

---

# Being a Bit Frequentist Improves Bayesian Neural Networks

---

Agustinus Kristiadi  
University of Tübingen

Matthias Hein  
University of Tübingen

Philipp Hennig  
University of Tübingen  
MPI for Intelligent Systems, Tübingen

## Abstract

Despite their compelling theoretical properties, Bayesian neural networks (BNNs) tend to perform worse than frequentist methods in classification-based uncertainty quantification (UQ) tasks such as out-of-distribution (OOD) detection. In this paper, based on empirical findings in prior works, we hypothesize that this issue is because even recent Bayesian methods have never considered OOD data in their training processes, even though this “OOD training” technique is an integral part of state-of-the-art frequentist UQ methods. To validate this, we treat OOD data as a first-class citizen in BNN training by exploring four different ways of incorporating OOD data into Bayesian inference. We show in extensive experiments that OOD-trained BNNs are competitive to recent frequentist baselines. This work thus provides strong baselines for future work in Bayesian UQ.

## 1 INTRODUCTION

Uncertainty quantification (UQ) allows learning systems to “know when they do not know”. It is important functionality, especially in safety-critical applications of AI (Amodei et al., 2016), where a system encountering a novel task (which should be associated with high uncertainty) should hesitate or notify a human supervisor. Both the Bayesian and frequentist deep learning communities address similar UQ functionality, but it appears that even recent Bayesian neural networks (BNNs, Osawa et al., 2019; Tomczak et al., 2020; Dusenberry et al., 2020; Izmailov et al.,

---

Proceedings of the 25<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).




Figure 1: Average confidence on uniform OOD test data (lower is better). LA, VB, DE, and OE stand for the Laplace approximation, variational Bayes, Deep Ensemble, and Outlier Exposure, respectively. All methods have similar accuracy and confidence on all the in-distribution test sets, but OE is the only method that achieves low confidence on these outliers.

2021a,b, etc.) tend to underperform compared to the state-of-the-art frequentist UQ methods (Hendrycks et al., 2019; Lee et al., 2018; Hein et al., 2019; Meinke and Hein, 2020; Bitterwolf et al., 2020, etc.). Figure 1 shows this observation: the frequentist Outlier Exposure method (Hendrycks et al., 2019) performs much better than BNNs and even Deep Ensemble (Lakshminarayanan et al., 2017), which has been considered as a strong baseline in Bayesian deep learning (Ovadia et al., 2019).

This paper thus seeks to answer the question of “how can we bring the performance of BNNs on par with that of recent frequentist UQ methods?” Our working hypothesis is that the disparity between them is *not* due to some fundamental advantage of the frequentist viewpoint. Rather, it is due to the more mundane practical fact that recent frequentist UQ methods leverage OOD data in their training process, via the so-called “OOD training” technique. The benefits of this technique are well-studied, both for improving generalization (Zhang and LeCun, 2017) and more recently, for OOD detection (Lee et al., 2018; Hein et al., 2019; Meinke and Hein, 2020; Bitterwolf et al.,

2020). But while OOD data have been used for tuning the hyperparameters of BNNs (Kristiadi et al., 2020a), it appears that even recently proposed deep Bayesian methods have not considered OOD training. A reason for this may be that it is unclear how one can incorporate OOD data in the Bayesian inference itself.

We explore four options—some are philosophically clean, others are heuristics—of incorporating OOD data into Bayesian inference. These methods are motivated by the assumptions that the data (i) have an extra “none class”, (ii) are entirely represented by “soft labels”, or (iii) have mixed “hard” and “soft labels”. Moreover, we also (iv) investigate an interpretation of the popular OE loss as a likelihood that can readily be used in Bayesian inference. We compare the four of them against strong baselines in various UQ tasks and show that BNNs equipped with these likelihoods can outperform a recent OOD-trained frequentist baseline. Our empirical findings thus validate the hypothesis that OOD training is the cause of BNNs underperformance. We hope that the proposed approaches, especially the simple yet best-performing “none-class” method, can serve as strong baselines in the Bayesian deep learning community.

## 2 PRELIMINARIES

We focus on classification tasks. Let  $F : \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}^c$  defined by  $(x, \theta) \mapsto F(x; \theta)$  be an  $\ell$ -layer,  $c$ -class NN with any activation function. Here,  $\mathbb{R}^n$ ,  $\mathbb{R}^d$ , and  $\mathbb{R}^c$  are the input, parameter, and output spaces of the network, respectively. Let  $P(X)$  and  $P(Y|X)$  be unknown probability distributions on  $\mathbb{R}^n$  and  $\{1, \dots, c\}$ , respectively. Given an i.i.d. dataset  $\mathcal{D} := \{(x^{(i)}, y^{(i)})\}_{i=1}^m$  sampled from the previous distributions, the standard training procedure for  $F$  amounts (from the Bayesian perspective) to *maximum a posteriori (MAP) estimation* under a given likelihood  $p(\mathcal{D}|\theta) := \prod_{i=1}^m p(y^{(i)}|x^{(i)}, \theta)$  and a prior  $p(\theta)$ .

Let  $\Delta^c$  be the  $(c - 1)$ -probability simplex. We define the *softmax* inverse-link function  $\sigma : \mathbb{R}^c \rightarrow \Delta^c$  by  $\sigma_k(z) := \exp(z_k) / \sum_{i=1}^c \exp(z_i)$  for each  $k = 1, \dots, c$ . A common choice of the likelihood function for  $c$ -class classification networks is the softmax-Categorical likelihood: Given a datum pair  $(x, y) \in \mathcal{D}$  and parameter  $\theta$ , the mapped output  $\sigma(F(x; \theta))$  can be interpreted as a probability vector and the *Categorical log-likelihood* over it can be defined by  $\log p_{\text{Cat}}(y|x, \theta) := \log \sigma_{y_k}(F(x; \theta))$ , i.e. it is simply the logarithm of the  $y_k$ -th component of the network’s softmax output. Considering all the data, we write  $\log p_{\text{Cat}}(\mathcal{D}|\theta) := \sum_{i=1}^m \log p_{\text{Cat}}(y^{(i)}|x^{(i)}, \theta)$ . Then, given a prior  $p(\theta)$ , the MAP estimate is a particular  $\theta$  satisfying  $\arg \max_{\theta} \log p_{\text{Cat}}(\mathcal{D}|\theta) + \log p(\theta) =$

$\arg \max_{\theta} \log p(\theta|\mathcal{D})$ . In the case of a zero-mean isotropic Gaussian prior, this corresponds to a maximum likelihood under a weight decay regularization.

While MAP-estimated networks are established to achieve high accuracy, they are overconfident (Nguyen et al., 2015). Bayesian methods promise to mitigate this issue (MacKay, 1992a; Wilson and Izmailov, 2020, etc.). The core idea of *Bayesian neural networks (BNNs)* is to infer the full posterior  $p(\theta|\mathcal{D})$  (instead of just a single estimate of  $\theta$  in the case of MAP estimation) and marginalizing it to obtain the predictive distribution  $p(y|x, \mathcal{D}) = \int \sigma(F(x; \theta)) p(\theta|\mathcal{D}) d\theta$ . Unfortunately, due to the nonlinearity of  $F$  in  $\theta$ , this exact posterior is intractable, and various approximations have been proposed. Prominent among them are variational Bayes (VB, Hinton and Van Camp, 1993; Graves, 2011; Blundell et al., 2015, etc.) and Laplace approximations (LAs, MacKay, 1992b; Ritter et al., 2018, etc.). Once an approximation  $q(\theta) \approx p(\theta|\mathcal{D})$  has been obtained, one can easily use it to approximate  $p(y|x, \mathcal{D})$  via Monte Carlo (MC) integration.

## 3 WHY BNNs UNDERPERFORM

Let  $U$  be the *data region*, i.e. a subset of the input space  $\mathbb{R}^n$  where the distribution  $P(X)$  assigns non-negligible mass. Suppose  $V := \mathbb{R}^n \setminus U$  is the remaining subset of the input space  $\mathbb{R}^n$  that has low mass under  $P(X)$ , i.e. it is the *OOD region*. It has recently been shown that any point estimate of  $F$  can induce an arbitrarily overconfident prediction on  $V$  (Hein et al., 2019; Nguyen et al., 2015). While Bayesian methods have been shown to “fix” this issue in the asymptotic regime (i.e. when the distance between a test point  $x \in V$  and the data region  $U$  tends to infinity; see Kristiadi et al., 2020a,b), no such guarantee has been shown for *non-asymptotic* regime which contains outliers that are relatively close to  $U$ . In fact, empirical evidence shows that BNNs often yield suboptimal results in this regime, as Fig. 1 shows.

In an adjacent field, the frequentist community has proposed a technique—referred to here as *OOD training*—to address this issue. The core idea is to “expose” the network to a particular kind of OOD data and let it generalize to unseen outliers. Suppose  $\mathcal{D}_{\text{out}} := \{\hat{x}_i \in V\}_{i=1}^{m_{\text{out}}}$  is a collection of  $m_{\text{out}}$  points sampled from some distribution on  $V$ . Then, one can incorporate these OOD samples into the standard MAP objective via an additional objective function  $\mathcal{L}$  that depends on the network and  $\mathcal{D}_{\text{out}}$ , but not the dataset  $\mathcal{D}$ . We thus do the following:

$$\arg \max_{\theta} \log p_{\text{Cat}}(\mathcal{D}|\theta) + \log p(\theta) + \mathcal{L}(\theta; \mathcal{D}_{\text{out}}). \quad (1)$$

For instance, Hendrycks et al. (2019) define  $\mathcal{L}$  to be

the negative cross-entropy between the softmax output of  $F$  under  $\mathcal{D}_{\text{out}}$  and the uniform discrete distribution. Intuitively, (1) tries to find a parameter vector of  $F$  that induces well-calibrated predictions both inside and outside of  $U$ . In particular, ideally, the network should retain the performance of the MAP estimate in  $U$ , while attaining the maximum entropy or uniform confidence prediction everywhere in  $V$ . Empirically, this frequentist robust training scheme obtains state-of-the-art performance in OOD detection benchmarks (Hendrycks et al., 2019; Meinke and Hein, 2020; Bitterwolf et al., 2020, etc.).

While some works have employed OOD data for tuning the hyperparameters of BNNs (Kristiadi et al., 2020a,b), ultimately OOD data are not considered as a first-class citizen in the Bayesian inference itself. Furthermore, while one can argue that theoretically, BNNs can automatically assign high uncertainty over  $V$  and thus robust to outliers, as we have previously discussed, empirical evidence suggests otherwise. Altogether, it thus now seems likely that indeed the fact that BNNs are not exposed to OOD data during training is a major factor contributing to the discrepancy in their UQ performance compared to that of the state-of-the-art frequentist methods.

## 4 OOD TRAINING FOR BNNs

Motivated by the hypothesis laid out in the previous section, our goal here is to come up with an OOD training scheme for standard BNN inference while retaining a reasonable Bayesian interpretation. To this end, we explore four different ways of incorporating OOD data in Bayesian inference by making different assumptions about the data and hence the likelihood, starting from the most philosophically clean to the most heuristic.<sup>1</sup>

### Method 1: Extra “None Class”

The most straightforward yet philosophically clean way to incorporate unlabeled OOD data is by adding an extra class, corresponding to the “none class”—also known as the “dustbin” or “garbage class” (Zhang and LeCun, 2017). That is, we redefine our network  $F$  as a function  $\mathbb{R}^n \times \mathbb{R}^{d+\hat{d}} \rightarrow \mathbb{R}^{c+1}$  where  $\hat{d}$  is the number of additional parameters in the last layer associated with the extra class. Note that this is different from the “background class” method (Zhang and LeCun, 2017; Wang and Aitchison, 2021) which assumes that the extra class is tied to the rest of the classes and thus does

<sup>1</sup>One might be tempted to treat  $\mathcal{L}$  in (1) as a log-prior. However not only does this mean that we have a (controversial) data-dependent prior, but also it introduces implementation issues, e.g. the KL-divergence term in VB’s objective cannot be computed easily anymore.

not have additional parameters. We choose to use the “dustbin class” since Zhang and LeCun (2017) showed that it is the better of the two.

Under this assumption, we only need to label all OOD data in  $\mathcal{D}_{\text{out}}$  with the class  $c+1$  and add them to the true dataset  $\mathcal{D}$ . That is, the new dataset is  $\tilde{\mathcal{D}} := \mathcal{D} \amalg \{(x_{\text{out}}^{(1)}, c+1), \dots, (x_{\text{out}}^{(m_{\text{out}})}, c+1)\}$ , where  $\amalg$  denotes disjoint union. Under this setting, we can directly use the Categorical likelihood, and thus, a BNN with this assumption has a sound Bayesian interpretation.

### Method 2: Soft Labels

In this method, we simply assume that the data have “soft labels”, i.e. the labels are treated as general probability vectors, instead of restricted to integer labels (Thiel, 2008).<sup>2</sup> Thus, we can assume that the target  $Y$  is a  $\Delta^c$ -valued random variable. Under this assumption, since one-hot vectors are also elements of  $\Delta^c$  (they represent the  $c$  corners of  $\Delta^c$ ), we do not have to redefine  $\mathcal{D}$  other than to one-hot encode the original integer labels.

Now let us turn our attention to the OOD training data. The fact that these data should be predicted with maximum entropy suggests that the suitable label for any  $x_{\text{out}} \in \mathcal{D}_{\text{out}}$  is the uniform probability vector  $u := (1/c, \dots, 1/c)$  of length  $c$ —the center of  $\Delta^c$ . Thus, we can redefine  $\mathcal{D}_{\text{out}}$  as the set  $\{(x_{\text{out}}^{(i)}, u)\}_{i=1}^m$ , and then define a new joint dataset  $\tilde{\mathcal{D}} := \mathcal{D} \amalg \mathcal{D}_{\text{out}}$  containing both the soft-labeled in- and out-distribution training data. Note that without the assumption that  $Y$  is a simplex-valued random variable, we cannot assign the label  $u$  to the OOD training data, and thus we cannot naturally convey our intuition that we should be maximally uncertain over OOD data.

Under the previous assumption, we have to adapt the likelihood. A straightforward choice for simplex-valued random variables is the Dirichlet likelihood  $p_{\text{Dir}}(y|x, \theta) := \text{Dir}(y|\alpha(F(x; \theta)))$  where we have made the dependence of  $\alpha$  to the network output  $F(x; \theta)$  explicit. So, we obtain the log-likelihood function

$$\log p_{\text{Dir}}(y|x, \theta) = \log \Gamma(\alpha_0) - \log \Gamma(\alpha_k(F(x; \theta))) + \sum_{k=1}^c (\alpha_k(F(x; \theta)) - 1) \log y_k, \quad (2)$$

where  $\alpha_0 := \sum_{k=1}^c \alpha_k(F(x; \theta))$  and  $\Gamma$  is the Gamma function. Therefore, the log-likelihood for  $\tilde{\mathcal{D}}$  is given by  $\log p_{\text{Dir}}(\tilde{\mathcal{D}}|x, \theta) = \sum_{i=1}^m \log p_{\text{Dir}}(y^{(i)}|x^{(i)}, \theta) + \sum_{i=1}^{m_{\text{out}}} \log p_{\text{Dir}}(u|x_{\text{out}}^{(i)}, \theta)$ , which can readily be used in a Bayesian inference.

<sup>2</sup>The term “soft label” here is different than “fuzzy label” (Kuncheva, 2000; El Gayar et al., 2006) where it is not constrained to sum to one.

One thing left to discuss is the definition of  $\alpha(F(x; \theta))$ . An option is to decompose it into the mean and precision (Minka, 2000). We do so by writing  $\alpha_k(F(x; \theta)) = \gamma \sigma_k(F(x; \theta))$  for each  $k = 1, \dots, c$ , where  $\gamma$  is the precision (treated as a hyperparameter) and the softmax output  $\sigma(F(x; \theta))$  represents the mean—which is valid since it is an element of  $\Delta^c$ . The benefits are two-fold: First, since we focus solely on the mean, it is easier for optimization (Minka, 2000). Indeed, we found that the alternatives, such as  $\alpha_k(F(x; \theta)) = \exp(F_k(x; \theta))$  yield worse results. Second, after training, we can use the softmax output of  $F$  as usual without additional steps, i.e. when making a prediction, we can treat the network as if it was trained using the standard softmax-Categorical likelihood.

### Method 3: Mixed Labels

There is a technical issue when using the Dirichlet likelihood for the in-distribution data: It is known that the Dirichlet likelihood does not work well with one-hot encoded vectors and that it is harder to optimize than the Categorical likelihood (Malinin and Gales, 2018). To see this, notice in (2) that the logarithm is applied to  $y_k$ , in contrast to  $\sigma_k(F(x; \theta))$  in the Categorical likelihood. If  $y$  is a one-hot encoded vector, this implies that for all but one  $k \in \{1, \dots, c\}$ , the expression  $\log y_k$  is undefined and thus the entire log-likelihood also is. While one can mitigate this issue via e.g. label smoothing (Malinin and Gales, 2018; Szegedy et al., 2016), ultimately we found that models with the Dirichlet likelihood generalize worse than their Categorical counterparts (more in Section 6). Fortunately, the Dirichlet log-likelihood (2) does not suffer from this issue when used for OOD data because their label  $u$  is the uniform probability vector—in particular, all components of  $u$  are strictly larger than zero.

Motivated by these observations, we combine the best of both worlds in the stability of the Categorical likelihood in modeling “hard” one-hot encoded labels (or equivalently, integer labels) and the flexibility of the Dirichlet likelihood in modeling soft labels. To this end, we assume that all the in-distribution data in  $\mathcal{D}$  have the standard integer labels, while all the OOD data in  $\mathcal{D}_{\text{out}}$  have soft labels. Then, assuming  $\tilde{\mathcal{D}} = \mathcal{D} \amalg \mathcal{D}_{\text{out}}$ , we define the following “mixed” log-likelihood:

$$\begin{aligned} \log p(\tilde{\mathcal{D}}|\theta) &:= \sum_{i=1}^m \log p_{\text{Cat}}(y^{(i)}|x^{(i)}, \theta) \\ &\quad + \sum_{i=1}^{m_{\text{out}}} \log p_{\text{Dir}}(u|x_{\text{out}}^{(i)}, \theta). \end{aligned}$$

The implicit assumption of this formulation is that, unlike the two previous methods, we have two distinct

generative processes for generating the labels of input points in  $U$  and  $V$ . Data in  $\mathcal{D}$  can thus have a different “data type” than data in  $\mathcal{D}_{\text{out}}$ . This method can therefore be interpreted as solving a multi-task or multi-modal learning problem.

### Method 4: Frequentist-Loss Likelihood

Considering the effectiveness of frequentist methods, it is thus tempting to give a direct Bayesian treatment upon them. But to do so, we first have to find a sound probabilistic justification of  $\mathcal{L}$  in (1) since not all loss functions can be interpreted as likelihood. We use the OE objective (Hendrycks et al., 2019) as a use case.

First, recall that OE’s OOD objective—the last term in (1)—is given by

$$\begin{aligned} \mathcal{L}_{\text{OE}}(\theta; \mathcal{D}_{\text{out}}) &:= - \mathbb{E}_{x_{\text{out}} \sim \mathcal{D}_{\text{out}}} (H(\sigma(F(x_{\text{out}}; \theta)), u)) \\ &= \frac{1}{c m_{\text{out}}} \sum_{i=1}^{m_{\text{out}}} \sum_{k=1}^c \log \sigma_k(F(x_{\text{out}}^{(i)}; \theta)), \end{aligned} \quad (3)$$

where  $u := (1/c, \dots, 1/c) \in \Delta^c$  is the uniform probability vector of length  $c$  and  $H$  is the functional measuring the *cross-entropy* between its two arguments. Our goal here is to interpret (3) as a log-likelihood function: we aim at finding a log-likelihood function  $\log p(\mathcal{D}_{\text{out}}|\theta)$  over  $\mathcal{D}_{\text{out}}$  that has the form of  $\mathcal{L}_{\text{OE}}$ . This is sufficient for defining the overall likelihood over  $\mathcal{D}$  and  $\mathcal{D}_{\text{out}}$  since given this function and assuming that these datasets are independent, we readily have a probabilistic interpretation of the log-likelihood terms in (1):  $\log p(\mathcal{D}, \mathcal{D}_{\text{out}}|\theta) = \log p_{\text{Cat}}(\mathcal{D}|\theta) + \log p(\mathcal{D}_{\text{out}}|\theta)$ .

We begin with the assumption that the Categorical likelihood is used to model both the in- and out-of-distribution data—in particular, we use the standard integer labels for both  $\mathcal{D}$  and  $\mathcal{D}_{\text{out}}$ . Now, recall that the OOD data ideally have the uniform confidence, that is, they are equally likely under all possible labels. But since we have assumed hard labels, we cannot use  $u$  directly as the label for  $x_{\text{out}} \in \mathcal{D}_{\text{out}}$ . To circumvent this, we redefine the OOD dataset  $\mathcal{D}_{\text{out}}$  by assigning all  $c$  possible labels to each  $x_{\text{out}}$ , based on the intuition that  $F$  should be “maximally confused” about the correct label of  $x_{\text{out}}$ .

$$\begin{aligned} \mathcal{D}_{\text{out}} &:= \{(x_{\text{out}}^{(1)}, 1), \dots, (x_{\text{out}}^{(1)}, K), \dots, \\ &\quad (x_{\text{out}}^{(m_{\text{out}})}, 1), \dots, (x_{\text{out}}^{(m_{\text{out}})}, K)\}. \end{aligned} \quad (4)$$

Thus, given  $m_{\text{out}}$  unlabeled OOD data, we have  $|\mathcal{D}_{\text{out}}| = c m_{\text{out}}$  OOD data points in our OOD training set. So, the negative log-Categorical likelihood over  $\mathcal{D}_{\text{out}}$  is given by

$$- \log p(\mathcal{D}_{\text{out}}|\theta) = \sum_{i=1}^{m_{\text{out}}} \sum_{k=1}^c \log \sigma_k(F(x_{\text{out}}^{(i)}; \theta)) \quad (5)$$



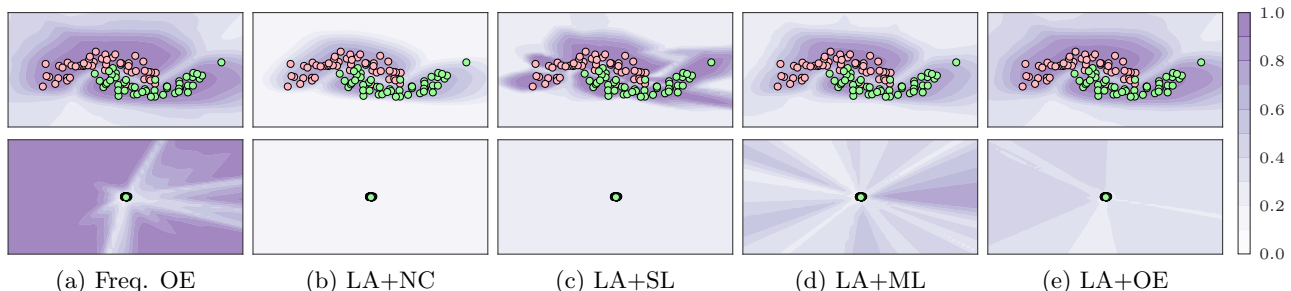


Figure 2: Confidence estimate of the frequentist OE baseline (a) and the Bayesian OOD-training methods discussed in Section 4 (b-e). For the BNNs, we use the Laplace approximation (LA). The suffixes “+NC”, “+SL”, “+ML”, and “+OE” refer to the “none class”, “soft labels”, “mixed labels”, and “OE likelihoods” methods, respectively. The bottom row shows the zoomed-out versions of the top one. The OOD training data is sampled from  $\text{Unif}[-6, 6]^2$ . While OE yields calibrated uncertainty near the training data (a), it is overconfident far from them. Meanwhile, Bayesian methods are less overconfident far away from the data (b-e).

Comparing this to (3), we identify that  $\log p(\mathcal{D}_{\text{out}}|\theta)$  is exactly  $\mathcal{L}_{\text{OE}}$ , up to a constant factor  $1/(cm_{\text{out}})$ , which can be thought of as a tempering factor to  $p(\mathcal{D}_{\text{out}}|\theta)$ . We have thus obtained the probabilistic interpretation of OE’s objective—this likelihood can then be soundly used in a Bayesian inference—albeit arising from applying a heuristic (4) to the data.

#### 4.1 Remark

Here, we consider the question of whether there is an inherent advantage of using OOD-trained BNNs instead of the standard OOD-trained network. One answer to this question is given by the recent finding that Bayesian methods naturally yield low uncertainty in regions far away from the data (Kristiadi et al., 2020a). In contrast, OOD-trained point-estimated networks do not enjoy such a guarantee by default and must resort to e.g. generative models (Meinke and Hein, 2020). We illustrate this observation synthetically in Fig. 2.

## 5 RELATED WORK

OOD training for BNNs has recently been used for tuning the hyperparameters of LAs (Kristiadi et al., 2020a, 2021). However, it appears that OOD training is not commonly utilized by BNNs in the Bayesian inference itself. Wang and Aitchison (2021), at the same time window as this work, also proposed OOD training for BNNs by justifying the presence of OOD training data as a consequence of the data curation process. Nevertheless, their method is different than all the methods proposed here, as discussed in Section 4, and only validated on a single BNN. Meanwhile, we explore four distinct methods and extensively validate them on various BNNs, see Section 6.

From an adjacent field, adversarial training for BNNs has recently been studied. In particular, Liu et al. (2019) specifically employ VB and modify the first term of the ELBO to take into account the worst-case perturbation of each data point, which can be thought of as a particular type of OOD data. Unlike theirs, our methods are for general OOD data and are agnostic to the approximate inference method.

Non-Bayesian Dirichlet-based models have recently been studied for UQ (Sensoy et al., 2018). Similar to our proposed “soft labels” and “mixed labels” likelihoods, Malinin and Gales (2018, 2019); Nandy et al. (2020) use the Dirichlet distribution as the output of a non-Bayesian network and employ OOD training via a custom, non-standard loss. Their methods’ Bayesian interpretation is therefore unclear. In contrast, for modeling soft labels, we simply use the standard Dirichlet log-likelihood function, which is well-studied in the context of generalized linear models (Gueorguieva et al., 2008). Our methods thus retain a clear Bayesian interpretation when used in BNNs.

## 6 EXPERIMENTS

### 6.1 Setup

**Baselines** We use the following strong, recent baselines to represent non-Bayesian methods:<sup>3</sup> (i) standard MAP-trained network (**MAP**), (ii) Deep Ensemble (**DE**, Lakshminarayanan et al., 2017), and (iii) Outlier Exposure (**OE**, Hendrycks et al., 2019). Note that DE and OE are among the established state-of-the-art frequentist UQ methods.

For standard Bayesian methods, i.e. those considering

<sup>3</sup>Deep Ensemble can also be seen as a Bayesian method, but it was originally proposed as a frequentist method.

Table 1: Test accuracy ( $\uparrow$ ) / ECE ( $\downarrow$ ), averaged over five prediction runs. Best values in each categories (separated by horizontal lines) are in bold.

	MNIST	F-MNIST	SVHN	CIFAR-10	CIFAR-100
MAP	99.4±0.0 / 6.4±0.0	92.4±0.0 / 13.9±0.0	97.4±0.0 / 8.9±0.0	94.8±0.0 / 10.0±0.0	76.7±0.0 / 14.3±0.0
DE	<b>99.5±0.0</b> / 8.6±0.0	<b>93.6±0.0</b> / 3.6±0.0	<b>97.6±0.0</b> / <b>3.5±0.0</b>	<b>95.7±0.0</b> / <b>4.5±0.0</b>	<b>80.0±0.0</b> / <b>1.9±0.0</b>
OE	99.4±0.0 / <b>5.3±0.0</b>	92.3±0.0 / 12.1±0.0	97.4±0.0 / 10.6±0.0	94.6±0.0 / 13.2±0.0	76.7±0.0 / 15.0±0.0
VB	<b>99.5±0.0</b> / 11.2±0.3	92.4±0.0 / 3.7±0.2	97.5±0.0 / 5.7±0.2	94.9±0.0 / 5.8±0.2	<b>75.4±0.0</b> / <b>8.3±0.0</b>
+NC	99.4±0.0 / 12.6±0.3	92.2±0.0 / 3.3±0.1	97.5±0.0 / <b>4.1±0.1</b>	94.4±0.0 / 5.5±0.1	74.1±0.0 / 10.7±0.1
+SL	<b>99.5±0.0</b> / 10.5±0.3	<b>93.1±0.0</b> / 6.3±0.1	<b>97.6±0.0</b> / 9.3±0.2	93.0±0.0 / 11.0±0.1	71.4±0.0 / 13.0±0.0
+ML	99.3±0.0 / 11.8±0.2	92.0±0.0 / <b>2.5±0.1</b>	<b>97.6±0.0</b> / 4.2±0.0	<b>95.0±0.0</b> / 4.9±0.2	<b>75.4±0.0</b> / 10.4±0.0
+OE	99.4±0.0 / <b>10.0±0.2</b>	92.3±0.0 / 3.0±0.2	<b>97.6±0.0</b> / 5.7±0.2	94.8±0.0 / <b>4.6±0.2</b>	74.2±0.0 / 8.9±0.0
LA	99.4±0.0 / 7.6±0.1	92.5±0.0 / 11.3±0.2	97.4±0.0 / 3.3±0.3	<b>94.8±0.0</b> / 7.5±0.3	76.6±0.1 / 8.3±0.1
+NC	99.4±0.0 / 5.4±0.7	92.4±0.0 / 8.5±0.3	97.3±0.0 / 4.6±0.2	94.0±0.0 / <b>6.6±0.3</b>	76.2±0.0 / 6.1±0.0
+SL	<b>99.7±0.0</b> / 12.1±1.1	<b>93.2±0.0</b> / <b>3.2±0.3</b>	<b>97.5±0.0</b> / 7.4±0.2	93.6±0.0 / 10.2±0.2	72.3±0.1 / 7.1±0.2
+ML	99.4±0.0 / 7.5±1.0	92.5±0.0 / 5.9±0.2	97.4±0.0 / <b>2.9±0.2</b>	<b>94.8±0.0</b> / 6.9±0.3	76.5±0.1 / <b>4.4±0.1</b>
+OE	99.4±0.0 / <b>4.8±0.7</b>	92.3±0.0 / 7.4±0.1	97.4±0.0 / 3.2±0.1	94.6±0.0 / 8.8±0.1	<b>76.7±0.1</b> / <b>4.4±0.1</b>

only  $\mathcal{D}$  in the inference, we use (iv) the all-layer diagonal Laplace approximation on top of the MAP network (**LA**) and (v) the last-layer mean-field VB (**VB**, Graves, 2011; Blundell et al., 2015). We mainly use only these simple Bayesian methods to validate that the proposed likelihood could make even these crudely approximated BNNs competitive to the strong baselines. Results with more advanced BNNs are in Tab. 5.

To represent our methods, we again use the same LA and VB but with the modifications proposed in Section 4. We denote these modified methods **LA+X** and **VB+X**, respectively. Here, **X** is the abbreviation for the proposed methods, i.e. **NC** for “none class” (Method 1), **SL** for “soft label” (Method 2), **ML** for “mixed label” (Method 3), and **OE** for “OE likelihood” (Method 4).

Finally, we use the LeNet and WideResNet-16-4 architectures, trained in the usual manner—see Appendix B. Source code is available at [https://github.com/wiseodd/bayesian\\_ood\\_training](https://github.com/wiseodd/bayesian_ood_training).

**Datasets** As the in-distribution datasets, we use: (i) MNIST, (ii) Fashion-MNIST (F-MNIST), (iii) SVHN, (iv) CIFAR-10, and (v) CIFAR-100. For each of them, we obtain a validation set of size 2000 by randomly splitting the test set. For methods requiring OOD training data, i.e. OE, LA+X, and VB+X, we use the 32×32 downsampled ImageNet dataset (Chrabaszcz et al., 2017) as an alternative to the 80M Tiny Images dataset used by Hendrycks et al. (2019); Meinke and Hein (2020), since the latter is not available anymore.<sup>4</sup> For OOD detection tasks, we use various *unseen* (i.e. not used for training or tuning) OOD test sets as used in (Meinke and Hein, 2020; Hein et al., 2019), both real-world (e.g. E-MNIST) and synthetic (e.g. uniform noise). For text classification, we use the Stan-

ford Sentiment Treebank (SST, Socher et al., 2013) and the TREC dataset (Voorhees, 2001). We detail of all OOD test sets in Appendix A. We furthermore test the methods in a dataset-shift robustness task using the corrupted CIFAR-10 (CIFAR-10-C) dataset (Ovadia et al., 2019; Hendrycks and Dietterich, 2019).

**Metrics** To measure OOD detection performance, we use the standard FPR95 metric, which measures the false positive rate at 95% true positive rate. Other metrics such as average confidence and area under the ROC curve are presented in the appendix. Meanwhile, to measure dataset-shift robustness and predictive performance, we use test accuracy and expected calibration error (ECE) with 15 bins (Naeini et al., 2015).

## 6.2 Generalization and Calibration

We present the generalization and calibration performance in Table 1. We note that generally, all methods discussed in Section 4 attain comparable accuracy to, and are better calibrated than the vanilla MAP/OE models. However, the “soft label” method tends to underperform in both accuracy and ECE—this can be seen clearly on CIFAR-100. This issue appears to be because of the numerical issue we have discussed in Section 4. Note that this issue seems to also plague other Dirichlet-based methods (Malinin and Gales, 2018, 2019). Overall, it appears that Bayesian OOD training with NC, ML, and OE is not harmful to the in-distribution performance—they are even more calibrated than the frequentist OE.

## 6.3 OOD Detection

We present the OOD detection results on image classification datasets in Table 2.<sup>5</sup> As indicated in Fig. 1,

<sup>4</sup><https://groups.csail.mit.edu/vision/TinyImages/>.

<sup>5</sup>Refer to Appendix C for the detailed, non-averaged results for Tables 2 to 5, along with additional metrics.

Table 2: OOD data detection in terms of FPR95. Values are averages over six OOD test sets and five prediction runs—lower is better. The best values of each category are in bold. Details are in Table 6 in the appendix.

Methods	MNIST	F-MNIST	SVHN	CIFAR-10	CIFAR-100
MAP	17.7	69.4	22.4	52.4	81.0
DE	10.6	61.4	10.1	32.3	73.3
OE	<b>5.4</b>	<b>16.2</b>	<b>2.1</b>	<b>22.8</b>	<b>54.0</b>
VB	25.7	63.3	22.0	36.5	77.6
+NC	7.5	15.0	<b>1.4</b>	<b>28.0</b>	<b>49.9</b>
+SL	<b>2.7</b>	<b>4.2</b>	1.8	40.4	62.3
+ML	7.4	19.6	<b>1.4</b>	29.1	50.2
+OE	6.8	22.4	1.5	29.8	53.3
LA	19.4	68.7	17.1	53.6	81.3
+NC	6.6	8.3	1.5	<b>20.1</b>	<b>47.4</b>
+SL	<b>2.2</b>	<b>4.1</b>	<b>1.0</b>	38.5	60.9
+ML	5.5	14.3	1.1	21.8	52.5
+OE	5.4	17.0	1.1	23.3	53.9

Table 3: OOD data detection on text classification tasks. Values are averages over five prediction runs and additionally, three OOD test sets for FPR95. Details are in the appendix (Tables 12 and 13).

Methods	ECE		FPR95	
	SST	TREC	SST	TREC
MAP	20.8	17.2	100.0	96.3
DE	<b>2.5</b>	10.6	100.0	24.2
OE	13.0	<b>9.4</b>	<b>0.0</b>	<b>0.0</b>
LA	21.0	17.3	100.0	96.4
+NC	17.9	18.6	<b>0.0</b>	<b>0.0</b>
+SL	17.5	10.4	95.3	0.8
+ML	<b>11.4</b>	11.5	84.6	<b>0.0</b>
+OE	12.8	<b>8.4</b>	<b>0.0</b>	<b>0.0</b>

OE is in general significantly better than even DE while retaining the computational efficiency of MAP. The vanilla Bayesian baselines, represented by VB and LA, achieve worse results than DE (and thus OE). But, when OOD training is employed to train these BNNs using the four methods we considered in Section 4, their performance improves. We observe that all Bayesian OOD training methods generally yield better results than DE and become competitive to OE. In particular, while the “soft label” method (SL) is best for “easy” datasets (MNIST, F-MNIST), we found that the simplest “none class” method (NC) achieves the best results in general.

In Table 3, we additionally show the results on text classification datasets. We found that the OOD training methods consistently improve both the calibration and OOD-detection performance of the vanilla Bayesian methods, making them on par with OE. As before, the “none class” method performs well in OOD detection. This is a reassuring result since NC is

 Table 4: Average ECE and FPR95 under models trained with synthetic noises as  $\mathcal{D}_{\text{out}}$ . Both values are averaged over five prediction runs, and additionally six OOD test sets for FPR95. See Tables 10 and 11 in the appendix for details.

Methods	SVHN	CIFAR-10	CIFAR-100
<b>ECE ↓</b>			
MAP	<b>8.9</b>	<b>10.0</b>	<b>14.3</b>
OE	<b>8.9</b>	11.5	16.1
LA	<b>3.3</b>	7.5	8.3
+NC	5.0	8.3	<b>3.8</b>
+SL	13.5	16.0	4.0
+ML	7.4	<b>7.2</b>	3.3
+OE	3.8	<b>7.2</b>	8.6
<b>FPR95 ↓</b>			
MAP	22.4	52.4	81.0
OE	<b>11.4</b>	<b>31.0</b>	<b>60.1</b>
LA	17.1	53.6	81.3
+NC	10.5	<b>26.4</b>	64.5
+SL	93.7	37.9	68.6
+ML	14.4	28.4	61.0
+OE	<b>10.1</b>	35.3	<b>56.4</b>

also the most philosophically clean (i.e. requires fewer heuristics) than the other three methods considered.

A common concern regarding OOD training is the choice of  $\mathcal{D}_{\text{out}}$ . As an attempt to address this, in Table 4 we provide results on OOD detection when the model is trained using a synthetic noise dataset. The noise dataset used here is the “smooth noise” dataset (Hein et al., 2019), obtained by permuting, blurring, and contrast-rescaling the original training dataset. We found that even with such a simple OOD dataset, we can still generally obtain better OOD detection results than OE, as shown by the FPR95 values. Moreover, the combination of Bayesian formalism and OOD




Figure 3: Dataset shift performance on the corrupted-CIFAR-10-C dataset, following Ovadia et al. (2019). Values are ECE—lower is better.

Table 5: Average ECE and FPR95 with more sophisticated base models. “C-10” stands for CIFAR-10, while “C-100” for CIFAR-100

Methods	ECE		FPR95	
	C-10	C-100	C-10	C-100
Flipout	10.9	19.8	65.0	85.4
+NC	<b>8.2</b>	<b>13.8</b>	<b>40.9</b>	<b>56.2</b>
CSGHMC	<b>1.7</b>	4.0	60.3	81.0
+NC	6.2	<b>2.4</b>	<b>25.0</b>	<b>43.0</b>
DE	<b>4.5</b>	1.9	32.3	73.3
+NC	4.8	<b>1.7</b>	<b>17.0</b>	<b>44.4</b>

training is beneficial in calibrating the in-distribution uncertainty: We found that using this  $\mathcal{D}_{\text{out}}$  to train OE yields worse-calibrated results than even the vanilla MAP model, as the ECE values show. In contrast, OOD-trained LA yields better ECE results in general.

Finally, we show that OOD training is beneficial for other BNNs. In Table 5, we consider two recent (all-layer) BNNs: a VB with the flipout estimator (Flipout, Wen et al., 2018) and the cyclical stochastic-gradient Hamiltonian Monte Carlo (CSGHMC, Zhang et al., 2020). Evidently, OOD training improves their OOD detection performance by a large margin. Moreover, OOD training also improves the performance of DE.

#### 6.4 Dataset-Shift Robustness

In this UQ task, OOD training is beneficial for both MAP and the vanilla Bayesian methods (VB, LA), making them competitive to the state-of-the-art DE’s performance in larger severity levels, see Fig. 3. Moreover, the OOD-trained VB and LA are in general more calibrated than OE, which shows the benefit of the Bayesian formalism vis-à-vis the point-estimated OE. This indicates that *both* being Bayesian and considering OOD data during training are beneficial.

Even though it is the best in OOD detection, here we observe that NC is less calibrated in terms of ECE than its counterparts. This might be due to the incompatibility of calibration metrics with the additional class: When the data are corrupted, they become closer to the OOD data, and thus NC tends to assign higher probability mass to the last class which does not correspond to any of the true classes (contrast this to other the approaches). Therefore, in this case, the confidence over the true class becomes necessarily lower—more so than the other approaches. Considering that calibration metrics depend on the confidence of the true class, the calibration of NC thus suffers. One way to overcome this issue is to make calibration metrics aware of the “none class”, e.g. by measuring calibration only on data that have low “none class” probability. We leave the investigation for future research.

#### 6.5 Costs

The additional costs associated with all the OOD-training methods presented here are negligible: Like other non-Bayesian OOD-training methods, the only overhead is the additional minibatch of OOD training data at each training iteration—the costs are similar to when considering a standard training procedure with double the minibatch size. Additionally for LA, in its Hessian computation, one effectively computes it with twice the number of the original data. However, this only needs to be done *once* post-training.

#### 6.6 Limitations

Our methods require a choice of OOD training dataset and it is unclear how can one obtain an OOD training set for specialized problems like medical analyses. Nevertheless, for image classification or natural language processing, we know that large-scale natural image/text OOD training sets are most useful (Hendrycks et al., 2019). Moreover, since our methods



are agnostic to the choice of OOD training data, any future advances in the choice of OOD data for OOD training from both the Bayesian and the frequentist communities can be applied to ours.

## 7 CONCLUSION

We raised an important observation regarding contemporary BNNs’ performance in uncertainty quantification, in particular in OOD detection tasks: BNNs tend to underperform compared to non-Bayesian UQ methods. We hypothesized that this issue is because recent frequentist UQ methods utilize an auxiliary OOD training set. To validate this hypothesis, we explored ways to incorporate OOD training data into BNNs while still maintaining a reasonable Bayesian interpretation. Our experimental results showed that using OOD data in approximate Bayesian inference significantly improved the performance of BNNs, making them competitive or even better than their non-Bayesian counterparts. In particular, we found that the most philosophically Bayesian-compatible way of OOD training—simply add an additional “none class”—performs best. We hope that the studied methods can be strong baselines for future work in the Bayesian deep learning community.

## Acknowledgments

The authors gratefully acknowledge financial support by the European Research Council through ERC StG Action 757275 / PANAMA; the DFG Cluster of Excellence “Machine Learning - New Perspectives for Science”, EXC 2064/1, project number 390727645; the German Federal Ministry of Education and Research (BMBF) through the Tübingen AI Center (FKZ: 01IS18039A); and funds from the Ministry of Science, Research and Arts of the State of Baden-Württemberg. AK is grateful to the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for support. AK is also grateful to Felix Dangel and Jonathan Schmidt for feedback.

## References

- D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané. Concrete Problems in AI Safety. *arXiv preprint arXiv:1606.06565*, 2016.
- J. Bitterwolf, A. Meinke, and M. Hein. Certifiably Adversarially Robust Detection of Out-of-Distribution Data. In *NeurIPS*, 2020.
- C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra. Weight Uncertainty in Neural Networks. In *ICML*, 2015.
- S. R. Bowman, G. Angeli, C. Potts, and C. D. Manning. A Large Annotated Corpus for Learning Natural Language Inference. In *EMNLP*, 2015.
- K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning Phrase Representations Using RNN Encoder-Decoder for Statistical Machine Translation. In *EMNLP*, 2014.
- P. Chrabaszcz, I. Loshchilov, and F. Hutter. A Downsampled Variant of ImageNet as an Alternative to the CIFAR Datasets. *arXiv preprint arXiv:1707.08819*, 2017.
- T. Clanuwat, M. Bober-Irizar, A. Kitamoto, A. Lamb, K. Yamamoto, and D. Ha. Deep Learning for Classical Japanese Literature. *arXiv preprint arXiv:1812.01718*, 2018.
- G. Cohen, S. Afshar, J. Tapson, and A. Van Schaik. EMNIST: Extending MNIST to Handwritten Letters. In *International Joint Conference on Neural Networks*, 2017.
- M. Dusenberry, G. Jerfel, Y. Wen, Y. Ma, J. Snoek, K. Heller, B. Lakshminarayanan, and D. Tran. Efficient and Scalable Bayesian Neural Nets with Rank-1 Factors. In *ICML*, 2020.
- N. El Gayar, F. Schwenker, and G. Palm. A Study of the Robustness of KNN Classifiers Trained Using Soft Labels. In *IAPR Workshop on Artificial Neural Networks in Pattern Recognition*, 2006.
- D. Elliott, S. Frank, K. Sima’an, and L. Specia. Multi30K: Multilingual English-German Image Descriptions. In *ACL Workshop on Vision and Language*, 2016.
- A. Graves. Practical Variational Inference for Neural Networks. In *NIPS*. 2011.
- R. Gueorguieva, R. Rosenheck, and D. Zelterman. Dirichlet Component Regression and Its Applications to Psychiatric Data. *Computational Statistics & Data Analysis*, 52(12), 2008.
- M. Hein, M. Andriushchenko, and J. Bitterwolf. Why Relu Networks Yield High-Confidence Predictions Far Away from the Training Data and How to Mitigate the Problem. In *CVPR*, 2019.
- D. Hendrycks and T. Dietterich. Benchmarking Neural Network Robustness to Common Corruptions and Perturbations. In *ICLR*, 2019.
- D. Hendrycks, M. Mazeika, and T. Dietterich. Deep Anomaly Detection with Outlier Exposure. In *ICLR*, 2019.
- G. E. Hinton and D. Van Camp. Keeping the Neural Networks Simple by Minimizing the Description Length of the Weights. In *COLT*, 1993.

- P. Izmailov, P. Nicholson, S. Lotfi, and A. G. Wilson. Dangers of Bayesian Model Averaging under Covariate Shift. In *NeurIPS*, 2021a.
- P. Izmailov, S. Vikram, M. D. Hoffman, and A. G. Wilson. What Are Bayesian Neural Network Posteriors Really Like? In *ICML*, 2021b.
- A. Kristiadi, M. Hein, and P. Hennig. Being Bayesian, Even Just a Bit, Fixes Overconfidence in ReLU Networks. In *ICML*, 2020a.
- A. Kristiadi, M. Hein, and P. Hennig. An Infinite-Feature Extension for Bayesian ReLU Nets That Fixes Their Asymptotic Overconfidence. *arXiv preprint arXiv:2010.02709*, 2020b.
- A. Kristiadi, M. Hein, and P. Hennig. Learnable Uncertainty under Laplace Approximations. In *UAI*, 2021.
- L. Kuncheva. *Fuzzy Classifier Design*. Springer Science & Business Media, 2000.
- B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and Scalable Predictive Uncertainty Estimation Using Deep Ensembles. In *NIPS*, 2017.
- K. Lee, H. Lee, K. Lee, and J. Shin. Training Confidence-Calibrated Classifiers for Detecting Out-of-Distribution Samples. In *ICLR*, 2018.
- X. Liu, Y. Li, C. Wu, and C.-J. Hsieh. Adv-BNN: Improved Adversarial Defense through Robust Bayesian Neural Network. In *ICLR*, 2019.
- I. Loshchilov and F. Hutter. SGDR: Stochastic Gradient Descent with Warm Restarts. In *ICLR*, 2017.
- D. J. MacKay. The evidence framework applied to classification networks. *Neural computation*, 1992a.
- D. J. MacKay. A Practical Bayesian Framework for Backpropagation Networks. *Neural Computation*, 4(3), 1992b.
- A. Malinin and M. Gales. Predictive uncertainty estimation via prior networks. In *NIPS*, 2018.
- A. Malinin and M. Gales. Reverse KL-Divergence Training of Prior Networks: Improved Uncertainty and Adversarial Robustness. In *NIPS*, 2019.
- A. Meinke and M. Hein. Towards Neural Networks that Provably Know when They don't Know. In *ICLR*, 2020.
- T. Minka. Estimating a Dirichlet distribution, 2000.
- M. P. Naeini, G. Cooper, and M. Hauskrecht. Obtaining Well Calibrated Probabilities using Bayesian Binning. In *AAAI*, 2015.
- J. Nandy, W. Hsu, and M. L. Lee. Towards Maximizing the Representation Gap between In-Domain & Out-of-Distribution Examples. In *NeurIPS*, 2020.
- A. Nguyen, J. Yosinski, and J. Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *CVPR*, 2015.
- K. Osawa, S. Swaroop, M. E. E. Khan, A. Jain, R. Eschenhagen, R. E. Turner, and R. Yokota. Practical Deep Learning with Bayesian Principles. In *NeurIPS*, 2019.
- Y. Ovadia, E. Fertig, J. Ren, Z. Nado, D. Sculley, S. Nowozin, J. Dillon, B. Lakshminarayanan, and J. Snoek. Can You Trust Your Model's Uncertainty? Evaluating Predictive Uncertainty under Dataset Shift. In *NeurIPS*, 2019.
- H. Ritter, A. Botev, and D. Barber. A Scalable Laplace Approximation for Neural Networks. In *ICLR*, 2018.
- M. Sensoy, L. Kaplan, and M. Kandemir. Evidential deep learning to quantify classification uncertainty. In *NIPS*, 2018.
- R. Socher, A. Perelygin, J. Wu, J. Chuang, C. D. Manning, A. Y. Ng, and C. Potts. Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. In *EMNLP*, 2013.
- C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the Inception Architecture for Computer Vision. In *CVPR*, 2016.
- C. Thiel. Classification on Soft Labels is Robust Against Label Noise. In *International Conference on Knowledge-Based and Intelligent Information and Engineering Systems*, 2008.
- M. Tomczak, S. Swaroop, and R. Turner. Efficient Low Rank Gaussian Variational Inference for Neural Networks. In *NeurIPS*, 2020.
- E. M. Voorhees. Overview of the TREC-9 Question Answering Track. In *Text REtrieval Conference (TREC)*, 2001.
- X. Wang and L. Aitchison. Bayesian OOD Detection with Aleatoric Uncertainty and Outlier Exposure. *arXiv preprint arXiv:2102.12959v2*, 2021.
- Y. Wen, P. Vicol, J. Ba, D. Tran, and R. Grosse. Flipout: Efficient Pseudo-Independent Weight Perturbations on Mini-Batches. In *ICLR*, 2018.
- A. G. Wilson and P. Izmailov. Bayesian Deep Learning and a Probabilistic Perspective of Generalization. In *NeurIPS*, 2020.
- F. Yu, A. Seff, Y. Zhang, S. Song, T. Funkhouser, and J. Xiao. LSUN: Construction of a Large-Scale Image Dataset Using Deep Learning with Humans in the Loop. *arXiv preprint arXiv:1506.03365*, 2015.
- S. Zagoruyko and N. Komodakis. Wide Residual Networks. In *BMVC*, 2016.

- G. Zhang, S. Sun, D. Duvenaud, and R. Grosse. Noisy Natural Gradient as Variational Inference. In *ICML*, 2018.
- R. Zhang, C. Li, J. Zhang, C. Chen, and A. G. Wilson. Cyclical Stochastic Gradient MCMC for Bayesian Deep Learning. In *ICLR*, 2020.
- X. Zhang and Y. LeCun. Universum Prescription: Regularization Using Unlabeled Data. In *AAAI*, 2017.

---

# Supplementary Material: Being a Bit Frequentist Improves Bayesian Neural Networks

---

## Appendix A OOD Test Sets

For image-based OOD detection tasks, we use the following test sets on top of MNIST, F-MNIST, SVHN, CIFAR-10, and CIFAR-100:

- E-MNIST: Contains handwritten letters (“a”-“z”)—same format as MNIST (Cohen et al., 2017).
- K-MNIST: Contains handwritten Hiragana scripts—same format as MNIST (Clanuwat et al., 2018).
- LSUN-CR: Contains real-world images of classrooms (Yu et al., 2015).
- CIFAR-Gr: Obtained by converting CIFAR-10 test images to grayscale.
- F-MNIST-3D: Obtained by converting single-channel F-MNIST images into three-channel images—all these three channels have identical values.
- UNIFORM: Obtained by drawing independent uniformly-distributed random pixel.
- SMOOTH: Obtained by permuting, smoothing, and contrast-rescaling the original (i.e. the respective in-distribution) test images (Hein et al., 2019).

Meanwhile, for text classification, we use the following OOD test set, following (Hendrycks et al., 2019):

- MULTI30K: Multilingual English-German image description dataset (Elliott et al., 2016).
- WMT16: Machine-translation dataset, available at <http://www.statmt.org/wmt14/translation-task.html>.
- SNLI: Collection of human-written English sentence pairs manually labeled for balanced classification with the labels entailment, contradiction, and neutral (Bowman et al., 2015).

Finally, for dataset-shift robustness tasks, we use the standard dataset:

- CIFAR-10-C: Contains 19 different perturbations—e.g. snow, motion blur, brightness rescaling—with 5 level of severity for a total of 95 distinct shifts (Hendrycks and Dietterich, 2019; Ovadia et al., 2019).

## Appendix B Training Details

**Non-Bayesian** For MNIST and F-MNIST, we use a five-layer LeNet architecture. Meanwhile, for SVHN, CIFAR-10, and CIFAR-100, we use WideResNet-16-4 (Zagoruyko and Komodakis, 2016). For all methods, the training procedures are as follows. For LeNet, we use Adam with initial learning rate of  $1 \times 10^{-3}$  and annealed it using the cosine decay method (Loshchilov and Hutter, 2017) along with weight decay of  $5 \times 10^{-4}$  for 100 epochs. We use a batch size of 128 for both in- and out-distribution batches, amounting to an effective batch size of 256 in the case of OOD training. The standard data augmentation pipeline (random crop and horizontal flip) is applied to both in-distribution and OOD data. For WideResNet-16-4, we use SGD instead with an initial learning rate of  $1 \times 10^{-1}$  and Nesterov momentum of 0.9 along with the dropout regularization with rate 0.3—all other hyperparameters are identical to LeNet. Finally, we use 5 ensemble members for DE.

**Bayesian** For both LA, VB, and their variants (i.e. LA+X and VB+X), we use the identical setup as in the non-Bayesian training above. Additionally, for LA and LA+X, we use the diagonal Fisher matrix as the approximate Hessian. Moreover, we tune prior variance by minimizing the validation Brier score. All predictions are done using 20 MC samples. For VB and VB+X, we use a diagonal Gaussian variational posterior for both the last-layer weight matrix and bias vector. Moreover, the prior is a zero-mean isotropic Gaussian with prior precision  $5 \times 10^{-4}$  (to emulate the choice of the weight decay in the non-Bayesian training). The trade-off hyperparameter  $\tau$  of the ELBO is set to the standard value of 0.1 (Osawa et al., 2019; Zhang et al., 2018). We do not use weight decay on the last layer since the regularization of its parameters is done by the KL-term of the ELBO. Lastly, we use 5 and 200 MC samples for computing the ELBO and for making predictions, respectively.

**Text Classification** The network used is a two-layer Gated Recurrent Unit (GRU, Cho et al., 2014) with 128 hidden units on each layer. The word-embedding dimension is 50 and the maximum vocabulary size is 10000. We put an affine layer on top of the last GRU output to translate the hidden units to output units. Both the LA and VB are applied only on this layer. We use a batch size of 64 and Adam optimizer with a learning rate of 0.01 without weight decay, except for LA in which case we use weight decay of  $5 \times 10^{-4}$ . The optimization is done for 5 epochs, following (Hendrycks et al., 2019).

## Appendix C Additional Results

The detailed, non-averaged results for the FPR95 metric are in Table 6. Furthermore, additional results with the area-under-ROC-curve (AUROC), area-under-precision-recall-curve (AUPRC), and mean confidence (MMC) metrics are in Tables 6 to 9, respectively. For the full results for models trained with the SMOOTH noise dataset as  $\mathcal{D}_{\text{out}}$  are in Tables 10 and 11. Furthermore, the full results of the NLP experiment is in Tables 12 and 13. Finally, detailed, non-averaged results for sophisticated models (Flipout and CSGHMC) are in Tables 14 and 15.

Table 6: OOD data detection in terms of FPR95. Lower is better. Values are averages over five prediction runs.

Datasets	MAP	OE	DE	VB					LA				
				Plain	NC	SL	ML	OE	Plain	NC	SL	ML	OE
<b>MNIST</b>													
F-MNIST	11.8±0.0	0.0±0.0	5.3±0.0	12.5±0.1	0.1±0.0	0.0±0.0	0.4±0.0	1.1±0.0	12.0±0.0	0.2±0.0	0.0±0.0	0.1±0.0	0.0±0.0
E-MNIST	35.6±0.0	26.4±0.0	30.4±0.0	34.5±0.1	34.7±0.1	14.3±0.1	34.2±0.1	31.4±0.1	35.8±0.1	30.6±0.0	12.6±0.1	26.8±0.1	26.7±0.1
K-MNIST	14.4±0.0	5.9±0.0	7.7±0.0	14.0±0.1	10.5±0.1	2.1±0.0	9.7±0.1	8.5±0.0	14.5±0.1	8.9±0.0	0.7±0.0	5.8±0.0	5.9±0.0
CIFAR-Gr	0.2±0.0	0.0±0.0	0.0±0.0	0.2±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.2±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0
Uniform	44.3±0.0	0.0±0.0	19.8±0.0	93.1±0.2	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	54.2±0.4	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0
Smooth	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0
<b>F-MNIST</b>													
MNIST	73.5±0.0	38.5±0.0	65.8±0.0	66.8±0.1	43.5±0.1	9.5±0.0	50.1±0.1	57.2±0.1	72.2±0.2	25.6±0.1	11.5±0.1	38.9±0.3	39.9±0.1
E-MNIST	73.6±0.0	21.0±0.0	58.6±0.0	68.1±0.1	18.7±0.0	5.0±0.0	34.0±0.0	40.6±0.0	72.2±0.2	6.0±0.0	4.6±0.1	14.7±0.1	23.1±0.1
K-MNIST	73.7±0.0	37.4±0.0	47.2±0.0	62.6±0.1	28.0±0.1	10.6±0.0	33.4±0.0	36.7±0.0	71.5±0.2	18.2±0.1	8.7±0.1	32.5±0.2	38.7±0.3
CIFAR-Gr	87.2±0.0	0.0±0.0	86.6±0.0	75.3±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	87.7±0.1	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0
Uniform	81.3±0.0	0.0±0.0	86.3±0.0	87.3±0.1	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	81.0±0.2	0.0±0.0	0.0±0.0	0.0±0.0	0.1±0.0
Smooth	26.8±0.0	0.0±0.0	24.2±0.0	19.6±0.0	0.0±0.0	0.0±0.0	0.2±0.0	0.1±0.0	27.3±0.1	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0
<b>SVHN</b>													
CIFAR-10	18.9±0.0	0.1±0.0	9.5±0.0	15.0±0.0	0.3±0.0	0.1±0.0	0.0±0.0	0.1±0.0	15.4±0.1	0.4±0.0	0.0±0.0	0.0±0.0	0.1±0.0
LSUN-CR	19.7±0.0	0.0±0.0	8.3±0.0	17.2±0.2	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	15.5±0.1	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0
CIFAR-100	21.8±0.0	0.2±0.0	11.6±0.0	18.1±0.0	0.5±0.0	0.5±0.0	0.1±0.0	0.2±0.0	17.6±0.1	0.6±0.0	0.2±0.0	0.2±0.0	0.1±0.0
FMNIST-3D	26.7±0.0	0.0±0.0	17.5±0.0	24.5±0.1	0.0±0.0	0.6±0.0	0.0±0.0	0.0±0.0	27.2±0.1	0.1±0.0	0.0±0.0	0.0±0.0	0.0±0.0
Uniform	30.0±0.0	0.0±0.0	6.4±0.0	48.2±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	17.0±0.1	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0
Smooth	17.3±0.0	12.0±0.0	6.9±0.0	9.1±0.0	7.7±0.0	9.5±0.0	8.3±0.0	8.4±0.0	10.1±0.1	8.1±0.1	5.9±0.0	6.6±0.1	6.4±0.0
<b>CIFAR-10</b>													
SVHN	34.5±0.0	10.0±0.0	33.9±0.0	33.5±0.0	30.6±0.1	59.4±0.0	18.3±0.1	33.9±0.1	35.5±0.1	12.7±0.2	47.2±0.3	8.7±0.1	10.8±0.0
LSUN-CR	53.3±0.0	28.0±0.0	44.0±0.0	49.4±0.4	25.9±0.2	43.7±0.1	36.8±0.1	34.8±0.2	53.8±0.6	17.5±0.3	41.2±0.9	30.1±0.5	28.4±0.4
CIFAR-100	61.2±0.0	57.8±0.0	52.5±0.0	58.4±0.1	58.5±0.1	63.3±0.0	56.8±0.1	57.1±0.1	61.4±0.1	59.6±0.2	62.2±0.2	60.4±0.2	57.9±0.1
FMNIST-3D	42.4±0.0	26.8±0.0	30.7±0.0	37.4±0.0	19.0±0.1	43.9±0.0	32.2±0.1	29.6±0.1	43.2±0.2	15.4±0.1	36.8±0.2	24.2±0.1	27.8±0.1
Uniform	87.7±0.0	0.0±0.0	0.0±0.0	13.8±0.1	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	92.8±0.1	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0
Smooth	35.1±0.0	14.2±0.0	32.9±0.0	26.4±0.0	34.0±0.2	31.9±0.0	30.3±0.0	23.1±0.1	34.9±0.1	15.5±0.2	43.6±0.2	7.5±0.0	14.9±0.1
<b>CIFAR-100</b>													
LSUN-CR	82.0±0.0	64.3±0.0	75.3±0.0	73.8±0.2	62.3±0.5	76.3±0.1	65.3±0.1	67.6±0.3	82.8±0.5	55.9±0.5	75.6±0.7	65.3±0.5	64.1±1.0
CIFAR-10	79.8±0.0	81.9±0.0	76.4±0.0	78.2±0.1	81.4±0.1	82.8±0.0	79.5±0.1	79.0±0.0	79.5±0.1	80.9±0.2	81.7±0.1	80.8±0.1	80.0±0.1
FMNIST-3D	65.8±0.0	58.5±0.0	61.8±0.0	57.1±0.1	41.0±0.2	72.0±0.1	51.7±0.1	56.0±0.1	66.1±0.1	58.6±0.3	69.0±0.2	59.2±0.1	59.3±0.3
Uniform	97.6±0.0	0.0±0.0	94.3±0.0	100.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	98.8±0.1	0.0±0.0	0.0±0.0	0.1±0.0	0.0±0.0
Smooth	79.5±0.0	65.2±0.0	58.7±0.0	79.1±0.0	64.8±0.1	80.2±0.1	54.4±0.1	64.0±0.0	79.2±0.1	41.6±0.2	78.0±0.1	57.1±0.3	66.2±0.1



Being a Bit Frequentist Improves Bayesian Neural Networks

Table 7: OOD data detection in terms of AUROC. Higher is better. Values are averages over five prediction runs.

Datasets	MAP	OE	DE	VB					LA					
				Plain	NC	SL	ML	OE	Plain	NC	SL	ML	OE	
<b>MNIST</b>														
F-MNIST	97.3±0.0	99.9±0.0	98.7±0.0	97.4±0.0	99.9±0.0	99.9±0.0	99.8±0.0	99.6±0.0	97.4±0.0	99.9±0.0	99.9±0.0	99.9±0.0	99.9±0.0	99.9±0.0
E-MNIST	89.1±0.0	93.7±0.0	90.4±0.0	89.9±0.1	90.4±0.0	95.7±0.0	91.1±0.0	92.1±0.1	89.1±0.0	91.2±0.0	94.9±0.0	93.3±0.0	93.6±0.0	
K-MNIST	96.9±0.0	98.5±0.0	98.1±0.0	96.9±0.0	97.8±0.0	98.8±0.0	98.0±0.0	98.2±0.0	96.9±0.0	97.9±0.0	99.2±0.0	98.5±0.0	98.5±0.0	
CIFAR-Gr	99.6±0.0	100.0±0.0	99.8±0.0	99.6±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	99.6±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	
Uniform	95.0±0.0	100.0±0.0	95.8±0.0	90.5±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	94.6±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	
Smooth	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	
<b>F-MNIST</b>														
MNIST	79.7±0.0	92.9±0.0	83.0±0.0	85.3±0.0	87.9±0.0	98.6±0.0	86.7±0.0	86.2±0.0	80.3±0.0	94.2±0.0	98.2±0.0	92.9±0.0	92.5±0.0	
E-MNIST	81.8±0.0	96.5±0.0	87.5±0.0	85.1±0.0	95.6±0.0	99.2±0.0	92.0±0.0	91.3±0.0	82.3±0.0	98.9±0.0	99.3±0.0	97.6±0.0	96.1±0.0	
K-MNIST	83.1±0.0	94.4±0.0	91.7±0.0	86.9±0.0	94.3±0.0	98.4±0.0	93.5±0.0	93.3±0.0	83.9±0.0	96.9±0.0	98.7±0.0	94.9±0.0	94.1±0.0	
CIFAR-Gr	82.2±0.0	100.0±0.0	83.6±0.0	87.5±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	81.4±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	
Uniform	85.5±0.0	100.0±0.0	85.7±0.0	85.8±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	85.3±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	
Smooth	95.7±0.0	100.0±0.0	96.4±0.0	97.2±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	95.5±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	
<b>SVHN</b>														
CIFAR-10	96.2±0.0	100.0±0.0	97.9±0.0	95.6±0.0	99.9±0.0	99.9±0.0	100.0±0.0	100.0±0.0	97.1±0.0	99.9±0.0	100.0±0.0	100.0±0.0	100.0±0.0	
LSUN-CR	95.7±0.0	100.0±0.0	97.7±0.0	95.9±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	97.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	
CIFAR-100	95.5±0.0	99.9±0.0	97.4±0.0	94.7±0.0	99.9±0.0	99.8±0.0	100.0±0.0	99.9±0.0	96.5±0.0	99.9±0.0	99.9±0.0	100.0±0.0	100.0±0.0	
FMNIST-3D	95.5±0.0	100.0±0.0	97.1±0.0	91.4±0.0	100.0±0.0	99.8±0.0	100.0±0.0	100.0±0.0	95.6±0.0	100.0±0.0	99.9±0.0	100.0±0.0	100.0±0.0	
Uniform	94.3±0.0	100.0±0.0	98.2±0.0	80.2±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	96.8±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	
Smooth	96.5±0.0	97.6±0.0	98.4±0.0	97.5±0.0	97.7±0.0	95.9±0.0	97.5±0.0	97.7±0.0	97.7±0.0	98.4±0.0	98.1±0.0	98.5±0.0	98.7±0.0	
<b>CIFAR-10</b>														
SVHN	95.6±0.0	98.2±0.0	95.6±0.0	95.7±0.0	95.8±0.0	89.0±0.0	97.2±0.0	95.6±0.0	95.5±0.0	97.8±0.0	92.7±0.0	98.7±0.0	98.1±0.0	
LSUN-CR	91.8±0.0	95.9±0.0	93.7±0.0	91.5±0.0	96.1±0.0	92.5±0.0	94.1±0.0	94.4±0.0	92.0±0.0	97.5±0.0	93.8±0.1	95.7±0.0	96.0±0.1	
CIFAR-100	89.8±0.0	90.1±0.0	91.3±0.0	88.6±0.0	88.2±0.0	85.9±0.0	88.3±0.0	88.9±0.0	89.9±0.0	89.8±0.0	86.2±0.0	89.7±0.0	90.0±0.0	
FMNIST-3D	94.4±0.0	96.2±0.0	95.8±0.0	94.5±0.0	97.2±0.0	92.9±0.0	95.1±0.0	95.9±0.0	94.3±0.0	97.7±0.0	94.0±0.0	96.6±0.0	96.1±0.0	
Uniform	93.0±0.0	100.0±0.0	99.5±0.0	97.6±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	92.2±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	
Smooth	94.3±0.0	97.6±0.0	95.6±0.0	96.0±0.0	95.3±0.0	94.8±0.0	95.4±0.0	96.2±0.0	94.6±0.0	97.6±0.0	93.9±0.0	98.8±0.0	97.6±0.0	
<b>CIFAR-100</b>														
LSUN-CR	78.4±0.0	85.3±0.0	83.8±0.0	81.3±0.1	87.9±0.0	80.1±0.0	85.6±0.0	85.3±0.1	78.8±0.1	89.1±0.0	82.0±0.3	86.3±0.1	85.9±0.2	
CIFAR-10	77.4±0.0	77.1±0.0	79.8±0.0	77.7±0.0	76.8±0.0	76.3±0.0	77.8±0.0	77.3±0.0	77.8±0.0	77.5±0.0	76.6±0.0	77.6±0.0	77.7±0.0	
FMNIST-3D	85.3±0.0	86.2±0.0	87.6±0.0	87.7±0.0	91.5±0.0	84.0±0.0	89.6±0.0	87.9±0.0	85.2±0.0	87.0±0.0	85.7±0.0	86.3±0.0	86.0±0.0	
Uniform	80.1±0.0	100.0±0.0	87.7±0.0	64.8±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	82.6±0.1	100.0±0.0	100.0±0.0	99.9±0.0	100.0±0.0	
Smooth	77.0±0.0	76.6±0.0	82.9±0.0	69.4±0.0	74.5±0.0	80.3±0.0	83.2±0.0	78.3±0.0	79.4±0.1	91.7±0.0	78.1±0.0	86.4±0.1	77.7±0.1	

Table 8: OOD data detection in terms of AUPRC. Higher is better. Values are averages over five prediction runs.

Datasets	MAP	OE	DE	VB					LA				
				Plain	NC	SL	ML	OE	Plain	NC	SL	ML	OE
<b>MNIST</b>													
F-MNIST	96.9±0.0	99.9±0.0	98.7±0.0	97.5±0.0	99.9±0.0	99.8±0.0	99.8±0.0	99.6±0.0	97.0±0.0	99.9±0.0	99.9±0.0	99.9±0.0	99.9±0.0
E-MNIST	74.2±0.0	86.7±0.0	77.2±0.0	76.3±0.2	77.8±0.1	83.3±0.0	80.6±0.2	82.5±0.2	74.1±0.1	79.6±0.0	78.8±0.1	85.5±0.0	86.4±0.0
K-MNIST	96.5±0.0	98.5±0.0	98.0±0.0	96.4±0.0	97.5±0.0	97.2±0.0	97.8±0.0	98.1±0.0	96.6±0.0	97.8±0.0	98.9±0.0	98.4±0.0	98.4±0.0
CIFAR-Gr	99.7±0.0	100.0±0.0	99.8±0.0	99.6±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	99.6±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0
Uniform	96.7±0.0	100.0±0.0	97.3±0.0	93.4±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	96.5±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0
Smooth	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0
<b>F-MNIST</b>													
MNIST	75.3±0.0	92.2±0.0	79.0±0.0	83.1±0.1	84.0±0.1	98.4±0.0	83.9±0.1	84.1±0.1	76.3±0.0	92.5±0.0	98.0±0.0	92.3±0.0	91.9±0.0
E-MNIST	66.9±0.0	92.7±0.0	76.4±0.0	74.0±0.3	88.7±0.1	98.2±0.0	82.8±0.0	82.5±0.1	67.8±0.0	96.8±0.0	98.4±0.0	94.8±0.0	92.0±0.0
K-MNIST	81.7±0.0	94.4±0.0	91.1±0.0	85.2±0.1	93.1±0.0	98.1±0.0	92.5±0.0	92.6±0.0	82.8±0.0	96.2±0.0	98.6±0.0	94.7±0.0	94.2±0.0
CIFAR-Gr	85.5±0.0	100.0±0.0	87.2±0.0	89.6±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	84.9±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0
Uniform	88.1±0.0	100.0±0.0	88.9±0.0	89.2±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	87.9±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0
Smooth	95.3±0.0	100.0±0.0	96.1±0.0	96.9±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	95.1±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0
<b>SVHN</b>													
CIFAR-10	98.3±0.0	100.0±0.0	99.1±0.0	96.9±0.0	99.9±0.0	99.9±0.0	100.0±0.0	100.0±0.0	98.8±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0
LSUN-CR	99.9±0.0	100.0±0.0	100.0±0.0	99.9±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0
CIFAR-100	97.7±0.0	100.0±0.0	98.8±0.0	96.4±0.0	99.9±0.0	99.9±0.0	100.0±0.0	100.0±0.0	98.3±0.0	99.9±0.0	100.0±0.0	100.0±0.0	100.0±0.0
FMNIST-3D	98.1±0.0	100.0±0.0	98.8±0.0	93.5±0.0	100.0±0.0	99.9±0.0	100.0±0.0	100.0±0.0	98.2±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0
Uniform	97.3±0.0	100.0±0.0	99.3±0.0	82.8±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	98.7±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0
Smooth	98.5±0.0	98.9±0.0	99.4±0.0	98.6±0.0	98.6±0.0	96.5±0.0	98.4±0.0	98.7±0.0	99.1±0.0	99.3±0.0	98.7±0.0	99.4±0.0	99.5±0.0
<b>CIFAR-10</b>													
SVHN	93.3±0.0	96.5±0.0	93.3±0.0	93.3±0.0	92.8±0.0	77.6±0.0	94.5±0.0	92.8±0.0	93.3±0.0	95.9±0.0	85.2±0.0	97.0±0.0	96.4±0.0
LSUN-CR	99.7±0.0	99.8±0.0	99.7±0.0	99.6±0.0	99.8±0.0	99.7±0.0	99.7±0.0	99.7±0.0	99.7±0.0	99.9±0.0	99.7±0.0	99.8±0.0	99.8±0.0
CIFAR-100	89.9±0.0	90.0±0.0	91.3±0.0	86.7±0.0	85.7±0.0	82.8±0.0	85.3±0.0	86.9±0.0	90.0±0.0	90.0±0.0	82.4±0.0	89.8±0.0	90.0±0.0
FMNIST-3D	95.1±0.0	96.4±0.0	96.1±0.0	94.7±0.0	97.0±0.0	92.8±0.0	95.0±0.0	96.1±0.0	95.0±0.0	97.6±0.0	93.4±0.0	96.7±0.0	96.3±0.0
Uniform	95.5±0.0	100.0±0.0	99.6±0.0	98.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	95.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0
Smooth	94.4±0.0	97.5±0.0	95.8±0.0	95.7±0.0	95.5±0.0	93.9±0.0	95.1±0.0	95.9±0.0	94.7±0.0	97.6±0.0	93.9±0.0	98.6±0.0	97.4±0.0
<b>CIFAR-100</b>													
LSUN-CR	99.0±0.0	99.4±0.0	99.3±0.0	99.0±0.0	99.5±0.0	99.1±0.0	99.3±0.0	99.3±0.0	99.0±0.0	99.5±0.0	99.2±0.0	99.4±0.0	99.4±0.0
CIFAR-10	77.2±0.0	77.0±0.0	79.3±0.0	76.8±0.0	77.1±0.0	75.4±0.0	77.4±0.0	77.1±0.0	77.3±0.0	77.1±0.0	75.8±0.0	77.2±0.0	77.1±0.0
FMNIST-3D	85.7±0.0	86.0±0.0	88.2±0.0	87.6±0.0	90.9±0.0	84.8±0.0	89.3±0.0	87.6±0.0	85.6±0.0	87.1±0.0	86.5±0.0	86.3±0.0	85.7±0.0
Uniform	84.6±0.0	100.0±0.0	91.5±0.0	73.3±0.0	100.0±0.0	100.0±0.0	100.0±0.0	100.0±0.0	87.1±0.1	100.0±0.0	100.0±0.0	99.9±0.0	100.0±0.0
Smooth	75.4±0.0	71.0±0.0	80.3±										

Table 9: Confidences in terms of MMC, averaged over five prediction runs. Lower is better for OOD datasets.

Datasets	VB					LA							
	MAP	OE	DE	Plain	NC	SL	ML	OE	Plain	NC	SL	ML	OE
<b>MNIST</b>	99.1	99.4	99.3	98.7	98.2	99.1	98.2	98.6	99.0	99.3	99.2	99.2	99.3
F-MNIST	66.3±0.0	22.0±0.0	64.9±0.0	70.4±0.0	6.1±0.0	20.4±0.0	21.2±0.0	27.3±0.0	65.2±0.0	7.7±0.0	21.0±0.0	20.7±0.0	22.2±0.0
E-MNIST	82.3±0.0	79.3±0.0	80.9±0.0	78.6±0.0	74.0±0.0	75.9±0.0	73.0±0.0	73.9±0.0	81.1±0.0	79.6±0.0	80.4±0.0	76.5±0.0	78.1±0.0
K-MNIST	73.3±0.0	65.7±0.0	69.7±0.0	67.6±0.0	51.8±0.0	61.2±0.0	50.3±0.0	52.4±0.0	71.8±0.0	64.3±0.0	66.3±0.0	62.3±0.0	64.3±0.0
CIFAR-Gr	48.0±0.0	10.0±0.0	43.1±0.0	45.2±0.0	0.0±0.0	10.0±0.0	10.1±0.0	10.1±0.0	47.3±0.0	0.0±0.0	10.1±0.0	10.4±0.0	10.2±0.0
Uniform	96.8±0.0	10.0±0.0	97.4±0.0	97.8±0.0	0.1±0.0	10.0±0.0	10.1±0.0	10.1±0.0	96.5±0.0	0.0±0.0	10.2±0.0	10.3±0.0	10.2±0.0
Smooth	12.9±0.0	10.1±0.0	12.6±0.0	12.7±0.0	0.3±0.0	21.6±0.0	10.2±0.0	10.2±0.0	12.8±0.0	0.5±0.0	20.5±0.0	10.1±0.0	10.1±0.0
<b>F-MNIST</b>	96.1	96.0	94.8	93.6	91.0	95.0	91.2	92.9	95.7	94.9	93.5	94.5	94.7
MNIST	82.8±0.0	60.1±0.0	74.3±0.0	70.9±0.0	50.0±0.0	32.9±0.0	57.2±0.0	63.7±0.0	80.9±0.0	33.8±0.1	33.9±0.0	55.9±0.1	57.9±0.0
E-MNIST	82.5±0.0	45.0±0.0	70.0±0.0	71.5±0.0	25.0±0.0	25.5±0.0	44.4±0.0	50.7±0.0	80.6±0.0	9.4±0.0	24.3±0.0	36.2±0.0	44.7±0.0
K-MNIST	82.5±0.0	60.1±0.0	64.0±0.0	68.4±0.0	37.0±0.0	33.5±0.0	44.8±0.0	48.0±0.0	80.0±0.0	28.1±0.0	32.0±0.0	52.6±0.0	57.3±0.0
CIFAR-Gr	89.1±0.0	10.0±0.0	84.6±0.0	73.1±0.0	0.0±0.0	10.0±0.0	10.1±0.0	10.2±0.0	88.6±0.0	0.0±0.0	10.8±0.0	10.3±0.0	10.3±0.0
Uniform	85.5±0.0	13.2±0.0	82.3±0.0	79.4±0.0	0.6±0.0	10.2±0.0	10.6±0.0	11.1±0.0	84.2±0.0	0.0±0.0	10.9±0.0	12.2±0.0	15.0±0.0
Smooth	48.0±0.0	10.5±0.0	44.0±0.0	42.0±0.0	1.0±0.0	11.0±0.0	11.9±0.0	10.7±0.0	47.4±0.0	0.3±0.0	11.6±0.0	10.4±0.0	10.8±0.0
<b>SVHN</b>	98.6	98.6	98.1	97.7	97.3	98.4	97.1	97.6	97.9	98.0	97.9	97.8	97.3
CIFAR-10	69.0±0.0	11.6±0.0	57.3±0.0	61.3±0.0	3.7±0.0	15.4±0.0	11.1±0.0	11.7±0.0	60.7±0.0	3.6±0.0	12.4±0.0	12.1±0.0	12.5±0.0
LSUN-CR	69.8±0.0	10.2±0.0	57.7±0.0	63.9±0.0	0.1±0.0	10.8±0.0	10.2±0.0	10.2±0.0	61.5±0.1	0.0±0.0	11.0±0.0	10.3±0.0	10.7±0.0
CIFAR-100	70.3±0.0	12.4±0.0	58.8±0.0	63.5±0.0	4.3±0.0	17.3±0.0	11.6±0.0	12.3±0.0	62.3±0.0	4.1±0.0	13.4±0.0	12.7±0.0	13.4±0.0
FMNIST-3D	73.6±0.0	11.0±0.0	62.5±0.0	67.8±0.0	1.5±0.0	17.6±0.0	10.8±0.0	11.5±0.0	68.0±0.0	1.8±0.0	15.5±0.0	10.9±0.0	12.0±0.0
Uniform	77.5±0.0	10.3±0.0	57.1±0.0	82.2±0.0	0.1±0.0	10.0±0.0	10.1±0.0	10.2±0.0	65.7±0.1	0.0±0.0	10.4±0.0	10.3±0.0	10.7±0.0
Smooth	68.8±0.0	52.1±0.0	51.3±0.0	55.5±0.0	42.3±0.0	62.7±0.0	45.6±0.0	44.0±0.0	58.3±0.1	43.9±0.0	44.9±0.0	44.6±0.1	40.8±0.0
<b>CIFAR-10</b>	96.7	96.9	95.9	95.7	94.5	96.0	95.1	95.5	96.4	94.9	95.8	96.1	96.1
SVHN	65.2±0.0	45.7±0.0	60.8±0.0	59.5±0.0	49.7±0.0	73.3±0.0	45.0±0.0	55.8±0.0	63.5±0.0	36.8±0.1	64.7±0.0	33.0±0.0	44.0±0.0
LSUN-CR	73.9±0.0	58.7±0.0	65.2±0.0	68.2±0.0	41.1±0.1	64.3±0.0	56.7±0.0	57.1±0.0	71.6±0.0	29.2±0.2	61.7±0.4	53.6±0.3	55.2±0.3
CIFAR-100	77.2±0.0	76.1±0.0	69.8±0.0	72.4±0.0	67.7±0.0	75.9±0.0	69.8±0.0	70.5±0.0	75.4±0.0	67.3±0.0	74.3±0.0	73.2±0.0	72.6±0.0
FMNIST-3D	67.7±0.0	55.9±0.0	58.5±0.0	60.4±0.0	32.7±0.0	64.3±0.0	54.5±0.0	52.3±0.0	66.0±0.0	27.4±0.0	57.6±0.0	48.5±0.0	53.1±0.1
Uniform	82.1±0.0	10.3±0.0	40.7±0.0	52.2±0.0	0.1±0.0	10.2±0.0	10.7±0.0	10.1±0.0	81.9±0.1	0.0±0.0	10.2±0.0	10.3±0.0	10.3±0.0
Smooth	64.3±0.0	45.2±0.0	57.2±0.0	54.0±0.0	51.4±0.0	53.3±0.0	53.0±0.0	49.3±0.0	62.5±0.0	32.6±0.0	59.5±0.1	26.4±0.0	43.5±0.0
<b>CIFAR-100</b>	84.7	85.1	80.9	69.0	66.2	80.1	67.1	67.1	81.6	80.3	77.0	79.1	79.4
SVHN	52.8±0.0	43.9±0.0	43.8±0.0	27.3±0.0	31.0±0.0	46.2±0.0	24.9±0.0	26.9±0.0	46.9±0.0	33.7±0.1	42.1±0.0	34.8±0.0	37.9±0.1
LSUN-CR	62.7±0.0	51.3±0.0	48.8±0.0	33.3±0.1	21.1±0.0	49.7±0.0	26.5±0.0	26.4±0.1	57.0±0.1	36.4±0.1	43.5±0.5	40.9±0.1	42.6±0.3
CIFAR-10	62.9±0.0	63.7±0.0	53.4±0.0	39.0±0.0	36.9±0.0	55.3±0.0	37.2±0.0	37.4±0.0	57.3±0.0	55.7±0.0	50.8±0.0	53.8±0.0	54.0±0.0
FMNIST-3D	51.8±0.0	48.2±0.0	42.5±0.0	24.8±0.0	15.9±0.0	44.5±0.0	20.6±0.0	22.6±0.0	47.1±0.0	40.7±0.0	38.4±0.0	40.3±0.0	41.4±0.0
Uniform	64.2±0.0	1.4±0.0	45.0±0.0	59.4±0.0	0.0±0.0	2.2±0.0	1.2±0.0	1.2±0.0	54.4±0.1	0.0±0.0	1.5±0.0	4.4±0.1	1.5±0.0
Smooth	61.7±0.0	58.9±0.0	47.3±0.0	49.3±0.0	38.4±0.0	50.6±0.0	28.7±0.0	34.8±0.0	54.8±0.1	30.4±0.1	49.1±0.0	40.2±0.1	51.5±0.1

 Table 10: Test accuracy ( $\uparrow$ ) / ECE ( $\downarrow$ ) of models trained with random noises (Hein et al., 2019) as  $\mathcal{D}_{\text{out}}$ , averaged over five prediction runs.

	MNIST	F-MNIST	SVHN	CIFAR-10	CIFAR-100
MAP	99.4±0.0 / 6.4±0.0	92.4±0.0 / 13.9±0.0	97.4±0.0 / 8.9±0.0	94.8±0.0 / 10.0±0.0	76.7±0.0 / 14.3±0.0
DE	99.5±0.0 / 8.6±0.0	93.6±0.0 / 3.6±0.0	97.6±0.0 / 3.5±0.0	95.7±0.0 / 4.5±0.0	80.0±0.0 / 1.9±0.0
OE	99.6±0.0 / 6.4±0.0	92.6±0.0 / 12.7±0.0	97.5±0.0 / 8.9±0.0	94.7±0.0 / 11.5±0.0	76.5±0.0 / 16.1±0.0
VB	99.5±0.0 / 11.2±0.3	92.4±0.0 / 3.7±0.2	97.5±0.0 / 5.7±0.2	94.9±0.0 / 5.8±0.2	75.4±0.0 / 8.3±0.0
+NC	99.4±0.0 / 10.6±0.1	92.3±0.0 / 3.0±0.1	97.4±0.0 / 4.2±0.2	94.9±0.0 / 5.1±0.1	74.0±0.1 / 8.8±0.1
+SL	99.6±0.0 / 12.4±0.1	93.2±0.0 / 12.4±0.1	97.3±0.0 / 12.8±0.0	91.5±0.0 / 18.5±0.1	1.0±0.1 / 0.1±0.0
+ML	99.4±0.0 / 11.6±0.2	92.1±0.0 / 2.4±0.1	97.6±0.0 / 3.3±0.1	95.1±0.0 / 3.6±0.2	75.2±0.0 / 9.8±0.0
+OE	99.5±0.0 / 10.2±0.2	92.5±0.0 / 3.2±0.1	97.5±0.0 / 5.0±0.1	94.9±0.0 / 6.8±0.2	74.1±0.0 / 8.0±0.0
LA	99.4±0.0 / 7.6±0.1	92.5±0.0 / 11.3±0.2	97.4±0.0 / 3.3±0.3	94.8±0.0 / 7.5±0.3	76.6±0.1 / 8.3±0.1
+NC	99.3±0.0 / 10.1±0.8	92.5±0.0 / 2.8±0.1	96.3±0.0 / 5.0±0.1	94.9±0.0 / 8.3±0.3	75.9±0.1 / 3.8±0.1
+SL	99.6±0.0 / 11.1±0.4	93.0±0.0 / 9.1±0.1	18.8±0.0 / 13.5±0.1	91.5±0.0 / 16.0±0.2	72.2±0.0 / 4.0±0.1
+ML	99.5±0.0 / 5.5±0.2	92.3±0.0 / 9.3±0.1	97.5±0.0 / 7.4±0.3	94.8±0.0 / 7.2±0.3	76.8±0.1 / 3.3±0.2
+OE	99.6±0.0 / 6.6±0.3	92.3±0.0 / 2.0±0.1	97.5±0.0 / 3.8±0.2	94.6±0.0 / 7.2±0.3	76.3±0.0 / 8.6±0.1

Table 11: OOD data detection under models trained with random noises (Hein et al., 2019) as  $\mathcal{D}_{\text{out}}$ . Values are FPR95, averaged over five prediction runs—lower is better.

Datasets	MAP	OE	DE	VB					LA				
				Plain	NC	SL	ML	OE	Plain	NC	SL	ML	OE
<b>MNIST</b>													
F-MNIST	11.8±0.0	6.8±0.0	5.3±0.0	12.5±0.1	6.5±0.0	0.7±0.0	11.9±0.1	10.3±0.1	12.0±0.0	8.2±0.0	0.0±0.0	6.3±0.0	6.8±0.0
E-MNIST	35.6±0.0	30.7±0.0	30.4±0.0	34.5±0.1	35.1±0.1	17.9±0.0	37.3±0.1	34.3±0.1	35.8±0.1	34.2±0.0	15.3±0.1	31.0±0.0	30.7±0.0
K-MNIST	14.4±0.0	7.8±0.0	7.7±0.0	14.0±0.1	14.5±0.1	1.1±0.0	15.8±0.1	14.0±0.1	14.5±0.1	10.6±0.0	0.7±0.0	8.5±0.0	7.8±0.0
CIFAR-Gr	0.2±0.0	0.0±0.0	0.0±0.0	0.2±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.1±0.0	0.2±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0
Uniform	44.3±0.0	0.7±0.0	19.8±0.0	93.1±0.2	0.0±0.0	0.0±0.0	1.9±0.1	1.8±0.0	54.2±0.4	0.7±0.0	0.0±0.0	0.6±0.0	0.8±0.0
<b>F-MNIST</b>													
MNIST	73.5±0.0	62.2±0.0	65.8±0.0	66.8±0.1	62.2±0.0	30.4±0.1	59.1±0.1	60.4±0.1	72.2±0.2	60.8±0.2	24.3±0.3	55.4±0.2	57.4±0.3
E-MNIST	73.6±0.0	50.2±0.0	58.6±0.0	68.1±0.1	43.9±0.0	25.7±0.0	54.7±0.1	54.6±0.1	72.2±0.2	44.2±0.1	22.5±0.2	39.6±0.2	48.3±0.3
K-MNIST	73.7±0.0	47.4±0.0	47.2±0.0	62.6±0.1	31.4±0.1	19.1±0.0	35.6±0.1	38.1±0.1	71.5±0.2	33.9±0.2	20.9±0.2	31.7±0.2	43.0±0.4
CIFAR-Gr	87.2±0.0	0.5±0.0	86.6±0.0	75.3±0.0	0.1±0.0	0.2±0.0	0.8±0.0	1.1±0.0	87.7±0.1	0.7±0.0	0.2±0.0	0.7±0.0	1.0±0.0
Uniform	81.3±0.0	26.0±0.0	86.3±0.0	87.3±0.1	47.1±0.2	0.0±0.0	0.1±0.0	4.9±0.0	81.0±0.2	43.4±0.7	0.0±0.0	22.0±0.2	38.1±0.9
<b>SVHN</b>													
CIFAR-10	18.9±0.0	13.8±0.0	9.5±0.0	15.0±0.0	13.0±0.1	16.5±0.0	8.4±0.0	11.5±0.0	15.4±0.1	8.4±0.0	94.8±1.6	14.9±0.1	11.4±0.1
LSUN-CR	19.7±0.0	9.0±0.0	8.3±0.0	17.2±0.2	10.5±0.1	8.8±0.1	5.4±0.1	9.1±0.1	15.5±0.1	8.3±0.3	95.4±4.0	12.6±0.2	8.2±0.1
CIFAR-100	21.8±0.0	15.6±0.0	11.6±0.0	18.1±0.0	14.8±0.1	17.9±0.0	10.2±0.0	12.4±0.0	17.6±0.1	11.7±0.0	93.9±0.7	16.6±0.1	13.4±0.1
FMNIST-3D	26.7±0.0	29.8±0.0	17.5±0.0	24.5±0.1	31.1±0.0	30.4±0.0	30.0±0.1	25.3±0.0	27.2±0.1	34.6±0.1	95.1±0.8	23.3±0.1	27.7±0.1
Uniform	30.0±0.0	0.0±0.0	6.4±0.0	48.2±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	17.0±0.1	0.0±0.0	90.2±0.6	19.0±0.1	0.0±0.0
<b>CIFAR-10</b>													
SVHN	34.5±0.0	7.3±0.0	33.9±0.0	33.5±0.0	11.1±0.0	20.6±0.0	11.0±0.0	9.7±0.0	35.5±0.1	6.6±0.0	26.0±0.3	7.5±0.0	8.3±0.1
LSUN-CR	53.3±0.0	49.0±0.0	44.0±0.0	49.4±0.4	46.7±0.1	61.7±0.0	45.7±0.2	47.6±0.3	53.8±0.6	48.9±0.5	54.9±0.4	51.9±0.3	48.3±0.5
CIFAR-100	61.2±0.0	58.2±0.0	52.5±0.0	58.4±0.1	57.7±0.1	71.1±0.0	56.3±0.1	56.6±0.1	61.4±0.1	59.3±0.2	70.7±0.3	57.6±0.1	59.0±0.3
FMNIST-3D	42.4±0.0	44.9±0.0	30.7±0.0	37.4±0.0	39.5±0.1	62.7±0.0	43.3±0.1	44.0±0.1	43.2±0.2	40.2±0.1	57.8±0.4	40.8±0.1	46.4±0.3
Uniform	87.7±0.0	26.7±0.0	0.0±0.0	13.8±0.1	100.0±0.0	57.3±0.1	98.0±0.0	100.0±0.0	92.8±0.1	3.5±0.1	17.7±0.4	12.9±0.3	49.8±0.9
<b>CIFAR-100</b>													
LSUN-CR	82.0±0.0	79.7±0.0	75.3±0.0	73.8±0.2	80.9±0.3	91.5±8.5	77.9±0.1	71.1±0.4	82.8±0.5	82.0±0.8	78.5±0.8	72.8±0.9	79.7±0.8
CIFAR-10	79.8±0.0	80.5±0.0	76.4±0.0	78.2±0.1	81.3±0.0	93.9±1.1	80.0±0.1	81.4±0.1	79.5±0.1	80.6±0.2	82.1±0.1	78.8±0.2	80.2±0.2
FMNIST-3D	65.8±0.0	66.9±0.0	61.8±0.0	57.1±0.1	69.3±0.1	93.6±1.1	63.1±0.1	61.9±0.1	66.1±0.1	71.2±0.2	82.0±0.1	64.4±0.4	67.9±0.2
Uniform	97.6±0.0	73.3±0.0	94.3±0.0	100.0±0.0	99.7±0.0	95.4±1.4	99.5±0.0	99.8±0.0	98.8±0.1	88.4±0.5	100.0±0.0	88.9±0.5	54.0±0.6

Table 12: Accuracy and ECE on text classification tasks, averaged over five prediction runs.

Methods	SST	TREC
MAP	78.1±0.0 / 20.8±0.0	76.0±0.0 / 17.2±0.0
DE	82.9±0.0 / 2.5±0.0	80.6±0.0 / 10.6±0.0
OE	78.6±0.0 / 13.0±0.0	68.8±0.0 / 9.4±0.0
LA	78.0±0.0 / 21.0±0.4	75.9±0.1 / 17.3±0.3
+NC	78.0±0.0 / 17.9±0.1	43.4±0.0 / 18.6±0.2
+DL	69.2±0.2 / 17.5±0.7	45.0±0.2 / 10.4±0.8
+ML	48.9±0.3 / 11.4±0.2	55.8±0.1 / 11.5±0.3
+OE	78.5±0.0 / 12.8±0.4	68.1±0.1 / 8.4±0.7

Table 13: OOD data detection on text classification tasks. Values are FPR95, averaged over five prediction runs—lower is better.

Datasets	MAP	OE	DE	LA				
				Plain	NC	SL	ML	OE
<b>SST</b>								
SNLI	100.0±0.0	0.0±0.0	100.0±0.0	100.0±0.0	0.0±0.0	97.0±0.3	89.6±0.7	0.0±0.0
Multi30k	100.0±0.0	0.0±0.0	100.0±0.0	100.0±0.0	0.0±0.0	99.5±0.0	83.5±1.5	0.0±0.0
WMT16	100.0±0.0	0.0±0.0	100.0±0.0	100.0±0.0	0.0±0.0	89.3±0.6	80.7±1.4	0.0±0.0
<b>TREC</b>								
SNLI	99.7±0.0	0.0±0.0	31.0±0.1	99.7±0.0	0.0±0.0	0.7±0.3	0.0±0.0	0.0±0.0
Multi30k	100.0±0.0	0.0±0.0	14.2±0.0	100.0±0.0	0.0±0.0	0.8±0.5	0.0±0.0	0.0±0.0
WMT16	89.2±0.0	0.0±0.0	27.3±0.0	89.3±0.0	0.0±0.0	0.8±0.7	0.0±0.0	0.0±0.0

Table 14: Test accuracy ( $\uparrow$ ) / ECE ( $\downarrow$ ) of more sophisticated base models, averaged over five prediction runs.

	<b>CIFAR-10</b>	<b>CIFAR-100</b>
Flipout	91.3 $\pm$ 0.0 / 10.9 $\pm$ 0.2	70.4 $\pm$ 0.1 / 19.8 $\pm$ 0.2
+NC	89.7 $\pm$ 0.1 / 8.2 $\pm$ 0.2	67.1 $\pm$ 0.1 / 13.8 $\pm$ 0.1
CSGHMC	93.9 $\pm$ 0.0 / 1.7 $\pm$ 0.0	74.0 $\pm$ 0.0 / 4.0 $\pm$ 0.0
+NC	92.2 $\pm$ 0.0 / 6.2 $\pm$ 0.0	71.6 $\pm$ 0.0 / 2.4 $\pm$ 0.0
DE	95.7 $\pm$ 0.0 / 4.5 $\pm$ 0.0	80.0 $\pm$ 0.0 / 1.9 $\pm$ 0.0
+NC	94.9 $\pm$ 0.0 / 4.8 $\pm$ 0.0	79.0 $\pm$ 0.0 / 1.7 $\pm$ 0.0

Table 15: OOD data detection with more sophisticated base models. Values are FPR95, averaged over five prediction runs.

<b>Datasets</b>	<b>Flipout</b>		<b>CSGHMC</b>		<b>DE</b>	
	<b>Plain</b>	<b>NC</b>	<b>Plain</b>	<b>NC</b>	<b>Plain</b>	<b>NC</b>
<b>CIFAR-10</b>						
SVHN	72.1 $\pm$ 0.3	39.6 $\pm$ 0.2	56.8 $\pm$ 0.0	16.4 $\pm$ 0.0	33.9 $\pm$ 0.0	8.1 $\pm$ 0.0
LSUN-CR	63.7 $\pm$ 1.0	37.5 $\pm$ 0.2	56.7 $\pm$ 0.0	24.0 $\pm$ 0.0	44.0 $\pm$ 0.0	18.3 $\pm$ 0.0
CIFAR-100	74.5 $\pm$ 0.2	70.4 $\pm$ 0.1	63.4 $\pm$ 0.0	63.1 $\pm$ 0.0	52.5 $\pm$ 0.0	51.7 $\pm$ 0.0
FMNIST-3D	65.0 $\pm$ 0.2	38.2 $\pm$ 0.1	51.0 $\pm$ 0.0	14.8 $\pm$ 0.0	30.7 $\pm$ 0.0	10.3 $\pm$ 0.0
Uniform	53.8 $\pm$ 0.6	0.0 $\pm$ 0.0	87.0 $\pm$ 0.0	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0
Smooth	61.1 $\pm$ 0.2	59.6 $\pm$ 0.2	47.1 $\pm$ 0.0	31.9 $\pm$ 0.0	32.9 $\pm$ 0.0	13.6 $\pm$ 0.0
<b>CIFAR-100</b>						
LSUN-CR	85.8 $\pm$ 0.6	55.6 $\pm$ 0.4	79.3 $\pm$ 0.0	38.0 $\pm$ 0.0	75.3 $\pm$ 0.0	54.0 $\pm$ 0.0
CIFAR-10	86.1 $\pm$ 0.3	87.0 $\pm$ 0.2	82.1 $\pm$ 0.0	84.2 $\pm$ 0.0	76.4 $\pm$ 0.0	78.5 $\pm$ 0.0
FMNIST-3D	73.4 $\pm$ 0.4	65.8 $\pm$ 0.2	67.0 $\pm$ 0.0	45.5 $\pm$ 0.0	61.8 $\pm$ 0.0	50.0 $\pm$ 0.0
Uniform	99.7 $\pm$ 0.0	0.0 $\pm$ 0.0	93.8 $\pm$ 0.0	0.0 $\pm$ 0.0	94.3 $\pm$ 0.0	0.0 $\pm$ 0.0
Smooth	82.0 $\pm$ 0.3	72.5 $\pm$ 0.3	83.0 $\pm$ 0.0	47.2 $\pm$ 0.0	58.7 $\pm$ 0.0	39.3 $\pm$ 0.0