

Noisy Annotation Refinement for Object Detection

Jiafeng Mao
mao@hal.t.u-tokyo.ac.jp

The University of Tokyo
Tokyo, Japan

Qing Yu
yu@hal.t.u-tokyo.ac.jp

Yoko Yamakata
yamakata@hal.t.u-tokyo.ac.jp

Kiyoharu Aizawa
aizawa@hal.t.u-tokyo.ac.jp

Abstract

Supervised training of object detectors requires well-annotated large-scale datasets, whose production is costly. Therefore, some efforts have been made to obtain annotations in economical ways, such as cloud sourcing. However, datasets obtained by these methods tend to contain noisy annotations such as inaccurate bounding boxes and incorrect class labels. In this paper, we propose a new problem setting of training object detectors on datasets with entangled noises of annotations of class labels and bounding boxes. Our proposed method efficiently decouples the entangled noises, corrects the noisy annotations, and subsequently trains the detector using the corrected annotations. We verified the effectiveness of our proposed method and compared it with state-of-the-art methods on noisy datasets with different noise levels. The experimental results show that our proposed method significantly outperforms state-of-the-art methods.

1 Introduction

Object-detection technologies using supervised learning have achieved rapid development in recent years. Most research results on object detection are based on reliable well-annotated datasets, such as ImageNet [1] and MS-COCO [2]. The production of a well-annotated dataset is labor intensive and extremely expensive in real-world applications, and some efforts [3, 4, 5, 6] have been made to reduce the cost of dataset production. However, the datasets obtained by these methods usually contain noisy annotations such as inaccurate bounding boxes and incorrect class labels. Owing to the interference of noise, the detectors trained on these datasets tend to perform worse than detectors trained on well-annotated datasets.

In this paper, we refer to datasets that contain incorrect annotations as noisy datasets. We propose a novel problem setting of training object detectors on noisy datasets with entangled noise. As multi-task learning models, object detectors are affected by incorrect annotations in a far more complex manner than CNN-based image classifiers. The entanglement of incorrect classification annotations and inaccurate localization annotations makes the analysis of complex effects challenging. For example, an incorrect localization annotation not only

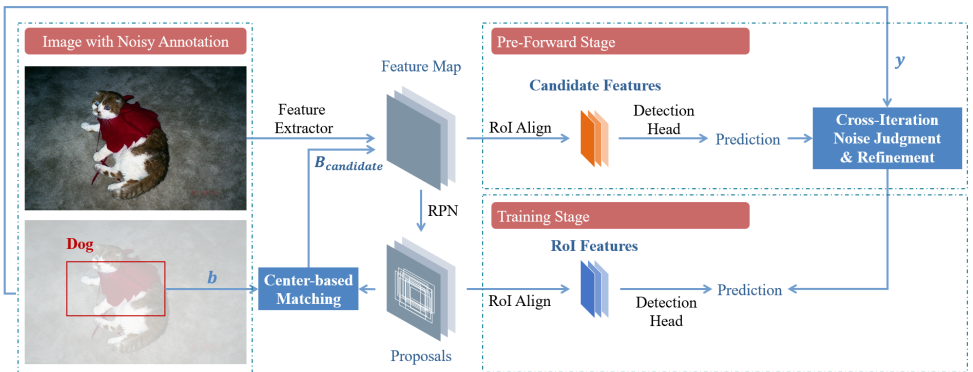


Figure 1: Architecture of the proposed framework. The center matching initially corrects the annotated bounding boxes and outputs the results grouped with the candidate regions, as detailed in Section 3.2. The pre-forward stage classifies and regresses the candidate regions, and the cross-iteration noise judgment identifies the noisy labels on the basis of the results; these are introduced in Sections 3.1 and 3.3, respectively.

can directly mislead the training of the regression head of the detector but also can affect the training of the classification head by providing misplaced region features. From the results of the experiments on artificial noisy datasets, we observe a huge gap in performance between the detectors trained on noisy datasets and those trained on clean datasets. To this end, we propose a framework based on Faster R-CNN to mitigate the negative impact of both types of noise by correcting incorrect training annotations online. The main idea of our proposal is to correct the two types of noise in the dataset by the prediction of the model as the training proceeds and subsequently using the corrected annotations to supervise the training of the model.

We summarize the contributions of this paper as follows:

- We propose a novel problem setting, which is training object detector on dataset with entangled noisy annotation.
- We comprehensively evaluate the impact of location noise and classification noise on object detector.
- We propose a framework and several novel methods to decouple the entangled noise. By reducing the noise level of the noisy training dataset, our framework can efficiently alleviate the negative impact of noisy annotations on object detector.
- We conduct experiments on artificial noisy datasets with different types and levels of noise; our proposal significantly outperforms state-of-the-art methods.

2 Related Work

For training DNNs with noisy data in classification task, Arpit *et al.* [1] suggested that DNNs first learn simple patterns and subsequently memorize noisy data. Furthermore, it was verified that noisy labels lead to high training losses [2]. Based on these results, memorization-based methods for image classification have been proposed. Co-teaching [19, 26] simultaneously trains two classifiers, and each classifier selects samples with relatively low losses

in each mini-batch to teach the peer classifier. The joint optimization framework [24] uses model prediction to update the labels in the dataset and uses a label distribution-based regularization term to prevent the model from irreversibly over-modifying the dataset. In addition, methods such as JoCoR [24] and symmetric cross-entropy loss [23] that use the predictions of the model to guide the training are proposed.

Unlike with the noisy label-supervised image classification, which has achieved many results [4, 10, 11, 20, 22, 27, 28], only a few studies have focused on the noisy annotation-supervised object detection. Mao *et al.* [15] observed a positive correlation between the noise level of the dataset and the training loss of the object detector, and consequently applied the joint optimization framework to the noisy annotation-supervised object-detection task, thereby significantly reducing the impact of noise on the model accuracy. However, their method only focuses on localization annotations and assumes that all classification labels are correct. NOTE-RCNN [6] addresses the noisy annotations using a pre-trained detector, whose training requires a small number of clean annotations. Li *et al.* [12] used a gradient-based approach, named CA-BBC, to correct noisy localization annotations and a pseudo-label-based approach to correct noisy classification labels. They updated all annotations in the training dataset indiscriminately. In contrast to their method, our method first discriminates noisy labels using a reliable method and subsequently applies pseudo-label correction to those labels judged to be incorrect on the basis of the accurate noise judgment.

3 Proposal

For simplicity, in the following explanation, we assume that there is only one object and its corresponding annotation is $\{b, y\}$ on each image x , where b denotes the annotated bounding box of the object and y denotes the class label of the object, and there is only one image in each iteration. In practice, the steps described below are applied to multiple objects in each input image.

3.1 Pipeline

Faster R-CNN In our proposed framework, we train the Faster R-CNN with an additional pre-forward stage to decouple the entangled localization annotation noise and classification noise. Faster R-CNN consists of the feature extractor \mathcal{E} , the region proposal network (RPN) \mathcal{R} , the classification head \mathcal{H}_c and the regression head \mathcal{H}_r . For each iteration during the training, the feature map f of the training image x is extracted by the feature extractor \mathcal{E} as follows:

$$f = \mathcal{E}(x). \quad (1)$$

The RPN \mathcal{R} takes f and generates region proposals P , which is denoted by

$$P = \mathcal{R}(f). \quad (2)$$

Then, P is subsequently sampled and annotated on the basis of the annotation $\{b, y\}$, and each region r in sampled P is used to extract the region features f_r as follows:

$$f_r = \text{RoIAlign}(f, r). \quad (3)$$

Finally, the classification head \mathcal{H}_c takes the region features and outputs the class-wise probability $s(r)$ of each region r , which is denoted by:

$$s(r) = \mathcal{H}_c(f_r). \quad (4)$$

At the same time, the regression head \mathcal{H}_r takes the region features as follows:

$$r_{regressed} = \mathcal{H}_r(r, f_r), \quad (5)$$

where $r_{regressed}$ is the regressed region of r , which is the refined region according to the offset prediction of the regression head.

In the general training of the Faster R-CNN, the feature extractor \mathcal{E} , the RPN \mathcal{R} , the classification head \mathcal{H}_c , and the regression head \mathcal{H}_r are optimized by SGD under the supervision of the annotated region proposals. To deal with the noisy annotation, we insert an additional pre-forward stage after the generation of the proposals P (Eq. 2) and before the training of the network, as shown in Figure 1.

Pre-forward Stage for Annotation Refinement First, to deal with the localization noise, we propose a center-matching correction to generate an initially corrected bounding box b^* and a set of candidate regions $B_{candidate}$ based on the region proposals P , which will be detailed in Section 3.2. Then, the region features of b^* and $B_{candidate}$ are calculated using the RoI Align [9] as follows:

$$f_{b^*} = \text{RoIAlign}(f, b^*), \quad (6)$$

$$F_{candidate} = \{\text{RoIAlign}(f, b) \mid b \in B_{candidate}\}. \quad (7)$$

Next, we feed f_{b^*} and $F_{candidate}$ to the classification head and the regression head as follows:

$$s(b^*) = \mathcal{H}_c(f_{b^*}), \quad (8)$$

$$S_{candidate} = \{\mathcal{H}_c(f) \mid f \in F_{candidate}\}, \quad (9)$$

$$B_{regressed} = \{\mathcal{H}_r(b, f) \mid (b, f) \in (B_{candidate}, F_{candidate})\}, \quad (10)$$

where $s(b^*)$ and $S_{candidate}$ denote the classification probability of the initially corrected bounding box $s(b^*)$ and candidate regions $B_{candidate}$, respectively. $B_{regressed}$ denotes the regressed candidate regions of $B_{candidate}$.

To identify the label noise, we apply cross-iteration noise judgment (as described in Section 3.3) on the basis of the classification loss $\mathcal{L}(b^*, y)$ for b^* , which is denoted by

$$\mathcal{L}(b^*, y) = \text{CrossEntropy}(s(b^*), y). \quad (11)$$

For each label y that is judged to be noisy, we use the class having the largest predicted probability as the pseudo-label y^* , which is denoted by

$$y^* = \arg \max_c s(c|b^*), \quad (12)$$

where c denotes the classes of the dataset. When

$$s(y^*|b^*) > 0.5, \quad (13)$$

we consider the pseudo-label y^* to be reliable and we update y by y^* . Otherwise, $\{y, b\}$ is discarded and not used for training in the current iteration. Once an annotation $\{y^*, b\}$ is used for training, the annotated bounding box b is further modified to b_m^* as follows:

$$b_m^* = (b^* + b_1^* + b_2^*)/3, \quad (14)$$

where $\{b_1^*, b_2^*\} \subset B_{regressed}$ denotes the regression results of $\{b_1, b_2\} \subset B_{candidate}$, where $\{b_1, b_2\}$ are the candidate regions with the highest and second-highest confidence in category y^* . We use the refined annotation $\{y^*, b_m^*\}$ to train the Faster R-CNN.

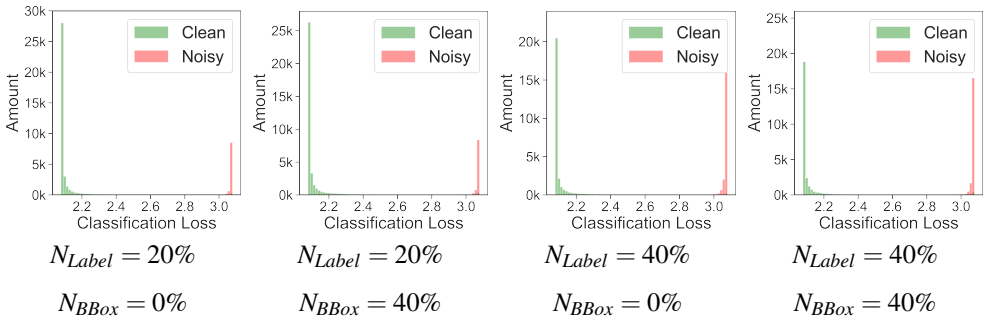


Figure 2: Distribution of the classification loss caused by noisy labels and clean labels. The loss values of the noisy labels are clearly separated from those of the clean labels.

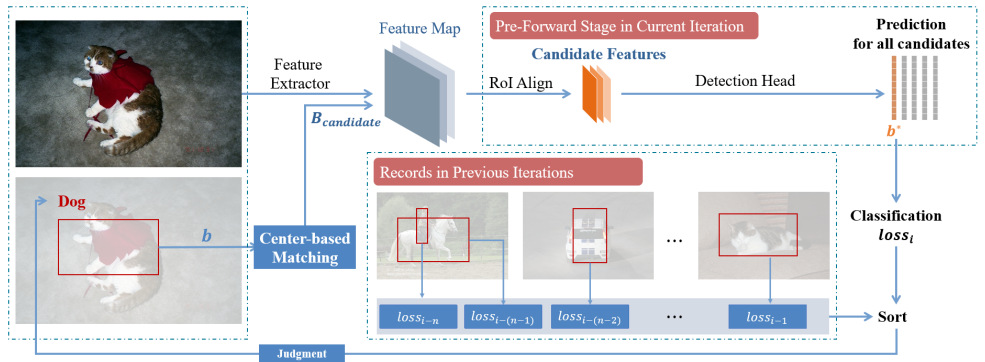


Figure 3: Architecture of the cross-iteration noise judgment. A dynamic threshold is calculated according to the loss values maintained in the previous iterations. The CINJ judges the credibility of the current labels by comparing their loss values with the threshold. The labels judged to be clean are maintained for training, whereas the labels judged to be noisy are modified, as described in Section 3.1.

3.2 Center-Matching Correction

Previous works [15, 21] verified the conclusion that the noise level of classification labels and location annotations is positively correlated with the corresponding training loss and proposed that modifying the annotations in the direction of reducing the loss value can reduce the noise level of the dataset. Inspired by their work, we also attempt to correct the training annotation with prediction and optimize the network parameters with SGD alternately.

Because the classification head \mathcal{H}_c and the regression head \mathcal{H}_r are dependent on the region proposals generated by RPN \mathcal{R} , we decide to find the correct region proposals to deal with localization noise first. Fortunately, we notice that both the training and the inference of the RPN are class-agnostic, which implies that we can ignore the class label noise and correct the location annotation by matching the proper proposals to each annotated bounding box.

The most straightforward method is to assign the proposals to the closest annotated bounding boxes based on their IoU. However, when the annotated bounding boxes are significantly distorted owing to noise, these bounding boxes cannot be matched to any proposals, even if the proposals are accurate. To solve this problem, we propose a center matching

strategy that is more effective than the IoU, as it can match a bounding box with nearby proposals, even if the annotated bounding box is severely distorted in shape and has low IoUs with the corresponding proposals.

We introduce the criterion that calculates the fitness between the bounding boxes and the proposals. Given an annotated bounding box b and a proposal $p \in P$, where P denotes all the proposals generated by the RPN, the fitness between b and p can be calculated as follows:

$$fit(b, p) = 1 - (\mathcal{D}(b, p) + \gamma \cdot \mathcal{C}(b, p)), \quad (15)$$

$fit(b, p)$ measures the similarity between the annotated bounding box b and the region proposal p . The center of the noisy annotated bounding boxes are usually still near the objects, even if their shapes are severely distorted. Therefore, we expect to match the region proposals and annotated bounding boxes on the basis of the distances of their centers. $\mathcal{D}(b, p)$ measures the relative distance between b and p . Let (x_c, y_c) , W , and H denote the coordinates of the center point, width, and height of the bounding box, respectively. The calculation of $\mathcal{D}(b, p)$ is formulated as follows:

$$\mathcal{D}(b, p) = \frac{\sqrt{(x_c^b - x_c^p)^2 + (y_c^b - y_c^p)^2}}{W_p + H_p}. \quad (16)$$

Matching based on distance makes it more difficult to find the corresponding proposals for large objects than for small objects. Therefore, we divide the Euclidean distance by the size of the object to mitigate the difference between objects of different sizes. For simplicity, the size of the objects are approximately represented by the half circumference of the proposals. $\mathcal{C}(b, p)$ denotes the size difference between bounding box b and proposal p as follows:

$$\mathcal{C}(b, p) = \left| \frac{W_b + H_b}{W_p + H_p} - 1 \right|. \quad (17)$$

$\mathcal{C}(b, p)$ is designed to prevent proposals from being assigned to incorrect overlapping bounding boxes. We expect this term to work only when the differences in the calculated $\mathcal{D}(b, p)$ are not significant. Therefore, the parameter γ is set to 0.1 to prevent this term from interfering with the distance determination in the usual case.

For the proposals P_h that have Top-100 objectness scores generated by the RPN, where $P_h \subset P$, we calculate the fitness between the proposals and the annotated bounding box b , and maintain the proposals with a large fitness as follows:

$$B_{candidate} = \{ p \mid fit(b, p) > 0.9, p \in P_h \}. \quad (18)$$

Subsequently, a few proposals in $B_{candidate}$ that have Top-10 objectness scores are maintained and the others are discarded to reduce redundancy. We use the proposal that has the highest objectness score p^* to perform class-agnostic refinement on each annotated bounding box b as follows:

$$b^* = \alpha \cdot p^* + (1 - \alpha) \cdot b \quad (19)$$

After all, the center-matching correction initially modifies the annotated bounding box b to b^* , which is used for classification noise judgment and label refinement. At the same time, other candidate regions $B_{candidate}$ are utilized for further localization refinement, as described in Section 3.1.



Figure 4: Visualization of the bounding box refinement during training. The red boxes denote the noisy annotations in the training dataset with the noisy parameters of $N_{Label} = 40\%$ and $N_{BBox} = 40\%$. The correction results of the first step are shown in yellow boxes, whereas the final correction results are shown in white boxes.

3.3 Cross-Iteration Noise Judgment

We noticed that the correctness of the classification labels is a dominant factor in the magnitude of the classification loss, as shown in Figure 2. This phenomenon motivates us to distinguish the incorrect labels from the correct labels, which allow us to refine the noisy labels and maintain the clean class labels. Thus, we propose a noise judgment method called cross-iteration noise judgment (CINJ), which uses an adaptive threshold to judge the classification labels of the current training iteration, as shown in Figure 3.

We maintain the classification loss $\mathcal{L}(b^*, y)$ for all the annotated bounding boxes in the pre-forward stage of each iteration. The CINJ adaptively discriminates the credibility of the current class labels by comparing the relative magnitude of the classification losses in the current iteration with the loss values maintained in the previous iterations.

Let $\mathcal{L}_i(b^*, y)$ denote the classification loss corresponding to an annotation $\{b, y\}$ in the i -th iteration. The CINJ mechanism maintains a queue $Q = [\mathcal{L}_{i-N}, \mathcal{L}_{i-N+1}, \dots, \mathcal{L}_{i-1}]$ of loss values, which contains N loss values in the latest dozens of iterations. For simplicity, we assume that there is only one image in each iteration. Each $\mathcal{L}_j \in Q$ uniquely corresponds to an annotation $\{b_j, y_j\}$ in the training dataset. Let $Q^* = [\mathcal{L}_1^*, \mathcal{L}_2^*, \dots, \mathcal{L}_N^*]$ denote the result of sorting Q in the ascending order. When $R\%$ of the classification labels in the dataset are incorrect and Q^* is sufficiently long, we can assume that approximately $R\%$ of the loss values in Q^* correspond to the incorrect labels. Because the noisy labels cause larger loss values than those of correct labels, as shown in Figure 2, the loss values corresponding to the $R\%$ incorrect labels are clustered at the back end of Q^* . Therefore, we use an acceptance rate $r = 1 - R\%$ to compute an adaptive threshold $\mathcal{T} = Q_{[r \cdot N]}^*$, and labels with loss values greater than this threshold \mathcal{T} are considered as noisy labels. As queue Q is updated after each iteration, \mathcal{T} dynamically changes.

4 Experiments

4.1 Dataset

We manually introduce different levels of symmetric classification label noise and uniform localization noise into the clean PASCAL VOC [5] and MS-COCO [13] datasets. We use two parameters, N_{Label} and N_{BBox} , to control the level of symmetric class label noise and uniform bounding box localization noise, respectively. In particular, for all artificial noisy datasets, N_{Label} of the classification labels are randomly chosen and replaced by other categories. The horizontal coordinates of all bounding boxes are perturbed by a number of pixels uniformly

Table 1: Comparison with state-of-the-art methods on the noisy PASCAL VOC dataset 07+12. The results (mAP@50) are evaluated on the clean test split of the VOC 2007 dataset. Bounding box refinement (BR) includes the initial bounding box correction in the center-matching correction (Eq. 19) and the final correction (Eq. 14). Label refinement (LR) includes the CINJ (Section 3.3) and the class label refinement (Eq. 12).

N_{BBox}	0%			20%				40%			
N_{Label}	20%	40%	60%	0%	20%	40%	60%	0%	20%	40%	60%
Baseline [1, 8]	74.2	70.3	65.8	75.5	73.1	69.6	63.2	58.9	55.9	53.7	47.9
Co-teaching [12, 19]	76.5	74.1	69.9	75.6	73.2	69.7	65.1	60.6	59.7	55.8	50.4
NOTE-RCNN [6, 12]	76.7	74.9	69.9	76.0	73.7	70.1	65.8	63.4	61.5	57.8	53.7
CA-BBC [12]	79.1	77.7	74.1	77.9	76.7	74.8	71.9	71.9	70.6	69.1	64.5
Baseline + BR	–	–	–	78.4	74.3	71.2	66.0	73.4	70.2	65.9	60.4
Baseline + LR	79.9	79.3	76.6	–	76.9	75.5	71.6	–	60.3	58.6	55.0
Proposal	79.9	79.3	76.6	78.4	77.2	76.4	72.8	73.4	71.7	70.3	66.6

Table 2: Comparison with state-of-the-art methods on the MS-COCO dataset. The results are evaluated on the clean validation split of the MS-COCO 2017.

Method	$N_l = 20\%, N_b = 20\%$		$N_l = 40\%, N_b = 40\%$	
	mAP@0.5	mAP@[.5,.95]	mAP@0.5	mAP@[.5,.95]
Baseline [1, 8]	47.9	23.9	29.7	10.3
Co-teaching [12, 19]	49.7	24.6	35.9	14.6
NOTE-RCNN [6, 12]	50.4	25.1	38.5	15.2
CA-BBC [12]	53.5	27.7	47.4	21.2
Baseline + BR	52.9	28.1	44.6	19.3
Baseline + LR	54.0	29.7	40.1	17.2
Proposal	54.3	30.2	47.5	21.7

drawn from $[-wN_{BBox}, +wN_{BBox}]$, and from $[-hN_{BBox}, +hN_{BBox}]$ for the vertical coordinates, where w and h denote the width and height of the corresponding box, respectively. By combining these two parameters, we obtain artificial noisy datasets with different levels of noise, where the classification noise and the localization noise are entangled together.

4.2 Implementation

We construct the proposed framework based on a Faster R-CNN with ResNet-50 [8] and FPN [12] as the feature extractor. All the models trained on the noisy PASCAL VOC datasets use an SGD optimizer with a learning rate of 0.01, a momentum of 0.9, and a weight decay of 1×10^{-4} . For the experiments on the MS-COCO dataset, the learning rate is changed to 5×10^{-3} , and the others remain unchanged. Our implementation is based on the mmdetection toolbox [2] and the training schedule basically follow the default training configuration provided by the toolbox. When training the Faster R-CNN using our proposal, we start the judgment and correction after the first epoch to obtain stable results. The length N for queue \mathcal{Q} is set to 128, and the acceptance rate r is set close to the percentage of correct class labels in the training dataset for all the results reported in Table 1. α is set to 0.2 and 0.4 for datasets with noise rates $N_{BBox} = 20\%$ and 40% , respectively. The quantitative sensitive analysis on

Table 3: Ablation experiments: mAP of the noisy localization annotation refinement.

Bounding Box Refinement (BR)		N_{BBox}	
Center-Matching (Eq. 19)	Refinement (Eq. 14)	20%	40%
		75.5	58.9
✓		77.3	68.4
✓	✓	78.4	73.4

Table 4: Ablation experiments: mAP of the noisy label refinement.

Label Refinement (LR)		N_{Label}		
CINJ	Refinement (Eq. 12)	20%	40%	60%
		74.2	70.3	65.8
✓		78.3	74.4	70.0
	✓	78.9	76.9	71.4
✓	✓	79.9	79.3	76.6

the hyperparameters are shown in the Supplementary Material.

4.3 Experiment Results

We apply our proposal on the PASCAL VOC 07+12 dataset with different levels of artificial noise, and the results are shown in Table 1. Our proposed method outperforms state-of-the-art methods by a significant margin on all noisy datasets with different noise levels. Figure 4 shows the results of the bounding box correction while training the detector on a dataset with $N_{Label} = 40\%$, $N_{BBox} = 40\%$. Our proposed method can significantly reduce localization noise even if the classification labels are unreliable. Furthermore, center-matching correction can correct the bounding boxes with high-level noise. The evaluation of the corrected training dataset is provided in Supplementary Material.

Table 2 shows the experimental results on the noisy MS-COCO dataset. Our proposal significantly outperforms state-of-the-art methods under the supervision of noisy datasets. The experimental results confirm that our proposal is effective in alleviating the negative impact caused by noisy annotations.

4.4 Ablation Study

In this section, we conduct ablation experiments on noisy Pascal VOC datasets to examine the effectiveness of each part of our proposal. As each of our modules explicitly targets a particular type of noise, we separately examine our methods on the datasets with the corresponding noise.

For the bounding box noise, the bounding box refinement (BR) involves two steps. The first step is performed by the center-matching correction and initially corrects the bounding boxes regardless of the class labels. The second step uses the class labels to choose candidates with top-2 confidence in the corresponding category and subsequently correct the corresponding bounding box using Eq. 14. Table 3 shows the effect of the two steps, both of which contribute to the performance of our proposal.

Table 5: mAP on datasets with pair classification label noise.

Method	Pair-Noise Rate		
	20%	30%	40%
Baseline	76.3	71.3	61.7
Baseline + LR	79.8	78.6	69.8

Table 6: mAP on datasets with Gaussian localization annotation noise.

Method	σ^2		
	0.2	0.25	0.3
Baseline	71.5	66.4	59.8
Baseline + BR	74.9	73.1	71.0

For the classification noise, our label refinement (LR) first performs the CINJ to judge the correctness of the labels and subsequently replace those labels judged to be noisy with pseudo-labels. Table 4 shows the effectiveness of both steps when dealing with noisy class labels. The results in the first row are obtained using the vanilla Faster R-CNN. The second row shows the results when only the CINJ is performed, and all labels judged to be noisy are discarded without modification. The third row shows the results when all class labels are replaced by high-confidence pseudo-labels without noise judgment. When the two steps work together, the proposed method achieves the best performance.

4.5 Evaluation on Other Types of Noise

We further construct more challenging noisy datasets to examine the effectiveness of our proposal. For the classification labels, we introduce pair noise, which randomly select $R\%$ samples from each class y and change their label indices to $y + 1$. For the localization annotations, we introduce Gaussian noise. The horizontal coordinates of all bounding boxes are perturbed by δw pixels and by δh for the vertical coordinates, where w and h denote the width and height of the corresponding box, respectively, and δ is drawn from a Gaussian distribution $\delta \sim N(0, \sigma^2)$. The precise definitions of the noises are provided in the Supplementary Material.

The results are shown in Tables 5 and 6. The experimental results confirm that our proposal can significantly alleviate the negative impact caused by noisy annotations, despite the type of noise.

5 Conclusion

In this study, we propose a novel problem setting of training object detector on datasets with entangled classification noise and localization annotation noise. We proposed an effective correction method based on center matching for localization annotation noise to match extremely noisy bounding boxes to the proposals generated by the RPN. We observed the behavior of the object detector under the supervision of noisy class labels and noticed that noisy labels led to significantly high values of classification loss. On the basis of this finding, we proposed a noise-judging method, called cross-iteration noise judgment, to identify incorrect training labels. We arrange the steps with an additional pre-forward stage to decouple the entangled noise. Our proposal significantly outperforms the baseline by reducing the noise level of the training dataset. Each proposed step is experimentally shown to be effective in alleviating the negative impact caused by incorrect annotations.

ACKNOWLEDGEMENTS

This work is partially supported by JST CREST 90249938 and JSPS KAKENHI 18H03254.

References

- [1] Devansh Arpit, Stanisław Jastrzębski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *ICML*, 2017.
- [2] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, et al. Mmdetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019.
- [3] Etienne David, Simon Madec, Pouria Sadeghi-Tehran, Helge Aasen, Bangyou Zheng, Shouyang Liu, Norbert Kirchgessner, Goro Ishikawa, Koichi Nagasawa, Minhajul A Badhon, et al. Global wheat head detection (gwhd) dataset: a large and diverse dataset of high-resolution rgb-labelled images to develop and benchmark wheat head detection methods. *Plant Phenomics*, 2020.
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 2009.
- [5] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010.
- [6] Jiyang Gao, Jiang Wang, Shengyang Dai, Li-Jia Li, and Ram Nevatia. Note-rcnn: Noise tolerant ensemble rcnn for semi-supervised object detection. In *ICCV*, 2019.
- [7] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [9] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017.
- [10] Ishan Jindal, Matthew Nokleby, and Xuewen Chen. Learning deep networks from noisy labels with dropout regularization. In *ICDM*, 2016.
- [11] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4. *International Journal of Computer Vision*, pages 1–26, 2020.
- [12] Junnan Li, Caiming Xiong, Richard Socher, and Steven Hoi. Towards noise-resistant object detection with noisy annotations. *arXiv preprint arXiv:2003.01285*, 2020.
- [13] 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 *ECCV*, 2014.
- [14] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017.

- [15] Jiafeng Mao, Qing Yu, and Kiyoharu Aizawa. Noisy localization annotation refinement for object detection. In *ICIP*, 2020.
- [16] Dim P Papadopoulos, Jasper RR Uijlings, Frank Keller, and Vittorio Ferrari. Extreme clicking for efficient object annotation. In *ICCV*, 2017.
- [17] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *CVPR*, 2017.
- [18] Shaoqing Ren, Kaiming He, Ross B Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015.
- [19] Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *NeurIPS*, 2018.
- [20] Sainbayar Sukhbaatar, Joan Bruna, Manohar Paluri, Lubomir Bourdev, and Rob Fergus. Training convolutional networks with noisy labels. *arXiv preprint arXiv:1406.2080*, 2014.
- [21] Daiki Tanaka, Daiki Ikami, Toshihiko Yamasaki, and Kiyoharu Aizawa. Joint optimization framework for learning with noisy labels. In *CVPR*, 2018.
- [22] Arash Vahdat. Toward robustness against label noise in training deep discriminative neural networks. In *NeurIPS*, 2017.
- [23] Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross entropy for robust learning with noisy labels. In *ICCV*, 2019.
- [24] Hongxin Wei, Lei Feng, Xiangyu Chen, and Bo An. Combating noisy labels by agreement: A joint training method with co-regularization. In *CVPR*, 2020.
- [25] Tong Xiao, Tian Xia, Yi Yang, Chang Huang, and Xiaogang Wang. Learning from massive noisy labeled data for image classification. In *CVPR*, 2015.
- [26] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor Tsang, and Masashi Sugiyama. How does disagreement help generalization against label corruption? In *ICML*, 2019.
- [27] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.
- [28] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.