
Towards Robust Object Detection: Bayesian RetinaNet for Homoscedastic Aleatoric Uncertainty Modeling

Abstract

According to recent studies, commonly used computer vision datasets contain about 4% of label errors. For example, the COCO dataset is known for its high level of noise in data labels, which limits its use for training robust neural deep architectures in a real-world scenario.

To model such a noise, in this paper we have proposed the homoscedastic aleatoric uncertainty estimation, and present a series of novel loss functions to address the problem of image object detection at scale.

Specifically, the proposed functions are based on Bayesian inference and we have incorporated them into the common community-adopted object detection deep learning architecture RetinaNet.

We have also shown that modeling of homoscedastic aleatoric uncertainty using our novel functions allows to increase the model interpretability and to improve the object detection performance being evaluated on the COCO dataset.

1 INTRODUCTION

Usually, training a predictive algorithm involves training a machine learning model on a labeled dataset from a scratch or using this dataset to fine-tune a model previously pre-trained on a large publicly available dataset such as ImageNet or MS COCO. However, a recent study [Curtis G. Northcutt, 2021] concluded that commonly used open datasets for computer vision tasks contain about 4% of errors in image labels. The MS COCO dataset for detection models benchmarking is also known for its noisy labels of both object classes and bounding boxes [Khetan et al., 2017, Vahdat, 2017]. At the same time, popular cross-entropy loss is considered to be sensitive to noisy labeling [Feng et al., 2020]. Moreover, the deeper the model, the more it adapts

to these labeling errors. This negatively affects not only the integrity of the contests on the corresponding datasets, but also a real-world scale, since these datasets are often used for model pre-training to solve various problems.

One way to account for the label errors is to estimate the aleatoric uncertainty, which reflects the noise level in the training data and can be used at the inference time [Hüllermeier and Waegeman, 2021]. The aleatoric uncertainty is divided into homoscedastic, i.e. constant for the data distribution in a particular task, and heteroscedastic, i.e. different for each data object [Kendall and Gal, 2017]. Despite the estimation of heteroscedastic uncertainty is more useful for computer vision problems in general [Kendall and Gal, 2017], its modeling requires changes in the neural network (NN) architecture. Moreover, its application in practice requires developing tools to postprocess prediction for a particular object with this uncertainty.

At the same time, the modeling of homoscedastic aleatoric uncertainty can be performed based on the modification of the loss functions rather than the architecture, which is less time-consuming. In addition, homoscedastic aleatoric modeling even improves the accuracy of solving the computer vision problems [Kendall et al., 2018]. Researchers Kendall et al. [2018] consider the application of modeling this type of uncertainty for multi-task NN architecture, solving semantic, instance segmentation, and depth regression problems. Quantification of aleatoric uncertainty can greatly increase model performance in the detection problem [Feng et al., 2019, Meyer et al., 2019]

Recently, Bayesian deep learning has been widely used in object detection [Bendale and Boulton, 2016, Harakeh et al., 2020, Kraus and Dietmayer, 2019, Miller et al., 2018, 2019, 2021, Postels et al., 2019]. However, all these works focus on epistemic uncertainty.

Fewer number of papers are devoted to the aleatoric uncertainty estimation [Kraus and Dietmayer, 2019, Le et al., 2018] including those on 3d object detection [Feng et al., 2018, 2019, Meyer et al., 2019] and one-stage detec-

tor [Kraus and Dietmayer, 2019, Le et al., 2018]. However, existing works do not study the application of homoscedastic aleatoric uncertainty modeling for the detection problem, although this can help isolate noise from data and improve model robustness. Moreover, as the detection is the multi-task problem (i.e. includes localization and classification tasks), the modeling can be performed without changes in the neural network architecture, using tools, developed by Kendall et al. [2018].

Being inspired by this, we aimed to answer the following research questions:

RQ1: Can homoscedastic aleatoric uncertainty modeling improve the detection accuracy based on deep neural networks?

RQ2: Can Bayesian approximation be effectively applied to modeling homoscedastic aleatoric uncertainty for existing detection models?

In order to answer them, we propose novel loss functions, whose optimization is equivalent to modeling homoscedastic aleatoric uncertainty for the joint localization and classification tasks. The paper contributions are the following:

1. A new loss function for the classification task for modeling the aleatoric uncertainty called **Bayesian Focal Loss**.
2. A new loss function for the localization task for modeling the aleatoric uncertainty called **Bayesian Smooth L_1 Loss**.

The proposed loss functions for modeling the homoscedastic aleatoric uncertainty can be applied to any NN detectors, which use cross-entropy or Focal loss and L_1 or Smooth L_1 loss, without changing their architecture and training pipeline. The uncertainty modeling can make existing detectors robust to noise in data labels and can improve detection accuracy as well.

2 RELATED WORK

2.1 BAYESIAN DEEP LEARNING FOR COMPUTER VISION

Recently, Kendall et al. [2018] suggested a tool for modeling homoscedastic aleatoric uncertainty to weigh multi-task losses. They considered three computer vision tasks: semantic segmentation, instance segmentation, and depth regression. The modeling required building a probabilistic model for both classification and regression tasks.

For the regression task, they defined a probabilistic model with a Gaussian likelihood, where the mean is given by the model output $f^W(x)$ with weights W on input x :

$$p(y|f^W(x)) = N(f^W(x), \sigma^2) \quad (1)$$

and the variance is given as an observation noise scalar σ , which captures homoscedastic aleatoric uncertainty.

Interpreting this Gaussian log likelihood maximization as objective, they obtained the modification of L_2 loss:

$$BL_2 = -\frac{1}{2\sigma^2} \|y - f^W(x)\|^2 - \log \sigma, \quad (2)$$

with BL_2 being the Bayesian L_2 loss. It is then maximized with respect to weights W and noise scalar σ .

For the classification task the likelihood appeared to be less trivial. Assume the model output $f^W(x)$ is scaled by $1/\sigma^2$ and then squashed through the Softmax activation function. Then, the likelihood is the following:

$$p(y|f^W(x), \sigma) = \text{Softmax} \left(\frac{1}{\sigma^2} f^W(x) \right), \quad (3)$$

which can be interpreted as a Boltzmann distribution with temperature σ .

The log likelihood is defined as:

$$\begin{aligned} \log p(y = c|f^W(x), \sigma) &= \frac{1}{\sigma^2} f_c^W(x) - \\ &- \log \sum_{c'} \exp \left(\frac{1}{\sigma^2} f_{c'}^W(x) \right), \end{aligned} \quad (4)$$

with $f_c^W(x)$ the element of $f^W(x)$ vector for a particular class c .

Using maximum likelihood inference for the multi-task neural network with output y_1 for the regression task and y_2 for the classification task the following minimization objective can be obtained:

$$\begin{aligned} L(W, \sigma_1, \sigma_2) &= \frac{1}{2\sigma_1^2} L_1(W) + \log \sigma_1 + \\ &+ \frac{1}{\sigma_2^2} L_2(W) + \log \sigma_2, \end{aligned} \quad (5)$$

where $L_1(W) = \|y_1 - f^W(x)\|^2$ is the Euclidean loss for y_1 and $L_2(W) = -\log(\text{Softmax}(y_2, f^W(x)))$ is the cross-entropy loss for y_2 . This loss is optimised with respect to W as well as σ_1 and σ_2 .

The main difficulty with the $L_2(W)$ loss is to release x in $f^W(x)$ from scaling factor $1/\sigma^2$. To achieve this, Kendall et al. [2018] performed the following: subtracted and added $\frac{1}{\sigma^2} \log(\sum_{c'} \exp(f_{c'}^W(x)))$ to Eq. 4, then used a simplifying assumption

$$\left[\sum_{c'} \exp(f_{c'}^W(x)) \right]^{\frac{1}{\sigma^2}} \approx \frac{1}{\sigma} \sum_{c'} \exp \left(\frac{1}{\sigma^2} f_{c'}^W(x) \right),$$

which becomes an equality when $\sigma \rightarrow 1$.

2.2 RETINANET DETECTOR

RetinaNet [Lin et al., 2017b] is a one-stage anchor-based neural network for object detection. This architecture is most famous by the proposed classification loss function, referred as Focal Loss. RetinaNet consists of four subnetworks:

- Backbone is a basic convolutional network that extracts features from the input image. Traditionally, the state-of-the-art networks are used as backbones, such as ResNet [He et al., 2016], VGG [Simonyan and Zisserman, 2015], EfficientNet [Tan and Le, 2019].
- Feature Pyramid Network (FPN) is a “neck” convolutional neural network proposed by Lin et al. [2017a]. It combines feature maps from different layers of the backbone network in a top-down pathway using lateral connection. This allows to solve a task (classification or regression) at different image resolutions and semantic scales.
- Localization subnetwork is a “head” subnetwork that extracts information from the FPN about the coordinates of objects in the image, solving the regression task. It trains based on the Smooth L_1 loss proposed by Girshick [2015].
- Classification subnetwork is a “head” subnetwork that extracts information about object classes from the FPN, solving the classification task. It trains based on the Focal loss.

For the bounding boxes regression, RetinaNet uses Smooth L_1 loss. This is a combination of L_1 and L_2 loss functions, which was initially inspired by Huber [1992]. Its formula is

$$\text{Smooth}_{L_1}(x) = \begin{cases} \frac{\beta^2}{2} \cdot \epsilon^2, & \text{if } \epsilon < \frac{1}{\beta^2} \\ \epsilon - \frac{1}{2\beta^2}, & \text{otherwise} \end{cases}, \quad (6)$$

with $1/\beta^2$ the threshold for switching from the L_1 to the L_2 loss function, and $\epsilon = \|y - f^W(x)\|$ with x the network input, its output $f^W(x)$, and y the ground truth coordinate of the object bounding box. The main difference from the L_2 loss function is that addition of L_1 case helps avoid over-penalizing outliers.

For the classification, Lin et al. [2017b] introduced Focal loss. Focal loss is proven to penalize the network better than the cross-entropy loss [Lin et al., 2017b] on hard negative examples. Its formula is

$$FL(p_t) = -\log p_t \cdot (1 - p_t)^\gamma, \quad (7)$$

$$\text{where } p_t = \begin{cases} p & y = 1 \\ 1 - p & \text{otherwise} \end{cases},$$

with $p = \text{Sigmoid}(f^W(x))$, y the ground truth class label of an object. The main difference of Focal loss from the

cross-entropy loss is the modulation factor $\gamma \in (0, +\infty)$ introduced to handle the problem of class imbalance, which is typical for object detection, since an object of interest usually occupies relatively little space in the image. Thus, Focal loss results in higher gradient values for higher error values and vice versa. This forces the network to focus on hard negative examples better, which are the objects of interest. The generalized RetinaNet loss function can then be written as $L = \alpha L_{class} + L_{reg}$, with L_{class} the classification Focal loss function, L_{reg} the regression (localization) Smooth L_1 loss function, α the balancing coefficient that adjusts the impact of the L_{class} term.

Although RetinaNet loss functions are quite effective, they do not allow to capture homoscedastic aleatoric uncertainty making RetinaNet sensitive to the noisy data. To overcome this issue, we propose the novel Focal and Smooth L_1 loss functions, which are able to model homoscedastic aleatoric uncertainty. We call our neural network, that utilizes them, *Bayesian RetinaNet*.

3 BAYESIAN RETINANET

In this section, we introduce the novel loss functions with homoscedastic uncertainty based on maximum likelihood estimation.

Let $f^W(x)$ denote the output of a neural network with weights W on input x and ϵ be the error that is the norm of difference between the ground truth value and our prediction:

$$\epsilon = \|y - f^W(x)\|.$$

3.1 BAYESIAN SMOOTH L_1 LOSS FOR HOMOSCEDASTIC ALEATORIC UNCERTAINTY

First, we introduce the novel likelihood for the localization task, which is to predict object coordinates. As localization is the regression task, we adopt the likelihood from Section 2.1 for the Smooth L_1 loss and define our likelihood as the combination of Gaussian and Laplace likelihoods:

$$p(y|f^W(x), \sigma, \alpha) = \begin{cases} p_G(y|f^W(x), \sigma), & \text{if } \epsilon < \frac{1}{\beta^2} \\ p_L(y|f^W(x), \alpha), & \text{otherwise} \end{cases}, \quad (8)$$

where p_G is Gaussian likelihood, p_L is Laplace likelihood with observation noise scalars σ and α , respectively.

As in maximum likelihood inference, here we maximise the log likelihood of the model. Thus, following the likelihood for regression in the case of L_2 loss [Kendall et al., 2018],

for Smooth L_1 it can be written as

$$\log p(y|f^W(x), \sigma, \alpha) \propto \begin{cases} -\frac{\epsilon^2}{2\sigma^2} - \log \sigma, & \text{if } \epsilon < \frac{1}{\beta^2} \\ -\alpha\epsilon + \log \alpha & \text{otherwise} \end{cases}, \quad (9)$$

where L_2 corresponds to Gaussian likelihood p_G , L_1 to Laplace likelihood p_L .

This leads to the following minimization objective $L(W, \sigma, \alpha)$:

$$\begin{aligned} L(W, \sigma, \alpha) &= -\log p(y|f^W(x), \sigma) \propto \\ &\propto \begin{cases} \frac{\epsilon^2}{2\sigma^2} + \log \sigma, & \text{if } \epsilon < \frac{1}{\beta^2} \\ \alpha\epsilon - \log \alpha & \text{otherwise} \end{cases}, \quad (10) \\ &= \begin{cases} \frac{1}{2\sigma^2} L_2(W) + \log \sigma, & \text{if } \epsilon < \frac{1}{\beta^2} \\ \alpha L_1(W) - \log \alpha & \text{otherwise} \end{cases} \end{aligned}$$

where we write $L_1(W) = \epsilon$ for the L_1 loss of y , write $L_2(W) = \epsilon^2$ for Euclidean loss of y .

The likelihood in Eq. 10 has two variances corresponding to L_1 and L_2 . However, this is inconvenient in practice because the ground truth bounding box coordinates are unknown in the real world model inference. Thus, to find the dependency between α and σ and also save the property of the likelihood, we solve the equation of density function of our likelihood:

$$\int_{-\infty}^{\infty} p(y|f^W(x), \sigma) dt = 1. \quad (11)$$

From this equation, we obtain the following dependency between variances:

$$\alpha = -\beta^2 \log \tau, \quad (12)$$

where $\tau = 1 - \text{erf}\left(\frac{1}{\beta^2 \sqrt{2\sigma^2}}\right)$ with erf the Gauss error function [Abramowitz et al., 1988].

Whether we place Eq. 12 into Eq. 10, we obtain Bayesian Smooth L_1 Loss:

$$\begin{aligned} BSmooth_{L1}(\epsilon) &= \\ &= \begin{cases} \frac{\epsilon^2}{2\sigma^2} + \log \sigma & \text{if } \epsilon < \frac{1}{\beta^2} \\ -\beta^2 \epsilon \log \tau - \log(-\beta^2 \log \tau) & \text{otherwise} \end{cases}. \quad (13) \end{aligned}$$

The first and second cases of Eq. 13 are not equal, when ϵ equals to $1/\beta^2$. The second case requires a small correction. To solve this issue, we smooth this function and obtain the

following loss function:

$$\begin{aligned} BSmooth_{L1}(\epsilon) &= \\ &= \begin{cases} \frac{\epsilon^2}{2\sigma^2} + \log \sigma & \text{if } \epsilon < \frac{1}{\beta^2} \\ -\beta^2 \epsilon \log \tau + \log \tau + \frac{1}{2\sigma^2 \beta^4} + \log \sigma & \text{otherwise} \end{cases}. \quad (14) \end{aligned}$$

Following Kendall et al. [2018] in experiments we train the network to predict the log variance, $s := \log \sigma^2$, which is more numerically stable than regressing the variance σ^2 directly to avoid division by zero. The proposed Bayesian Smooth L_1 loss function plot is presented in Fig. 1 in comparison with the original Smooth L_1 loss. The proposed loss function penalizes the neural network better than the original Smooth L_1 loss: for less noisy data, it penalizes the neural network more for large prediction errors. For noisier data, it penalizes the neural network more uniformly, less “trusting” the data labels.




Figure 1: The proposed Bayesian Smooth L_1 loss function for different estimates of aleatoric uncertainty. At $\sigma = 0.5$, $\sigma = 2.0$, compared to the original Smooth L_1 loss (blue line).

3.2 BAYESIAN FOCAL LOSS FOR HOMOSCEDASTIC ALEATORIC UNCERTAINTY

Now we introduce the novel likelihood function for the classification task, which is the modified Focal loss. In RetinaNet, the classification activation function is logistic, which is more convenient for datasets with non-mutually exclusive classes. Thus, for the classification task, the likelihood can be defined as:

$$p(y|f^W(x), \sigma) = \text{Sigmoid}\left(\frac{1}{\sigma^2} f^W(x)\right) \quad (15)$$

with a positive noise scalar σ , which reflects homoscedastic uncertainty. This likelihood can also be interpreted as the Boltzmann distribution where the input is scaled by σ^2 . We aim to maximise the likelihood. For the classification task it can be effectively done using Focal loss, which behaves the same way as log likelihood. Focal loss likelihood can be defined as:

$$BFL(p(y|f^W(x), \sigma)) = - \left(\frac{1}{\sigma} (1 - p_t)^{\sigma-2} \right)^\gamma \times \left(\frac{1}{\sigma^2} \log p_t - \log \sigma \right) \quad (16)$$

where BFL is Bayesian Focal loss and

$$p_t = \begin{cases} p & y = 1 \\ 1 - p & \text{otherwise} \end{cases}.$$

To obtain BFL from the original Focal loss, the main issue is to release $f^W(x)$ in logistic function from the scaling factor $1/\sigma^2$. To solve this issue and obtain the new form of likelihood, the following transitions are used: subtraction and addition of $\frac{1}{\sigma^2} \log(1 + \exp(f^W(x)))$ term; simplifying assumptions

$$\begin{aligned} \frac{1}{\sigma} \left[1 + \exp \left(\frac{1}{\sigma^2} f^W(x) \right) \right] &\approx \\ &\approx [1 + \exp(f^W(x))] \frac{1}{\sigma^2}, \end{aligned} \quad (17)$$

and

$$\begin{aligned} \frac{1}{\sigma} \left[1 + \exp \left(-\frac{1}{\sigma^2} f^W(x) \right) \right] &\approx \\ &\approx [1 + \exp(-f^W(x))] \frac{1}{\sigma^2}, \end{aligned} \quad (18)$$

which become an equality when $\sigma \rightarrow 1$.

Bayesian Focal loss is equal to the original Focal Loss, when $\sigma = 1$, or $p(y|f^W(x)) = \text{Sigmoid}(f^W(x))$.

In our experiments, we train the network to predict the log variance, $s := \log \sigma^2$ to preserve the numerical stability. The proposed Bayesian Focal loss function plot is presented in Fig. 2 in comparison with the original Focal loss. Our loss function penalizes the neural network better than the original Focal loss ($\sigma=1$ in the figure): for less noisy data, it penalizes the neural network less for well-classified examples and more for large prediction errors. For noisier data, it penalizes the neural network more uniformly, less “trusting” the data labels.

3.3 MULTI-TASK LIKELIHOOD FOR BAYESIAN RETINANET

For the multi-task Bayesian RetinaNet with output y_1 for a localization task and y_2 for a classification task, we obtain




Figure 2: The proposed Bayesian Focal loss function for different estimates of aleatoric uncertainty at $\sigma = 0.7$, $\sigma = 1.0$, $\sigma = 2.0$. At $\sigma = 1.0$ function is equal to the original Focal loss.

the following minimization objective:

$$L(W, \sigma_1, \sigma_2) = BSmooth_{L1}(f^W(x), \sigma_1) + \alpha \cdot BFL(f^W(x), \sigma_2) \quad (19)$$

where $BSmooth_{L1}(f^W(x), \sigma_1)$ is the Bayesian Smooth L_1 loss for y_1 , $BFL(f^W(x), \sigma_2)$ is the Bayesian Focal loss for y_2 , α is the balancing coefficient that adjusts the impact of the $BFL(f^W(x), \sigma_2)$ term. This multi-task loss is optimised with respect to W as well as σ_1 and σ_2 .

Unlike in [Kendall et al., 2018], our multi-task objective does not allow to weigh losses by tuning σ_1 and σ_2 . It only allows to learn these noise scalars and thus capture homoscedastic uncertainty.

4 EXPERIMENTS AND RESULTS

In our experiments, as the backbone for RetinaNet and Bayesian RetinaNet we used only ResNet-50 due to the memory limitation. The architecture of Bayesian RetinaNet was the same as the original RetinaNet model. The changes were made only for losses, which were replaced with the developed objectives. For both models, we used image scale equal to 800.

4.1 DATASET

We evaluated our loss functions on the COCO 2017 dataset [Lin et al., 2014]. This dataset is known for being quite noisy [Khetan et al., 2017, Vahdat, 2017], because it was scrapped from the Flickr image database. The dataset consists of more than 330,000 images, with 220,000 labelled images and more than 1.5 million objects in total. All objects are presented in the wild. The COCO dataset contains

80 object classes. Commonly, images contain objects of multiple classes, but about 10% contain a single class only. All objects are annotated with bounding box coordinates and classes, which are stored in the JSON format.

4.2 EVALUATION

Experiments were conducted on a single NVIDIA Titan RTX GPU with 24GB of VRAM. The original implementation of the RetinaNet model was taken from the detectron2 [Wu et al., 2019] library, based on the pytorch [Paszke et al., 2019] framework.

First, we trained the original RetinaNet model using Adam [Kingma and Ba, 2015] optimizer with an initial learning rate of 0.00001. The learning rate scheduler with warmup was used.

Next, we trained our model, which is Bayesian RetinaNet, using Adam optimizer with an initial learning rate of 0.00001. The learning rate scheduler with warmup was also used. We initialized $s_1 = \log \sigma_1^2$ for the localization task with 1.0, $s_2 = \log \sigma_2^2$ for the classification task with 0.0. Both models training took 900,000 iterations, which is about 3 days on average.

For evaluation, we used a standard script from the cocoapi [Lin et al., 2014] library. Models were evaluated on the val and test-dev splits of the MS COCO 2017 dataset. The primary metric of the COCO is mean average precision (mAP).

4.3 RESULTS ANALYSIS

Tables 1 and 2 show the results of comparing metrics for our model and the original RetinaNet-ResNet-50 model. The original model achieved 35.9% mAP on the val set and 35.7% mAP on the test-dev set, as reported in paper [Lin et al., 2017b].

As can be seen, our model provides an increase of 1.6% for the main mAP metric on the test-dev set and increase of 1% on the val set. This result seems to confirm the hypothesis that modeling aleatoric uncertainty can improve the accuracy of the detection problem solving, which answers the **RQ1**. We can conclude that our proposed losses penalize the neural network better than the original losses of RetinaNet. The estimations of aleatoric uncertainties obtained during the training were 0.124 for the regression task and 0.8 for the classification task. These values correlate with the fact that the COCO dataset has noisy labels.

While all the average precision metrics obtained by Bayesian RetinaNet are higher or equal compared to the baseline, average recall metrics on the val set are better only in 2 of 5 cases. The reason of such an effect is that our loss functions penalize the model more for false positive errors, while the

Table 1: Comparison of RetinaNet trained with original loss functions and Bayesian RetinaNet trained with proposed loss functions, which model homoscedastic aleatoric uncertainty, on the val set of the MS COCO dataset. Here, mAP is mean average precision presented for different IoU thresholds and object sizes (small, medium, large), mAR is mean average recall presented for different numbers of detections per image and object sizes.

Metric	RetinaNet	Bayesian RetinaNet (our)
mAP	35.9%	36.8%
mAP50	54.2%	55.8%
mAP75	38.4%	39.3%
mAPs	20.6%	20.6%
mAPm	38.9%	40.8%
mAPl	46.2%	47.0%
mARmax1	31.8%	32.0%
mARmax10	51.6%	51.3%
mARmax100	54.8%	54.4%
mARs	35.8%	34.2%
mARm	58.5%	58.7%
mARl	69.2%	69.0%

Table 2: Comparison of RetinaNet trained with original loss functions and Bayesian RetinaNet trained with proposed loss functions, which model homoscedastic aleatoric uncertainty, on the test-dev set of the MS COCO dataset. Here, mAP is mean average precision presented for different IoU thresholds and object sizes (small, medium, large).

Metric	RetinaNet	Bayesian RetinaNet (our)
mAP	35.7%	37.3%
mAP50	55.0%	56.2%
mAP75	38.5%	40.0%
mAPs	18.9%	21.0%
mAPm	38.9%	39.6%
mAPl	46.3%	46.8%

true positive rate increases less significantly. This fact is consistent with the functions plots: for example, Bayesian Focal loss provides higher gradient values for bigger errors than the original Focal loss.

The proposed loss functions are easy for incorporating to the existing neural networks, that utilize cross-entropy/Focal loss for the classification and L_1 /Smooth L_1 loss for the localization tasks solving. Thus, in future our losses can be scaled and applied for SpineNet [Du et al., 2020], ATSS [Zhang et al., 2020] and other current state-of-the-art detection models. This answers **RQ2**. Modeling homoscedastic aleatoric uncertainty can advance the neural network detectors robustness, help them better generalize to the real-world scenarios and achieve higher performance.

5 CONCLUSION AND DISCUSSION

In this work, we have proposed the novel loss functions for the detection problem (i.e. joint classification and localization), namely *Bayesian Focal loss* and *Bayesian Smooth L_1 loss* functions. The proposed functions are able to model homoscedastic aleatoric uncertainty during model training and do not require the architecture changes.

The proposed losses were studied using the COCO 2017 dataset based on the RetinaNet-ResNet-50 model. As a result of the study, an increase of 1.6% by the mAP metric on the test-dev set was achieved. The obtained result confirms the hypothesis that modeling homoscedastic aleatoric uncertainty improves the accuracy of the detection problem solution. The values of aleatoric uncertainties obtained using our losses were 0.8 for the classification task and 0.124 for the regression task.

In future work, we plan to apply the proposed loss functions to other models, which are based on the RetinaNet architecture, for example, SpineNet [Du et al., 2020], ATSS [Zhang et al., 2020]. We also plan to evaluate the developed functions on other datasets with known noise values to prove that the uncertainties estimates correlate with these values. Furthermore, it is interesting to apply the developed loss functions to model heteroscedastic aleatoric uncertainty, which can advance the detection accuracy and increase the interpretability of detection per object.

References

- Milton Abramowitz, Irene A Stegun, and Robert H Romer. Handbook of mathematical functions with formulas, graphs, and mathematical tables, 1988.
- Abhijit Bendale and Terrance E Boult. Towards open set deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1563–1572, 2016.
- Jonas Mueller Curtis G. Northcutt, Anish Athalye. Pervasive label errors in test sets destabilize machine learning benchmarks. *ICLR 2021 RobustML and Weakly Supervised Learning Workshops*, 2021.
- Xianzhi Du, Tsung-Yi Lin, Pengchong Jin, Golnaz Ghiasi, Mingxing Tan, Yin Cui, Quoc V. Le, and Xiaodan Song. Spinenet: Learning scale-permuted backbone for recognition and localization. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 11589–11598. IEEE, 2020.
- Di Feng, Lars Rosenbaum, and Klaus Dietmayer. Towards safe autonomous driving: Capture uncertainty in the deep neural network for lidar 3d vehicle detection. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 3266–3273. IEEE, 2018.
- Di Feng, Lars Rosenbaum, Fabian Timm, and Klaus Dietmayer. Leveraging heteroscedastic aleatoric uncertainties for robust real-time lidar 3d object detection. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 1280–1287. IEEE, 2019.
- Lei Feng, Senlin Shu, Zhuoyi Lin, Fengmao Lv, Li Li, and Bo An. Can cross entropy loss be robust to label noise. In *Proceedings of the 29th International Joint Conferences on Artificial Intelligence*, pages 2206–2212, 2020.
- Ross B. Girshick. Fast R-CNN. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 1440–1448. IEEE Computer Society, 2015.
- Ali Harakeh, Michael Smart, and Steven L Waslander. Bayesod: A bayesian approach for uncertainty estimation in deep object detectors. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 87–93. IEEE, 2020.
- 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.
- Peter J Huber. Robust estimation of a location parameter. In *Breakthroughs in statistics*, pages 492–518. Springer, 1992.
- Eyke Hüllermeier and Willem Waegeman. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110(3): 457–506, 2021.
- Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5574–5584, 2017.
- Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 7482–7491. IEEE Computer Society, 2018.
- Ashish Khetan, Zachary C Lipton, and Anima Anandkumar. Learning from noisy singly-labeled data. *arXiv preprint arXiv:1712.04577*, 2017.

- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- Florian Kraus and Klaus Dietmayer. Uncertainty estimation in one-stage object detection. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 53–60. IEEE, 2019.
- Michael Truong Le, Frederik Diehl, Thomas Brunner, and Alois Knol. Uncertainty estimation for deep neural object detectors in safety-critical applications. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 3873–3878. IEEE, 2018.
- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In David J. Fleet, Tomás Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V*, volume 8693 of *Lecture Notes in Computer Science*, pages 740–755. Springer, 2014.
- Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 936–944. IEEE Computer Society, 2017a.
- Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 2999–3007. IEEE Computer Society, 2017b.
- Gregory P Meyer, Ankit Laddha, Eric Kee, Carlos Vallespi-Gonzalez, and Carl K Wellington. Lasernet: An efficient probabilistic 3d object detector for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12677–12686, 2019.
- Dimity Miller, Lachlan Nicholson, Feras Dayoub, and Niko Sünderhauf. Dropout sampling for robust object detection in open-set conditions. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3243–3249. IEEE, 2018.
- Dimity Miller, Feras Dayoub, Michael Milford, and Niko Sünderhauf. Evaluating merging strategies for sampling-based uncertainty techniques in object detection. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2348–2354. IEEE, 2019.
- Dimity Miller, Niko Sünderhauf, Michael Milford, and Feras Dayoub. Uncertainty for identifying open-set errors in visual object detection. *arXiv preprint arXiv:2104.01328*, 2021.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 8024–8035, 2019.
- Janis Postels, Francesco Ferroni, Huseyin Coskun, Nassir Navab, and Federico Tombari. Sampling-free epistemic uncertainty estimation using approximated variance propagation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2931–2940, 2019.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 6105–6114. PMLR, 2019.
- Arash Vahdat. Toward robustness against label noise in training deep discriminative neural networks. *arXiv preprint arXiv:1706.00038*, 2017.
- Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z. Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 9756–9765. IEEE, 2020.