

Multiclass Confidence and Localization Calibration for Object Detection

Bimsara Pathiraja, Malitha Gunawardhana, Muhammad Haris Khan
Mohamed bin Zayed University of Artificial Intelligence, UAE

{bimsara.pathiraja,malitha.gunawardhana,muhammad.haris}@mbzuai.ac.ae

Abstract

Albeit achieving high predictive accuracy across many challenging computer vision problems, recent studies suggest that deep neural networks (DNNs) tend to make overconfident predictions, rendering them poorly calibrated. Most of the existing attempts for improving DNN calibration are limited to classification tasks and restricted to calibrating in-domain predictions. Surprisingly, very little to no attempts have been made in studying the calibration of object detection methods, which occupy a pivotal space in vision-based security-sensitive, and safety-critical applications. In this paper, we propose a new train-time technique for calibrating modern object detection methods. It is capable of jointly calibrating multiclass confidence and box localization by leveraging their predictive uncertainties. We perform extensive experiments on several in-domain and out-of-domain detection benchmarks. Results demonstrate that our proposed train-time calibration method consistently outperforms several baselines in reducing calibration error for both in-domain and out-of-domain predictions. Our code and models are available at <https://github.com/bimsarapathiraja/MCCL>

1. Introduction

Deep neural networks (DNNs) are the backbone of many top-performing systems due to their high predictive performance across several challenging domains, including computer vision [16,17,41,45,52] and natural language processing [5,7]. However, some recent works [14,15,38,47] report that DNNs are susceptible to making overconfident predictions, which leaves them miscalibrated. This not only spurs a mistrust in their predictions, but more importantly, could lead to disastrous consequences in several safety-critical applications, such as healthcare diagnosis [8,43], self-driving cars [13], and legal research tools [50]. For instance, in self-driving cars, if the perception component wrongly detects a stop sign as a speed limit sign with high confidence, it can potentially lead to disastrous outcomes.

Several strategies have been proposed in the recent past for improving model calibration. A simple calibration technique is a post-processing step that re-scales the outputs of a trained model using parameters which are learnt on a hold-out portion of the training set [14]. Despite being easy to implement, these post-processing approaches are restrictive. They assume the availability of a hold-out set, which is not always possible in many real-world settings. Another route to reducing calibration error is train-time calibration techniques, which intervene at the training time by involving all model parameters. Typically train-time calibration methods feature an auxiliary loss term that is added to the application-specific loss function to regularize predictions [18,27,33,35].

We note that almost all prior efforts towards improving model calibration target the task of visual image classification. Surprisingly, little to no noticeable attempts have been made in studying the calibration of visual object detection models. Visual object detection methods account for a major and critical part of many vision-based decision-making systems. Moreover, most of the current calibration techniques only aim at reducing calibration error for in-domain predictions. However, in many realistic settings, it is likely that, after model deployment, the incoming data distribution could continuously change from the training data distribution. In essence, the model should be well-calibrated for both in-domain and out-of-domain predictions.

To this end, in this paper, we aim to study the calibration of (modern) deep learning-based object detection methods. In this pursuit, we observe that, (a) object detection methods are intrinsically miscalibrated, (b) besides displaying noticeable calibration errors for in-domain predictions, they are also poorly calibrated for out-of-domain predictions and, (c) finally, the current calibration techniques for classification are sub-optimal for object detection (Figure 1). Towards improving the calibration performance of object detection methods, inspired by the train-time calibration route, we propose a new train-time calibration approach aims at jointly calibrating the predictive multiclass confidence and bounding box localization.




Figure 1. DNN-based object detectors are inherently miscalibrated for both in-domain and out-of-domain predictions. Also, calibration methods for image classification are sub-optimal for object detection. Our proposed train-time calibration method for object detection is capable of reducing the calibration error (D-ECE%) of DNN-based detectors in both in-domain and out-of-domain scenarios.

Contributions: (1) We study the relatively unexplored direction of calibrating modern object detectors and observe that they are intrinsically miscalibrated in both in-domain and out-of-domain predictions. Also, the existing calibration techniques for classification are sub-optimal for calibrating object detectors. (2) We propose a new train-time calibration method for detection, at the core of which is an auxiliary loss term, which attempts to jointly calibrate multiclass confidences and bounding box localization. We leverage predictive uncertainty in multiclass confidences and bounding box localization. (3) Our auxiliary loss term is differentiable, operates on minibatches, and can be utilized with other task-specific loss functions. (4) We perform extensive experiments on challenging datasets, featuring several in-domain and out-of-domain scenarios. Our train-time calibration method consistently reduces the calibration error across DNN-based object detection paradigms, including FCOS [45] and Deformable DETR [52], both in in-domain and out-of-domain predictions.

2. Related works

Post-processing calibration methods: A simple approach to calibration is a post-processing step, which re-scales the outputs of a trained model using some parameters that are learned on the hold-out portion of the training set. Temperature scaling (TS), which is an adaptation of Platt scaling [40], is a prominent example. It divides the logits (pre-softmax activations) from a trained network with a fixed temperature parameter ($T > 0$) that is learned using a hold-out validation set. An obvious limitation of TS is that it decreases the confidence of the whole (confidence) vector, including the confidence of the correct class. Beyond using a single temperature parameter (T), some works use a matrix (M) to transform the logits. The matrix (M) is also learnt using a hold-out validation set. Dirichlet calibration (DC) employed Dirichlet distributions to generalize the Beta-calibration [26] method, originally proposed for

binary classification, to a multi-class setting. DC is realized as an extra layer in a neural network whose input is log-transformed class probabilities. The work of [3] proposed a differentiable approximation of expected calibration error (ECE) and utilizes it in a meta-learning framework to obtain well-calibrated models. Islam et al. [23] achieved class-distribution-aware calibration using temperature scaling (TS) and label smoothing (LS) [44] for long-tailed visual recognition. Majority of the aforementioned work address in-domain calibration. Recently, [46] proposed to gradually perturb the hold-out validation set for simulating out-of-domain prior to learning the temperature parameter (T). Despite being easy-to-implement and effective, TS methods require a hold-out validation set, which is not readily available in many realistic scenarios.

Train-time calibration techniques: Another approach to improving model calibration are train-time calibration techniques. Brier score is considered one of the earliest attempts for calibrating binary probabilistic forecast [4]. Some recent works report that models trained with negative log-likelihood (NLL) are prone to making overconfident predictions. A dominant class in train-time methods typically propose an auxiliary loss term that is used in conjunction with NLL. For instance, [39] utilized the Shannon entropy to penalize overconfident predictions. Similarly, Muller et al. [36] showed that label smoothing [44] also improves calibration. Recently, [33] introduced a margin into the label smoothing technique to obtain well-calibrated models. While re-visiting focal loss (FL) [32], [35] demonstrated that it is capable of implicitly calibrating DNNs. Liang et al. [30] incorporated the difference between confidence and accuracy (DCA) as an auxiliary loss term with the Cross-Entropy loss to achieve model calibration. Likewise, [27] developed MMCE loss for calibrating DNNs, which is formulated using a reproducible kernel in Hilbert space [12]. Most of these methods only calibrate the confidence of the predicted label ignoring the confidences of non-predicted

classes. Recently, [18] proposed an auxiliary loss term for calibrating the whole confidence vector.

Probabilistic and non-probabilistic methods: Many probabilistic approaches stem from Bayesian formalism [1], which assumes a prior distribution over the neural network (NN) parameters, and training data is leveraged to obtain the posterior distribution over the NN parameters. This posterior is then used to estimate the predictive uncertainty. The exact Bayesian inference is computationally intractable, consequently, we can see approximate inference methods, including variational inference [2,34], and stochastic expectation propagation [20]. A non-probabilistic approach is ensemble learning that can be used to quantify uncertainty; it uses the empirical variance of the network predictions. Ensembles can be created with the differences in model hyperparameters [47], random initialization of weights and random shuffling of training data [29], dataset shift [38], and Monte Carlo (MC) dropout [10,51]. In this work, we propose to use MC dropout [10] to quantify predictive uncertainty both in class confidences and the bounding localization. It allows creating a distribution over both outputs from a typical DNN-based object detector. The naive implementation of MC dropout can incur high computational cost for large datasets and network architectures during model training. So, we resort to an efficient implementation of MC dropout that greatly reduces this computational overhead.

We note that, almost all prior work for addressing calibration is targeted at classification task [14, 18, 26, 27, 33], and no noticeable study has been published that strives to improve the calibration of object detection methods, especially for out-of-domain predictions. In this paper, we explore the problem of calibrating object detectors and observe that they are inherently miscalibrated for both in-domain and out-domain predictions. To this end, we propose a train-time calibration method aimed at jointly calibrating multiclass confidence and bounding box localization.

3. Method

3.1. Defining and Measuring Calibration

Calibration for classification: A perfectly calibrated model for (image) classification outputs class confidences that match with the predictive accuracy. If the accuracy is less than the confidence, then the model is overconfident and if the accuracy is higher than the confidence, then the model is underconfident. Let $\mathcal{D} = \langle (\mathbf{x}_i, y_i^*) \rangle_{i=1}^N$ denote a dataset consisting of N examples drawn from a joint distribution $\mathcal{D}(\mathcal{X}, \mathcal{Y})$, where \mathcal{X} is an input space and \mathcal{Y} is the label space. For each sample $\mathbf{x}_i \in \mathcal{X}$, $y_i^* \in \mathcal{Y} = \{1, 2, \dots, K\}$ is the corresponding ground truth class label. Let $\mathbf{s} \in \mathbb{R}^K$ be the vector containing the predicted confidences of all K classes, and $s_i[y]$ be the confidence predicted for a class y

on a given input example \mathbf{x}_i . The model is said to be perfectly calibrated when, for each sample $(\mathbf{x}, y) \in \mathcal{D}$:

$$\mathbb{P}(y = y^* | \mathbf{s}[y] = s) = s \quad (1)$$

where $\mathbb{P}(y = y^* | \mathbf{s}[y] = s)$ is the accuracy for each confidence scores in \mathbf{s} .

Calibration for object detection: Contrary to classification, in object detection, the dataset contains the ground-truth annotations for each object in an image, specifically the object localization information and the associated object categories. Let $\mathbf{b}^* \in \mathcal{B} = [0, 1]^4$ be the bounding box annotation of the object and y^* be the corresponding class label. The prediction from an object detection model consists of a class label \hat{y} , with a confidence score \hat{s} and a bounding box $\hat{\mathbf{b}}$. Unlike classification, for object detection, precision is used instead of accuracy for calibration. Therefore, an object detector is perfectly calibrated when [28]:

$$\mathbb{P}(m = 1 | \hat{s} = s, \hat{y} = y, \hat{\mathbf{b}} = \mathbf{b}) = s \quad (2)$$

$$\forall s \in [0, 1], y \in \mathcal{Y}, \mathbf{b} \in [0, 1]^4$$

where $m = 1$ denotes a correctly classified prediction i.e. whose \hat{y} matches with the y^* and the Intersection-over-Union (IoU) between $\hat{\mathbf{b}}$ and \mathbf{b}^* is greater than a certain threshold γ . Thus, $\mathbb{P}(m = 1)$ amounts to approximating $\mathbb{P}(\hat{y} = y^*, \hat{\mathbf{b}} = \mathbf{b}^*)$ with a certain IoU threshold γ .

Measuring miscalibration for classification and object detection: For classification, the expected calibration error (ECE) is used to measure the miscalibration of a model. The ECE measures the expected deviation of the predictive accuracy from the estimated confidence [14, 28, 37]:

$$\mathbb{E}_{\hat{s}} [|\mathbb{P}(\hat{y} = y | \hat{s} = s) - s|] \quad (3)$$

As \hat{s} is a continuous random variable, the ECE is approximated by binning the confidence space of \hat{s} into N equally spaced bins. Therefore, ECE is approximated by [37]:

$$\text{ECE} = \sum_{n=1}^N \frac{|I(n)|}{|\mathcal{D}|} \cdot |\text{acc}(n) - \text{conf}(n)| \quad (4)$$

where $|I(n)|$ is the number of examples in the n^{th} bin, and $|\mathcal{D}|$ is the total number of examples. $\text{acc}(n)$ and $\text{conf}(n)$ denote the average accuracy and average confidence in the n^{th} bin, respectively. Although the ECE measure can be used for measuring miscalibration of object detectors, it fails to reflect the calibration improvement when additional box coordinates are used for calibration since the ECE considers confidence of each example independent of the box properties to apply binning and to calculate an average precision. In this work, we use location-dependent calibration, termed as detection ECE (D-ECE). It is defined as the expected deviation of the observed precision with respect to the given box properties.

$$\mathbb{E}_{\hat{s}\hat{\mathbf{b}}} \left[\left| \mathbb{P}(m = 1 | \hat{s} = s, \hat{y} = y, \hat{\mathbf{b}} = \mathbf{b}) - s \right| \right] \quad (5)$$

Similar to ECE, the multidimensional D-ECE is calculated by partitioning both the confidence and box property spaces in each dimension k into N_k equally spaced bins. Thus, D-ECE is given by [28]:

$$\text{D-ECE}_k = \sum_{n=1}^{N_{total}} \frac{|I(n)|}{|\mathcal{D}|} \cdot |\text{prec}(n) - \text{conf}(n)| \quad (6)$$

where N_{total} is the total number of bins. $\text{prec}(n)$ and $\text{conf}(n)$ denote the average precision and confidence in each bin, respectively.

3.2. Proposed train-time calibration: MCCL

This section describes our new train-time calibration method at the core of which is an auxiliary loss function. This auxiliary loss formulation aims at jointly calibrating the multiclass confidence and bounding box localization. It is based on the fact that, the modern object detectors (based on DNNs) predict a confidence vector along with the bounding box parameters. The two key quantities to our loss function are (1) the predictive certainty in class logits and the bounding box localization and, (2) the class-wise confidence after computing class-wise logits mean (termed mean logits based class-wise confidence hereafter) and mean bounding box localization. The predictive certainty in class-wise logits is used in-tandem with the mean logits based class-wise confidence to calibrate the multi-class confidence scores. While, the predictive certainty in the bounding box prediction is used to calibrate the bounding box localization. Instead of inputting the class-wise logits and predicted bounding box parameters to the classification loss and regression loss in task-specific detection losses, we input the class-wise mean logits and mean bounding box parameters, respectively. We first describe how to compute the mean logits based class-wise confidence, mean bounding box parameters, and the certainty in both class logits and bounding box localization.

Quantifying means and certainties: For the n^{th} positive location, we aim to quantify the mean logits based class-wise confidence $\bar{\mathbf{s}}_n \in \mathbb{R}^K$ and class-wise certainty in logits $\mathbf{c}_n \in \mathbb{R}^K$ as well as the mean bounding box parameters $\{\mu_n\}_{j=1}^J$ and certainty in bounding box localization \mathbf{g}_n , where J is the number of bounding box parameters. Given an input sample (image), we perform N stochastic forward passes by applying the Monte-Carlo (MC) dropout [10]. It generates a distribution over class logits and bounding box localization. Assuming one-stage object detector (e.g., [45]), we insert a dropout layer before the classification layer and the regression layer. Let $\mathbf{z}_n \in \mathbb{R}^{N \times K}$ and $\mathbf{r}_n \in \mathbb{R}^{N \times J}$ encode the distributions over class-wise logit scores and bounding box parameters, respectively, corresponding to n^{th} positive location obtained after performing N , MC forward passes.

We obtain the mean logits based class-wise confidence $\bar{\mathbf{s}}_n \in \mathbb{R}^K$ by first taking the mean along the first dimension of \mathbf{z}_n to get class-wise mean logits and then applying the softmax. To obtain class-wise certainty \mathbf{c}_n , we first estimate the uncertainty $\mathbf{d}_n \in \mathbb{R}^K$ by computing the variance along the first dimension of \mathbf{z}_n . Then, we apply \tanh over \mathbf{d}_n and subtract it from 1 as: $\mathbf{c}_n = 1 - \tanh(\mathbf{d}_n)$, where \tanh is used to scale the uncertainty $\mathbf{d}_n \in [0, \text{inf})$ between 0 and 1.

Similarly, we estimate the certainty \mathbf{g}_n in the (predicted) bounding box parameters for the n^{th} positive location. Let $\{\sigma_n^2\}_{j=1}^J$ and $\{\mu_n\}_{j=1}^J$ be the vectors (J is the number of bbox parameters) comprised of variances and means of predicted bounding box parameters distribution \mathbf{r}_n . These variances and the means are computed along the first dimension of \mathbf{r}_n . Also, let $\mu_{n,com}$ denote the combined mean, computed as $\mu_{n,com} = \frac{1}{J} \sum_{j=1}^J \mu_{n,j}$. Then, we estimate the (joint) uncertainty u_n as:

$$u_n = \frac{1}{J} \sum_{j=1}^J [\sigma_{n,j}^2 + (\mu_{n,j} - \mu_{n,com})^2]. \quad (7)$$

The certainty \mathbf{g}_n in the n^{th} positive bounding box localization is then computed as: $\mathbf{g}_n = 1 - \tanh(u_n)$.

We leverage these estimated mean logits based class-wise confidence, class-wise certainty and the certainty in bounding box localization to formulate the two components of our auxiliary loss: multi-class confidence calibration (MCC), and localization calibration (LC). For MCC, we compute the difference between the fused mean confidence and certainty with the accuracy. For LC, we calculate the deviation between the predicted bounding box overlap and the predictive certainty of the bounding box. Both quantities are computed over the mini-batch during training.

Multi-class confidence calibration (MCC): To achieve multi-class confidence calibration, we leverage the mean logits based class-wise confidence and class-wise certainty and fuse them by computing class-wise mean. The resulting vector is termed as the multiclass fusion of mean confidence and certainty. Then, we calculate the absolute difference between the fused vector and the accuracy as:

$$\mathcal{L}_{MCC} = \frac{1}{K} \sum_{k=1}^K \left| \frac{1}{M} \sum_{l=1}^{N_b} \sum_{n=1}^{N_{pos}} \mathbf{v}_{l,n}[k] - \frac{1}{M} \sum_{l=1}^{N_b} \sum_{n=1}^{N_{pos}} \mathbf{q}_{l,n}[k] \right| \quad (8)$$

where $M = N_b \times N_{pos}$. N_b is the number of samples in the minibatch and N_{pos} represents the number of positive locations. $\mathbf{q}_{l,n}[k] = 1$ if k is the ground truth class of the bounding box predicted for the n^{th} location in the l^{th} sample. $\mathbf{v}_{l,n}[k] = (\bar{\mathbf{s}}_{l,n}[k] + \mathbf{c}_{l,n}[k])/2$, where $\bar{\mathbf{s}}_{l,n}[k]$ and $\mathbf{c}_{l,n}[k]$ are the mean confidence and the certainty, respectively, for the class k of the n^{th} positive location in the l^{th} sample. The \mathcal{L}_{MCC} is capable of calibrating the confidence of both the predicted label and non-predicted labels. It penalizes the model if, for a given class k , the fusion (of mean

logits based class-wise confidence and certainty in class-wise logits) across minibatch deviates from the average occurrence of this class across minibatch.

Localization calibration (LC): We calibrate the localization component by leveraging the certainty in bounding box prediction. Next, we compute the absolute difference between the bounding box overlap (with the ground truth) and the certainty in the bounding box prediction.

$$\mathcal{L}_{LC} = \frac{1}{N_b} \sum_{l=1}^{N_b} \frac{1}{N_{pos}^l} \sum_{n=1}^{N_{pos}^l} \left| \left[\text{IoU}(\hat{\mathbf{b}}_{n,l}, \mathbf{b}_{n,l}^*) - g_{n,l} \right] \right| \quad (9)$$

where N_{pos}^l denotes the number of positive bounding box regions in the l^{th} sample, and $g_{n,l}$ is the certainty in the n^{th} positive bounding box prediction from l^{th} sample. Both \mathcal{L}_{MCC} and \mathcal{L}_{LC} operate over the mini-batches, and we combine them to get our new auxiliary loss term $\mathcal{L}_{MCCL-aux}$:

$$\mathcal{L}_{MCCL-aux} = \mathcal{L}_{MCC} + \beta \mathcal{L}_{LC} \quad (10)$$

where β is a hyperparameter to control the relative contribution of \mathcal{L}_{LC} to the overall loss $\mathcal{L}_{MCCL-aux}$.

4. Experiments

Datasets: To evaluate the in-domain calibration performance, we use the following five datasets: Sim10K [24], KITTI [11], Cityscapes (CS) [6], COCO [31], and PASCAL VOC(2012) [9]. **Sim10K** [24] contains synthetic images of the car category, and offers 10K images which are split into 8K for training, 1K for validation and 1K for testing. **Cityscapes** [6] is an urban driving scene dataset and consists of 8 object categories. It has 2975 training images and 500 validation images, which are used for evaluation. **KITTI** [11] is similar to Cityscapes as it contains images of road scenes with a wide view of the area, except that KITTI images were captured with a different camera setup. Following prior works, we consider car class for experiments. We use train2017 version of **MS-COCO** [31] and it offers 118K training images, 5K validation images, and 41K test images. **PASCAL VOC 2012** [9] consists of 5,717 training and 5,823 validation images, and provides bounding box annotations for 20 classes. For evaluating out-of-domain calibration performance, we use Sim10K to CS, KITTI to CS, CS to Foggy-CS, COCO to Cor-COCO, CS to BDD100K [49], VOC to Clipart1k [22], VOC to Watercolor2k [22], and VOC to Comic2k [22]. **Foggy Cityscapes** (CS-F) [42] dataset is developed using Cityscapes dataset [6] by simulating foggy weather leveraging the depth maps in Cityscapes with three levels of foggy weather. **Cor-COCO** is a corrupted version of MS-COCO val2017 dataset for out-of-domain evaluation, and is constructed by introducing random corruptions with severity levels defined in [19]. **Clipart1k** [22] contains 1K images,

which are split into 800 for training and 200 for validation, and shares 20 object categories with PASCAL VOC. Both **Comic2k** [22] and **Watercolor2k** [22] are comprised of 1K training images and 1K test images, and share 6 categories with Pascal VOC. **BDD100k** [48] offers 70K training images, 20K test images and 10K validation images. We use validation set for out-of-domain evaluation.

Implementation Details: For all experiments, we use Tesla V100 GPUs. For COCO experiments, we use 8 GPUs and follow training configurations reported in [45]. For experiments on all other datasets, we utilize 4 GPUs and follow training configurations listed in [21]. We chose β in Equation (10) from $\{0.01, 1\}$. For further training details, we refer to the supplementary material.

Evaluation metrics: We use D-ECE metric defined in Equation (6) at IoU of 0.5 to measure calibration performance. Note that, in addition to classification scores, it takes into account the calibration of center-x, center-y, width, and height of the predicted box. For reporting detection performance, we use mAP and AP@0.5 metrics.

Baselines: We evaluate our train-time calibration method against models trained with task-specific losses of a CNN-based object detector, namely FCOS [45], and ViT-based object detector, namely Deformable DETR [52]. We then compare with the temperature scaling post-hoc method and further with the recently proposed auxiliary loss functions for classification, including MDCA [18] and AvUC [25].

4.1. Results

In-domain experiments: We compare the in-domain performance on five challenging datasets with the models trained with task-specific loss of FCOS in Table 1. The results reveal that our train-time calibration method (MCCL) consistently improves the calibration performance of the task-specific losses. Notably, when added to the task-specific loss of FCOS, our MCCL reduces the D-ECE by 5.86% and 1.76% in VOC and CS datasets, respectively.

Out-of-domain experiments: Table 2 and Table 3 report out-of-domain performance on eight challenging shifts. We see that our MCCL is capable of consistently improving the calibration performance in all shift scenarios. We notice a major decrease in D-ECE of 2.91% in Sim10K to CS shift. Similarly, we observe a reduction in D-ECE by a visible margin of 2.47% for CS to CS-foggy (CS-F).

Comparison with post-hoc method: We choose temperature scaling (TS) as post-hoc calibration for comparison. The temperature parameter T is optimized using a hold-out validation set to re-scale the logits of the trained model (FCOS). Table 5 compares the performance of TS with our method (MCCL) on COCO, Sim10K, CS and COCO corrupted datasets. We note that TS performs inferior to our method and to baseline. This could be because when there are multiple dense prediction maps, as in FCOS, it is likely

In-domain performance										
Methods	Sim10K		KITTI		CS		COCO		VOC	
	D-ECE	AP@0.5	D-ECE	AP@0.5	D-ECE	AP@0.5	D-ECE	AP@0.5	D-ECE	mAP
Baseline (FCOS)	12.90	87.45	9.54	94.54	9.40	70.48	15.42	54.91	11.88	59.68
Ours (MCCL)	11.18	86.47	7.79	93.76	7.64	70.22	14.94	54.85	6.02	59.17

Table 1. In-domain calibration performance (in D-ECE%) on five challenging datasets, including Sim10K, KITTI, Cityscapes (CS), COCO and VOC. Best results are in bold.

Out-of-domain performance										
Methods	Sim10K → CS		KITTI → CS		CS → CS-F		COCO → Cor-COCO		CS → BDD100K	
	D-ECE	AP@0.5	D-ECE	AP@0.5	D-ECE	AP@0.5	D-ECE	AP@0.5	D-ECE	AP@0.5
Baseline (FCOS)	9.51	45.18	7.53	38.11	11.18	19.81	15.90	30.01	18.82	14.18
Ours (MCCL)	6.60	44.30	6.43	38.73	8.97	19.54	14.45	29.96	16.12	14.20

Table 2. Out-of-domain calibration performance (in D-ECE%) on five challenging domain shifts.

Out-of-domain performance						
Methods	VOC → clipart		VOC → watercolor		VOC → comic	
	D-ECE	mAP	D-ECE	mAP	D-ECE	mAP
Baseline (FCOS)	1.54	14.57	3.23	24.23	3.75	9.89
Ours (MCCL)	1.06	13.71	2.33	28.70	2.84	11.50

Table 3. Out-of-domain calibration performance on three challenging domain drifts.

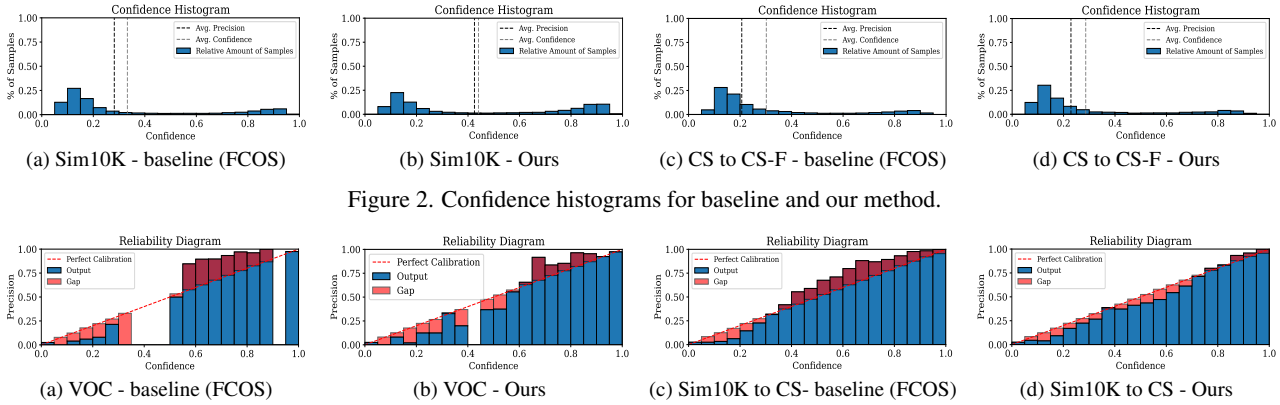


Figure 2. Confidence histograms for baseline and our method.

Figure 3. Reliability diagrams for baseline and our method.

that a single temperature parameter T will not be optimal for the corresponding logit vectors.

Test accuracy/precision: We note that in addition to consistently reducing D-ECE, our MCCL also preserves the mAP or AP@0.5 in almost all cases. In the in-domain experiments (Table 1), the maximum reduction in AP@0.5 is only 0.98% in the Sim10K dataset. In the out-of-domain experiments (Table 2 & Table 3), it mostly remains same in KITTI to CS, CS to BDD100K, VOC to watercolor, and VOC to comic shifts.

Overcoming under/overconfidence: We plot confidence histogram (Figure 2) and reliability diagrams (Figure 3) to illustrate the effectiveness of our method in mitigating over-

confidence or underconfidence. In confidence histograms (Figure 2) from Sim10K in-domain and CS to CS-F out-of-domain datasets, the average confidence is greater than the average precision which indicates the overconfident model. Our method reduces this gap in both scenarios compared to the baseline (FCOS) method and alleviates the overconfidence of the baseline. Similarly, the reliability diagrams (Figure 3) for VOC in-domain and Sim10K to CS domain shifts reveal that our method can mitigate both underconfident and overconfident predictions by a visible margin.

Confidence values of incorrect detections: We analyse the confidence of our method in case of incorrect predictions (Figure 4). Compared to baseline, our method is capable

Comparison with Deformable DeTR (Baseline)				
Dataset	Baseline		Ours	
	D-ECE	AP@0.5	D-ECE	AP@0.5
Sim10K	7.51	89.85	6.36	89.63
KITTI	6.31	96.76	3.87	96.55
CS	9.74	71.03	9.69	70.63
COCO	7.92	62.93	7.89	62.95
VOC	6.13	66.83	5.65	65.67
Sim10K → CS	7.28	49.65	6.79	51.55
KITTI → CS	12.93	29.93	12.57	29.62
CS → CS-F	9.54	25.73	9.39	25.39
COCO → Cor-COCO	6.77	35.51	6.71	35.02
VOC → Clipart	4.32	16.40	3.55	18.11
VOC → Watercolor	6.60	27.23	6.56	26.51
VOC → Comic	6.49	9.41	6.31	9.54

Table 4. Comparison of calibration performance of models trained with ViT-based object detector (Deformable DeTR) method and after integrating our method (MCCL).

of reducing the confidence of incorrect predictions over the whole spectrum of confidence range.

With another baseline: Table 4 reports results with ViT-based object detector, namely Deformable DETR [52]. Compared to FCOS, the Deformable DETR, is already a relatively strong baseline in calibration error. We observe that our MCCL reduces the calibration error (D-ECE) for both in-domain and out-of-domain predictions. The major improvement (2.44% reduction in D-ECE) in calibration performance is observable for KITTI in-domain predictions.




Figure 4. Histogram of confidence values for incorrect predictions by baseline (FCOS) and our method in COCO dataset.

4.2. Ablation study

Impact of each component in MCCL: We report the result of ablation experiments for validating the performance contribution of different components in our method (MCCL) (Table 5). Moreover, we report the calibration performance of two train-time calibration losses for image classification: MDCA [18] and AvUC [25]. We can observe the following trends from Table 5. The calibration performance of our MCCL is not due to providing only the class-wise mean logits and mean bounding box parameters to classification




Figure 5. D-ECE convergence for baseline, the classification and localization components of our MCCL, and MCCL.

loss and regression loss of detection-specific loss, respectively, (ours w/o \mathcal{L}_{LC} & \mathcal{L}_{MCC}). Both \mathcal{L}_{MCC} and \mathcal{L}_{LC} are integral components of our method (MCCL). They are complementary to each other and their proposed combination is vital to delivering the best calibration performance. For instance, in Sim10K to CS shift, the proposed combination of \mathcal{L}_{MCC} and \mathcal{L}_{LC} achieves a significant reduction in D-ECE compared to MCC and LC alone. Further, the classification-based calibration losses are sub-optimal for calibrating object detection methods.

D-ECE convergence: Figure 5 compares the convergence of D-ECE for baseline, the two components (classification and localization) of our method (MCCL), and MCCL. Although our MCCL and its two constituents does not directly optimize the D-ECE metric, they provide improved D-ECE convergence compared to the baseline.

Impact on location-dependent calibration: Figure 6 and Figure 7 depict that miscalibration error (D-ECE) relies highly on the relative object location (c_x, c_y) and/or its relative width and height (w, h). Moreover, it tends to increase as we approach image boundaries. Figure 6 plots the precision, confidence and D-ECE over individual parameters i.e. c_x . Figure 7 plots 2D calibration heatmaps over object location and width/height, where each location in a heatmap represents D-ECE. Both figures show that, compared to baseline, besides other locations, our MCCL can decrease D-ECE at image boundaries. Figure 6 also shows that, compared to baseline, our MCCL allows the adaptation of confidence score at all image locations differently by adjusting the shape of confidence curve accordingly.

MCDO overhead & its Tradeoff analysis: Table 6 reveals that, in our implementation, upon increasing Monte-Carlo dropout (MCDO) passes $N=\{3, 5, 10, 15\}$, there is a *little overhead* in time cost over $N=1$. Table 7 shows the impact of varying the number of MC dropout passes (N) on calibration performance. Upon increasing the N , we see improved calibration, especially in OOD scenario.

Methods	Ablation study							
	COCO → COCO		Sim10K → Sim10K		Sim10K → CS		COCO → Cor-COCO	
	D-ECE	AP@0.5	D-ECE	AP@0.5	D-ECE	AP@0.5	D-ECE	AP@0.5
Baseline (FCOS)	15.42	54.91	12.90	88.01	9.51	45.18	15.90	30.01
Post-hoc (TS)	17.34	54.77	17.99	29.97	14.66	87.29	24.03	45.91
AvUC [25]	15.17	54.73	13.24	87.91	10.64	39.39	15.78	29.75
MDCA [18]	15.25	54.44	13.99	87.95	11.14	39.51	15.33	30.07
Ours (w/o \mathcal{L}_{LC} & \mathcal{L}_{MCC})	15.26	54.25	<u>13.12</u>	86.34	8.87	42.89	15.28	29.67
Ours (w/o \mathcal{L}_{LC})	<u>15.00</u>	54.11	13.00	87.86	<u>8.63</u>	45.31	<u>14.52</u>	29.73
Ours (w/o \mathcal{L}_{MCC})	15.12	54.40	<u>12.86</u>	87.24	9.12	41.14	15.54	29.61
Ours (MCCL)	14.94	54.85	11.18	86.47	6.60	44.29	14.45	29.96

Table 5. Ablations in MCCL and comparison of MCCL with TS, and classification-based train-time losses: MDCA [18] and AvUC [25].

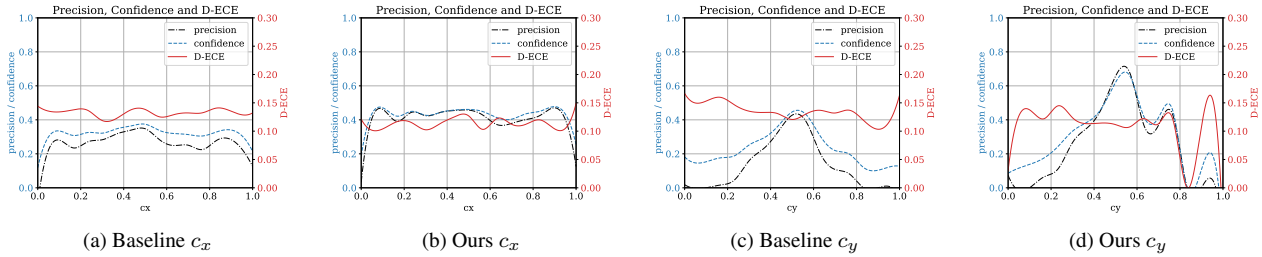


Figure 6. Calibration precision, confidence, and ECE with IOU @0.5 of (a) baseline (FCOS) relative to center-x (c_x), (b) our method relative to center-x (c_x) (c) baseline (FCOS) relative to center-y (c_y), and (d) our method relative to center-y (c_y).

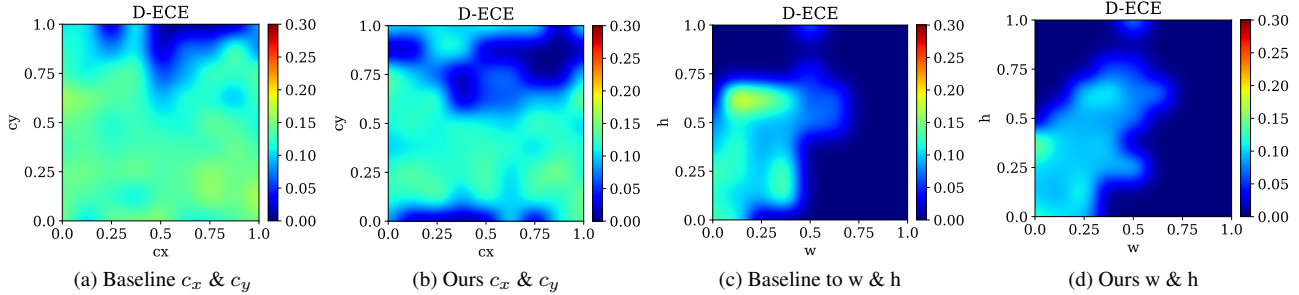


Figure 7. Calibration heatmap of (a,b) baseline (FCOS) and our method over center-x (c_x) and center-y (c_y) with IOU @0.5, (c,d) baseline (FCOS) and our method over width (w) and height (h) with IOU @0.5.

Time cost with MC dropout					
N	N=1	N=3	N=5	N=10	N=15
Time per iteration(s)	0.143	0.197	0.243	0.350	0.463
Increment	-	0.372	0.694	1.442	2.230

Table 6. Increment in time per iteration (in secs) over N=1 upon increasing MC forward passes.

Calibration performance with MC dropout						
Method	Metric	baseline	N=3	N=5	N=10	N=15
COCO	D-ECE	16.57	16.27	16.25	16.19	16.12
	AP@0.5	52.34	51.00	51.49	51.23	51.89
COCO-Corr.	D-ECE	17.27	16.18	16.21	15.87	16.17
	AP@0.5	29.25	28.24	28.59	28.31	28.80

Table 7. Impact on D-ECE after increasing N in MC dropout.

5. Conclusion

Very little to no attempts have been made towards studying the calibration of object detectors. In this paper, we explored this direction and presented a new train-time technique for calibrating DNN-based object detection methods. At the core of our method is an auxiliary loss which aims

at jointly calibrating multiclass confidence and box localization after leveraging their predictive uncertainties. Extensive experiments reveal that our method can consistently reduce the calibration error of object detectors from two different DNN-based object detection paradigms for both in-domain and out-of-domain detections.

References

- [1] José M Bernardo and Adrian FM Smith. *Bayesian theory*, volume 405. John Wiley & Sons, 2009. [3](#)
- [2] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International conference on machine learning*, pages 1613–1622. PMLR, 2015. [3](#)
- [3] Ondrej Bohdal, Yongxin Yang, and Timothy Hospedales. Meta-calibration: Meta-learning of model calibration using differentiable expected calibration error. *arXiv preprint arXiv:2106.09613*, 2021. [2](#)
- [4] Glenn W Brier et al. Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78(1):1–3, 1950. [2](#)
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. [1](#)
- [6] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [5](#)
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. [1](#)
- [8] Michael W Dusenberry, Dustin Tran, Edward Choi, Jonas Kemp, Jeremy Nixon, Ghassen Jerfel, Katherine Heller, and Andrew M Dai. Analyzing the role of model uncertainty for electronic health records. In *Proceedings of the ACM Conference on Health, Inference, and Learning*, pages 204–213, 2020. [1](#)
- [9] M. Everingham, S. M. A. Eslami, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111(1):98–136, Jan. 2015. [5](#)
- [10] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016. [3, 4](#)
- [11] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 3354–3361. IEEE, 2012. [5](#)
- [12] Arthur Gretton. Introduction to rkhs, and some simple kernel algorithms. *Adv. Top. Mach. Learn. Lecture Conducted from University College London*, 16:5–3, 2013. [2](#)
- [13] Sorin Grigorescu, Bogdan Trasnea, Tiberiu Cocias, and Gigel Macesanu. A survey of deep learning techniques for autonomous driving. *Journal of Field Robotics*, 37(3):362–386, 2020. [1](#)
- [14] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International Conference on Machine Learning*, pages 1321–1330. PMLR, 2017. [1, 3](#)
- [15] Marton Havasi, Rodolphe Jenatton, Stanislav Fort, Jeremiah Zhe Liu, Jasper Snoek, Balaji Lakshminarayanan, Andrew M Dai, and Dustin Tran. Training independent subnetworks for robust prediction. *arXiv preprint arXiv:2010.06610*, 2020. [1](#)
- [16] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask r-cnn. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. [1](#)
- [17] 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. [1](#)
- [18] Ramya Hebbalaguppe, Jatin Prakash, Neelabh Madan, and Chetan Arora. A stitch in time saves nine: A train-time regularizing loss for improved neural network calibration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16081–16090, June 2022. [1, 3, 5, 7, 8](#)
- [19] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *International Conference on Learning Representations (ICLR)*, 2019. [5](#)
- [20] José Miguel Hernández-Lobato and Ryan Adams. Probabilistic backpropagation for scalable learning of bayesian neural networks. In *International conference on machine learning*, pages 1861–1869. PMLR, 2015. [3](#)
- [21] Cheng-Chun Hsu, Yi-Hsuan Tsai, Yen-Yu Lin, and Ming-Hsuan Yang. Every pixel matters: Center-aware feature alignment for domain adaptive object detector. In *European Conference on Computer Vision*, pages 733–748. Springer, 2020. [5](#)
- [22] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [5](#)
- [23] Mobarakol Islam, Lalithkumar Seenivasan, Hongliang Ren, and Ben Glocker. Class-distribution-aware calibration for long-tailed visual recognition. *arXiv preprint arXiv:2109.05263*, 2021. [2](#)
- [24] Matthew Johnson-Roberson, Charles Barto, Rounak Mehta, Sharath Nittur Sridhar, Karl Rosaen, and Ram Vasudevan. Driving in the matrix: Can virtual worlds replace human-generated annotations for real world tasks? In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 746–753. IEEE, 2017. [5](#)
- [25] Ranganath Krishnan and Omesh Tickoo. Improving model calibration with accuracy versus uncertainty optimization. *Advances in Neural Information Processing Systems*, 2020. [5, 7, 8](#)
- [26] Meelis Kull, Telmo Silva Filho, and Peter Flach. Beta calibration: a well-founded and easily implemented improvement on logistic calibration for binary classifiers. In *Artificial Intelligence and Statistics*, pages 623–631. PMLR, 2017. [2, 3](#)

- [27] Aviral Kumar, Sunita Sarawagi, and Ujjwal Jain. Trainable calibration measures for neural networks from kernel mean embeddings. In *International Conference on Machine Learning*, pages 2805–2814. PMLR, 2018. 1, 2, 3
- [28] Fabian Küppers, Jan Kronenberger, Amirhossein Shantia, and Anselm Haselhoff. Multivariate confidence calibration for object detection. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020. 3, 4
- [29] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017. 3
- [30] Gongbo Liang, Yu Zhang, Xiaoqin Wang, and Nathan Jacobs. Improved trainable calibration method for neural networks on medical imaging classification. In *British Machine Vision Conference (BMVC)*, 2020. 2
- [31] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014. 5
- [32] Tsung-Yi Lin, Priyank Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *IEEE transactions on pattern analysis and machine intelligence*, 2018. 2
- [33] Bingyuan Liu, Ismail Ben Ayed, Adrian Galdran, and Jose Dolz. The devil is in the margin: Margin-based label smoothing for network calibration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 80–88, June 2022. 1, 2, 3
- [34] Christos Louizos and Max Welling. Structured and efficient variational deep learning with matrix gaussian posteriors. In *International conference on machine learning*, pages 1708–1716. PMLR, 2016. 3
- [35] Jishnu Mukhoti, Viveka Kulharia, Amartya Sanyal, Stuart Golodetz, Philip Torr, and Puneet Dokania. Calibrating deep neural networks using focal loss. *Advances in Neural Information Processing Systems*, 33:15288–15299, 2020. 1, 2
- [36] Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? *Advances in neural information processing systems*, 32, 2019. 2
- [37] Mahdi Pakdaman Naeini, Gregory Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015. 3
- [38] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems*, 32, 2019. 1, 3
- [39] Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. Regularizing neural networks by penalizing confident output distributions. *arXiv preprint arXiv:1701.06548*, 2017. 2
- [40] John Platt et al. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3):61–74, 1999. 2
- [41] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 1
- [42] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision*, 126(9):973–992, 2018. 5
- [43] Monika Sharma, Oindrila Saha, Anand Sriraman, Ramya Hebbalaguppe, Lovekesh Vig, and Shirish Karande. Crowdsourcing for chromosome segmentation and deep classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 34–41, 2017. 1
- [44] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016. 2
- [45] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 9627–9636, 2019. 1, 2, 4, 5
- [46] Christian Tomani, Sebastian Gruber, Muhammed Ebrar Erdem, Daniel Cremers, and Florian Buettner. Post-hoc uncertainty calibration for domain drift scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10124–10132, 2021. 2
- [47] Florian Wenzel, Jasper Snoek, Dustin Tran, and Rodolphe Jenatton. Hyperparameter ensembles for robustness and uncertainty quantification. *Advances in Neural Information Processing Systems*, 33:6514–6527, 2020. 1, 3
- [48] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2636–2645, 2020. 5
- [49] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving video database with scalable annotation tooling. *arXiv preprint arXiv:1805.04687*, 2(5):6, 2018. 5
- [50] Ronald Yu and Gabriele Spina Ali. What’s inside the black box? ai challenges for lawyers and researchers. *Legal Information Management*, 19(1):2–13, 2019. 1
- [51] Zhilu Zhang, Adrian V Dalca, and Mert R Sabuncu. Confidence calibration for convolutional neural networks using structured dropout. *arXiv preprint arXiv:1906.09551*, 2019. 3
- [52] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable {detr}: Deformable transformers for end-to-end object detection. In *International Conference on Learning Representations*, 2021. 1, 2, 5, 7