

From Black-box to White-box: Examining Confidence Calibration under different Conditions

Franziska Schwaiger¹, Maximilian Henne¹, Fabian Küppers², Felipe Schmoeller Roza¹,
Karsten Roscher¹, Anselm Haselhoff²

¹Fraunhofer Institute for Cognitive Systems IKS, Munich, Germany

²Ruhr West University of Applied Sciences, Bottrop, Germany

{franziska.schwaiger, maximilian.henne, felippe.schmoeller.da.roza}@iks.fraunhofer.de

{fabian.kueppers, anselm.haselhoff}@hs-ruhrwest.de

Abstract

Confidence calibration is a major concern when applying artificial neural networks in safety-critical applications. Since most research in this area has focused on classification in the past, confidence calibration in the scope of object detection has gained more attention only recently. Based on previous work, we study the miscalibration of object detection models with respect to image location and box scale. Our main contribution is to additionally consider the impact of box selection methods like non-maximum suppression to calibration. We investigate the default intrinsic calibration of object detection models and how it is affected by these post-processing techniques. For this purpose, we distinguish between black-box calibration with non-maximum suppression and white-box calibration with raw network outputs. Our experiments reveal that post-processing highly affects confidence calibration. We show that non-maximum suppression has the potential to degrade initially well-calibrated predictions, leading to overconfident and thus miscalibrated models.

1 Introduction

Modern deep neural networks achieve remarkable results on various tasks but it is a well-known issue that these networks fail to provide reliable estimates about the correctness of predictions in many cases (Niculescu-Mizil and Caruana 2005; Guo et al. 2017). A network outputs a score attached to each prediction that can be interpreted as the probability of correctness. Such a model is *well-calibrated* if the observed accuracy matches the estimated confidence scores. However, recent work has shown that these confidence scores neither represent the actual observed accuracy in classification (Niculescu-Mizil and Caruana 2005; Naeini, Cooper, and Hauskrecht 2015; Guo et al. 2017) nor the observed precision in object detection (Küppers et al. 2020). Calibrated confidence estimates integrated in safety-critical applications like autonomous driving can provide valuable additional information with respect to situational awareness and can reduce the risk of hazards resulting from functional insufficiencies by decreasing the space of unknown unsafe scenarios which is a critical part for the safety of the intended functionality (SOTIF ISO/PAS 21448).

Copyright © 2021 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

In the past, most research in this area has focused on classification (Naeini, Cooper, and Hauskrecht 2015; Kull, Silva Filho, and Flach 2017; Guo et al. 2017; Seo, Seo, and Han 2019; Mukhoti et al. 2020), whereas calibration in object detection has recently gained more attention (Neumann, Zisserman, and Vedaldi 2018; Feng et al. 2019; Küppers et al. 2020). Object detection is a joint task of classification and regression of the predictions' position and scale. Recent work has shown that the regression branch of object detection models also affects confidence calibration (Küppers et al. 2020). However, the observable detections of a model are commonly processed by non-maximum suppression (NMS) and/or thresholded by a certain confidence score. In this work, our goal is to investigate the influence of such post-processing techniques on the model calibration. For this purpose, we adapt common object detection models and examine their miscalibration before NMS on the one hand (white-box scenario). In this way, we have access to the raw predictions of a network and are thus able to examine the network's calibration properties by default. On the other hand, we further apply NMS with increasing intersection over union (IoU) thresholds (black-box scenario), which varies the number of boxes that are suppressed. Changing the parameters of the NMS enables us to examine to what extent the models are intrinsically calibrated and how this is affected by postprocessing techniques. An illustrative representation for the problem setting is demonstrated in Fig. 1. Furthermore, we use a Faster R-CNN architecture (Ren et al. 2015) that uses the cross entropy loss during training and compare it to a RetinaNet (Lin et al. 2017) that uses a focal loss. It is already known that models trained with focal loss produce much less confident predictions (Mukhoti et al. 2020). This enables us to further investigate the effect of post-processing methods by comparing the default calibration properties of both model architectures with and without NMS.

This work is structured as follows: we give a review of the current state-of-the-art research in confidence calibration in Section 2. We further give a definition of white-box and black-box calibration and a description of our calibration targets in Section 3. In Section 4, our experimental results are demonstrated and in Section 5 we discuss our findings.




Figure 1: Typically a non-maximum suppression (NMS) is applied to all detections of a detection model to fuse and reduce redundant bounding boxes. In our work, we investigate how NMS affects confidence calibration. Thus, we study the difference in calibration before NMS (white-box) and afterward (black-box).

2 Related Work

Numerous methods have been developed in the past to address the miscalibration of neural networks. One of the first representatives of post-processing calibration methods has been histogram binning (Zadrozny and Elkan 2001), isotonic regression (Zadrozny and Elkan 2002), Bayesian binning (Naeini, Cooper, and Hauskrecht 2015), and Platt scaling (Platt 1999), whereas more recently temperature scaling (Guo et al. 2017), beta calibration (Kull, Silva Filho, and Flach 2017), and Dirichlet calibration (Kull et al. 2019) have been developed to tackle miscalibration in the scope of classification. In object detection models, dealing with miscalibration presents a different set of challenges and was first addressed by (Neumann, Zisserman, and Vedaldi 2018), who proposed an additional model output to be utilized as a regularizing temperature applied to the remaining logits. Recently, (Küppers et al. 2020) have studied the effect of position and scale of detected objects to miscalibration and concluded that calibration also depends on the regression output of a detection model. They further provide a framework to include position and scale information into a calibration mapping.

For the task of classification, a common way to measure miscalibration is to adapt the *expected calibration error* (ECE), proposed by (Naeini, Cooper, and Hauskrecht 2015), which uses a binning scheme to measure the gap between observed frequency and average confidence. (Kumar, Liang, and Ma 2019) show in their work that the common ECE underestimates the true calibration error in some cases and provide a differentiable upper bound called *maximum mean calibration error* (MMCE) that can also be used during model training as a second regularization term. For measuring miscalibration in object detection tasks, an extension of the ECE called *detection expected calibration error* (D-ECE) was proposed by (Küppers et al. 2020), consisting of a multidimensional binning scheme to assess the miscalibration over all predicted features of an object detection model.

Since the standard cross-entropy loss is prone to favor overly confident predictions, further research directions in-

vestigate how to directly obtain well-calibrated models after training. Besides the previously mentioned MMCE, the authors in (Pereyra et al. 2017) introduce a regularization term penalizing high confident predictions. In contrast, (Müller, Kornblith, and Hinton 2019) show that label smoothing yields good probabilities after training. Recently, (Mukhoti et al. 2020) investigate the effects of focal loss, originally proposed as a loss term for RetinaNet (Lin et al. 2017), on confidence calibration. They show that using focal loss in conjunction with an adaptive parameter significantly improves the confidence calibration of classification models. We also observe that using focal loss prevents overconfident predictions in our experiments on the RetinaNet with standard hyperparameters (Lin et al. 2017). While well-known object detection models like Faster-RCNN (Ren et al. 2015) commonly tend to output overconfident predictions, the probability scores of a RetinaNet rather underestimate the observed frequency.

3 Defining Confidence Calibration for Object Detection Models

In this section, we describe the definition of black-box and white-box calibration. The idea behind this distinction is to analyze the impact of bounding-box postprocessing on calibration.

An object detector takes an image as input x and outputs predictions in form of a class label $y \in \mathcal{Y}$ with corresponding confidence score $p \in [0, 1]$ and bounding box $r = (c_x, c_y, h, w) \in \mathcal{R}^J$, with (c_x, c_y) being the center position, (h, w) the box height and width and J the size of the used box encoding. The authors in (Küppers et al. 2020) propose a confidence calibration that not only considers the confidence score p but also includes the box information r . The performance of a detector is thus evaluated by matching its predictions $(\hat{y}, \hat{p}, \hat{r})$ with the ground-truth annotations, where $m = 1$ denotes a matched box and $m = 0$ a mismatch. More formally, perfect calibration in the scope

of object detection is defined by

$$\underbrace{\mathbb{P}(M = 1 | \hat{P} = p, \hat{Y} = y, \hat{R} = r)}_{\text{precision given } p, y, r} = \underbrace{p}_{\text{confidence}}, \quad (1)$$

$$\forall p \in [0, 1], y \in \mathcal{Y}, r \in \mathcal{R}^J.$$

As the detections rarely match the ground-truth perfectly, true positives (TP, $m = 1$) and false positives (FP, $m = 0$) are obtained by comparing the IoU to a fixed threshold τ . TP and FP correspond to boxes with $\text{IoU} \geq \tau$ and $\text{IoU} < \tau$, respectively. The process of inference is commonly followed by a non-maximum suppression since an object detection model outputs a huge amount of mostly less confident and redundant detections. On the one hand, we can consider the definition of calibration given by Equation 1 to the raw outputs of a detector without any post-processing. We denote this case as the white-box calibration case for the following of this paper. On the other hand, we can also view the NMS as part of the detector and treat the output of the NMS as our desired calibration target. This is denoted as black-box calibration.

In (Küppers et al. 2020), the detection expected calibration error (D-ECE) is defined as an extension of the commonly used ECE (Naeini, Cooper, and Hauskrecht 2015) for object detection tasks. The D-ECE also includes the box information r by partitioning the space of each variable k into N_k equally spaced bins. The total amount of bins is given by $N_{total} = \prod_{k=1}^K N_k$ and the D-ECE is defined as

$$\text{D-ECE}_K = \sum_{n=1}^{N_{total}} \frac{|I(n)|}{|D|} \cdot |\text{prec}(n) - \text{conf}(n)|, \quad (2)$$

where $I(n)$ is the set of all samples in a single bin and $|D|$ the total amount of samples, while $\text{prec}(n)$ and $\text{conf}(n)$ denote the average precision and confidence within each bin, respectively. We use this metric to measure miscalibration in both cases: For white-box, we consider all possible box predictions whereas for black-box only the winning boxes after NMS are considered. This is explained in more detail in the following section.

4 Experimental Evaluation

In order to analyze the confidence calibration under different conditions, we use the COCO 2017 validation dataset (Lin et al. 2014) with a random split of 70% and 30% for training and testing the calibration, respectively.

Evaluation Protocol

We perform both black-box and white-box calibration by following the evaluation protocol of (Küppers et al. 2020) and use their provided calibration framework. The final calibration results are obtained as an average over 20 independent training and testing results. For inference, we use a pre-trained RetinaNet (Lin et al. 2017) and a Faster R-CNN (Ren et al. 2015) model provided by the *Detectron2* framework (Wu et al. 2019). While the classification branch of the former model is trained by cross entropy loss, the latter one

uses a focal loss that enables to focus on hard examples during training with low confidence. On the other hand, good predictions with high confidence are less weighted during training that in turn leads to less confident predictions (Lin et al. 2017; Mukhoti et al. 2020). Our experiments are restricted to the predictions of class *person*.

To study the effect of non-maximum suppression, we apply different IoU thresholds to merge boxes denoted by $\text{NMS}@\{0.5, 0.75, 0.9\}$. In the white-box case without NMS, we use the raw predictions for measuring and performing calibration on the one hand. On the other hand, we further adopt *top-k* box selection where only k bounding boxes with the highest confidence are kept using $k = 1000$. This is the common case during inference to reduce low confidential and mostly redundant predictions. Following (Küppers et al. 2020), the predictions of all models are obtained by inference with a probability threshold of 0.3 which means discarding all predictions with a confidence score less than this threshold. As the relative amount of predictions per image with low confidence score is significantly higher than the relative amount of the remaining predictions, this probability threshold ensures that the D-ECE is not dominated by these low confidence samples.

For confidence calibration, we use multivariate histogram binning (Zadrozny and Elkan 2001; Küppers et al. 2020) for calibration as a fast and reliable calibration method. We also evaluate several setups with different subsets of box information to evaluate the effect of the used feature set. We either use the confidence only, also including the box centers (\hat{c}_x, \hat{c}_y) or box scales (\hat{h}, \hat{w}), or we use all features for measuring and performing calibration. For the histogram-based calibration, we use 15 bins for confidence only, $N_k = 5$ bins for $(\hat{p}, \hat{c}_x, \hat{c}_y)$ and $(\hat{p}, \hat{h}, \hat{w})$, and $N_k = 3$ when using all available features. In contrast, for D-ECE computation we use 20 bins for confidence only, $N_k = 8$ bins for $(\hat{p}, \hat{c}_x, \hat{c}_y)$ and $(\hat{p}, \hat{h}, \hat{w})$, and $N_k = 5$ when using all available information. We increase the robustness of the D-ECE calculation by also neglecting bins with less than 8 samples (Küppers et al. 2020).

Results

In Tables 2 and 3, the results for black-box and white-box calibration for RetinaNet and Faster R-CNN are presented, respectively. Three different IoU threshold values of $\tau = \{0.5, 0.6, 0.75\}$ are considered to match predictions with ground-truth annotations. In the tables, each cell presents the D-ECE for the baseline (without calibration) and the corresponding D-ECE after histogram-based calibration (HB). The Tables 2 and 3 show the results of the black-box models with varying strength of NMS as well as the calibration results for the white-box case without NMS. The D-ECE is evaluated with different additional box information: The first column shows the confidence only calibration, the second and third columns the calibration with box centers and box scales, and the last columns show the results for the calibration with all box information considered. The best D-ECE scores are highlighted for each set of features and IoU value across all variants.

	(\hat{p})	(\hat{p}, c_x, c_y)	(\hat{p}, h, w)	full
NMS@0.5	0	1	29	528
NMS@0.75	0	0	20	485
NMS@0.9	0	0	12	435
Without NMS	0	0	9	414
Baseline	20	256	256	1024

Table 1: Amount of neglected bins within D-ECE calculation of the three black-box models and the white-box model for Faster R-CNN (Ren et al. 2015; Wu et al. 2019). A similar amount of bins is also neglected during the examinations for RetinaNet (Lin et al. 2017; Wu et al. 2019).

For Faster R-CNN, we observe that the white-box model calibrates consistently better by default than the black-box models in most cases. In contrast, we observe the opposite behavior for the RetinaNet model. Therefore, we further study the calibration properties of those networks by inspecting their reliability diagrams shown in Fig. 3 for the black-box and white-box cases. The RetinaNet white-box model without NMS offers underconfident predictions which is a known property of models trained by focal loss (Lin et al. 2017). After NMS, a particular behavior can be observed in Fig. 3e with overconfident predictions in the low confidence interval ($\hat{p} < 0.5$) and underconfident predictions in the high confidence interval ($\hat{p} > 0.5$). Also, when comparing the calibrated results shown in Fig. 3b and 3f, it is evident that the calibration for the white-box model leads to a better D-ECE score. In contrast, Faster R-CNN outputs reasonably well calibrated predictions before NMS but is highly overconfident after NMS. Again, we observe that the white-box D-ECE score is much better compared to the black-box model **after** calibration has been applied.

We also study the effect of position-dependent miscalibration as in (Küppers et al. 2020), shown in Fig. 4. We compare the white-box and black-box models before and after calibration for each object detector. These figures allow to analyze if calibration is influenced by the position of predicted bounding boxes. All images show a tendency of higher miscalibration close to the borders. That may be caused by the difficulty of detecting objects correctly which are cropped out of the frame. However, this is of minor relevance considering that most of the positional discrepancies are mitigated after calibration in all cases.

As shown in Tables 2 and 3, the calibration for the white-box model performs better than the calibration for the black-box model for the first and second columns. The opposite happens when including the box scales into the computation of the D-ECE. Here, the black-box model with NMS@0.5 provides the best results. A possible explanation for this observation could be, that by increasing the NMS value, the number of samples also increases from 4,229 and 4,496 to 117,292 and 37,355 for RetinaNet and Faster R-CNN, respectively. As expected, the more we go in the white-box direction, the less predictions are discarded. Having more samples for the miscalibration computation also means that there are possibly more samples within each bin leading to a more robust miscalibration estimation (Kumar, Liang, and

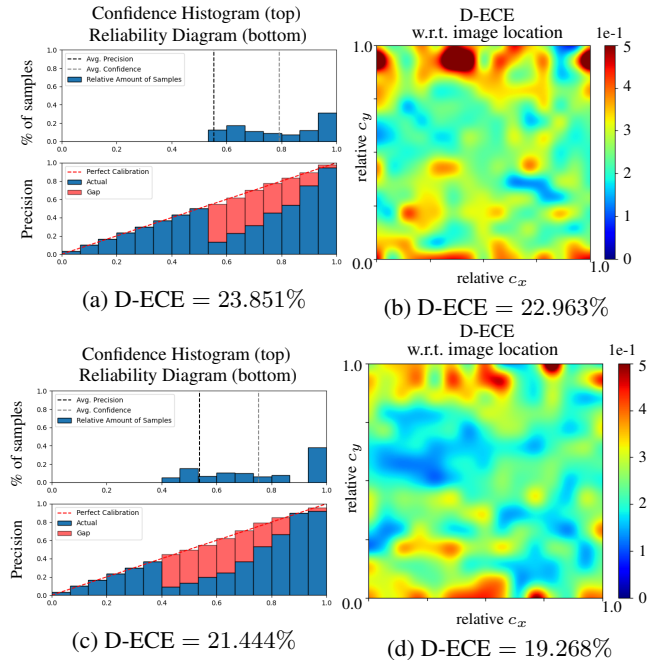


Figure 2: Confidence histogram and reliability diagram (left) and position-dependent heatmap (right) for RetinaNet (Lin et al. 2017; Wu et al. 2019) (top row) and Faster RCNN (Ren et al. 2015; Wu et al. 2019) (bottom row) after white-box calibration and then further application of non-maximum suppression with NMS@0.5.

Ma 2019). As previously mentioned, bins with less than 8 samples are neglected for the computation of the D-ECE. The total amount of neglected bins for each configuration is illustrated in Table 1. Especially using all available information for calibration (full case), more and more bins are left out when going from white-box (bottom) to black-box (top) resulting in less bins contributing to the miscalibration score.

A critical question arises how to integrate white-box calibration into the object detection pipeline. As demonstrated in the previous results, NMS has a significant impact in the calibration affecting the precision as well as the confidence scores of the detections. It has been shown that NMS has the potential to degrade the calibration results. Therefore, we investigate the calibration properties of the detection models that are processed by a NMS with histogram-based calibration beforehand. The results are shown in Fig. 2: It can be seen that calibration before NMS leads to higher miscalibration as the confidence is calibrated before NMS as well. However, as NMS also affects the precision, the detection model gets too overconfident in both cases. In order to preserve good calibrations from the white-box method, alternative box suppression methods should be investigated. One option would be to integrate the confidence calibration with the box merging strategies compared by (Roza et al. 2020), such as *weighted box fusion* and *variance voting* and test how such methods influence the model calibration.

	(\hat{p})	(\hat{p}, c_x, c_y)	(\hat{p}, h, w)	full		(\hat{p})	(\hat{p}, c_x, c_y)	(\hat{p}, h, w)	full
IoU@0.5					IoU@0.5				
Baseline	16.200	12.004	14.963	13.478	Baseline	15.448	15.750	14.246	12.586
HB	1.636	6.335	2.775	5.391	HB	1.388	6.123	4.252	6.665
IoU@0.6					IoU@0.6				
Baseline	20.862	15.303	18.743	16.546	Baseline	3.435	7.486	6.710	7.064
HB	1.673	6.091	2.991	5.743	HB	1.441	6.273	4.192	6.444
IoU@0.75					IoU@0.75				
Baseline	31.659	24.684	28.864	25.765	Baseline	20.980	20.840	20.041	17.504
HB	1.436	6.095	3.110	5.704	HB	1.227	4.847	3.315	4.974

(a) Black-box calibration with NMS@0.5, $|D| = 4, 229$.(b) Black-box calibration NMS@0.75, $|D| = 7, 923$

	(\hat{p})	(\hat{p}, c_x, c_y)	(\hat{p}, h, w)	full		(\hat{p})	(\hat{p}, c_x, c_y)	(\hat{p}, h, w)	full
IoU@0.5					IoU@0.5				
Baseline	30.748	30.672	30.436	29.427	Baseline	28.027	28.127	28.014	28.176
HB	1.212	5.290	3.671	6.686	HB	0.855	4.947	2.895	6.114
IoU@0.6					IoU@0.6				
Baseline	21.773	21.954	21.612	21.350	Baseline	23.097	23.290	23.118	23.482
HB	1.195	5.717	3.981	7.647	HB	1.033	5.331	3.306	6.912
IoU@0.75					IoU@0.75				
Baseline	3.057	6.907	8.489	10.143	Baseline	8.487	10.190	10.190	11.992
HB	1.367	5.675	4.468	7.847	HB	1.132	5.892	4.207	8.266

(c) Black-box calibration with NMS@0.9, $|D| = 20, 005$ (d) White-box calibration without NMS, $|D| = 117, 292$

Table 2: D-ECE results (%) for RetinaNet (Lin et al. 2017; Wu et al. 2019) for different IoU scores. Each column shows the baseline D-ECE and the calibrated one using histogram-based (HB) calibration with different subsets using either confidence only \hat{p} , including the box centers (c_x, c_y) , box scales (h, w) or using all features. Note that comparing the D-ECE scores of columns to each other is not applicable since different subsets of data have been used for D-ECE measurement and calibration.

	(\hat{p})	(\hat{p}, c_x, c_y)	(\hat{p}, h, w)	full		(\hat{p})	(\hat{p}, c_x, c_y)	(\hat{p}, h, w)	full
IoU@0.5					IoU@0.5				
Baseline	7.781	9.060	7.829	7.168	Baseline	7.597	9.927	10.828	9.804
HB	1.789	6.186	2.947	4.960	HB	1.523	6.968	3.778	6.134
IoU@0.6					IoU@0.6				
Baseline	9.370	10.041	9.033	7.810	Baseline	16.100	15.226	16.417	14.933
HB	1.564	6.075	3.105	5.142	HB	1.343	6.323	3.490	5.610
IoU@0.75					IoU@0.75				
Baseline	31.659	24.684	28.864	25.765	Baseline	34.634	32.535	31.861	27.883
HB	1.436	6.095	3.110	5.704	HB	1.123	4.878	3.018	4.996

(a) Black-box calibration with NMS@0.5, $|D| = 4, 496$.(b) Black-box calibration with NMS@0.75, $|D| = 7, 231$.

	(\hat{p})	(\hat{p}, c_x, c_y)	(\hat{p}, h, w)	full		(\hat{p})	(\hat{p}, c_x, c_y)	(\hat{p}, h, w)	full
IoU@0.5					IoU@0.5				
Baseline	7.323	10.431	10.042	10.318	Baseline	6.914	9.619	8.638	10.061
HB	1.354	6.697	4.062	7.121	HB	1.038	5.234	3.206	6.239
IoU@0.6					IoU@0.6				
Baseline	7.499	10.328	11.622	11.630	Baseline	4.592	7.720	8.540	9.548
HB	1.184	6.383	4.050	7.141	HB	1.099	5.523	3.603	6.959
IoU@0.75					IoU@0.75				
Baseline	25.689	25.539	25.792	25.002	Baseline	13.067	13.883	15.658	16.462
HB	1.139	5.478	4.126	6.908	HB	0.999	5.996	4.505	8.652

(c) Black-box calibration with NMS@0.9, $|D| = 17, 742$.(d) White-box calibration without NMS, $|D| = 37, 355$.

Table 3: D-ECE results (%) for Faster R-CNN (Ren et al. 2015; Wu et al. 2019) before and after histogram-based (HB) calibration using different IoU thresholds for NMS. The structure of this table is comparable to Tab. 2.

5 Conclusion

In this paper, we analyzed the influence of box suppression methods on confidence calibration for object detection models. To do so, we adapt models without box suppression methods denoted as white-box models, contrasting to the black-box approach commonly suggested. We performed histogram-based calibration for both black-box and white-box scenarios on the COCO dataset. We found that the initial calibration of detection models is highly impacted by NMS. Additionally, we observed that calibration also depends on the architecture of the object detection model. For RetinaNet, the model predictions are underconfident before applying NMS whereas, for Faster R-CNN, the white-box model outputs quite well calibrated detections that become overconfident after NMS.

Knowing that the miscalibration not only depends on the classification outputs but also on the regression output for the bounding boxes, we performed histogram-based calibration using different subsets of the output data. For the confidence only and (\hat{p}, c_x, c_y) case, the white-box model outperforms the black-box models while the black-box models present slightly better results on the other scenarios.

While the white-box calibration has given good results, the most effective integration of white-box calibration methods in existing object detectors utilizing NMS remains as an open issue. As shown by the results in this paper, the NMS layer affects the results by giving different calibration profiles before and after the suppression. Corroborating with further results presented in this paper, the calibrated detections obtained by the white-box models deteriorated after NMS for both RetinaNet and Faster R-CNN. However, we think this problem can be solved by using other suppression methods which consider a larger set of the overall better calibrated boxes than NMS.

For future work we suggest alternative applications to the standard NMS method to verify if they can lead to better calibrated object detectors. One option would be to integrate the confidence calibration with box merging strategies compared by (Roza et al. 2020), such as *box averaging*, *weighted box fusion* or *variance voting*.

Acknowledgements

This work was funded by the Bavarian Ministry for Economic Affairs, Regional Development and Energy as part of a project to support the thematic development of the Institute for Cognitive Systems and within the Intel Collaborative Research Institute Safe Automated Vehicles.

References

- Feng, D.; Rosenbaum, L.; Gläser, C.; Timm, F.; and Dietmayer, K. 2019. Can We Trust You? On Calibration of a Probabilistic Object Detector for Autonomous Driving. *ArXiv* abs/1909.12358.
- Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On Calibration of Modern Neural Networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, 1321–1330. PMLR.
- Kull, M.; Nieto, M. P.; Kängsepp, M.; Silva Filho, T.; Song, H.; and Flach, P. 2019. Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with Dirichlet calibration. In *Advances in Neural Information Processing Systems*, 12316–12326.
- Kull, M.; Silva Filho, T.; and Flach, P. 2017. Beta calibration: a well-founded and easily implemented improvement on logistic calibration for binary classifiers. In *Artificial Intelligence and Statistics*, 623–631.
- Kumar, A.; Liang, P. S.; and Ma, T. 2019. Verified Uncertainty Calibration. In *Advances in Neural Information Processing Systems* 32, 3792–3803. Curran Associates, Inc. URL <http://papers.nips.cc/paper/8635-verified-uncertainty-calibration.pdf>.
- Küppers, F.; Kronenberger, J.; Shantia, A.; and Haselhoff, A. 2020. Multivariate Confidence Calibration for Object Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 326–327.
- Lin, T.; Maire, M.; Belongie, S. J.; Bourdev, L. D.; Girshick, R. B.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft COCO: Common Objects in Context. *CoRR* abs/1405.0312. URL <http://arxiv.org/abs/1405.0312>.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal Loss for Dense Object Detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2980–2988.
- Mukhoti, J.; Kulharia, V.; Sanyal, A.; Golodetz, S.; Torr, P. H.; and Dokania, P. K. 2020. Calibrating Deep Neural Networks using Focal Loss. In *Advances in Neural Information Processing Systems*.
- Müller, R.; Kornblith, S.; and Hinton, G. E. 2019. When does label smoothing help? In *Advances in Neural Information Processing Systems*, 4694–4703.
- Naeini, M.; Cooper, G.; and Hauskrecht, M. 2015. Obtaining Well Calibrated Probabilities Using Bayesian Binning. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, 2901–2907.
- Neumann, L.; Zisserman, A.; and Vedaldi, A. 2018. Relaxed Softmax: Efficient Confidence Auto-Calibration for Safe Pedestrian Detection. In *Workshop on Machine Learning for Intelligent Transportation Systems (NIPS)*.
- Niculescu-Mizil, A.; and Caruana, R. 2005. Predicting Good Probabilities with Supervised Learning. In *Proceedings of the 22nd International Conference on Machine Learning*, 625–632.
- Pereyra, G.; Tucker, G.; Chorowski, J.; Łukasz Kaiser; and Hinton, G. 2017. Regularizing Neural Networks by Penalizing Confident Output Distributions. *CoRR*.
- Platt, J. 1999. Probabilistic Outputs for Support Vector Machines and Comparisons to Regularized Likelihood Methods. *Advances in Large Margin Classifiers* 61–74.
- Ren, S.; He, K.; Girshick, R. B.; and Sun, J. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *CoRR* abs/1506.01497. URL <http://arxiv.org/abs/1506.01497>.
- Roza, F. S.; Henne, M.; Roscher, K.; and Gunnemann, S. 2020. Assessing Box Merging Strategies and Uncertainty

Estimation Methods in Multimodel Object Detection. In *Beyond mAP: Reassessing the Evaluation of Object Detectors @ECCV*, –.

Seo, S.; Seo, P. H.; and Han, B. 2019. Learning for Single-Shot Confidence Calibration in Deep Neural Networks Through Stochastic Inferences. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Wu, Y.; Kirillov, A.; Massa, F.; Lo, W.-Y.; and Girshick, R. 2019. Detectron2. <https://github.com/facebookresearch/detectron2>.

Zadrozny, B.; and Elkan, C. 2001. Obtaining Calibrated Probability Estimates from Decision Trees and Naive Bayesian Classifiers. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML)*, 609–616.

Zadrozny, B.; and Elkan, C. 2002. Transforming Classifier Scores into Accurate Multiclass Probability Estimates. In *Proceedings of the Eighth International Conference on Knowledge Discovery and Data Mining, July 23-26, 2002, Edmonton, Alberta, Canada*, 694–699. doi:10.1145/775047.775151. URL <https://doi.org/10.1145/775047.775151>.