

# Beyond Classification: Definition and Density-based Estimation of Calibration in Object Detection

Teodora Popordanoska  
 ESAT-PSI, KU Leuven

teodora.popordanoska@kuleuven.be

Aleksei Tiulpin  
 HST Research Unit, University of Oulu

aleksei.tiulpin@oulu.fi

Matthew B. Blaschko  
 ESAT-PSI, KU Leuven

matthew.blaschko@esat.kuleuven.be

## Abstract

Despite their impressive predictive performance in various computer vision tasks, deep neural networks (DNNs) tend to make overly confident predictions, which hinders their widespread use in safety-critical applications. While there have been recent attempts to calibrate DNNs, most of these efforts have primarily been focused on classification tasks, thus neglecting DNN-based object detectors. Although several recent works addressed calibration for object detection and proposed differentiable penalties, none of them are consistent estimators of established concepts in calibration. In this work, we tackle the challenge of defining and estimating calibration error specifically for this task. In particular, we adapt the definition of classification calibration error to handle the nuances associated with object detection, and predictions in structured output spaces more generally. Furthermore, we propose a consistent and differentiable estimator of the detection calibration error, utilizing kernel density estimation. Our experiments demonstrate the effectiveness of our estimator against competing train-time and post-hoc calibration methods, while maintaining similar detection performance.

## 1. Introduction

Calibration is a property of a model, which directly translates to the ability to estimate its own predictive uncertainty, thereby facilitating safe and responsible deployment. Intuitively, a well-calibrated model produces confidence scores that accurately reflect the uncertainty associated with its predictions. In deep neural networks (DNNs), which are driving most of the current progress in machine learning and computer vision, this has become an emerging concern, as they can yield wrong predictions with high con-




Figure 1. Calibration error of RetinaNet on Pascal VOC. The model without calibration (Base) is compared with post-hoc (temperature scaling; TS) and train-time calibration methods (TCD and Ours). Our KDE-based estimator effectively reduces the calibration error. The error bars represent 95% CI.

fidence [8, 28]. Unreliable uncertainty estimates produced by such models can lead to erroneous decision-making, which is especially risky in fields like autonomous driving [40] and medicine [6].

The field of model calibration typically studies three sub-problems: *estimating* calibration error, *regularization* during training, and *post-hoc* calibration. While in recent years there have been multiple advances in all of these domains, most of the existing literature studies calibration in the context of classification. Calibration in structured prediction problems, such as object detection (OD), have received substantially less attention, despite their increased integration in many safety-critical applications. While some attempts have been made to define, measure, enforce, and adjust calibration [15, 24, 25, 27], we argue that the field of OD lacks a solid mathematical foundation in both definition of calibration and estimation of calibration errors.

In spite of its importance, defining calibration for object

detection is a challenging problem due to the variability of DNN-based detectors. In particular, there are ambiguities related to how many bounding boxes are returned, how to select a detection confidence threshold below which detections are rejected, what is the intersection over union (IoU) threshold for considering a “correct” detection, and so forth. There have been a few recent attempts to define calibration by either replacing accuracy with precision in the definition of calibration for classification [15, 25, 29], or by requiring that both classification and localization performances jointly match the confidence score [27].

In this paper, we propose a general framework that unifies existing definitions of calibration in OD, and is flexible to parametrize the notion of a “correct” detection. Beyond defining calibration, we derive a *consistent* and *differentiable* estimator of calibration error in OD, relying on a kernel density estimator (KDE) recently proposed for classification [31], and recommended among the best practices [21] for assessing calibration of classifiers. Due to its consistency and differentiability, our estimator of calibration error can be used not only as a reliable tool to assess calibration, but also in calibration-regularized training of popular detectors. We perform extensive experiments on MS COCO [18], Cityscapes [4], and PASCAL VOC [7], and demonstrate that our estimator, when incorporated as an auxiliary loss during training, consistently reduces the calibration error across several object detectors. Furthermore, we show that finetuning with our loss is an effective way to improve calibration, without sacrificing accuracy and without the need to re-train the models from scratch.

In summary, our contributions are:

1. We propose a general unifying definition of calibration in OD, which addresses the nuances related to assessing the “correctness” of a detection.
2. We develop a consistent and differentiable estimator of calibration error based on KDE [31] for OD.
3. We perform rigorous empirical analysis of calibration of popular object detectors on several datasets.
4. We demonstrate that our estimator can be used as an auxiliary train-time loss, which effectively reduces the calibration error, while maintaining similar average precision (AP).

## 2. Related work

### 2.1. Calibration in classification

Most of the prior work on developing techniques for improving calibration in the vision domain target the image classification task. Post-hoc methods [8, 20, 41, 42] rescale the output of a trained model using parameters learned on a validation set. The most successful method of this group is temperature scaling [8], where a single parameter  $T$  is learned, usually by minimizing NLL, to scale the

logits before applying the softmax function. Train-time calibration methods [10, 14, 17, 22, 23, 31] typically incorporate calibration-related auxiliary loss directly into the training process of a neural network. Even though these methods may introduce computational overhead, they do not require a hold-out validation set and have demonstrated superior performance in handling domain shift [12], compared to post-hoc methods.

### 2.2. Calibration in object detection

Recent works have shown that miscalibration is a concern not only in classification tasks, but also in the domain of object detection [15, 29]. Pathiraja et al. [29] showed that existing calibration methods for classification are not as effective for calibrating object detectors. As a result, there have been a few efforts to establish the definition of calibration [15, 27] and to devise techniques to mitigate the issue. Oksuz et al. [27] propose a binning-based metric called Localisation-aware Expected Calibration Error (LaECE), and use it as part of existing post-hoc calibration approaches like linear regression, histogram binning [42] and isotonic regression [41]. In the realm of train-time calibration methods, several auxiliary loss terms have been proposed to be used in addition to the detection-specific loss. For instance, TCD [24] and MCCL [29] were proposed to jointly calibrate the class-wise confidences and localization performance. Another recently introduced auxiliary loss is BPC [25], which is based on a heuristic that maximizes confidence scores for accurate predictions while minimizing scores for inaccurate predictions. Harakeh and Waslander [9] propose an energy score to train probabilistic detectors, which is empirically shown to improve calibration.

However, none of the proposed binning-based metrics and trainable auxiliary losses fulfill both requirements of being consistent and differentiable estimators of calibration error in the context of object detection. In this paper, we derive a novel estimator for calibration error in object detection with those desired properties.

## 3. Methods

### 3.1. Defining calibration

Several notions of different strength exist for multi-class calibration, including top-label, class-wise, and canonical calibration [8, 36, 37]. However, following recent trends of training the classification branch of object detectors in a multi-label setting, i.e. by using  $K$  independent binary classifiers [19, 32, 35], we also favor the approach of defining calibration by considering  $K$  binary object detectors, as the mutual-exclusion principle of multi-class classification does not hold (e.g. the area of an image occupied by a person may overlap with that of a bicycle).

Intuitively, calibration requires that the confidence score

of a prediction should be aligned with some notion of correctness. In classification tasks, a correct prediction is determined by the match between the predicted and ground truth labels. However, in object detection, assessing correctness is more nuanced, since it involves evaluating the degree of overlap between two sets of pixels (e.g. bounding boxes). To account for this, we will introduce a similarity measure (e.g. IoU) and a so-called link function, so that we can define a family of notions of calibration that also consider the degree of correctness of a given prediction. The formal definitions are given below.

**Binary classification** Let  $X \in \mathcal{X}$  and  $Y \in \mathcal{Y} = \{0, 1\}$  be random variables denoting the input and target, given by a joint distribution  $P$ . Let  $f$  be a neural network with  $f(X) = \hat{S}$ , where  $\hat{S} \in [0, 1]$  denotes the model’s confidence that the label is 1. Then  $f$  is said to satisfy *binary calibration* [2, 13, 41] if:

$$\mathbb{P}(Y = 1 \mid \hat{S} = s) = s, \quad \forall s \in [0, 1]. \quad (1)$$

**Object detection** Given a set of images  $\{X_i\}_{1 \leq i \leq n}$ , the goal of an object detector is to predict bounding boxes, class labels and confidence scores for the objects in  $\{X_i\}$ . Let  $g_k$  be a binary object detector for class  $k$  with  $g_k(X_i) = \{(\hat{S}_{ij}, \hat{B}_{ij})\}_{1 \leq j \leq m_{ik}}$ , where  $m_{ik}$  boxes are predicted for image  $X_i$  and class  $k$ ,  $\hat{B}_{ij} \in \mathcal{B} = \mathbb{R}^4$  denotes a predicted bounding box,<sup>1</sup> and  $\hat{S}_{ij} \in [0, 1]$  its corresponding score.

Let  $L : \mathcal{B} \times \mathcal{B} \rightarrow [0, 1]$  denote a similarity measure, for example IoU, Dice score, Hamming similarity etc. Let  $\psi : [0, 1] \rightarrow [0, 1]$  denote a monotonic function, which we will refer to as a *link function*. Some choices for the link function in the context of object detectors may include an identity function  $\psi(L) = L$ , or a piece-wise linear ramp function parametrized with  $\alpha \leq \beta \in [0, 1]$  as:

$$\psi(L) = \begin{cases} 0 & L \leq \alpha \\ \frac{L-\alpha}{\beta-\alpha} & \alpha < L < \beta \\ 1 & L \geq \beta \end{cases}. \quad (2)$$

Two special cases of this function that could be of interest are a step (threshold) function ( $0 < \alpha = \beta \leq 1$ ), and a piece-wise linear (hinge) function ( $\alpha = 0.5, \beta = 1$ ).

Let  $Z := \psi(L(\hat{B}, B^*))$  be a random variable that denotes the “correctness“ of a detection, and  $B^*$  be the ground-truth box that  $\hat{B}$  matches with. If  $\psi$  is a threshold function then  $Z \in \{0, 1\}$  denotes whether the predicted bounding box matched a ground truth box with  $L \geq \beta$ , and if  $\psi$  is an identity function then  $Z \in [0, 1]$  represents the degree of “correctness“ of the detection, e.g. the IoU between

<sup>1</sup>We note that in general this can be any set of pixels, not necessarily limited to the ones defined by a bounding box. We adopt this notation for convenience, but it does not change the mathematical principles involved.

predicted and ground truth box. Then we define *calibration for object detection* as:

$$\mathbb{P}(Z \mid \hat{S} = s) = s, \quad \forall s \in [0, 1]. \quad (3)$$

The introduction of a link function  $\psi$ , provides a concise way to relate the notion of calibration to existing literature. For example, by letting  $\psi$  be a threshold function and taking  $L = \text{IoU}$ , we recover the definition of calibration used in [15, 24, 25, 29], i.e.,  $\mathbb{P}(Z = 1 \mid \hat{S} = s) = s, \forall s \in [0, 1]$ , with  $Z = 1$  denoting an accurate detection in the sense that  $\text{IoU}(\hat{B}, B^*) \geq \beta$ . Similarly, by using an identity function for  $\psi$ , we obtain the definition proposed in [27, Equation (3)], i.e.,  $\mathbb{P}(Z = \text{IoU}(\hat{B}, B^*) \mid \hat{S} = s) = s, \forall s \in [0, 1]$ . Note that in both cases the requirement that the class prediction matches the ground truth class is simplified since we’re considering a binary detector and all predictions share the same class as the ground truth.

### 3.2. Measuring calibration

A predictive function trained with risk minimization on a continuous domain will be exactly calibrated with probability 0. We therefore need to measure a notion of calibration error. Analogous to the definition of calibration error for classification [8], we define the  $L_1$  *calibration error for a binary object detector*  $g_k$  for the  $k$ th class as:

$$\text{CE}(g_k) = \mathbb{E} \left[ \left| \mathbb{E}[Z \mid \hat{S} = s] - s \right| \right]. \quad (4)$$

**Estimating calibration error** In practice, we rely on finite samples to estimate the quantities of interest. One of the desired properties of such estimators is **consistency**, i.e., the estimator  $\widehat{\text{CE}}(g_k)$  should converge in probability to the true value  $\text{CE}(g_k)$  [11]:  $\text{plim}_{n \rightarrow \infty} \widehat{\text{CE}}(g_k) = \text{CE}(g_k)$ . Another desired property for calibration error estimators is **differentiability**, as it allows the estimator to be directly optimized alongside the task-specific loss function. Before presenting our estimator, which possesses both properties, we will introduce existing estimators and discuss their shortcomings.

In classification, one of the most widely used estimators of calibration error is the  $L_1$  binned estimator, commonly referred to as ECE [8]. Recent studies have proposed extensions of this estimator to the detection task, called detection ECE (D-ECE) [15] and Localization-aware ECE [27] (LaECE), to be used as finite sample estimators of the notion of calibration as defined by a threshold link and an identity link, respectively. Analogous to ECE, both D-ECE and LaECE partition the predictions into  $M$  equally-spaced bins. D-ECE is computed by taking a weighted average of the difference between precision and confidence in each bin:

$$\text{D-ECE} = \sum_{m=1}^M \frac{|\mathcal{D}_m|}{|\mathcal{D}|} |\text{prec}(m) - \text{conf}(m)|, \quad (5)$$

where  $\mathcal{D}_m$  is the set of detections in  $m^{\text{th}}$  bin,  $|\mathcal{D}|$  is the total number of detections,  $\text{prec}(m)$  denotes the precision and  $\text{conf}(m)$  the average confidence over samples in the  $m^{\text{th}}$  bin. LaECE is computed as an average of  $k$  per-class calibration errors obtained as:

$$\text{LaECE}^k = \sum_{m=1}^M \frac{|\mathcal{D}_m^k|}{|\mathcal{D}^k|} \left| \text{prec}^k(m) \times \text{IoU}^k(m) - \text{conf}^k(m) \right|, \quad (6)$$

Several works have discussed the numerous flaws of binning estimators [1, 5, 36, 37], such as sensitivity to the binning scheme, asymptotic inconsistency in many cases, and curse of dimensionality. Moreover, these estimators are not differentiable and require additional approximations to be integrated as part of the optimization process.

As a remedy, Popordanoska et al. [31], Zhang et al. [43] recently proposed kernel density-based [34] estimators of calibration error, which are consistent, differentiable, and have more favorable statistical and computational properties. Therefore, we extend the estimator of [31] to the detection task. That is, we wish to find an estimator:

$$\widehat{\text{CE}}(g_k) = \frac{1}{w} \sum_{v=1}^w \left| \mathbb{E}[Z \mid \hat{S} = s_v] - s_v \right|, \quad (7)$$

where  $w$  denotes the number of detections across images.

We will focus on deriving an estimator of the conditional expectation, depending on the definition of calibration resulting from using two distinct link functions  $\psi$  in Equation (3). In case  $\psi$  is a threshold function,  $Z$  is a discrete random variable, and the estimator obtained from this link function is identical to the one proposed in [31]. However, in case  $\psi$  is a continuous function such as the identity,  $Z$  is a continuous random variable, for which we derive the estimator below.

Let  $Z$  and  $\hat{S}$  be continuous random variables with joint density  $f_{Z, \hat{S}}(z, \hat{s})$ . We wish to find an estimator for the conditional expectation:

$$\mathbb{E}[Z \mid \hat{S} = s] = \int_0^1 z p_{Z|\hat{S}}(z \mid s) dz \quad (8)$$

$$= \frac{1}{p_{\hat{S}}(s)} \int_0^1 z p_{Z, \hat{S}}(z, s) dz. \quad (9)$$

We derive an estimator for  $p_{Z, \hat{S}}(z, s)$  using KDE:

$$p_{Z, \hat{S}}(z, s) \approx \frac{1}{w} \sum_{u=1}^w \underbrace{k_s(\hat{S}, s_u) k_z(Z, z_u)}_{k(Z, \hat{S}, z_u, s_u)}. \quad (10)$$

where  $k_s$  and  $k_z$  denote any consistent kernels over their respective domains. The resulting density estimate remains consistent [38, Theorem 6]. Therefore, for the integral in

Equation (9), we have:

$$\int_0^1 z p_{Z, \hat{S}}(z, s) dz \approx \int_0^1 z \frac{1}{w} \sum_{u=1}^w k_s(\hat{S}, s_u) k_z(Z, z_u) dz \quad (11)$$

$$= \frac{1}{w} \sum_{u=1}^w k_s(\hat{S}, s_u) \int_0^1 z k_z(Z, z_u) dz. \quad (12)$$

The final integral  $\int_0^1 z k_z(Z, z_u) dz$  converges to  $z_u$  as the bandwidth of  $k_z$  approaches zero, the density estimate approaches the true density, and the kernel  $k_z$  approaches a Dirac delta function centered at  $z_u$ . Finally, the resulting estimator can be written as:

$$\widehat{\mathbb{E}}[Z \mid \hat{S}] = \frac{\sum_{u=1}^w k(\hat{S}, s_u) z_u}{\sum_{u=1}^w k(\hat{S}, s_u)}. \quad (13)$$

Plugging this back into Equation (7), and generalizing  $z_u$  with a link function, we obtain the following *estimator of  $L_1$  calibration error for a binary object detector  $g_k$* , a given link function  $\psi$  and a similarity measure  $L$ :

$$\widehat{\text{CE}}(g_k) = \frac{1}{w} \sum_{v=1}^w \left| \frac{\sum_{u \neq v} k(s_v, s_u) \psi(L(b_u, b_u^*))}{\sum_{u \neq v} k(s_v, s_u)} - s_v \right|. \quad (14)$$

The estimator has values  $\in [0, 1]$ , it is consistent, asymptotically unbiased [31], and differentiable almost everywhere, as desired.

### 3.3. Train-time calibration

A common approach to designing calibration methods is to include a calibration-related auxiliary loss to the task-specific loss to be minimized during training [24, 25, 29]. In this section, we will show that existing auxiliary losses are not consistent estimators of calibration error. Instead, we propose to use our estimator derived in the previous section.

Munir et al. [24] proposed an auxiliary loss  $\text{TCD} = \frac{1}{2}(d_{cls} + d_{det})$ . Adapting the notation from [24, Equations (8) and (9)], we have that  $d_{cls} = \mathbb{E}[|\hat{S} - Y|]$  and  $d_{det} = \mathbb{E}[|\text{IoU} - \hat{S}|]$ . We will derive a relationship between the classification term  $d_{cls}$  and the classification calibration error, as well as between the localization term  $d_{det}$  and the detection calibration error. The calibration error for binary classification is given by  $\text{CE}_{cls} := \mathbb{E}[|\mathbb{E}[Y = 1 | \hat{S}] - \hat{S}|]$ , whereas  $\text{CE}_{det} := \mathbb{E}[|\mathbb{E}[\text{IoU} | \hat{S}] - \hat{S}|]$  defines the calibration error for detection when  $\psi$  is an identity function.

**Proposition 3.1.** *The classification calibration error upper bounds  $d_{cls}$ :  $\text{CE}_{cls} \geq d_{cls}$ .*

**Proposition 3.2.** *The detection calibration error is upper bounded by  $d_{det}$ :  $\text{CE}_{det} \leq d_{det}$ .*



The proofs are given in appendix A. Thus, through Jensen’s inequality, we showed that neither of these terms is a consistent estimator for any established concept of calibration error. We note that the auxiliary loss components defined in [29, Equations (8) and (9)] have the same functional form as  $d_{cls}$  and  $d_{det}$ . BPC [25, Equation (9)] is based on a heuristic to maximize the confidence scores for accurate predictions and minimize scores for inaccurate ones. Therefore, none of the existing auxiliary losses are based on a consistent estimator of calibration error. Consequently, a principled way to minimize calibration error during training is to integrate our KDE-based estimator  $\widehat{CE}$ , as the only consistent and differentiable estimator of calibration error for object detection, with the task specific-loss  $\mathcal{L}_{det}$  as:

$$\mathcal{L} = \mathcal{L}_{det} + \lambda \widehat{CE}, \quad (15)$$

where  $\lambda$  is a regularization parameter set by cross-validation.

## 4. Experiments

In this section, we investigate the performance of our estimator, both as a metric to evaluate calibration error, and as an auxiliary loss function used for train-time calibration.

### 4.1. Setup

**Datasets** We use the following three datasets:

1. MS COCO [18] is a widely used object detection dataset that consists of more than 330000 object instances belonging to 80 object categories. It contains 118K training images and 5K validation images.
2. Cityscapes [4] is an urban driving scene dataset, consisting of 2975 training and 500 validation images split into 8 object categories.
3. PASCAL VOC [7] has 20 classes, which are a subset of those in MS COCO. Following other works, for training we combine the VOC 2007 and VOC 2012 `trainval` sets, resulting in a total of 16551 images and more than 40000 objects.

We create separate validation sets with three different seeds, by splitting the original train set into new train and validation sets with 90:10 ratio. As the labels for the test set of COCO and Cityscapes are not publicly available, the original `val` set is used for reporting the results, and for VOC we evaluate on the VOC 2007 `test` set.

**Object detectors** We use three popular object detectors with different architectures, which have achieved strong performance in this task:

1. Faster R-CNN (F-RCNN) [33] is a popular two-stage detector with softmax classifier. The two-stage approach involves first generating a sparse set of object proposals, then refining them to accurately locate and

classify the objects in the image. Typically they perform better compared to one-stage detectors, but they are slower and require more computational resources.

2. RetinaNet [19] is a one-stage detector with sigmoid classifiers. Using this approach, objects are detected with a single pass through the network. RetinaNet applies a set of predefined anchor boxes to each feature level and predicts the objectness score and class probabilities for each anchor. In spite of their higher computation efficiency, one-stage detectors may have lower accuracy and struggle with small objects.
3. FCOS [35] is another common one-stage detector with sigmoid classifiers. Compared to RetinaNet, FCOS predicts the object bounding boxes and their class probabilities without requiring an anchor-based mechanism. Both FCOS and RetinaNet use focal loss, which aims to address the class imbalance problem, and it is also known to help with calibration.

**Baselines** Every object detector we consider has a combination of classification and regression loss, which we refer to as task-specific loss, and represents the first baseline we compare to. Due to its strong performance, we include **TS** as a representative of the post-hoc calibration methods. As a train-time baseline method, we compare with the recently proposed auxiliary loss function **TCD** [24], which is added to the task-specific loss.

**Metrics** For reporting detection performance we use the COCO-defined metrics. Namely, average precision AP denotes an average across categories and 10 (.50:.05:.95) IoU overlap thresholds used for matching. In some experiments, we also report AP@0.5, AP@0.75, and AP for small, medium and large objects, as defined in the COCO challenge [18]. For measuring calibration error, we use our KDE-based estimator  $\widehat{CE}$  (Equation (14)) with different link functions throughout the experiments. We use the same breakdown across object sizes and IoU overlaps as the one used for AP. In addition, we report the D-ECE metric (Equation (5)) used in [24, 25, 29], both at IoU@0.5 and averaged over 10 thresholds (.50:.05:.95). In our synthetic experiment, we also compare with LaECE [27].

**Implementation** Traditionally, detections are obtained in two steps: generating predictions and post-processing. The latter step performs further processing to filter out redundant or low-confidence detections, through procedures like Non-Maximum Suppression (NMS) and top- $k$  selection, where typically  $k = 100$  for COCO. Choosing the optimal value for  $k$ , or a threshold  $\gamma$  below which detections are rejected is an open area of research [26, 27]. For simplicity, in our experiments we set the score threshold  $\gamma = 0.5$  for our main results, and show an ablation study of

the effect of this parameter on the reported metrics. Our implementation for evaluating CE on a test set is based on the official API for assessing detection performance on COCO, i.e. we utilize the same information about matching predicted and ground truth boxes and IoU overlaps as for evaluating AP. We use a Beta kernel [3] and set the bandwidth parameter with a leave-one-out maximum likelihood procedure. For D-ECE we adopt the implementation from Küppers et al. [15]. We rely on Detectron2 [39] for the implementation and training configuration of the detectors. The code and trained models are available at <https://github.com/tpopordanoska/calibration-object-detection>.

## 4.2. Results

Here we empirically showcase the utilities of our proposed estimator and compare it with related work. The main purpose of an estimator of calibration error is to assess the extent to which a notion of calibration is violated. Therefore, we start our experiments by comparing the different variants of our KDE-based estimator (obtained with different choices for the link function) with the binning-based estimators D-ECE and LaECE. Subsequently, we provide a comprehensive evaluation of  $\widehat{CE}$  and the standard AP metric, at different IoU thresholds and object sizes. Finally, we demonstrate the differentiability of our estimator by incorporating it as an auxiliary loss function during training, both from scratch and for fine-tuning. We note that the estimator could also be integrated as part of a post-hoc method.

**Comparison of metrics** D-ECE measures the notion of calibration obtained by letting  $\psi$  be a threshold function, whereas LaECE is used for assessing calibration as obtained through an identity function. We compare these binning-based estimators to the corresponding KDE-based estimators on a synthetic binary problem, where the ground truth calibration error is known. Similar to [31], we apply temperature scaling with  $t_1 = 0.6$  to a uniform sample of numbers in  $[0, 1]$ . The purpose of this step is to ensure that the confidence scores are concentrated around 0 and 1, as in a realistic scenario. Then, we sample the labels, denoting  $\mathbb{1}[\text{IoU}(\hat{B}, B^*) \geq \beta]$ , according to that distribution, and therefore obtain perfect calibration. In order to simulate miscalibration, we perform another temperature scaling with  $t_2 = 0.6$ . D-ECE is calculated using 20 equally-spaced bins, as in [15]. For LaECE we use 25 bins, following [27], and compute the weighted difference between  $t_1$  and  $t_2$  scaled scores in each bin. As shown in Figure 2, D-ECE and  $\widehat{CE}$  have similar performance and both yield good CE estimates with a few thousand points. The comparison with LaECE can be found in Appendix B.

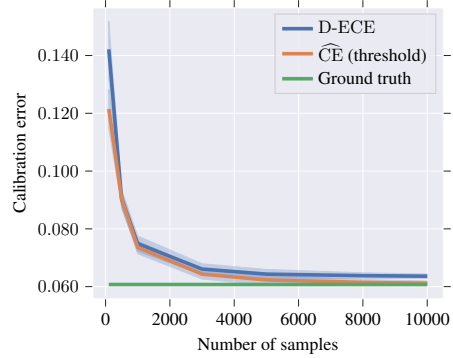


Figure 2. Comparison of D-ECE vs.  $\widehat{CE}$  (threshold link) as a function of the number of points used for the estimation.

**Evaluating a model zoo** We conduct a rigorous analysis to determine the calibration of F-RCNN, RetinaNet and FCOS on popular benchmark datasets like COCO, Cityscapes and Pascal VOC. Analogous to the evaluation on the COCO challenge, we estimate calibration error across different sizes of objects, and for different thresholds for the IoU overlap required to determine the matching between predicted and ground truth box. For this evaluation, we use a threshold link function. The results are summarized in Tables 1 and 2. The first thing to note is that modern object detectors are intrinsically miscalibrated. When comparing the different architectures, an observation can be made that RetinaNet and FCOS have lower calibration errors than F-RCNN. This may be attributed to the use of focal loss for the classification branch of RetinaNet and FCOS, which has been empirically shown to improve calibration. Moreover, different metrics reveal complementary information about the calibration of the detector. For instance,  $CE_{50}$  on Pascal VOC shows similar performance of F-RCNN and FCOS, whereas  $CE_{75}$  reveals that RetinaNet and FCOS have much better performance than F-RCNN for a calibration error that requires an IoU overlap of 0.75 for a correct detection. Similarly, both RetinaNet and FCOS demonstrate similar calibration errors when detecting small objects. However, FCOS exhibits superior calibration scores for medium-sized objects, despite both detectors achieving the same  $AP_M$ .

**Calibration-regularized training** Another notion of calibration that might be of interest is when the link  $\psi$  is an identity function. Intuitively, this type of calibration requires that the predicted score corresponds to the IoU overlap with a ground truth box. Using this setup for  $\widehat{CE}$ , we performed extensive experiments to compare the performance of our estimator with a post-hoc method (TS), and with a train-time loss (TCD). The temperature is chosen on a validation set by minimizing NLL. Table 3 reports the AP,  $\widehat{CE}$ , and D-ECE with IoU overlap of 0.5, for RetinaNet and

Table 1. Detection performance of object detectors on three popular datasets.

Dataset	Model	AP	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>S</sub>	AP <sub>M</sub>	AP <sub>L</sub>
COCO	F-RCNN	36.11 $\pm$ 0.10	53.35 $\pm$ 0.15	40.04 $\pm$ 0.12	18.88 $\pm$ 0.10	39.40 $\pm$ 0.11	47.98 $\pm$ 0.08
	RetinaNet	30.83 $\pm$ 0.12	43.33 $\pm$ 0.22	34.03 $\pm$ 0.07	13.59 $\pm$ 0.28	34.13 $\pm$ 0.13	43.16 $\pm$ 0.24
	FCOS	34.02 $\pm$ 0.06	48.84 $\pm$ 0.04	37.37 $\pm$ 0.10	17.29 $\pm$ 0.17	37.84 $\pm$ 0.17	45.15 $\pm$ 0.17
Cityscapes	F-RCNN	35.36 $\pm$ 0.65	54.46 $\pm$ 0.82	37.90 $\pm$ 0.79	11.54 $\pm$ 0.31	35.34 $\pm$ 0.51	56.62 $\pm$ 0.48
	RetinaNet	34.60 $\pm$ 0.23	52.60 $\pm$ 0.81	36.41 $\pm$ 0.35	10.39 $\pm$ 0.15	35.56 $\pm$ 0.15	57.47 $\pm$ 0.39
	FCOS	34.81 $\pm$ 0.08	52.31 $\pm$ 0.29	36.62 $\pm$ 0.23	10.66 $\pm$ 0.60	32.84 $\pm$ 0.44	56.90 $\pm$ 0.81
Pascal VOC	F-RCNN	52.96 $\pm$ 0.05	75.88 $\pm$ 0.08	59.67 $\pm$ 0.05	18.48 $\pm$ 0.39	40.90 $\pm$ 0.11	61.43 $\pm$ 0.08
	RetinaNet	53.20 $\pm$ 0.02	73.36 $\pm$ 0.02	58.75 $\pm$ 0.11	16.39 $\pm$ 0.42	40.19 $\pm$ 0.26	62.06 $\pm$ 0.02
	FCOS	52.13 $\pm$ 0.02	73.50 $\pm$ 0.07	57.63 $\pm$ 0.04	19.23 $\pm$ 0.23	40.52 $\pm$ 0.12	60.28 $\pm$ 0.06

Table 2. Calibration performance ( $CE \times 100$ ) of object detectors on three popular datasets. The calibration error is broken down into variants depending on the IoU threshold and the size of the objects, same as for AP.

Dataset	Model	$\widehat{CE}$	$\widehat{CE}_{50}$	$\widehat{CE}_{75}$	$\widehat{CE}_S$	$\widehat{CE}_M$	$\widehat{CE}_L$
COCO	F-RCNN	37.33 $\pm$ 0.09	20.48 $\pm$ 0.10	32.56 $\pm$ 0.15	35.72 $\pm$ 0.51	38.94 $\pm$ 0.10	40.81 $\pm$ 0.12
	RetinaNet	21.89 $\pm$ 0.36	12.03 $\pm$ 0.19	14.70 $\pm$ 0.41	26.73 $\pm$ 0.52	26.04 $\pm$ 0.13	27.58 $\pm$ 0.17
	FCOS	24.40 $\pm$ 0.13	16.55 $\pm$ 0.21	19.78 $\pm$ 0.17	26.34 $\pm$ 0.39	26.33 $\pm$ 0.24	29.68 $\pm$ 0.03
Cityscapes	F-RCNN	37.45 $\pm$ 0.46	17.00 $\pm$ 0.46	32.55 $\pm$ 0.48	40.85 $\pm$ 5.08	41.15 $\pm$ 1.57	38.22 $\pm$ 0.78
	RetinaNet	32.22 $\pm$ 0.29	12.91 $\pm$ 0.95	28.38 $\pm$ 0.74	29.99 $\pm$ 0.57	38.16 $\pm$ 1.27	35.91 $\pm$ 0.46
	FCOS	26.29 $\pm$ 0.98	13.91 $\pm$ 1.28	21.09 $\pm$ 1.57	28.44 $\pm$ 1.66	32.71 $\pm$ 0.78	28.13 $\pm$ 0.93
Pascal VOC	F-RCNN	34.88 $\pm$ 0.10	17.03 $\pm$ 0.16	28.75 $\pm$ 0.17	47.67 $\pm$ 0.91	41.99 $\pm$ 0.33	30.84 $\pm$ 0.07
	RetinaNet	24.11 $\pm$ 0.09	7.95 $\pm$ 0.19	17.89 $\pm$ 0.14	36.31 $\pm$ 1.22	31.53 $\pm$ 0.06	21.73 $\pm$ 0.11
	FCOS	22.32 $\pm$ 0.05	15.77 $\pm$ 0.10	15.55 $\pm$ 0.11	36.84 $\pm$ 0.65	27.12 $\pm$ 0.04	20.97 $\pm$ 0.02

FCOS, trained on Cityscapes.

Table 3. Comparison of detection (AP) and calibration performance ( $CE \times 100$ ) of models trained with RetinaNet and FCOS on Cityscapes. Models with no calibration, with post-hoc (TS) and train-time (TCD) methods are compared with our approach.

Model	AP $\uparrow$	$\widehat{CE} \downarrow$	D-ECE <sub>50</sub> $\downarrow$
RetinaNet	34.60 $\pm$ 0.23	23.25 $\pm$ 0.39	12.54 $\pm$ 0.66
RetinaNet + TS	34.60 $\pm$ 0.23	18.67 $\pm$ 2.07	11.05 $\pm$ 1.25
RetinaNet + TCD	33.94 $\pm$ 0.81	25.01 $\pm$ 1.13	13.04 $\pm$ 0.47
RetinaNet + $\widehat{CE}$	32.59 $\pm$ 0.71	17.33 $\pm$ 3.15	11.07 $\pm$ 0.22
FCOS	34.81 $\pm$ 0.08	14.43 $\pm$ 1.31	13.23 $\pm$ 1.18
FCOS + TS	33.62 $\pm$ 0.44	13.33 $\pm$ 2.16	11.96 $\pm$ 0.90
FCOS + TCD	35.57 $\pm$ 0.23	16.85 $\pm$ 1.27	13.68 $\pm$ 0.86
FCOS + $\widehat{CE}$	33.73 $\pm$ 0.57	13.07 $\pm$ 0.83	17.45 $\pm$ 0.61

**Fine-tuning** Since there already exist a variety of trained object detectors on COCO, we demonstrate that the benefits of adding our estimator as an auxiliary loss can be observed also for fine-tuning pre-trained models even for a few epochs. Table 4 shows the reduction in CE, achieved by fine-tuning for three epochs, by minimizing the loss defined in Equation (15) on COCO dataset.

Table 4. Performance (AP and  $CE \times 100$ ) of models without calibration and fine-tuned models using our auxiliary loss on COCO.

Model	AP $\uparrow$	$\widehat{CE} \downarrow$
F-RCNN	36.11 $\pm$ 0.10	31.76 $\pm$ 0.05
F-RCNN + $\widehat{CE}$	34.72 $\pm$ 0.09	26.91 $\pm$ 0.14
RetinaNet	30.83 $\pm$ 0.12	21.89 $\pm$ 0.36
RetinaNet + $\widehat{CE}$	30.10 $\pm$ 0.06	9.72 $\pm$ 0.18
FCOS	34.02 $\pm$ 0.06	24.40 $\pm$ 0.13
FCOS + $\widehat{CE}$	33.41 $\pm$ 0.12	15.32 $\pm$ 0.17

**Ablation study** We investigate the effect of the regularization parameter  $\lambda$ , the test score threshold  $\gamma$ , and the number of epochs used for fine-tuning on the reported metrics. In Figure 3 we present fine-tuned F-RCNN detectors on Cityscapes for different values of the  $\lambda$  parameter. We notice that increasing the weight of calibration regularization leads to noticeable reduction in calibration error. Table 5 shows the evaluated performance on test set using different thresholds  $\gamma$ . From Table 6 we can observe that fine-tuning even for a few epochs can lead to improvements in calibration error, without sacrificing AP. In Appendix B we conduct further experiments for the effect of  $\lambda$ , both for fine-tuned models and trained from scratch on the three datasets.




Figure 3. Effect of  $\lambda$  on AP and  $\widehat{CE}$ . The points represent fine-tuned Faster-RCNN detectors on Cityscapes for three epochs. The number next to the point denotes the value of  $\lambda$ .

Table 5. Effect of test score threshold  $\gamma$  on detection performance and calibration of Faster-RCNN trained on Cityscapes.

Model	AP	$\widehat{CE}$
F-RCNN ( $\gamma = 0.1$ )	38.34 $\pm$ 0.18	22.61 $\pm$ 1.11
F-RCNN ( $\gamma = 0.2$ )	37.53 $\pm$ 0.26	26.26 $\pm$ 0.78
F-RCNN ( $\gamma = 0.3$ )	36.92 $\pm$ 0.31	27.86 $\pm$ 0.80
F-RCNN ( $\gamma = 0.4$ )	36.15 $\pm$ 0.43	28.36 $\pm$ 0.83
F-RCNN ( $\gamma = 0.5$ )	35.36 $\pm$ 0.65	28.48 $\pm$ 0.68
F-RCNN ( $\gamma = 0.6$ )	34.31 $\pm$ 0.75	28.19 $\pm$ 0.90
F-RCNN ( $\gamma = 0.7$ )	33.26 $\pm$ 0.96	27.48 $\pm$ 0.99

Table 6. Effect of number of epochs for fine-tuning a Faster-RCNN network on Cityscapes, with  $\lambda = 1$  and  $\gamma = 0.5$ .

Model	AP	$\widehat{CE}$
F-RCNN	35.36 $\pm$ 0.65	28.48 $\pm$ 0.68
F-RCNN + $\widehat{CE}$ (1x)	34.39 $\pm$ 0.19	25.88 $\pm$ 1.40
F-RCNN + $\widehat{CE}$ (2x)	34.45 $\pm$ 0.22	25.48 $\pm$ 0.60
F-RCNN + $\widehat{CE}$ (3x)	35.06 $\pm$ 0.51	25.76 $\pm$ 0.85
F-RCNN + $\widehat{CE}$ (4x)	35.67 $\pm$ 0.41	26.72 $\pm$ 0.56
F-RCNN + $\widehat{CE}$ (5x)	35.80 $\pm$ 0.40	26.19 $\pm$ 1.46

## 5. Discussion

In this paper, we tackled the challenge of defining calibration and estimating calibration error for object detection. Specifically, calibration error for OD is a quantity that depends on the choice of a measure of “correctness” of a prediction, which we have decomposed into a similarity measure  $L$  and a link function  $\psi$ , which together determine the notion of calibration that is relevant for a particular scenario. Beyond the definition, we also proposed a consistent and differentiable estimator of a calibration error for OD, which can be used for common one-stage, two-stage, anchor-based and anchor-free DNN-based detectors.

The empirical results showed that our estimator performs equally well for estimating calibration error as existing binning-based versions, while offering superior statistical properties. Moreover, due to its differentiability, it can be directly included as part of both post-hoc and train-time calibration setups. In particular, when integrated into the task-specific loss during training, our estimator achieves substantial improvements in calibration error, while maintaining similar levels of AP. We have also shown that a reduction in calibration error can be achieved by fine-tuning with our auxiliary loss for a few epochs, thus eliminating the need to train object detectors from scratch.

This work also has some limitations. The main limitation is the intrinsic  $\mathcal{O}(n^2)$  [31] complexity of KDE. When using the estimator of calibration error as an auxiliary loss for one-stage object detectors, which generate a considerable number of candidate detections during training, the complexity becomes particularly challenging. Nevertheless, our synthetic experiments have revealed that the calibration error can be accurately estimated using only a few thousand data points. As a result, during training with RetinaNet and FCOS, we opt to randomly sample a subset of detections and evaluate the calibration error based on this reduced set of data points. Our empirical results verify that this is an effective strategy to conduct calibration-regularized training. Another limitation of this paper is that we did not include a transformer-based model, such as DINO [16], in our experiments. However, the presented mathematical principles are model-agnostic and our estimator can be trivially extended to this family of detectors as well. Finally, it is worth noting that in certain scenarios, calibration regularized training results in a slight reduction of the AP metric. Exploring the relationship between CE, AP, and the test score threshold, along with investigating the advantage that calibration brings to downstream tasks that utilize object detectors, is an interesting direction for future work.

To conclude, our study puts forward a *novel and principled* view on calibration in OD, and proposes a new mathematical framework, which enables both estimation of calibration error, and calibration regularized training for object detection. Considering the critical role of calibration in enhancing overall system robustness, our method is highly relevant across diverse computer vision applications, and especially in risk-sensitive scenarios.

## Acknowledgements

This research received funding from the Flemish Government (AI Research Program) and the Research Foundation - Flanders (FWO) through project number G0G2921N. The publication was also supported by funding from the Academy of Finland (Profif6 336449 funding program). The authors wish to acknowledge CSC – IT Center for Science, Finland, for generous computational resources.



## References

- [1] Arsenii Ashukha, Alexander Lyzhov, Dmitry Molchanov, and Dmitry Vetrov. Pitfalls of in-domain uncertainty estimation and ensembling in deep learning. In *International Conference on Learning Representations*, 2020. 4
- [2] Jochen Bröcker. Reliability, sufficiency, and the decomposition of proper scores. *Quarterly Journal of the Royal Meteorological Society*, 135(643):1512–1519, Jul 2009. 3
- [3] Song Xi Chen. Beta kernel estimators for density functions. *Computational Statistics & Data Analysis*, 31:131–145, 1999. 6
- [4] 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. 2, 5
- [5] Yukun Ding, Jinglan Liu, Jinjun Xiong, and Yiyu Shi. Revisiting the evaluation of uncertainty estimation and its application to explore model complexity-uncertainty trade-off. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 22–31, 2020. 4
- [6] 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
- [7] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html>. 2, 5
- [8] 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, 2, 3
- [9] Ali Harakeh and Steven L. Waslander. Estimating and evaluating regression predictive uncertainty in deep object detectors. In *International Conference on Learning Representations*, 2021. 2
- [10] 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*, pages 16081–16090, 2022. 2
- [11] Ildar Abdulovich Ibragimov and Rafail Zalmanovich Has’ Minskii. *Statistical estimation: asymptotic theory*, volume 16. Springer Science & Business Media, 2013. 3
- [12] Archit Karandikar, Nicholas Cain, Dustin Tran, Balaji Lakshminarayanan, Jonathon Shlens, Michael Curtis Mozer, and Rebecca Roelofs. Soft calibration objectives for neural networks. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. 2
- [13] Ananya Kumar, Percy S Liang, and Tengyu Ma. Verified uncertainty calibration. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 3792–3803. 2019. 3
- [14] Aviral Kumar, Sunita Sarawagi, and Ujjwal Jain. Trainable calibration measures for neural networks from kernel mean embeddings. In *ICML*, 2018. 2
- [15] 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. 1, 2, 3, 6
- [16] Feng Li, Hao Zhang, Huaizhe Xu, Shilong Liu, Lei Zhang, Lionel M Ni, and Heung-Yeung Shum. Mask dino: Towards a unified transformer-based framework for object detection and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3041–3050, 2023. 8
- [17] 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*, 2020. 2
- [18] 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 *European Conference on Computer Vision*, 2014. 2, 5
- [19] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42:318–327, 2017. 2, 5
- [20] Xingchen Ma and Matthew B. Blaschko. Meta-cal: Well-controlled post-hoc calibration by ranking. In *International Conference on Machine Learning*, 2021. 2
- [21] Lena Maier-Hein, Bjoern Menze, et al. Metrics reloaded: Pitfalls and recommendations for image analysis validation. *arXiv.org*. (2206.01653), 2022. 2
- [22] 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. 2
- [23] Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? *Advances in neural information processing systems*, 32, 2019. 2

- [24] Muhammad Akhtar Munir, Muhammad Haris Khan, M Sarfraz, and Mohsen Ali. Towards improving calibration in object detection under domain shift. *Advances in Neural Information Processing Systems*, 35:38706–38718, 2022. 1, 2, 3, 4, 5
- [25] Muhammad Akhtar Munir, Muhammad Haris Khan, Salman Khan, and Fahad Shahbaz Khan. Bridging precision and confidence: A train-time loss for calibrating object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11474–11483, 2023. 1, 2, 3, 4, 5
- [26] Kemal Oksuz, Baris Can Cam, Emre Akbas, and Sinan Kalkan. Localization recall precision (lrp): A new performance metric for object detection. In *Proceedings of the European conference on computer vision (ECCV)*, pages 504–519, 2018. 5
- [27] Kemal Oksuz, Tom Joy, and Puneet K. Dokania. Towards building self-aware object detectors via reliable uncertainty quantification and calibration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9263–9274, June 2023. 1, 2, 3, 5, 6
- [28] 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
- [29] Bimsara Pathiraja, Malitha Gunawardhana, and Muhammad Haris Khan. Multiclass confidence and localization calibration for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19734–19743, June 2023. 2, 3, 4, 5
- [30] Teodora Popordanoska, Jeroen Bertels, Dirk Vandermeulen, Frederik Maes, and Matthew B. Blaschko. On the relationship between calibrated predictors and unbiased volume estimation. In *Medical Image Computing and Computer-Assisted Intervention*. 2021. 11
- [31] Teodora Popordanoska, Raphael Sayer, and Matthew B. Blaschko. A consistent and differentiable lp canonical calibration error estimator. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. 2, 4, 6, 8
- [32] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 2
- [33] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39:1137–1149, 2015. 5
- [34] B. W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman & Hall, 1986. 4
- [35] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9626–9635, 2019. 2, 5
- [36] Juozas Vaicenavicius, David Widmann, Carl Andersson, Fredrik Lindsten, Jacob Roll, and Thomas Schön. Evaluating model calibration in classification. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3459–3467. PMLR, 2019. 2, 4
- [37] David Widmann, Fredrik Lindsten, and Dave Zachariah. Calibration tests in multi-class classification: A unifying framework. *Advances in Neural Information Processing Systems*, 32, 2019. 2, 4
- [38] Dominik Wied and Rafael Weißbach. Consistency of the kernel density estimator - a survey. *Statistical Papers*, 53(1): 1–21, 2010. 4
- [39] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 6
- [40] Ekim Yurtsever, Jacob Lambert, Alexander Carballo, and Kazuya Takeda. A survey of autonomous driving: Common practices and emerging technologies. *IEEE access*, 8:58443–58469, 2020. 1
- [41] B. Zadrozny and C. Elkan. Transforming classifier scores into accurate multiclass probability estimates. *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, 2002. 2, 3
- [42] Bianca Zadrozny and Charles Elkan. Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers. *ICML*, 1, 05 2001. 2
- [43] Jize Zhang, Bhavya Kailkhura, and T. Yong-Jin Han. Mix-n-match: Ensemble and compositional methods for uncertainty calibration in deep learning. In *International Conference on Machine Learning*, 2020. 4