

# Active Learning for Deep Object Detection via Probabilistic Modeling

Jiwoong Choi<sup>1,3</sup>, Ismail Elezi<sup>2,3</sup>, Hyuk-Jae Lee<sup>1</sup>, Clement Farabet<sup>3</sup>, and Jose M. Alvarez<sup>3</sup>  
<sup>1</sup>Seoul National University, <sup>2</sup>Technical University of Munich, <sup>3</sup>NVIDIA

{jwchoi, hjlee}@capp.snu.ac.kr, ismail.elezi@tum.de, {cfarabet, josea}@nvidia.com

## Abstract

Active learning aims to reduce labeling costs by selecting only the most informative samples on a dataset. Few existing works have addressed active learning for object detection. Most of these methods are based on multiple models or are straightforward extensions of classification methods, hence estimate an image’s informativeness using only the classification head. In this paper, we propose a novel deep active learning approach for object detection. Our approach relies on mixture density networks that estimate a probabilistic distribution for each localization and classification head’s output. We explicitly estimate the aleatoric and epistemic uncertainty in a single forward pass of a single model. Our method uses a scoring function that aggregates these two types of uncertainties for both heads to obtain every image’s informativeness score. We demonstrate the efficacy of our approach in PASCAL VOC and MSCOCO datasets. Our approach outperforms single-model based methods and performs on par with multi-model based methods at a fraction of the computing cost. Code is available at <https://github.com/NVlabs/AL-MDN>.

## 1. Introduction

The performance of deep detection networks is dependent on the size of the labeled data [31, 32]. Motivated by this, researchers have explored smart strategies to select the most informative samples in the dataset for labeling, known as active learning [35]. Typically, this is done by devising a scoring function that computes the network’s uncertainty, selecting to label the samples for which the network is least confident with regard to its predictions [2, 4, 40].

In general, the predictive uncertainty is decomposed into aleatoric and epistemic uncertainty [15, 20]. The former refers to the inherent noise in the data, such as sensor noise, and can be attributed to occlusions or lack of visual features [10, 24]. The latter refers to the uncertainty caused by the lack of knowledge of the model and is inversely proportional to the density of training data [38]. Modeling and distinguishing these two types of uncertainty is very impor-

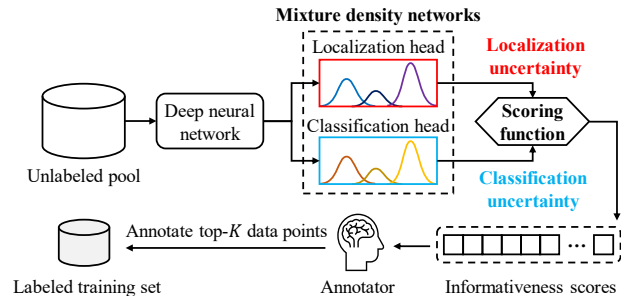


Figure 1: Our approach predicts aleatoric and epistemic uncertainties for both the localization and classification heads in a single forward pass of a single model. We propose an scoring function that aggregates epistemic and aleatoric uncertainties from both heads into single value. Then, those data points with the top-K scores are sent for labeling.

tant in active learning, as it allows the deep learning models to know about their limitations [21, 38], *i.e.*, recognize suspicious predictions in the sample (*aleatoric uncertainty*) and recognize samples that do not resemble the training set (*epistemic uncertainty*). To compute these types of uncertainty, researchers use multi-model based approaches, such as ensembles [2] or Monte Carlo (MC) dropout [13]. These methods reach good results but come with several limitations [11, 16]. In particular, by being multi-models, they require a much higher computing cost, and in the case of ensembles, they also increase the number of the network’s parameters [2]. Additionally, they rely only on classification uncertainty, totally ignoring the localization uncertainty.

In this paper, we propose a novel active learning approach for deep object detection. Our approach uses a single model with a single forward pass, significantly reducing the computing cost compared to multiple model-based methods. Despite this, our method reaches high accuracy. To manage so, our method utilizes both localization and classification-based aleatoric and epistemic uncertainties. As shown in Fig. 1, we base our method on a mixture density networks [3] that learns a Gaussian mixture model (GMM) for each of the network’s outputs, *i.e.*, localization and classification, to compute both aleatoric and

epistemic uncertainties. To efficiently train the network, we propose a loss function that serves as a regularizer for inconsistent data, leading to more robust models. Our method estimates every image’s informativeness score by aggregating all of the localization and classification-based uncertainties for every object contained in the image. We empirically show that leveraging both types of uncertainty coming from classification and localization heads is a critical factor for improving the accuracy. We demonstrate the benefits of our approach on PASCAL VOC [9] and MS-COCO [30] in a single-stage architecture such as SSD [31], and show generalized performance in a two-stage architecture such as Faster-RCNN [32]. Our approach consistently outperforms single-model based methods, and compared to methods using multi-models, our approach yields a similar accuracy while significantly reducing the computing cost.

In summary, our **contributions** are the following:

- We propose a novel deep active learning method for object detection that leverages the aleatoric and epistemic uncertainties, by considering both the localization and classification information. Our method is efficient and uses a single forward pass in a single model.
- We propose a novel loss to train the GMM-based object detection network that leads to overall performance improvements in the network.
- We demonstrate the effectiveness of our approach using different models on two different datasets.

## 2. Related Work

**Deep active learning for object detection** has recently acquired interest. The work of [16] trains an ensemble [2] of neural networks and then selects the samples with the highest score defined by some acquisition function, *i.e.*, entropy [36] or mutual information [5]. Concurrent work [11] explores similar directions, but by approximating the uncertainty via MC-dropout [12, 26]. The work of [1] presents a method of calculating pixel scores and using them for selecting informative samples. Another approach [33] proposes a *query by committee* paradigm to choose the set of images to be queried. The work of [34] uses the feature space to select representative samples in the dataset, reaching good performance in object detection [40]. A different solution was given by [23] where the authors define two different scores: *localization tightness* which is the overlapping ratio between the region proposal and the final prediction; and *localization stability* that is based on the variation of predicted object locations when input images are corrupted by noise. In all cases, images with the highest scores are chosen to be labeled. The state-of-the-art (SOTA) method of [40] offers a heuristic but elegant solution while

outperforming the other single model-based methods. During the training, the method learns to predict the target loss for each sample. During the active learning stage, it chooses to label the samples with the highest predicted loss.

Most of the above-mentioned methods [11, 16, 23] require multiple models or multiple forward passes to calculate the image’s informativeness score, resulting in a high computational cost. In addition, all those studies, despite focusing in active learning for object detection, either rely on heuristic methods to estimate localization uncertainty [23, 40], or cannot estimate it at all [1, 11, 16, 33, 34]. Therefore, while giving promising directions, they are less than satisfactory in terms of accuracy and computing cost. In contrast to those methods, our approach estimates and leverages both the localization and classification uncertainties to reach high accuracy, while using a single forward pass of a single model, significantly reducing the computational cost.

**Mixture density networks** have been recently used for several deep learning tasks. The approach of [8] focuses on the regression task for the steering angle. The works of [18, 39] attempt to solve a multimodal regression task. The work of [41] focuses on density estimation, while the work of [7] attempts to explore the supervised learning problem with corrupted data. However, previous studies do not consider the classification task, which is an essential part of object detection [8, 18, 39]. Additionally, all these studies do not take into account both types of uncertainty coming from the bounding box regression and the classification tasks [7, 8, 18, 39, 41]. Moreover, none of those studies address the problem of active learning for object detection. In contrast, our approach estimates and leverages both the aleatoric and epistemic uncertainties for both tasks in the context of active learning for object detection.

## 3. Active Learning for Object Detection

The key novelty of our approach is designing the output layers of the neural network to predict a probability distribution, instead of predicting a single value for each output of the network (see Fig. 2a). To this end, we propose to make use of a mixture density network where the output of the network consists of the parameters of a GMM: the mean  $\mu^k$ , the variance  $\Sigma^k$ , and the mixture weight  $\pi^k$  for the  $k$ -th component of the GMM. Given these parameters, we can estimate the aleatoric  $u_{al}$  and epistemic  $u_{ep}$  uncertainties [8]:

$$u_{al} = \sum_{k=1}^K \pi^k \Sigma^k, \quad u_{ep} = \sum_{k=1}^K \pi^k \|\mu^k - \sum_{i=1}^K \pi^i \mu^i\|^2, \quad (1)$$

where  $K$  is the number of components in the GMM.

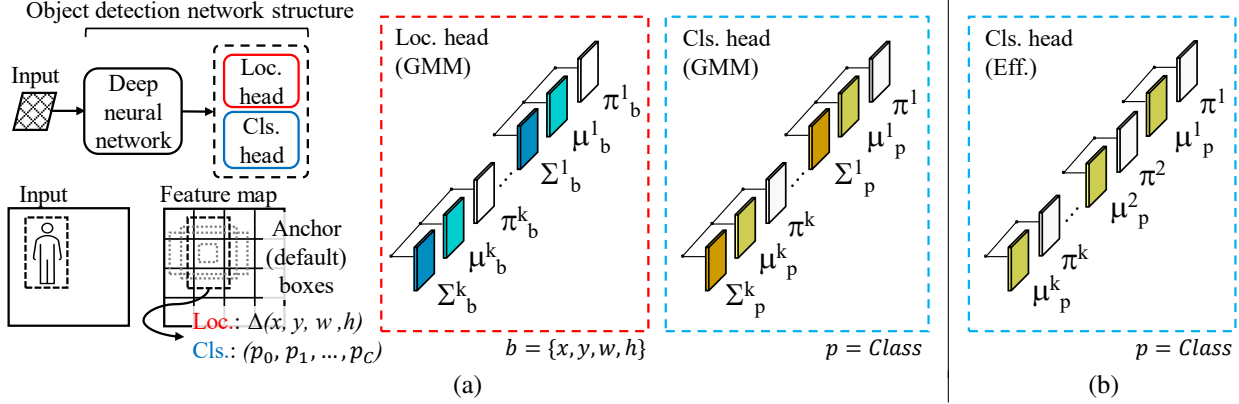


Figure 2: An overview of the proposed object detection network. The main difference with conventional object detectors [31, 32] is in the localization and classification heads (branches). a) Instead of having deterministic outputs, our approach learns the parameters of  $K$ -components GMM for each of the outputs: coordinates of the bounding box in the localization head and the class density distribution in the classification (confidence) head (see Section 3.1). b) A classification head that improves the efficiency by eliminating variance parameters from GMM’s classification head (see Section 3.2).

### 3.1. Object detection with probabilistic modeling

To introduce our approach, we first focus on the localization task and then extend it to the classification task. As we will show later in our experiments, our method is applicable to both single-stage and two-stage object detectors.

**Localization:** In object detection, a bounding box  $b$  is defined by its coordinates for the center ( $x$  and  $y$ ), its width ( $w$ ), and its height ( $h$ ). In our work, instead of predicting a deterministic value, our mixture model predicts 3 groups of parameters for each bounding box: the mean ( $\hat{\mu}_x$ ,  $\hat{\mu}_y$ ,  $\hat{\mu}_w$ , and  $\hat{\mu}_h$ ), the variance ( $\hat{\Sigma}_x$ ,  $\hat{\Sigma}_y$ ,  $\hat{\Sigma}_w$ , and  $\hat{\Sigma}_h$ ), and the weights of the mixture ( $\hat{\pi}_x$ ,  $\hat{\pi}_y$ ,  $\hat{\pi}_w$ , and  $\hat{\pi}_h$ ).

Let  $\{\hat{\pi}_b^k, \hat{\mu}_b^k, \hat{\Sigma}_b^k\}_{k=1}^K$ ,  $b \in \{x, y, w, h\}$  be the bounding box outputs obtained using our network. The parameters of a GMM with  $K$  models for each coordinate of the bounding box are obtained as follows:

$$\pi_b^k = \frac{e^{\hat{\pi}_b^k}}{\sum_{j=1}^K e^{\hat{\pi}_b^j}}, \quad \mu_b^k = \hat{\mu}_b^k, \quad \Sigma_b^k = \sigma(\hat{\Sigma}_b^k), \quad (2)$$

where  $\pi$  is the mixture weight for each component,  $\mu$  is the predicted value for each bounding box coordinate, and  $\Sigma$  is the variance for each coordinate representing its aleatoric uncertainty. As suggested in [8], we use a softmax function to keep  $\pi$  in probability space and use a sigmoid function to satisfy the positiveness constraint of the variance,  $\Sigma_b^k \geq 0$ .

**Localization loss:** The conventional bounding box regression loss, the smooth L1 loss [14], only considers the coordinates of the predicted bounding box and ground-truth (GT) box. Therefore, it cannot take into account the ambiguity (*aleatoric uncertainty*) of the bounding box. For training the mixture density network for localization, we propose a localization loss based on the negative log-likelihood loss.

Our loss regresses the parameters of the GMM to the offsets of the center ( $x$ ,  $y$ ), width ( $w$ ), and height ( $h$ ) of the anchor (default) box ( $d$ ) for positive matches:

$$L_{loc}(\lambda, l, g) = - \sum_{i \in Pos} \sum_b \lambda_G^{ij} \log \left( \sum_{k=1}^K \pi_b^{ik} \mathcal{N}(g_b^j | \mu_b^{ik}, \Sigma_b^{ik}) + \varepsilon \right),$$

$$\lambda_G^{ij} = \begin{cases} 1, & \text{if } IoU > 0.5, \\ 0, & \text{otherwise.} \end{cases}, \quad \hat{g}_x^j = \frac{(g_x^j - d_x^i)}{d_w^i}, \quad \hat{g}_y^j = \frac{(g_y^j - d_y^i)}{d_h^i},$$

$$\hat{g}_w^j = \log\left(\frac{g_w^j}{d_w^i}\right), \quad \hat{g}_h^j = \log\left(\frac{g_h^j}{d_h^i}\right), \quad (3)$$

where  $l$  is GMM parameters of bounding box ( $\pi_b^{ik}$ ,  $\mu_b^{ik}$ , and  $\Sigma_b^{ik}$ ),  $N$  is the number of matched anchor boxes (called *positive matches*),  $K$  is the number of mixtures,  $\lambda_G^{ij}$  is an indicator for matching the  $i$ -th anchor box  $d_b^i$  to the  $j$ -th GT box of category  $G$ , and  $\hat{g}_b^j$  is the  $j$ -th GT box. In experiments, we set  $\varepsilon = 10^{-9}$  for the numerical stability of the logarithm function.

**Classification:** We now focus on the classification head of the object detector. We model the output of every class as a GMM (see Fig. 2a). Our approach estimates the mean  $\hat{\mu}_p^k$  and variance  $\hat{\Sigma}_p^k$  for each class, and the weights of the mixture  $\hat{\pi}^k$  for each component of the GMM. We process the parameters of the GMM following Eq. 2, and obtain the class probability distribution for the  $k$ -th mixture by using the reparameterization trick [25] of applying Gaussian noise and variance  $\Sigma_p^k$  to  $\mu_p^k$  [24]:

$$\hat{c}_p^k = \mu_p^k + \sqrt{\Sigma_p^k} \gamma, \quad \gamma \sim \mathcal{N}(0, 1), \quad (4)$$

where  $\gamma$  is the auxiliary noise variable and has the same size as  $\mu_p^k$  and  $\Sigma_p^k$ .

**Classification loss:** For training the mixture density network for classification, we propose a loss function that takes into account the IoU of the anchor box compared to the GT box and considers the hard negative mining. More precisely, we formulate the classification loss as a combination of two terms  $L_{cl}^{Pos}$  and  $L_{cl}^{Neg}$  representing the contribution of positive and negative matches:

$$\begin{aligned} L_{cl}^{Pos}(\lambda, c) &= - \sum_{i \in Pos} \lambda_G^{ij} \sum_{k=1}^K \pi^{ik} (\hat{c}_G^j - \log \sum_{p=0}^C e^{\hat{c}_p^{ik}}) \\ L_{cl}^{Neg}(c) &= - \sum_{i \in Neg} \sum_{k=1}^K \pi^{ik} (\hat{c}_0^i - \log \sum_{p=0}^C e^{\hat{c}_p^{ik}}), \end{aligned} \quad (5)$$

where  $N$  is the number of positive matches,  $K$  is the number of mixtures,  $C$  is the number of classes, with 0 representing the background class  $\hat{c}_0^i$ ,  $\hat{c}_G^j$  is the GT class for the  $j$ -th GT box,  $\hat{c}_p^{ik}$  is the result calculated by Eq. 4,  $\lambda_G^{ij}$  is the same as used in Eq. 3, and  $M$  is the hard negative mining ratio. Instead of using all the negative matches, we sort them using the proposed mixture classification loss and choose top  $M \times N$  as final negative matches for training. In experiments, we set  $M$  to 3 as suggested in [31].

**Final loss:** We define the overall loss to train the object detector using mixture density network as:

$$L = \begin{cases} \frac{1}{N} (L_{loc}(\lambda, l, g)/\eta + L_{cl}^{Pos}(\lambda, c) + L_{cl}^{Neg}(c)), & \text{if } N > 0. \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

where  $N$  is the number of positive matches. In experiments, we set  $\eta$  to 2 as suggested in [6].

At inference, we can compute the coordinates of the bounding box  $R_b$  and the confidence score for each class  $P_i$  by summing the components of the mixture model as follows:

$$\begin{aligned} \text{Localization: } R_b &= \sum_{k=1}^K \pi_b^k \mu_b^k, \\ \text{Classification: } P_i &= \sum_{k=1}^K \pi^k \frac{e^{\mu_i^k}}{\sum_{j=0}^C e^{\mu_j^k}}. \end{aligned} \quad (7)$$

### 3.2. Improving parameter efficiency

In order to predict a probability distribution of the output values, our approach involves modifying the last layer of the network and therefore incurs an increase in the number of parameters, especially in the classification head. More precisely, for an output feature map of size  $F \times F$ , with  $C$  classes,  $D$  anchor boxes, and each bounding box defined using 4 coordinates, the number of parameters in the new layer added to estimate a  $K$ -component GMM with 3 parameters is  $F \times F \times D \times (4 \times 3 \times K)$  for the localization and  $F \times F \times D \times (C \times 2 \times K + K)$  for classification. We see that the number of parameters in the classification head is proportional to the number of classes.

In this section, we focus on improving the efficiency of the algorithm by reducing the number of parameters in the classification head. To this end, as shown in Fig. 2b, we relax the problem of estimating the variance  $\Sigma_p$ , in order to reduce the number of parameters with  $F \times F \times D \times (C \times K + K)$ . Instead, we obtain class probabilities as  $\hat{c}_p^k = \text{Softmax}(\mu_p^k)$ , and use them to estimate the aleatoric uncertainty as follows:

$$u_{al} = \sum_{k=1}^K \pi^k (\text{diag}(\hat{c}_p^k) - (\hat{c}_p^k)^{\otimes 2}), \quad (8)$$

where  $\text{diag}(q)$  is a diagonal matrix with the elements of the vector  $q$  and  $q^{\otimes 2} = qq^T$ . In this case,  $u_{al}$  is  $C \times C$  matrix where the value of each diagonal element can be interpreted as a class-specific aleatoric uncertainty [27].

Finally, we modify the classification loss for training the model with improved parameter efficiency as follows:

$$\begin{aligned} L_{cl}^{Pos}(\lambda, c) &= - \sum_{i \in Pos} \lambda_G^{ij} \sum_{k=1}^K \pi^{ik} (\hat{c}_G^j - \log \sum_{p=0}^C e^{\hat{\mu}_p^{ik}}) \\ L_{cl}^{Neg}(c) &= - \sum_{i \in Neg} \sum_{k=1}^K \pi^{ik} (\hat{c}_0^i - \log \sum_{p=0}^C e^{\hat{\mu}_p^{ik}}), \end{aligned} \quad (9)$$

where all parameters are same as Eq. 5, except for class probability  $\hat{\mu}_p^{ik}$ .

### 3.3. Scoring function

The scoring function in active learning provides a single value per image indicating its informativeness. Our scoring function estimates the informativeness of an image by aggregating all the aleatoric and epistemic uncertainty values for each detected object in the image.

Specifically, let  $U = \{u^{ij}\}$  be the set of uncertainty values (aleatoric or epistemic) of a group of images where  $u^{ij}$  is the uncertainty for the  $j$ -th object in the  $i$ -th image. For localization,  $u^{ij}$  is the maximum value over the 4 bounding box outputs. We first normalize these values using z-score normalization ( $\tilde{u}^{ij} = (u^{ij} - \mu_U)/\sigma_U$ ) to compensate for the fact that the values for the coordinates of the bounding box are unbounded and each uncertainty of an image might have a different range of values. We then assign to each image the maximum uncertainty over the detected objects  $u^i = \max_j \tilde{u}^{ij}$ . We empirically find that taking the maximum over the coordinates and the objects performs better than by taking the average.

Using the algorithm described above, we obtain four different normalized uncertainty values for each image: epistemic and aleatoric for classification and localization,  $\mathbf{u} = \{u_{ep_c}^i, u_{al_c}^i, u_{ep_b}^i, u_{al_b}^i\}$ , respectively. The remaining part is to aggregate these scores into a single one. We explore different combinations of scoring functions that aggregate these uncertainties, including sum or taking the maximum,



Method	Head	(a) VOC07		(b) MS-COCO	
		IoU>0.5	IoU>0.75	IoU>0.5	IoU>0.75
SSD	-	69.29±0.51	43.36±1.24	25.63±0.40	11.93±0.60
SGM	Loc	70.20±0.27	45.39±0.23	27.20±0.08	12.70±0.16
MDN	Loc	70.09±0.22	46.01±0.27	27.67±0.12	13.53±0.05
SGM	Cl	69.95±0.41	44.25±0.26	27.23±0.12	12.50±0.08
MDN	Cl	70.47±0.17	44.47±0.06	27.33±0.09	12.67±0.09
<i>Ours<sub>gmm</sub></i>	Loc+Cl	70.19±0.36	46.11±0.38	27.70±0.08	13.57±0.19
<i>Ours<sub>eff</sub></i>	Loc+Cl	70.45±0.06	46.18±0.26	27.33±0.04	13.33±0.12

Table 1: mAP (in %) of different instances of our approach compared to the original SSD network. *SGM* and *MDN* refer to single and multiple Gaussian models, and we apply those to localization (Loc), classification (Cl), and their combination (Loc+Cl).

like other active learning studies [16, 33]. As we will show in our experiments, taking the maximum over them achieves the highest results.

## 4. Experiments

In this section, we demonstrate the benefits of our approach. We first study the impact of using probabilistic modeling for the object detector and then analyze the proposed scoring function and relevant SOTA approaches in the context of active learning.

**Datasets:** We use PASCAL VOC [9] and MS-COCO [30] datasets. For PASCAL VOC, that contains 20 object categories, we use VOC07 (VOC2007) *trainval* and VOC07+12 *trainval* (union of VOC2007 and VOC2012) for training and evaluate our results on VOC07 *test*. For MS-COCO, that contains 80 object categories, we use MS-COCO *train2014* for training and evaluate our results on *val2017*.

**Experimental settings:** We employ Single Shot MultiBox Detector (SSD) [31], which is widely used in active learning studies [33, 40], with a VGG-16 backbone [37]. We train our models for 120,000 iterations using SGD with a batch size of 32 and a maximum learning rate of 0.001. We use a learning rate warm-up strategy for the first 1,000 iterations and divide the learning rate by 10 after 80,000 and 100,000 iterations. We set the number of Gaussian mixtures to 4, and in the supplementary materials, we provide an ablation study on the number of mixtures. Unless specified otherwise, we report the performance using the average and standard deviation of mAP for three independent trials.

### 4.1. Object detection with probabilistic modeling

We first analyze the impact of using our proposed probabilistic modeling for object detection on PASCAL VOC and MS-COCO. For MS-COCO, we use a random subset of 5,000 training images from *train2014*. We compare the accuracy of our GMM *Ours<sub>gmm</sub>* and the model with improved parameter efficiency *Ours<sub>eff</sub>* to the SSD [31] and several network configurations either using single or multiple Gaussians for the classification or localization heads.

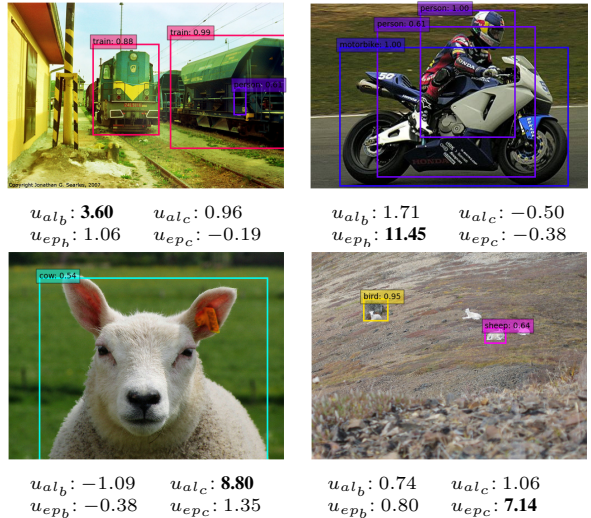


Figure 3: Examples of aleatoric and epistemic uncertainties for inaccurate detections, see more examples in the supplementary material. Starting from top-left image and going in clockwise direction: Person is a false positive; Person bounding box is not correct; A sheep is misclassified as a bird; A sheep is misclassified as a cow.

In Tab. 1a and Tab. 1b, we summarize the results of this experiment performed on VOC07 and MS-COCO, respectively. As shown, all networks that include probabilistic modeling outperform the SSD on both datasets. This is because of the regularization effect of the proposed loss function, which has a loss attenuation due to aleatoric uncertainty [6]. As a result, we obtain models that are robust to noisy data. Considering both the normal (IoU>0.5) and the strict metric (IoU>0.75), *Ours<sub>gmm</sub>* and *Ours<sub>eff</sub>* outperform all other variations on VOC07. On MS-COCO, *Ours<sub>gmm</sub>* outperforms all other instances and the baseline, while *Ours<sub>eff</sub>* reaches competitive results. We expect that the amount of noisy data in MS-COCO is larger than that of PASCAL VOC because MS-COCO has more diverse data. As shown in Eq. 9, there is no aleatoric uncertainty in *Ours<sub>eff</sub>*'s classification loss and therefore, we argue that the regularization by aleatoric uncertainty has a greater effect in MS-COCO.

In Fig. 3, we present representative examples of uncertainty scores for several images where the detector fails to detect the object. As shown, each uncertainty value (bold numbers in Fig. 3) provides a different insight into some particular failure. Localization uncertainties are related to the accuracy of the bounding box prediction, whereas classification uncertainties are related to the accuracy of the category prediction. Interestingly, in these examples, even if the predictions are wrong, uncertainty values seem to be uncorrelated suggesting each uncertainty could predict inaccurate results independently. From these results, we can

Aggregation function	mAP in % (# images)		
	1st (2k)	2nd (3k)	3rd (4k)
<i>random sampling</i>	62.43±0.10	66.36±0.13	68.47±0.09
$u_{alb}$	62.43±0.10	67.06±0.18	68.84±0.18
$u_{epb}$	62.43±0.10	66.75±0.26	69.01±0.17
$u_{alc}$	62.43±0.10	67.09±0.09	68.75±0.08
$u_{epc}$	62.43±0.10	66.51±0.12	68.95±0.13
$\sum_{j \in \{alb, epb\}} u_j$	62.43±0.10	67.01±0.10	68.58±0.29
$\sum_{j \in \{alc, epc\}} u_j$	62.43±0.10	67.07±0.27	69.03±0.20
$\sum_{j \in \{alb, alc\}} u_j$	62.43±0.10	66.96±0.08	68.92±0.23
$\sum_{j \in \{epb, epc\}} u_j$	62.43±0.10	66.49±0.14	68.62±0.24
$\sum_{j \in \{alb, epb, alc, epc\}} u_j$	62.43±0.10	67.04±0.28	69.09±0.30
$\max_{j \in \{alb, epb\}} u_j$	62.43±0.10	66.82±0.21	68.95±0.22
$\max_{j \in \{alc, epc\}} u_j$	62.43±0.10	66.87±0.14	68.99±0.31
$\max_{j \in \{alb, alc\}} u_j$	62.43±0.10	67.18±0.10	69.06±0.25
$\max_{j \in \{epb, epc\}} u_j$	62.43±0.10	66.72±0.10	68.99±0.21
$\max_{j \in \{alb, epb, alc, epc\}} u_j$	62.43±0.10	<b>67.32±0.12</b>	<b>69.43±0.11</b>

Table 2: **VOC07**: Comparison of scoring aggregation functions for active learning based on the aleatoric uncertainty, epistemic uncertainty, and their combination of each task.

conclude that the proposed approach not only computes uncertainty in a single forward pass of a single model but also boosts the performance of the detection network. As shown in the next experiment, combining these values will improve the data selection process during active learning.

## 4.2. Active learning evaluation

We now focus on evaluating the performance of our active learning on PASCAL VOC and MS-COCO datasets. We use an initial set of 2,000 for VOC07, 1,000 for VOC07+12 as suggested by [40], and 5,000 samples in MS-COCO as suggested by [23]. Then, during the active learning stage, for each unlabeled image, we apply non-maximum suppression and we compute the uncertainties for each of the *surviving* objects. The scoring function aggregates these uncertainties using the maximum or sum to provide the final informativeness score for the image. We score the set of unlabeled images and select the 1,000 images [40] with the highest score. Then, we add them to the labeled training set and repeat this process for several active learning cycles. For every active learning iteration, we train the model from scratch, using ImageNet pretrained weight.

**Scoring aggregation function:** We compare the active learning performance obtained using different functions to aggregate the aleatoric and epistemic uncertainties of both classification and localization heads. In particular, we compare seven different instances of our approach with random sampling: 1) Only the aleatoric or epistemic uncertainty on each task; 2) The sum of aleatoric and epistemic uncertainty on the localization or classification head; 3) The sum of aleatoric or epistemic uncertainty on the localization and classification; 4) The sum of aleatoric and epistemic uncertainties for both localization and classification; 5) The maximum of aleatoric and epistemic uncertainty on the localiza-

	Localization		Classification	
	Aleatoric	Epistemic	Aleatoric	Epistemic
$u_{alb}$	100	48	6	11
$u_{epb}$	48	100	7	14
$u_{alc}$	6	7	100	33
$u_{epc}$	11	14	33	100

Table 3: Overlapping ratio (in %) of selected images as a function of the type of uncertainty used.

	mAP in % (# images)			Number of para. ( $\times 10^6$ )	Forward time (sec)
	1st (2k)	2nd (3k)	3rd (4k)		
Random [31]	62.43±0.10	66.36±0.13	68.47±0.09	52.35	0.031
Entropy [33]	62.43±0.10	66.85±0.12	68.70±0.18	52.35	0.031
Core-set [34]	62.43±0.10	66.57±0.20	68.57±0.26	52.35	0.031
LLAL [40]	62.47±0.16	67.02±0.11	68.90±0.15	52.71	0.036
MC-dropout [11]	62.43±0.19	67.10±0.07	69.39±0.09	52.35	0.689
Ensemble [16]	62.43±0.10	67.11±0.26	69.26±0.14	157.05	0.093
<i>Ours<sub>gmm</sub></i>	62.43±0.10	67.32±0.12	69.43±0.11	52.35	0.031
<i>Ours<sub>eff</sub></i>	62.91±0.16	67.61±0.17	69.66±0.17	41.12	0.029

Table 4: **VOC07**: Comparison of mAP and computing cost of active learning with most relevant approaches. *Para.* and *sec* refer to *parameters* and *seconds*, respectively.

tion or classification head; 6) The maximum of aleatoric or epistemic uncertainty on the localization and classification, and 7) The maximum value of these four uncertainties. The results for this comparison are shown in Tab. 2. Our approach using the maximum value of aleatoric and epistemic uncertainties of both localization and classification tasks consistently outperforms all the other aggregation functions on each active learning iterations. More concretely, the maximum value of all uncertainties for both tasks shows better data selection performance in active learning than others. Based on these results, we use the maximum value of all uncertainties as a scoring function during active learning to compare with other active learning studies.

In Tab. 3, we summarize the overlap in the selection as a function of the uncertainty measure. The overlapping ratio using both uncertainties is 48% for localization and 33% for classification. More importantly, when we consider both uncertainties on localization and classification together, the overlapping ratio decreases to barely 14%. This suggests that uncertainty measures obtained for localization and classification are diversified and their combination improves the image selection process.

**Comparison to SOTA on VOC07:** In Tab. 4, we summarize the active learning results and computing cost of our method compared to relevant active learning approaches in the literature. In order to compare the computing cost, we provide the number of parameters and the forward time of each method. In general, a fast forward (backward) step and a small model size leads to a lower training cost and data sampling time during active learning [19, 22].

To focus on the active learning, we reproduce all numbers by applying each sampling method to the proposed GMM architecture where the output is a mixture distribu-

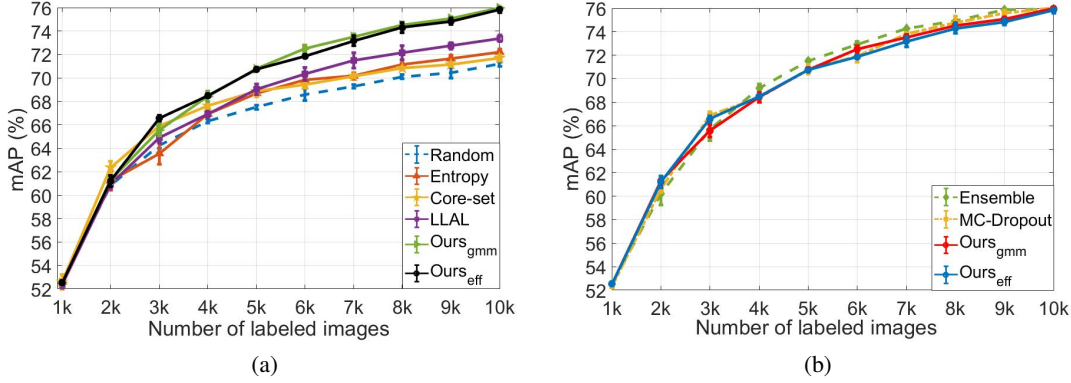


Figure 4: **VOC07+12:** a) Comparison to published works using a single model for scoring. Numbers are taken from [40]; b) Comparison to multiple model-based methods, ensemble and MC-dropout. Details of the numbers to reproduce the plot are in the supplementary material.




Figure 5: **VOC07+12:** Computing cost comparison to baseline and methods using multiple models; Model parameters in millions (M) and the forward time in seconds (sec). See the supplementary material for numerical details.

tion (*i.e.*, the same model as  $Ours_{gmm}$ ). For ensembles, we follow [2], building an ensemble of three independent models. For MC-dropout, we add dropout layers with  $p = 0.1$  to the six convolutional layers composing the extra feature layers in SSD. We compute the image scores using 25 forward passes [2]. For these two methods and entropy-based method, we follow the most common approach in the literature and estimate the final image score as the average entropy on the classification head [16]. For core-set [34], we follow [40], using features of fully connected layer-7 in VGG-16. For LLAL [40], we implement the *learning loss prediction module* on the proposed GMM architecture. As a baseline, we use random sampling on our GMM architecture. Note, we train all methods with exactly the same hyperparameters as mentioned in the *experimental settings*. As shown in Tab. 4, both instances of our approach consistently outperform all the other single-model based methods [31, 33, 34, 40] in every active learning iteration. Compared to multi-model based methods [11, 16], our proposed methods show higher accuracy while requiring a significantly lower computing cost. These results demonstrate that despite having a lower computational cost, our proposed method improves the active learning sampling performance compared to previous works.

**Comparison to SOTA on VOC07+12:** We now compare our approach to existing single-model based approaches on VOC07+12. Here, we consider the SOTA results reported in [40] including LLAL [40] and core-set [34], in addition to entropy [33] and random sampling. We use the same open-source and setting used in [40] for a fair comparison. To solely focus on active learning, we compare the performance based on the same baseline with [40], *i.e.*, SSD. To do this, we train the SSD using a dataset sampled by our proposed scoring function in the proposed GMM architecture and the architecture with improved parameter efficiency. To verify the influence of the initial training set for comparison, we run 5 independent trials with different seeds for the initial choice of the labeled set. We then obtain an average mAP of 0.5246 with a standard deviation of 0.003 that suggests little variations when experiments use a different initial subset of images. As shown in Fig. 4a, our method outperforms all the other single-model based methods. In the last active learning iteration, our approach achieves 0.7598 mAP which is 2.6 percent points higher than the score achieved by LLAL [40] (0.7338 mAP), thus showing a high-performance improvement in active learning based on a single model.

Finally, we compare our approach with methods using multiple models, *i.e.*, ensembles [16] and MC-dropout [11]. For ensembles and MC-dropout, we follow the same design mentioned in Tab. 4 and apply it to the SSD. In Fig. 4b and Fig. 5, we present the accuracy and computational cost comparison of these methods. As shown in Fig. 4b, in terms of the accuracy, our approach performs on par with MC-dropout and ensembles. However, our method uses a single forward pass of a single model to estimate the uncertainties, which is more efficient than ensembles and MC-dropout based methods. With respect to the number of parameters, MC-dropout has the same number of parameters as SSD since dropout layers do not add any new parameters, but it requires multiple forward passes. Our approach adds extra

	mAP in % (# images)			Number of para. ( $\times 10^6$ )	Forward time (sec)
	1st (5k)	2nd (6k)	3rd (7k)		
Random [31]	27.70 $\pm$ 0.08	28.70 $\pm$ 0.13	29.83 $\pm$ 0.04	116.51	0.152
Entropy [33]	27.70 $\pm$ 0.08	28.93 $\pm$ 0.11	29.89 $\pm$ 0.09	116.51	0.152
Core-set [34]	27.70 $\pm$ 0.08	28.99 $\pm$ 0.01	29.93 $\pm$ 0.06	116.51	0.152
LLAL [40]	27.71 $\pm$ 0.03	28.71 $\pm$ 0.06	29.53 $\pm$ 0.15	116.87	0.194
MC-dropout [11]	27.70 $\pm$ 0.10	29.20 $\pm$ 0.09	30.30 $\pm$ 0.08	116.51	3.718
Ensemble [16]	27.70 $\pm$ 0.08	29.03 $\pm$ 0.07	30.02 $\pm$ 0.06	349.53	0.456
<i>Ours<sub>gmm</sub></i>	27.70 $\pm$ 0.08	29.28 $\pm$ 0.05	30.51 $\pm$ 0.12	116.51	0.152
<i>Ours<sub>eff</sub></i>	27.33 $\pm$ 0.04	29.06 $\pm$ 0.08	30.02 $\pm$ 0.05	73.20	0.141

Table 5: **MS-COCO**: Comparison of mAP and computing cost of active learning with most relevant methods. *Para.* and *sec* refer to *parameters* and *seconds*, respectively.

		Baseline [32]	<i>Ours<sub>gmm</sub></i>	<i>Ours<sub>eff</sub></i>
mAP	IoU>0.5	75.31 $\pm$ 0.22	75.90 $\pm$ 0.09	75.80 $\pm$ 0.15
(%)	IoU>0.75	48.70 $\pm$ 0.11	49.36 $\pm$ 0.07	49.83 $\pm$ 0.30
# of parameters (M)		41.17	42.23	41.61
Forward time (sec)		0.059	0.062	0.060

Table 6: **VOC07**: Performance comparison of our mixture models based on Faster-RCNN and the original Faster-RCNN as a baseline [32].

parameters for the estimation of two types of uncertainty to the last layer of each head and therefore, the number of parameters is larger than in SSD. In ensemble-based methods, the number of parameters is proportional to the number of SSD models in the ensemble [28]. As shown in Fig. 5, our method requires significantly less computing cost than MC-dropout and ensemble-based method. In summary, our method provides the best trade-off between accuracy and computing cost for active learning.

**Comparison to SOTA on MS-COCO:** In Tab. 5, we summarize the active learning performance and computing cost of our approaches compared to active learning methods in the literature. To solely focus on the active learning, we reproduce all numbers by applying each sampling method to the proposed GMM architecture (*i.e.*, the same model as *Ours<sub>gmm</sub>*). For all methods, we follow the same settings as on Tab. 4. As shown, both instances of our approach consistently outperform all the other single-model based methods [31, 33, 34, 40] in each active learning cycle. In particular, LLAL [40] shows similar accuracy to random sampling on MS-COCO, because it does not take into account the large diversity of the data and the large number of classes present in the dataset. However, our approach also shows high accuracy on MS-COCO. Compared to multiple model-based methods [11, 16], both instances of our approach require a much less computing cost while *Ours<sub>gmm</sub>* outperforms those methods, and *Ours<sub>eff</sub>* shows competitive results with a much lower computational cost. These results suggest that our approach generalizes to larger dataset that have a larger number of classes.

### 4.3. Scalability and dataset transferability

Our method is not limited to single-stage detectors. Here, in a first experiment, we show how our method can be

Model	Backbone	Random selection	mAP in %	
			<i>Ours<sub>gmm</sub></i> selection	<i>Ours<sub>eff</sub></i> selection
SSD [31]	VGG-16 [37]	67.77 $\pm$ 0.12	68.71 $\pm$ 0.18	68.48 $\pm$ 0.31
	Resnet-34 [17]	65.53 $\pm$ 0.17	67.00 $\pm$ 0.14	67.20 $\pm$ 0.13
	Resnet-50 [17]	64.28 $\pm$ 0.39	65.73 $\pm$ 0.32	65.81 $\pm$ 0.21
Faster-RCNN [32]	Resnet-50-FPN [29]	72.93 $\pm$ 0.41	73.60 $\pm$ 0.18	75.45 $\pm$ 0.30

Table 7: **VOC07**: Transferability of a dataset created using the proposed scoring function and the mixture-based density models. As shown, datasets acquired using our method not only boost the performance of models using a different backbone but also the performance of two-stage detector such as Faster-RCNN.

applied to a two-stage detector such as Faster-RCNN [32] with FPN [29]. For this experiment, we use the same PASCAL VOC dataset as in Tab. 1a. In Tab. 6, we show the summary of the accuracy and the computing cost of our mixture models based on Faster-RCNN and the original Faster-RCNN as a baseline. As shown, both versions of our approach outperform the original model with up to 1.13 mAP improvement. Importantly, in this case, our approach is applied to the output layer of the detection network after region proposal in Faster-RCNN, therefore there is a negligible increase in computing cost and latency because the computation does not include the number of anchor boxes.

Finally, we study the transferability of actively acquired datasets. We compare the performance of SSD using different backbones such as Resnet-34 and Resnet-50 [17], and Faster-RCNN [32] detector trained using our actively sampled dataset. We perform the experiments in the actively sampled dataset from the last active learning cycle in Tab. 4. For completeness, we also report the accuracy obtained using random sampling. We summarize the results of this experiment in Tab. 7. As shown, networks trained using the actively sampled dataset outperform those trained using random sampling with up to 2.52 mAP improvement. Conclusively, our method not only scales to other object detection networks but also datasets actively acquired using our approach can be used to train other architectures.

## 5. Conclusions

We have proposed a novel deep active-learning approach for object detection. Our approach relies on mixture density networks to estimate, in a single forward pass of a single model, two types of uncertainty for both localization and classification tasks, and leverages them in the scoring function. Our proposed probabilistic modeling and scoring function achieve outstanding performance gains in accuracy and computing cost. We present a wide range of experiments on two publicly available datasets, PASCAL VOC and MS-COCO. Besides, our results suggest that our approach scales to new models with different architectures.



## References

- [1] Hamed Habibi Aghdam, Abel Gonzalez-Garcia, Antonio M. López, and Joost van de Weijer. Active learning for deep detection neural networks. In *International Conference on Computer Vision (ICCV)*, 2019.
- [2] William H. Beluch, Tim Genewein, Andreas Nürnberger, and Jan M. Köhler. The power of ensembles for active learning in image classification. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [3] Christopher M Bishop. Mixture density networks. 1994.
- [4] Kashyap Chitta, Jose M. Alvarez, Elmar Haussmann, and Clement Farabet. Less is more: An exploration of data redundancy with active dataset subsampling. *arXiv preprint arXiv:1811.03542*, 2019.
- [5] Kashyap Chitta, Jose M. Alvarez, and Adam Lesnikowski. Large-Scale Visual Active Learning with Deep Probabilistic Ensembles. *arXiv preprint arXiv:1811.03575*, 2018.
- [6] Jiwoong Choi, Dayoung Chun, Hyun Kim, and Hyuk-Jae Lee. Gaussian yolov3: An accurate and fast object detector using localization uncertainty for autonomous driving. In *International Conference on Computer Vision (ICCV)*, 2019.
- [7] Sungjoon Choi, Sanghoon Hong, Kyungjae Lee, and Sungbin Lim. Task agnostic robust learning on corrupt outputs by correlation-guided mixture density networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [8] Sungjoon Choi, Kyungjae Lee, Sungbin Lim, and Songhwa Oh. Uncertainty-aware learning from demonstration using mixture density networks with sampling-free variance modeling. In *International Conference on Robotics and Automation (ICRA)*, 2018.
- [9] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John M. Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *International Journal in Computer Vision (IJCV)*, 88(2):303–338, 2010.
- [10] Di Feng, Lars Rosenbaum, and Klaus Dietmayer. Towards safe autonomous driving: Capture uncertainty in the deep neural network for lidar 3d vehicle detection. In *International Conference on Intelligent Transportation Systems (ITSC)*, 2018.
- [11] Di Feng, Xiao Wei, Lars Rosenbaum, Atsuto Maki, and Klaus Dietmayer. Deep active learning for efficient training of a lidar 3d object detector. In *IEEE Intelligent Vehicles Symposium (IV)*, 2019.
- [12] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning (ICML)*, 2016.
- [13] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *International Conference on Machine Learning (ICML)*, 2017.
- [14] Ross Girshick. Fast r-cnn. In *International Conference on Computer Vision (ICCV)*, 2015.
- [15] Ali Harakeh, Michael Smart, and Steven L Waslander. Bayesod: A bayesian approach for uncertainty estimation in deep object detectors. In *International Conference on Robotics and Automation (ICRA)*, 2020.
- [16] Elmar Haussmann, Michele Fenzi, Kashyap Chitta, Jan Ivanecy, Hanson Xu, Donna Roy, Akshita Mittel, Nicolas Koumchatzky, Clement Farabet, and Jose M Alvarez. Scalable active learning for object detection. In *IEEE Intelligent Vehicles Symposium (IV)*, 2020.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [18] Yihui He and Jianren Wang. Deep multivariate mixture of gaussians for object detection under occlusion. *arXiv preprint arXiv:1911.10614*, 2019.
- [19] Tyler Highlander and Andres Rodriguez. Very efficient training of convolutional neural networks using fast fourier transform and overlap-and-add. *arXiv preprint arXiv:1601.06815*, 2016.
- [20] Stephen C. Hora. Aleatory and epistemic uncertainty in probability elicitation with an example from hazardous waste management. *Reliability Engineering and System Safety*, 54:217–223, 1996.
- [21] Eyke Hüllermeier and Willem Waegeman. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *arXiv preprint arXiv:1910.09457*, 2019.
- [22] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50× fewer parameters and <0.5mb model size. *arXiv preprint arXiv:1602.07360*, 2016.
- [23] Chieh-Chi Kao, Teng-Yok Lee, Pradeep Sen, and Ming-Yu Liu. Localization-aware active learning for object detection. In *Asian Conference on Computer Vision (ACCV)*, 2018.
- [24] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in Neural Information Processing Systems (NeurIPS)*. 2017.
- [25] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [26] Andreas Kirsch, Joost van Amersfoort, and Yarin Gal. Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [27] Yongchan Kwon, Joong-Ho Won, Beom Joon Kim, and Myunghee Cho Paik. Uncertainty quantification using bayesian neural networks in classification: Application to biomedical image segmentation. *Computational Statistics & Data Analysis*, 142:106816, 2020.
- [28] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [29] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [30] Tsung-Yi Lin, Michael Maire, Serge 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 (ECCV)*, 2014.

- [31] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European Conference on Computer Vision (ECCV)*, 2016.
- [32] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.
- [33] Soumya Roy, Asim Unmesh, and Vinay P. Namboodiri. Deep active learning for object detection. In *British Machine Vision Conference (BMVC)*, 2018.
- [34] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *International Conference on Learning Representations (ICLR)*, 2018.
- [35] Burr Settles. *Active Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. 2012.
- [36] Claude E. Shannon. A mathematical theory of communication. *Mobile Computing and Communications Review*, 5(1):3–55, 2001.
- [37] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, 2015.
- [38] Natasa Tagasovska and David Lopez-Paz. Single-model uncertainties for deep learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [39] Ali Varamesh and Tinne Tuytelaars. Mixture dense regression for object detection and human pose estimation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [40] Donggeun Yoo and In So Kweon. Learning loss for active learning. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [41] Jaeyoung Yoo, Geonseok Seo, and Nojun Kwak. Mixture-model-based bounding box density estimation for object detection. *arXiv preprint arXiv:1911.12721*, 2019.