

Figure 4: (a) Test performance with different sharpening temperature T for soft labels (Eq. 5). T controls the entropy of the label distribution. (b) Divergence between the two detection heads measured by their classifier discrepancy (Eq. 1), averaged over 1000 iterations (mini-batches). The two heads can stay sufficiently diverged due to different initialization and different proposal sampling during training. Training data for both figures is PASCAL VOC with 40% label noise and 40% bbox noise.

Ablation on soft label temperature. In Figure 4(a), we vary the sharpening temperature T , which controls the entropy of the soft label distribution (Eq. 5). A lower temperature encourages the model to produce more confident predictions, whereas a higher temperature emphasizes more on the uncertainty of the label. $T \rightarrow 0$ produces one-hot labels, which corresponds to the hard bootstrapping method (Reed et al., 2015), whereas $T = 1$ corresponds to the soft bootstrapping method. We observe that acceptable values of T ranges around 0.2 ~ 0.6. We use $T = 0.4$ in all experiments.

Dual-head divergence analysis.

In Figure 4(b), we show that the two detection heads can stay sufficiently diverged during the entire training process, due to different parameter initialization and different region proposal sampling. We measure divergence using classifier discrepancy (Eq. 1). Keeping the two detection heads diverged is important for several reasons. First, the effectiveness of the proposed CA-BBC requires disagreement between the two classifiers. Second, diverged detection heads can filter different kinds of noise, which results in more effective dual-head noise correction. Furthermore, diverged dual heads can lead to better ensemble performance during test.

Noise correction examples.

In Figure 5 we show example images with annotations before and after the proposed two-step noise correction method. Noisy GT labels and bounding boxes are shown in red whereas the corrected labels and boxes are shown in green. The proposed method can effectively clean the annotation noise by 1) regressing the bounding boxes to more accurate locations and 2) assigning high soft label values to the correct classes.

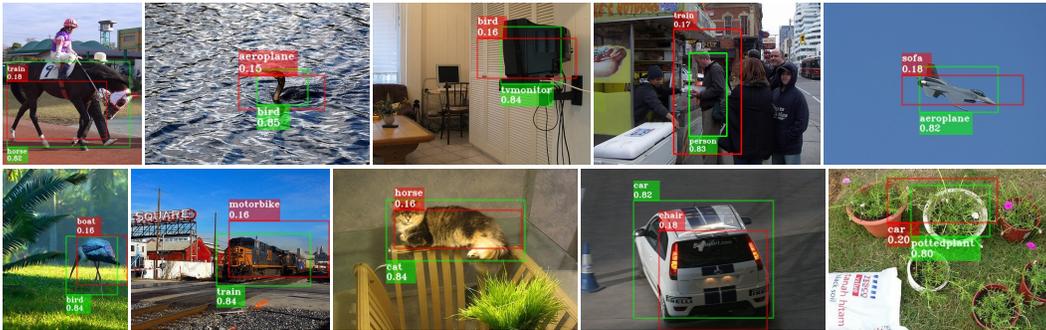


Figure 5: Examples of dual-head noise correction on PASCAL VOC dataset with 40% label noise and 40% bounding box noise (best viewed in color with zoom in). Noisy GT labels and GT bounding boxes are shown in red, whereas the corrected labels and bounding boxes are shown in green. Numbers under the class names indicate corresponding values in the soft labels.