13.8	21.2	3.2
	256, 10	13.1	19.2	3.1
STPD (ours)	64, 10	13.8	19.2	2.7
	128, 10	13.4	19.2	2.9
	256, 10	13.1	17.3	2.1

Then, the detection results of fine-tuning with the FA-INRIA using the pre-trained detectors on each dataset are shown in Figure 3, illustrating that the STPD pre-trained model is able to detect people that the WSPD pre-trained model misses.



Figure 3. Comparison of detection results of WSPD and STPD.

4. Analysis and Discussion

4.1. The Relationship between the Bias in the Quantity of Data and the Miss Rate

In the aforementioned results, we successfully generated a pseudo bounding box containing a person from the Places365 dataset. In Figure 1, we present a visualization of the location of a person's bounding box that was predicted during the process of self-training. This method was implemented based on the success of self-training in object detection [14] and was found to reduce the miss rate for adults and children, respectively. Moreover, it is effective in collecting data on pedestrians regardless of their age attributes, and not only on children for whom the number of data is small. If the bias in the quantity of data between age attributes is the primary cause of the disparity in detection performance, then it is only the bounding boxes for children that need to be more efficiently collected. However, manual annotation is very costly and impractical. Therefore, we applied data augmentation to the children's bounding boxes in the FA-INRIA training data to investigate the effect on the miss rate for adults and children. In our work, we tried to augment the children's bounding boxes by applying horizontal flip.

Table 5 shows the detection performance when data augmentation is applied to the children's bounding boxes. It can be seen that when the batch size is 256, the miss rate for both attributes decreases. However, when the batch size is 64 or 128, the miss rate for children does not change, while the miss rate decreases for adults. These results indicate that applying data augmentation is effective in improving the overall detection performance. On the other hand, when we focus on the standard deviation, we must not forget that the disparity in detection performance between age attributes is expanding. First and foremost, a "person" can be an adult or a child. If the detection performance for adults is improved solely by increasing the data of children, we would consider that the bias in the quantity of data between classes is not directly relevant.

Table 5. The impact of applying data augmentation (horizontal flip) only to the bounding boxes of the children in the training data. The results show that applying data augmentation is effective in improving the overall detection performance. On the other hand, it may increase the disparity in detection performance among age attributes.

Batch Size, Epochs	MR_{adult}		MR_{child}		MR_{std}	
	w/o Aug.	w/ Aug.	w/o Aug.	w/ Aug.	w/o Aug.	w/ Aug.
64, 10	13.8	12.6	19.2	19.2	2.7	3.3
128, 10	13.4	12.1	19.2	19.2	2.9	3.6
256, 10	13.1	10.7	17.3	15.4	2.1	2.4

4.2. The Relationship between the Size of a Person's Bounding Box and the Miss Rate

Detecting small objects is a difficult task in object and person detection research because of the limited information that can be obtained from a bounding box with a small image size. It is clear that children have smaller bodies than adults. Therefore, the bounding boxes of children tend to be smaller than those of adults. Thus, we thought it would be important to investigate the size of bounding boxes in the FA-INRIA.

Figure 4 presents the distribution of the size of the bounding boxes for adults and children. Adults are shown in red and children are shown in blue. This distribution indicates that most of the bounding boxes that exceed the size of 600 pixels \times 300 pixels in height and width, respectively, are for adults. In other words, the difference in the size distribution of the bounding boxes may be one of the factors affecting the disparity in the miss rate. Figure 5 also shows the distribution of the size of the bounding boxes in the image for the FA-INRIA (test set): the bounding boxes that could be detected are shown in red and the missed bounding boxes are shown in blue. As you can see in these figures, most of the missed bounding boxes are biased towards the smaller image size. In other words, in order to further mitigate the disparity in the miss rate, it is necessary to use detectors that can detect small persons.

In this paper, we investigated the effect of changing the image size of the input on the miss rate of each attribute. The SSD resizes the input image to a set size regardless of the size of the original image. This process is likely to result in the missing details of the image. In order to detect small bounding boxes, we thought that increasing the size of the input image would suppress the missing information. We examined three patterns of input image sizes: (i) 150 pixels \times 150 pixels; (ii) 300 pixels \times 300 pixels; and (iii) 600 pixels \times 600 pixels. The default size for the SSD is 300 pixels \times 300 pixels. For more accurate validation, we also used a sub-dataset with the same number of bounding boxes for adults and children in the training data.




Figure 4. Distribution of bounding boxes for adults and children in the FA-INRIA Person Dataset. Children's bounding boxes tend to be relatively smaller than those of adults.




Figure 5. Whether bounding boxes can be detected in test data (red: detected, blue: missed).

Table 6 shows the miss rate when the input image size of the SSD is changed. It can be seen that increasing the size of the input image is a major factor in reducing the miss rate. On the other hand, when the input size is small (150 pixels \times 150 pixels), the miss rate for children is very poor. We consider that this is because image information is also missing due to the relatively smaller bounding box. As shown in Figure 4, children's bounding boxes are more difficult to detect when the input size is small because children have a relatively higher proportion of small bounding boxes than adults. Based on this result and Figures 4 and 5, we conclude that the unbalanced distribution of the bounding box sizes is one of the main reasons for the disparity in detection performance between adults and children.

Table 6. The effect of changing the input size of the image to the SSD on the detection performance for each age attribute. The results show that increasing the input size decreases the miss rate. In addition, children are more strongly affected by changes in the size of the input. We conclude that the bias in the size of the bounding box is a major factor in the disparity in detection performance.

Batch	MR_{adult}			MR_{child}			MR_{std}		
Size, Epochs	Input size of the image (pixels \times pixels)								
	150	300	600	150	300	600	150	300	600
64, 10	14.9	14.1	14.4	17.3	15.4	15.4	1.2	0.7	0.5
128, 10	15.2	14.6	13.9	21.2	17.3	15.4	3.0	1.4	0.8
256, 10	14.1	14.7	13.4	21.2	17.3	15.4	3.6	1.3	1.0

4.3. Appearance Difference

We considered two aspects: the bias in the quantity of data between classes and the size of the bounding boxes. However, as shown in Figure 5, we can see that some people are not detected even though the bounding box is relatively large. Moreover, as mentioned in Section 4.1, we found that the bias in the quantity of data between classes is most likely not

relevant. These results suggest that there might be other factors that generate disparities in detection performance between age attributes. Subsequently, we hypothesized that there would be apparent differences between the distributions of bounding box sizes of adults and children as they differ significantly in size.

Figure 6 shows the compression of the image features using t-SNE and the visualization of the distribution. It is difficult to imagine that there is a disparity in detection performance based on the appearance of the distribution which is not clearly divided by age attribute and is evenly distributed. This result supports the fact that applying data augmentation to the children's bounding boxes was more effective in improving the detection rate for adults than for children. Since there is no apparent difference between adults and children, we reiterate that we do not need to consider the bias in the quantity of data between classes to reduce the miss rate for children.




Figure 6. Data visualization of bounding boxes using t-SNE (blue: adults, red: children). There is no apparent significant difference between the bounding boxes of children and adults. As mentioned in Section 4.1, when data augmentation was applied to children's bounding boxes, the miss rate was strongly affected for adults but not for children. This data visualization supports the consideration that the bias in the quantity of data between classes has little to do with the disparity in detection performance.

5. Conclusions

In this paper, we investigated and examined various perspectives on the causes of disparity in the detection performance between adults and children in the task of pedestrian detection. As a first experiment, we confirmed that self-training extends the pre-training model and improves the overall detection performance. Then, we found that applying data augmentation to the bounding boxes of children—for whom there is less data available than for adults—significantly improves the detection performance for adults but not children. We also visualized the feature distribution of the bounding boxes using t-SNE and found that there was no apparent difference between adults and children. These results indicate that it is not necessary to consider the bias in the quantity of data in terms of age attributes in pedestrian detection.

On the other hand, when we looked at the size of the bounding boxes in our FA-INRIA, we observed that the distribution was biased toward a smaller size for children than for adults. In addition, we found that changing the input size of the image fed to the detector had a significant impact on detection performance for children. In other words, we concluded that the disparity in the size of the bounding boxes was a major factor in the

disparity in detection performance among age attributes. In the future, focusing on the detection of small bounding boxes will help mitigate the bias between attributes.

Author Contributions: Main contribution: S.K., K.W. and R.Y.; supervision: Y.A., A.N. and H.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Brandao, M. Age and Gender Bias in Pedestrian Detection Algorithms. *arXiv* **2019**, arXiv:1906.10490.
2. Minoguchi, M.; Okayama, K.; Satoh, Y.; Kataoka, H. Weakly Supervised Dataset Collection for Robust Person Detection. *arXiv* **2020**, arXiv:2003.12263.
3. Dalal, N.; Triggs, B. Histograms of oriented gradients for human detection. In Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05), San Diego, CA, USA, 20–25 June 2005.
4. Girshick, R.; Donahue, J.; Darrell, T.; Malik, J. Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Columbus, OH, USA, 23–28 June 2014.
5. Girshick, R. Fast R-CNN. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Santiago, Chile, 7–13 December 2015.
6. Ren, S.; He, K.; Girshick, R.; Sun, J. Faster R-CNN: towards real-time object detection with region proposal networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **2016**, *39*, 1137–1149. [[CrossRef](#)] [[PubMed](#)]
7. Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.-Y.; Berg, A.C. SSD: Single Shot Multibox Detector. In Proceedings of the European Conference on Computer Vision (ECCV), Amsterdam, The Netherlands, 11–14 October 2016.
8. Redmon, J.; Divvala, S.; Girshick, R.; Farhadi, A. You Only Look Once: Unified, Real-Time Object Detection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 7–30 June 2016.
9. Zhao, Q.; Sheng, T.; Wang, Y.; Tang, Z.; Chen, Y.; Cai, L.; Ling, H. M2det: A single-shot Object Detector Based on Multi-level Feature Pyramid Network. *Proc. AAAI Conf. Artif. Intell.* **2019**, *33*, 9259–9266. [[CrossRef](#)]
10. Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollar, P. Focal Loss for Dense Object Detection. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Venice, Italy, 22–29 October 2017.
11. He, K.; Gkioxari, G.; Dollar, P.; Girshick, R. Mask R-CNN. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), Venice, Italy, 22–29 October 2017.
12. Wilson, B.; Hoffman, J.; Morgenstern, J. Predictive Inequity in Object Detection. *arXiv* **2019**, arXiv:1902.11097.
13. Zhou, B.; Lapedriza, A.; Khosla, A.; Oliva, A.; Torralba, A. Places: A 10 million Image Database for Scene Recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* **2017**, *40*, 1452–1464. [[CrossRef](#)]
14. Zoph, B.; Ghiasi, G.; Lin, T.-Y.; Cui, Y.; Liu, H.; Cubuk, E.D.; Le, Q. Rethinking Pre-training and Self-training. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 3833–3845.