

# Bayesian posterior approximation with stochastic ensembles

Oleksandr Balabanov<sup>1</sup>, Bernhard Mehlig<sup>2</sup>, Hampus Linander<sup>2,3</sup>

<sup>1</sup>Stockholm University <sup>2</sup>University of Gothenburg <sup>3</sup>Chalmers University of Technology

oleksandr.balabanov@fysik.su.se, bernhard.mehlig@physics.gu.se, hampus.linander@gu.se

Code: [github.com/oleksandr-balabanov/stochastic-ensembles](https://github.com/oleksandr-balabanov/stochastic-ensembles)

## Abstract

*We introduce ensembles of stochastic neural networks to approximate the Bayesian posterior, combining stochastic methods such as dropout with deep ensembles. The stochastic ensembles are formulated as families of distributions and trained to approximate the Bayesian posterior with variational inference. We implement stochastic ensembles based on Monte Carlo dropout, DropConnect and a novel non-parametric version of dropout and evaluate them on a toy problem and CIFAR image classification. For both tasks, we test the quality of the posteriors directly against Hamiltonian Monte Carlo simulations. Our results show that stochastic ensembles provide more accurate posterior estimates than other popular baselines for Bayesian inference.*

## 1. Introduction

Bayesian neural networks provide a principled way of reasoning about model selection and assessment with posterior distributions of model parameters [6, 8, 35, 49]. Although the analytical Bayesian posteriors can answer questions about model parameter uncertainty, the immense computational challenge for commonly used neural network architectures makes them practically infeasible<sup>1</sup>. Instead, we are forced to resign to approximation methods that makes a trade-off between posterior accuracy and computational complexity [27].

A prominent method to approximate the Bayesian posterior is deep ensembles [20, 31, 55], that can be shown to correspond to a variational inference approximation with a delta distribution family [24]. This method is implemented by simply training an ensemble of models and treating them as samples from the model posterior. In the variational inference formulation, this corresponds to approximating the posterior by sampling from many sets of maximum a posteriori parameters.

<sup>1</sup>There are examples of closed-form solutions for small architectures [56].

To further reduce the computational effort in evaluating the approximate posterior, stochastic methods such as Monte Carlo dropout [47] and DropConnect [51] inference have also been used extensively [13, 14, 40]. They benefit from computationally cheaper inference by virtue of sampling stochastically from a single model. Formulated as a variational approximation to the posterior, dropout samples from a family of parameter distributions where parameters can be randomly set to zero. Although this particular family of distributions might seem unnatural [11], it turns out that the stochastic property can help to find more robust regions of the parameter space, a fact well-known from a long history of using dropout as a regularization method.

Recently, there has been great progress towards understanding the analytical posteriors of larger neural networks by means of direct Markov Chain Monte Carlo sampling of the parameter posterior [26]. Through impressive computational efforts, the posteriors for models as large as ResNet-20 have been sampled using Hamiltonian Monte Carlo (HMC) simulations. The sampled posterior has been shown to provide more accurate predictions that are surprisingly sensitive to data distribution shift as compared to standard maximum likelihood estimation training procedures. These HMC computations have made it possible to compare approximation methods such as dropout inference and deep ensembles directly to the Bayesian posterior. Ensembling and stochastic methods such as dropout have been used successfully to find posterior approximations in many settings, but without a Bayesian formulation that includes both ensembling and stochastic methods it is difficult to understand if and how the two approaches can complement each other.

Recent work have shown that uncertainty quantification is subjective to neural network architectures, and that the accuracy of posterior approximations depend non-trivially on model architecture and dataset complexity [33]. To find computationally efficient methods that can accurately approximate the Bayesian posterior, for different data domains and architectures, is therefore an important goal with practical implications for applications that require accurate uncertainty quantification to assess network predictions [1, 17].

In this paper we combine deep ensembles and stochastic regularization methods into stochastic ensembles of neural networks. We formulate the stochastic ensemble construction within the Bayesian variational formalism, where multiple stochastic methods such as regular Monte Carlo dropout, DropConnect, and others are combined with ensembling into one general variational ansatz. We then conduct a series of tests using a simple toy model (synthetic data) and CIFAR (image classification), where stochastic deep ensembles are found to provide more accurate posteriors than MultiSWA [55] and regular deep ensembles in a number of settings. In particular, for CIFAR we use a neural network architecture evaluated by Izmailov et al. [26] in their large-scale experiments, allowing us to make a direct comparison of the ensemble methods to the “ground truth” HMC posterior.

## 2. Related Work

Bayesian inference for neural networks has a long history [6, 8, 35, 49]. A variety of methods have been developed for estimating the Bayesian posterior, including Laplace approximation [36], Markov Chain Monte Carlo methods [42, 46, 52], variational Bayesian methods [2, 4, 18, 22, 28, 34, 57] and Monte Carlo dropout [13, 14]. Deep ensembles [20, 31] have shown surprising abilities to approximate the Bayesian posterior well. They were found to empirically outperform other scalable methods [19, 43] and important conceptual aspects of their excellent performance were recently explained [10, 12, 55]. Forced diversity between ensemble members, motivated from a Bayesian perspective, has been shown to improve accuracy and out-of-distribution detection [9]. The dependence of the posterior on priors in weight and function space for deep ensembles has also been investigated [9, 50]. Various methods for improving deep ensembles have also recently been proposed [5, 7, 29, 39, 41, 53, 54, 58].

One of the widely used ensemble methods is MultiSWAG [55], where an ensemble of networks trained with Stochastic Weight Averaging Gaussian [25, 38, 55] is used for approximating the posterior. These networks were shown to have better generalization properties than networks trained using conventional SGD protocols [25, 38]. MultiSWAG was also found to provide posteriors that are more robust to distributional shifts than regular deep ensembles [55].

Our variational inference formulation is based on the works of Gal et al. [13, 14] for dropout and Hoffmann et al. [24] for deep ensembles. We combine these two methods and formulate one general Bayesian variational ansatz. We also note that our independent rederivation of the loss from Ref. [24] uncovered a new loss contribution promoting member diversity. A similar loss term was studied in Ref. [9] that was argued to be seen as Wasserstein gradient

descent on the Kullback-Leibler divergence. Here we derive it explicitly from Bayesian variational inference.

The numeric tests on CIFAR image classification follow closely the impressive effort of Izmailov et al. [26] reporting full-batch HMC computations on CIFAR datasets. We reproduced the exact network architecture and used their reported results for making a quantitative comparison of different ensemble types against HMC, indicating that our stochastic ensembles produce better posterior approximations than the non-stochastic baselines.

## 3. Our contribution

- We introduce stochastic ensembles. They are formulated using Bayesian variational inference where deep ensembles and stochastic methods such as dropout are combined into one general variational family of distributions. We show the theoretical advantage in using stochastic over regular (non-stochastic) ensembles and test different ensemble methods on a toy problem and CIFAR image classification.
- We show that for image classification on CIFAR, our stochastic ensembles based on Monte Carlo dropout produce the most accurate posterior approximations compared to other ensemble methods: They are found to be closer to HMC than MultiSWA and regular deep ensembles in terms of log-likelihood loss, accuracy, calibration, agreement, variance, out-of-distribution detection, predictive entropy, and robustness to distributional shifts.
- The posteriors produced by the Monte Carlo dropout ensembles for CIFAR image classification are also seen to outperform various stochastic gradient Monte Carlo (SGMCMC) methods: They are closer to the HMC posteriors and require less inference samples at test time.
- We introduce a novel stochastic model with a non-parametric version of dropout by jointly training an ensemble of models with exchange of parameters. It was found to outperform other methods for our simple toy task in terms of the agreement and variance computed in respect to the HMC posterior. It was also found to provide more accurate uncertainty estimates.

## 4. Bayesian deep learning

In modelling we aim to predict a label  $y^*$  for a given input  $x^*$  and provide uncertainty estimates of the prediction. The output should be in alignment with our observation, i.e. it should be conditioned on the training (observed) data  $\mathcal{D} = \{x_i, y_i\}_{i=1}^N$ . This is captured by the predictive probability distribution  $p(y^*|x^*, \mathcal{D})$ , which provides a distribution over possible output values  $y^*$ .



Consider a model with parameters  $\theta$ . The predictive distribution  $p(y^*|x^*, \mathcal{D})$  can then be represented as follows

$$p(y^*|x^*, \mathcal{D}) = \int d\theta \underbrace{p(y^*|x^*, \theta)}_{\text{aleatoric}} \underbrace{p(\theta|\mathcal{D})}_{\text{epistemic}}. \quad (1)$$

The posterior distribution  $p(\theta|\mathcal{D})$  describes epistemic uncertainty, the type of uncertainty that is by definition associated with the model parameters  $\theta$ . The term  $p(y^*|x^*, \theta)$  corresponds to aleatoric (data) uncertainty [16].

**Approximate methods.** In practice,  $p(\theta|\mathcal{D})$  is unknown and hard to compute: Direct methods for computing  $p(\theta|\mathcal{D})$  exist but they are intractable for models with large number of parameters. Approximate Bayesian inference aims to find a distribution  $q(\theta)$  that is as close to the true posterior  $p(\theta|\mathcal{D})$  as possible. To date, the golden standard method for approximate Bayesian inference is HMC [42]. It is based on performing Metropolis-Hastings updates using Hamiltonian dynamics associated with the potential energy  $U(\theta) = -\log p(y|x, \theta)p(\theta)$  with  $\mathcal{D} = (x, y)$ . This method is accurate but numerically tractable only for small models and datasets, severely limiting its practical applicability.

Variational inference [3, 4, 18] is less precise than HMC, but more numerically efficient. The idea is to look at a family of distributions  $q_\omega(\theta)$ , where  $\omega$  parameterizes the family, and find a distribution that approximates the posterior  $p(\theta|\mathcal{D})$  well. One common approach for finding the best family member  $q_\omega(\theta)$  is to numerically minimize the Kullback-Leibler (KL) divergence  $\text{KL}(q_\omega(\theta) || p(\theta|\mathcal{D}))$  [16], a measure quantifying how much information is lost when  $q_\omega(\theta)$  is approximated by  $p(\theta|\mathcal{D})$ :

$$\begin{aligned} \text{KL}(q_\omega(\theta) || p(\theta|\mathcal{D})) &= \int d\theta q_\omega(\theta) \log \frac{q_\omega(\theta)}{p(\theta|\mathcal{D})} \\ &= \text{KL}(q_\omega(\theta) || p(\theta)) - \mathbb{E}_{q_\omega(\theta)}[\log p(y|\theta, x)] + C, \end{aligned} \quad (2)$$

where the training data is denoted by  $\mathcal{D} = (x, y)$  and  $C$  is a  $\theta$ -independent term. Here we applied Bayes' theorem  $p(\theta|x, y) = p(y|\theta, x)p(\theta)/p(y|x)$ . By minimizing the loss function in Eq. (2) one obtains an approximation for the posterior  $p(\theta|\mathcal{D})$ . Note that the cost function consists of two terms, (1) the KL divergence against the prior  $p(\theta)$  and (2) the expected negative log likelihood (ENLL), that is a standard loss for classification tasks. Minimization of the ENLL loss in Eq. (2) is usually achieved via stochastic-gradient-based methods implemented on the Monte Carlo sampled model.

**Deep ensembles.** Conventional maximum likelihood training of a neural network can be naturally described within the Bayesian variational inference picture [18]. Deep ensembles can be explicitly reformulated within the

Bayesian setting as well [24] and are associated with ansatz

$$q_\omega(\theta) = \frac{1}{K} \sum_{k=1}^K \mathcal{N}(\theta; \omega_k, \sigma^2 I_{\dim[\theta]}), \quad (3)$$

where  $K$  is the number of networks in the ensemble,  $\omega$  consists of  $K$  independent sets  $\omega_k$  of parameters  $\theta$ ,  $k = 1, \dots, K$ . Conventional deep ensembles are associated with the case of infinitesimal (machine-precision)  $\sigma$  so that the normal distributions in Eq. (3) are sharply peaked and approach the delta function distributions  $\delta(\theta - \omega_k)$ .

The ENLL contribution to the cost function in Eq. (2) simplifies to a sum of independent loss terms associated with each ensemble member, weighted by factor  $1/K$ . In this case the KL divergence against the standard normally-distributed prior  $p(\theta) = \mathcal{N}(\theta; 0, \lambda^{-1} I_{\dim[\theta]})$  reduces to [24]

$$\begin{aligned} \text{KL}(q_\omega(\theta) || p(\theta)) &= \frac{1}{2} \dim[\theta] (\lambda \sigma^2 - \log \sigma^2 - 1 - \log \lambda) \\ &+ \underbrace{\frac{1}{2K} \sum_{k=1}^K \lambda \|\omega_k\|^2}_{\text{L2 regularization}} - \underbrace{\log K}_{\text{KL loss reduction due to ensembling}} + \underbrace{\text{RF}}_{\text{repulsive force}}, \end{aligned} \quad (4)$$

where the repulsive-force correction RF becomes negligibly small as  $\sigma$  approaches zero and was dropped in Ref. [24]. In our SM [48] we show that the repulsive-force correction can be approximated by the following upper bound

$$\text{RF} \leq \frac{1}{K} \sum_{k=1}^K \sum_{k' \neq k} \exp[-\|\omega_k - \omega_{k'}\|^2 / (8\sigma^2)]. \quad (5)$$

The L2 regularization term in Eq. (4) corresponds to our choice of a Gaussian prior distribution for the model parameters, other terms come from the entropy contribution  $H(q_\omega(\theta))$  to  $\text{KL}(q_\omega(\theta) || p(\theta))$  that is independent of the prior choice. The KL loss reduction due to ensembling is represented here by the  $[-\log K]$  term that decreases with the ensemble size. Interestingly, Eq. (4) also contains a repulsive force in weight space promoting the ensemble members to be diverse. Diversity among the ensemble members is believed to be a key element for reaching good performance [9, 12].

## 5. Stochastic Ensembles

**Posterior landscape.** To motivate our construction we start with an intuitive perspective. Regular deep ensembles are believed to be effective at approximating the posterior because of their ability to probe multiple modes of the posterior landscape [12, 24]. This is in contrast to less effective single network stochastic methods that only sample from a single mode. MultiSWAG [55] improves deep ensembles

by virtue of Stochastic Weight Averaging protocol that finds flatter parameter regions in the basins of attraction uncovered by deep ensembles. The flatter regions are believed to more accurately represent the posterior [25, 38, 55]. Note, however, that the basins of attraction sampled by MultiSWAG may not be optimal and there can exist wider basins that may be less probable for regular training protocols to end up in. Stochastic ensembles are built from trained networks regularized by stochastic methods such as dropout. They are also multimodal but they only find basins of attraction that are robust to stochastic perturbations and sampling from wide, robust parameter regions is essential for constructing an efficient posterior approximation, cf. the outlined arguments for MultiSWAG.

The length scale of parameter regions probed by a stochastic method is defined by its hyperparameters: dropout networks with larger drop rates, for example, will look for larger robust parameter regions. There will always be a trade off between the region size and its quality in terms of the robustness. To avoid the need to manually tune the hyper parameters one can aim towards finding a non-parametric stochastic method capable of probing parameter regions on different scales. This motivated us to introduce a non-parametric version of dropout described in the next section. We also note that different stochastic methods probe the parameter space in different ways. For example dropout probes random hyperplane projections in the parameter space.

**Bayesian description.** A well known example of stochastic regularization methods is Monte Carlo dropout [11]. For concreteness let us consider a single fully-connected layer with conventional (Bernoulli) dropout applied after some activation. The layer’s input-output transformation, say from the node layer  $x_l$  to the node layer  $x_{l+1}$ , can be expressed as follows

$$x_{l+1} = \sigma(W_{l+1,l} x_l + B_{l+1}) \circ z_{l+1}, \quad (6)$$

where  $W_{l+1,l}$  and  $B_{l+1}$  are the layer’s weights and biases respectively,  $\sigma(x)$  is the activation, and  $z_{l+1,n} \sim \mathcal{B}(p_{l+1,n})$  are Bernoulli-distributed variables, namely  $z_{l+1,n} = 1$  with probability  $p_{l+1,n}$  and  $z_{l+1,n} = 0$  with probability  $(1 - p_{l+1,n})$  with  $n = 1, \dots, N$ . Here  $N$  is the layer’s output dimension and  $\circ$  denotes the element-wise (Hadamard) product.

The dropout input-output transformation in Eq. (6) can be equivalently described as Monte Carlo sampling from a family of Bernoulli functions. Recently it was shown that networks regularized by Bernoulli methods can be viewed as variational inference [13–15]. We combine these networks into an ensemble and obtain a generalization of the deep ensemble variational ansatz in Eq. (3) that we call stochastic ensemble. For the example of a fully-connected

dropout layer in Eq. (6), the stochastic ensemble ansatz is given by

$$q_\omega(\theta_l) = \frac{1}{K} \sum_{k=1}^K \prod_{n=1}^N \hat{q}_{\omega_{l,n,k}}(\theta_{l,n}), \quad (7)$$

with

$$\hat{q}_{\omega_{l,n,k}}(\theta_{l,n}) = p_{l+1,n}^{(1)} \mathcal{N}(\theta_{l,n}; \omega_{l,n,k}^{(1)}, \sigma^2 I_{\dim[\theta_{l,n}]}) + p_{l+1,n}^{(2)} \mathcal{N}(\theta_{l,n}; \omega_{l,n,k}^{(2)}, \sigma^2 I_{\dim[\theta_{l,n}]}), \quad (8)$$

where  $n$  labels the layer’s output nodes  $x_{l+1}$ ,  $\theta_{l,n}$  consists of all weights and bias connecting the layer’s input  $x_l$  to the layer’s output node  $x_{l+1,n}$ , and  $\omega_{l,n,k}^{(i)}$  for  $i = 1, 2$  are independent realizations of the network’s parameters  $\theta_{l,n}$ ,  $p_{l,n}^{(i)}$  the Bernoulli probabilities for the two realizations, and  $k = 1, \dots, K$ .

As in the case of regular deep ensembles, the ENLL loss reduces to a sum of independent single network losses weighted by  $1/K$ . Each such loss term and its gradients can be handled using conventional stochastic methods such as Monte Carlo sampling [14]. The KL divergence against the normal prior reduces to (see our SM [48]):

$$\begin{aligned} \text{KL}(q_\omega(\theta) || p(\theta)) &= \frac{1}{2} \dim[\theta] (\lambda \sigma^2 - \log \sigma^2 - 1 - \log \lambda) \\ &+ \frac{1}{2K} \sum_{k=1}^K \sum_{n=1}^N \sum_{i=1}^2 \lambda p_{l+1,n}^{(i)} \|\omega_{l,n,k}^{(i)}\|^2 - \log K \\ &+ \underbrace{\sum_{n=1}^N \sum_{i=1}^2 p_{l+1,n}^{(i)} \log p_{l+1,n}^{(i)}}_{\text{KL loss reduction due to stochasticity}} + \text{RF}_2. \end{aligned} \quad (9)$$

The difference to regular deep ensembles in Eq. (4) is one additional term that describes the KL loss reduction corresponding to using stochastic over non-stochastic deep networks [14]. The repulsive force  $\text{RF}_2$  is of a more complex form than in the regular deep ensemble case, see our SM [48] for the derivation of the corresponding upper bound. The bound contains an accumulation of the RF contributions coming from every possible stochastic parameter realizations, weighted by the probabilities of obtaining these configurations. Note that it approaches zero for infinitesimal (machine-precision)  $\sigma$ .

**Stochastic ensemble realizations.** The ansatz in Eqs. (7) and (8) is equivalent to the dropout ensemble by assuming the limit of small  $\sigma$ , and taking  $\omega_{l,n,k}^{(2)}$  to be always zero. Other types of stochastic ensembles, such as DropConnect ensembles, as well as other varieties of network layers, beyond the fully-connected layers, can

be formulated similarly to Eq. (7). In our experiments we study the dropout and DropConnect realizations of stochastic ensembles, and denote them by SE1 and SE2 respectively. We also look at a realization of Eqs. (7) and (8) where two sets of trainable parameters are exchanged with equal probability  $p_{l+1,n}^{(1)} = p_{l+1,n}^{(2)} = 1/2$ . This provides a non-parametric method (in the sense of no tunable probabilities), referred to as SE3, that can be applied to any layer type with just minor changes, including the output layer.

**Limitations of the KL loss.** The KL loss in Eqs. (4) and (9) suggest that ensembles of stochastic networks are expected to more accurately approximate the posterior than regular deep ensembles, especially for large models because the loss reduction due to stochasticity in Eq. (9) scales with  $N$ . The KL divergence, however, is not a flawless measure. It is known to ignore parameter regions where the distribution  $q_{\omega}(\theta)$  is small, and therefore may not penalize the discrepancy with some important parts of the posterior. An interesting improvement to the KL loss was studied in Ref. [32].

We also stress that the total KL loss that is minimized during the training is not directly correlated with the performance metrics such as accuracy and calibration on the validation set. Smaller KL loss is expected to lead to better quantitative estimates, however, for a given task the quality of these estimates may depend on the data and considered metrics.

## 6. Numerical Results

### 6.1. Toy model

To quantitatively compare different ensemble methods we first evaluate them on a toy classification problem. The advantage of using a simple toy model is that the posterior can be straightforwardly sampled using HMC. We implement the NUTS extension [23] of HMC and directly evaluate performance of ensemble methods by comparing their outcomes with the corresponding HMC posterior. The HMC posterior is constructed by stacking 4 independent HMC chains, each of 2000 samples. The details on our HMC (NUTS) implementation are given in the SM [48].

We consider three training datasets shown in Fig. 1 (a-c), where the 2D data is classified into two classes. The datasets represent different class separability levels, helping us to make a more comprehensive comparison of the ensemble methods. We employ a simple feed forward network with two hidden layers of 10 neurons, ReLU activations, and softmax output. The trained models are then tested on two uniformly sampled data domains,  $\mathcal{D}_{in} = [-1, 1]^2$  (in-domain) and  $\mathcal{D}_{out} = [-10, 10]^2$  (out-of-domain). We consider three types of stochastic




Figure 1. (a-c) Three different datasets of training data and corresponding to them HMC uncertainty estimates, entropy and mutual information.

ensembles in our experiments: Monte Carlo dropout (SE1), DropConnect (SE2), and non-parametric dropout (SE3). More details on the model, training procedure and ensembles are provided in our SM [48].

**Uncertainty estimates.** In Fig. 1 we also depict two important uncertainty estimates evaluated using the HMC posterior. The total predictive uncertainty can be estimated by computing the predictive posterior entropy that formally quantifies how much information on average is contained in the output [45]. The second quantity is the (average) mutual Shannon information contained between network parameters and test data sample conditioned on the training dataset. It quantifies how the posterior distribution changes as we include a new data point to the training dataset [37]. The mutual information describes only the epistemic part of the total uncertainty and it is an important quantity used in many deep learning applications, for example, in active learning [44]. For more details on these two quantities, and our implementation, see the SM [48].

The toy classification problem is a good illustration for the concepts of aleatoric (data) uncertainty and epistemic (model) uncertainty. For the boundary regions between the two classes, the prediction is expected to exhibit high uncertainty that is mostly of aleatoric origin. Our HMC results in Fig. 1 are in agreement with the expectation: The total entropy correctly identifies regions of high uncertainty, both the boundary regions (high aleatoric uncertainty) and certain regions that were not used in the training (high

	Entropy 1e-3	MI 1e-3	Agr 1e-2	Var 1e-2
Toy-a (in-domain / out-of-domain)				
Regular	0.51 1.18	0.42 1.16	98.3 95.2	2.76 7.49
MultiSWA	0.71 1.32	0.48 1.31	97.30 94.01	4.00 8.46
NP Dropout (SE3)	<b>0.46</b> <b>1.14</b>	<b>0.39</b> <b>1.13</b>	<b>98.4</b> <b>95.5</b>	<b>2.59</b> <b>7.00</b>
Toy-b (in-domain / out-of-domain)				
Regular	0.30 0.89	0.19 0.87	<b>98.9</b> 95.4	1.61 5.44
MultiSWA	0.49 1.06	0.21 1.03	97.7 94.4	2.89 6.74
NP Dropout (SE3)	<b>0.25</b> <b>0.73</b>	<b>0.16</b> <b>0.71</b>	<b>98.9</b> <b>95.8</b>	<b>1.41</b> <b>4.51</b>
Toy-c (in-domain / out-of-domain)				
Regular	0.23 0.86	0.12 0.83	<b>99.1</b> <b>96.8</b>	<b>1.19</b> 4.39
MultiSWA	0.32 0.90	0.10 0.86	98.1 95.5	2.01 5.01
NP Dropout (SE3)	<b>0.22</b> <b>0.76</b>	<b>0.11</b> <b>0.73</b>	99.0 96.7	1.20 <b>3.86</b>

Table 1. Quantitative comparison of predictions obtained from HMC, regular ensemble, MultiSWA, and non-parametric dropout ensemble SE3. The tests are done using data from  $\mathcal{D}_{\text{in}} = [-1, 1]^2$  (in-domain, top rows) and  $\mathcal{D}_{\text{out}} = [-10, 10]^2$  (out-of-domain, bottom rows). We considered all three different toy datasets from Fig. 1. The considered metrics are agreement, variance, mean absolute difference of entropy and mutual information estimates computed in respect to the HMC runs. All variances are orders of magnitude smaller than quoted results given the large ensembles.

epistemic uncertainty). The mutual information is seen to be large only in the regions of high epistemic uncertainty. The aleatoric (epistemic) uncertainty is seen to increase (decrease) with the class mixing level.

**Performance of the ensemble methods.** We quantitatively evaluate performance of the ensemble methods by looking at the agreement and variance between the ensemble posteriors and “ground truth” HMC posterior. For an explicit definition of these two quantities see our SM [48] or Ref. [26]. We also report baseline results for the regular (non-stochastic) deep ensembles and MultiSWA ensemble method [55]. To quantify the quality of uncertainty estimates we also look at the mean absolute differences computed between the HMC and ensemble estimates. All three classification datasets depicted in Fig. 1 (a-c) are considered. For each dataset, we compute metrics on data from

$\mathcal{D}_{\text{in}}$  (in-domain) and  $\mathcal{D}_{\text{out}}$  (out-of-domain). The size of each ensemble is 1024. Here we always take only one Monte Carlo inference per trained ensemble member, in this way forcing all types of ensembles to use the same resources at test time. In accordance with this constraint, we implement the test-time efficient MultiSWA instead of MultiSWAG [55], see our SM [48] for implementation details.

The dropout (SE1) and DropConnect (SE2) ensembles were found to not perform as good as the regular and non-parametric dropout (SE3) ensembles, highlighting importance of correctly picking the stochastic type for a particular task. For our toy problems we could anticipate this behaviour by looking at the training KL loss exhibiting high fluctuations for SE1 and SE2. The KL loss of SE3 and non-stochastic networks did not fluctuate as much and resulted in smaller training KL loss values.

We present results corresponding to the non-parametric dropout (SE3), MultiSWA and regular ensembles in Table 1. We observe that SE3 outperforms the other two methods, in most places with large margin. The average improvement with respect to the regular ensemble is approximately 10 percents in each of the metrics. Interestingly, the improvement differs between the different versions of the classification dataset and drops with the mixing between the training datasets. We attribute this to the difficulty of stochastic methods to fit the boundary between mixed classes. We also note that for the case with the largest mixing, SE3 and regular ensembles are indistinguishable in agreement and variance for in-domain data. Surprisingly, MultiSWA does not perform better than the regular ensemble baseline.

## 6.2. CIFAR

Let us consider now more realistic data and focus on the classification of images. The CIFAR dataset [30] consists of 60000 low-resolution (32 by 32) RGB images classified into 10 (CIFAR-10) or 100 (CIFAR-100) classes. Complexity of this image dataset makes the full-batch HMC computations highly nontrivial and computationally demanding. Only recently the full-batch HMC method has been evaluated on CIFAR for large residual networks [26]. In this section we train the much less demanding stochastic and non-stochastic ensembles on the CIFAR dataset and compare our results to the HMC data from Ref. [26]. We employ the same network architecture: A ResNet-20-FRN residual network of depth 20 with batch normalization layers replaced with filter response normalization (FRN). The HMC posteriors are directly loaded from the publicly available resource [26].

Bayesian variational inference fixes some aspects of the training procedure such as weight decay and dispense with the need for early stopping. Our training procedure is therefore different from the implementation in Ref. [26] leading to different results for regular ensembles. Three



	Acc	Loss	ECE	Agr 1e-2	Var 1e-2	ODD
CIFAR-10						
HMC 3-chains	90.68	0.307	0.059	100	0.0	85.3
Regular	88.82 $\pm 0.03$	0.339 $\pm 0.001$	0.028 $\pm 0.001$	92.5 $\pm 0.1$	10.0 $\pm 0.2$	83.7 $\pm 0.2$
Multi SWA	88.54 $\pm 0.15$	0.374 $\pm 0.024$	0.057 $\pm 0.019$	92.4 $\pm 0.1$	11.6 $\pm 1.4$	83.1 $\pm 1.4$
Dropout (SE1)	<b>90.82</b> $\pm 0.03$	<b>0.301</b> $\pm 0.001$	<b>0.057</b> $\pm 0.001$	<b>94.1</b> $\pm 0.1$	<b>8.0</b> $\pm 0.1$	<b>85.5</b> $\pm 0.1$
CIFAR-100						
HMC 3-chain	67.39	1.205	0.131	100	0.0	72.5
Regular	62.84 $\pm 0.16$	1.477 $\pm 0.001$	0.157 $\pm 0.002$	71.4 $\pm 0.1$	29.8 $\pm 0.1$	71.5 $\pm 0.1$
Multi SWA	62.32 $\pm 0.4$	1.511 $\pm 0.017$	0.164 $\pm 0.001$	70.7 $\pm 0.3$	30.7 $\pm 0.6$	71.6 $\pm 0.6$
Dropout (SE1)	<b>68.49</b> $\pm 0.18$	<b>1.167</b> $\pm 0.001$	<b>0.124</b> $\pm 0.002$	<b>77.2</b> $\pm 0.1$	<b>22.2</b> $\pm 0.1$	<b>73.5</b> $\pm 0.1$

Table 2. Prediction accuracy (acc), test log-likelihood loss (loss), expected calibration error (ECE), agreement (agr), variance (var) and out-of-domain detection (ODD) for different inference methods trained on the CIFAR datasets. The HMC results are computed using the published HMC chains from Ref. [26]. Variances are calculated over independent runs.

types of stochastic ensembles are used in our CIFAR tests: Monte Carlo dropout (SE1), DropConnect (SE2), and non-parametric dropout (SE3). The drop rates of SE1 and SE2 were tuned to target the HMC results, rather than targeting best performance in terms of accuracy and loss. Each ensemble consists of 50 networks and we use strictly one Monte Carlo inference per ensemble member at test time, i.e. in total 50 inferences per ensemble. This restricts the methods to use the same resources. We also considered increasing the number of inferences per member, obtaining only a minor change in the performance. We report these results in our SM [48]. We also trained multiple realizations of the MultiSWA protocol corresponding to different hyperparameter sets. In the following we only present the closest MultiSWA realization to HMC. For the details on our hyperparameter selection see the SM [48].

**Stochastic ensembles.** We found our Monte Carlo dropout ensembles (SE1) to produce more accurate posteriors than DropConnect (SE2) and non-parametric dropout (SE3) ensembles. We therefore present only results corresponding to SE1. We could anticipate SE1 performing better than SE2 and SE3 from the training KL loss analysis.




Figure 2. Histogram plots of predictive uncertainty computed using different ensemble methods. The HMC baselines are shown in orange in every subplot for better visualisation contrast.




Figure 3. Calibration curves associated with different ensembles trained and evaluated on CIFAR.

**Test performance metrics.** In Table 2 we present the test accuracy, test loss, expected calibration error (ECE), agreement, variance, and out-of-distribution detection (ODD) evaluated for different types of ensembles. For metric definitions see our SM [48]. From Table 2 we read off that the Monte Carlo dropout ensembles (SE1) provide more accurate posteriors than the baseline MultiSWA and regular deep ensembles.

**Predictive entropy.** In Fig. 2 we present the predictive entropy computed for different ensembles trained on CIFAR. SE1 is seen to match more accurately the HMC distribution of entropy values than regular deep ensembles. We also note that MultiSWA and SE1 perform equally well on CIFAR-10 but SE1 is more accurate on CIFAR-100.

**Calibration curve.** An accurate approximation of Bayesian posterior is expected to be equally calibrated as HMC. In Fig. 3 we depict calibration curves associated with different ensemble methods. All methods are found to be relatively poorly calibrated. The SE1 ensembles are seen to be the closest to HMC.

**Robustness to distribution shift.** To test robustness to distribution shifts we evaluate the performance of different methods on the CIFAR-10-C and CIFAR-100-C

	Acc	Loss	ECE	Ag 1e-2	Var 1e-2
CIFAR-10-C (mean over corruptions)					
HMC 3-chains	71.05	0.879	0.079	100	0.0
Regular	78.48 $\pm 0.21$	0.647 $\pm 0.006$	0.042 $\pm 0.001$	79.8 $\pm 0.3$	20.6 $\pm 0.4$
Multi SWA	77.55 $\pm 0.36$	0.690 $\pm 0.021$	0.059 $\pm 0.012$	80.3 $\pm 0.1$	20.6 $\pm 0.6$
Dropout (SE1)	<b>76.32</b> $\pm 0.29$	<b>0.713</b> $\pm 0.002$	<b>0.064</b> $\pm 0.001$	<b>83.1</b> $\pm 0.1$	<b>16.4</b> $\pm 0.1$
CIFAR-100-C (mean over corruptions)					
HMC 3-chain	42.10	2.610	0.113	100	0.0
Regular	49.33 $\pm 0.13$	2.067 $\pm 0.009$	0.115 $\pm 0.001$	53.3 $\pm 0.3$	39.7 $\pm 0.2$
Multi SWA	<b>48.00</b> $\pm 0.70$	2.132 $\pm 0.032$	0.115 $\pm 0.002$	53.6 $\pm 0.1$	39.7 $\pm 0.1$
Dropout (SE1)	48.35 $\pm 0.05$	<b>2.115</b> $\pm 0.003$	<b>0.089</b> $\pm 0.001$	<b>61.1</b> $\pm 0.2$	<b>32.0</b> $\pm 0.1$

Table 3. Performance rates corresponding to different ensemble methods evaluated on CIFAR-C. The CIFAR test datasets are corrupted in 16 different ways at various intensities on the scale of 1 to 5. We average the results over all  $16 \cdot 5$  corrupted datasets and report variances over independent runs. The HMC data was computed using the published HMC chains from Ref. [26].

datasets [21]. These datasets contain CIFAR images altered using different types of corruptions with varying intensities. We use the same 16 corruptions as in Ref. [26]. For each method and corruption we compute the accuracy, loss, ECE, agreement, and variance, and then average over all the corrupted datasets. The corresponding data is presented in Table 3, where SE1 is seen to again be closer to HMC in terms of the agreement and variance. In Fig. 4 we plot the accuracy and loss resolved over different corruption types and intensities. SE1 is found to be more sensitive to the distribution shifts than the regular ensembles and MultiSWA but still not as much as HMC [26].

**Comparison to stochastic gradient Monte Carlo (SGMCMC) approximate posteriors.** Our SE1 is closer to HMC than published results [26] for SGMCMC methods in terms of accuracy, loss, ECE, agreement, variance. SE1 also provides closer posteriors to HMC for CIFAR-10-C evaluations, while also being more computationally efficient.

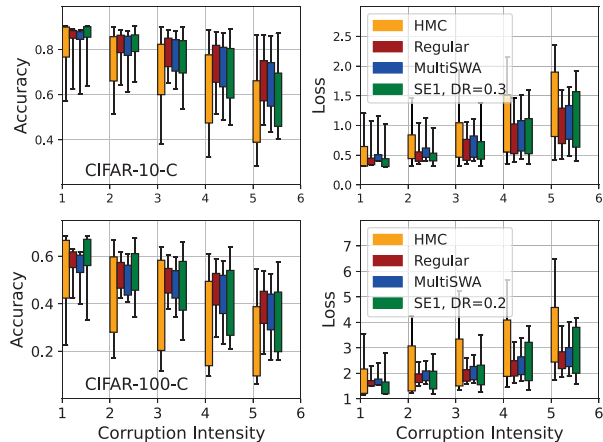


Figure 4. Evaluation on CIFAR-C. The CIFAR test sets are corrupted in 16 different ways at various intensities on the scale of 1 to 5. The error bars depict min and max of the corresponding metrics over each corruption, with the boxes indicating quartiles.

## 7. Conclusions

We introduced a variational ansatz for stochastic ensembles of neural networks. Common Bayesian posterior approximations such as regular deep ensembles, Monte Carlo dropout and DropConnect can all be formulated as special cases of these distributional families. Our theoretical and numerical results suggest that the ensembles of stochastic neural networks can provide more accurate approximations to the Bayesian posteriors than other baseline methods.

Using our variational ansatz we formulated a new type (SE3) of non-parametric stochastic method that can be applied to any layer type with only minor changes to incorporate parameter sharing between ensemble members. For a simple toy model, we showed with HMC sampling of the Bayesian posterior that our SE3 ensembles provide closer posteriors to HMC than all other considered ensemble methods.

We evaluated accuracy, loss, ECE, ODD, entropy, agreement, variance and robustness to distribution shift for a ResNet-20-FRN architecture using different types of ensembles in image classification tasks on CIFAR. We found that our dropout stochastic ensembles (SE1) are closer to the HMC posteriors than all other methods considered.

In each of our numeric tests, all the stochastic methods were manually tuned except non-parametric dropout. To understand better what hyperparameters and stochastic method to use for achieving the most accurate posterior approximations would be highly beneficial.

**Acknowledgements.** OB was supported by the Knut and Alice Wallenberg Foundation and the Swedish Research Council (grant 2017-05162). HL and BM were supported by Vetenskapsrådet (grants 2017-3865 and 2021-4452).

## References

- [1] Moloud Abdar, Farhad Pourpanah, Sadiq Hussain, Dana Rezazadegan, Li Liu, Mohammad Ghavamzadeh, Paul Fieguth, Xiaochun Cao, Abbas Khosravi, U. Rajendra Acharya, Vladimir Makarenkov, and Saeid Nahavandi. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information Fusion*, 76:243–297, 2021.
- [2] David Barber and Christopher M. Bishop. Ensemble learning in bayesian neural networks. In *Neural Networks and Machine Learning*, pages 215–237. Springer, 1998.
- [3] David Barber and Charles M. Bishop. Ensemble learning in bayesian neural networks. 1998.
- [4] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ICMML'15, page 1613–1622. JMLR.org, 2015.
- [5] Anh Bui, Trung Le, He Zhao, Paul Montague, Olivier deVel, Tamas Abraham, and Dinh Phung. Improving Ensemble Robustness by Collaboratively Promoting and Demoting Adversarial Robustness. *arXiv e-prints*, page arXiv:2009.09612, Sept. 2020.
- [6] Wray L. Buntine and Andreas S. Weigend. Bayesian back-propagation. *Complex Syst.*, 5, 1991.
- [7] Asa Cooper Stickland and Iain Murray. Diverse Ensembles Improve Calibration. *arXiv e-prints*, page arXiv:2007.04206, July 2020.
- [8] John S. Denker, Daniel B. Schwartz, Ben S. Wittner, Sara A. Solla, Richard E. Howard, Lawrence D. Jackel, and John J. Hopfield. Large automatic learning, rule extraction, and generalization. *Complex Syst.*, 1, 1987.
- [9] Francesco D' Angelo and Vincent Fortuin. Repulsive deep ensembles are bayesian. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 3451–3465. Curran Associates, Inc., 2021.
- [10] Vikranth Dwaracherla, Zheng Wen, Ian Osband, Xiuyuan Lu, Seyed Mohammad Asghari, and Benjamin Van Roy. Ensembles for Uncertainty Estimation: Benefits of Prior Functions and Bootstrapping. *arXiv e-prints*, page arXiv:2206.03633, June 2022.
- [11] Loic Le Folgoc, Vasileios Baltatzis, Sujal Desai, Anand Devaraj, Sam Ellis, Octavio E. Martinez Manzanera, Arjun Nair, Huaqi Qiu, Julia Schnabel, and Ben Glocker. Is mc dropout bayesian?, 2021.
- [12] Stanislav Fort, Huiyi Hu, and Balaji Lakshminarayanan. Deep Ensembles: A Loss Landscape Perspective. *arXiv e-prints*, page arXiv:1912.02757, Dec. 2019.
- [13] Yarin Gal and Zoubin Ghahramani. Bayesian Convolutional Neural Networks with Bernoulli Approximate Variational Inference. *arXiv e-prints*, page arXiv:1506.02158, June 2015.
- [14] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1050–1059, New York, New York, USA, 20–22 Jun 2016. PMLR.
- [15] Yarin Gal, Jiri Hron, and Alex Kendall. Concrete dropout. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [16] Jakob Gawlikowski, Cedrique Rovile Njjeutcheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, Muhammad Shahzad, Wen Yang, Richard Bamler, and Xiao Xiang Zhu. A Survey of Uncertainty in Deep Neural Networks. *arXiv e-prints*, page arXiv:2107.03342, July 2021.
- [17] Jakob Gawlikowski, Cedrique Rovile Njjeutcheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, Muhammad Shahzad, Wen Yang, Richard Bamler, and Xiao Xiang Zhu. A survey of uncertainty in deep neural networks, 2022.
- [18] Alex Graves. Practical variational inference for neural networks. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011.
- [19] Fredrik K. Gustafsson, Martin Danelljan, and Thomas B. Schön. Evaluating Scalable Bayesian Deep Learning Methods for Robust Computer Vision. *arXiv e-prints*, page arXiv:1906.01620, June 2019.
- [20] L. K. Hansen and P. Salamon. Neural network ensembles. *IEEE Trans. Pattern Anal. Mach. Intell.*, 12(10):993–1001, oct 1990.
- [21] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*, 2019.
- [22] Geoffrey E. Hinton and Drew van Camp. Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the Sixth Annual Conference on Computational Learning Theory*, COLT '93, page 5–13, New York, NY, USA, 1993. Association for Computing Machinery.
- [23] Matthew D. Hoffman and Andrew Gelman. The no-urn sampler: Adaptively setting path lengths in hamiltonian monte carlo, 2011.
- [24] Lara Hoffmann and Clemens Elster. Deep Ensembles from a Bayesian Perspective. *arXiv e-prints*, page arXiv:2105.13283, May 2021.
- [25] Pavel Izmailov, Dmitrii Podoprikin, T. Garipov, Dmitry P. Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. *ArXiv*, abs/1803.05407, 2018.
- [26] Pavel Izmailov, Sharad Vikram, Matthew D. Hoffman, and Andrew Gordon Wilson. What Are Bayesian Neural Network Posteriors Really Like? *arXiv e-prints*, page arXiv:2104.14421, Apr. 2021.
- [27] Laurent Valentin Jospin, Hamid Laga, Farid Boussaid, Wray Buntine, and Mohammed Bennamoun. Hands-on bayesian

- neural networks—a tutorial for deep learning users. *IEEE Computational Intelligence Magazine*, 17(2):29–48, 2022.
- [28] Diederik P. Kingma, Tim Salimans, and Max Welling. Variational Dropout and the Local Reparameterization Trick. *arXiv e-prints*, page arXiv:1506.02557, June 2015.
- [29] Lucas Kook, Andrea Götschi, Philipp FM Baumann, Torsten Hothorn, and Beate Sick. Deep interpretable ensembles. *arXiv e-prints*, page arXiv:2205.12729, May 2022.
- [30] Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009.
- [31] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. *arXiv e-prints*, page arXiv:1612.01474, Dec. 2016.
- [32] Yingzhen Li and Yarin Gal. Dropout Inference in Bayesian Neural Networks with Alpha-divergences. *arXiv e-prints*, page arXiv:1703.02914, Mar. 2017.
- [33] H. Linander, O. Balabanov, H. Yang, and B. Mehlig. Looking at the posterior: on the origin of uncertainty in neural-network classification, 2022.
- [34] Christos Louizos and Max Welling. Multiplicative normalizing flows for variational Bayesian neural networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2218–2227. PMLR, 06–11 Aug 2017.
- [35] David MacKay. Bayesian model comparison and backprop nets. In J. Moody, S. Hanson, and R.P. Lippmann, editors, *Advances in Neural Information Processing Systems*, volume 4. Morgan-Kaufmann, 1991.
- [36] D. MacKay. *Bayesian Methods for Adaptive Models*. PhD thesis, California Institute of Technology, 1992.
- [37] D. J. C. Mackay. Information-based objective functions for active data selection. *Neural Computation*, 4(2):550–604, 1992.
- [38] Wesley J Maddox, Pavel Izmailov, Timur Garipov, Dmitry P Vetrov, and Andrew Gordon Wilson. A simple baseline for bayesian uncertainty in deep learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [39] Hendrik Alexander Mehrkens, Camila González, and Anirban Mukhopadhyay. Improving robustness and calibration in ensembles with diversity regularization. *arXiv preprint arXiv:2201.10908*, 2022.
- [40] Aryan Mobiny, Hien V. Nguyen, Supratik Moulik, Naveen Garg, and Carol C. Wu. DropConnect Is Effective in Modeling Uncertainty of Bayesian Deep Networks. *arXiv e-prints*, page arXiv:1906.04569, June 2019.
- [41] Giung Nam, Jongmin Yoon, Yoonho Lee, and Juho Lee. Diversity Matters When Learning From Ensembles. *arXiv e-prints*, page arXiv:2110.14149, Oct. 2021.
- [42] Radford M. Neal. *Bayesian Learning for Neural Networks*. Springer-Verlag, Berlin, Heidelberg, 1996.
- [43] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, D. Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [44] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Brij B. Gupta, Xiaojiang Chen, and Xin Wang. A survey of deep active learning. *ACM Comput. Surv.*, 54(9), oct 2021.
- [45] Claude Elwood Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423, 1948.
- [46] Jost Tobias Springenberg, Aaron Klein, Stefan Falkner, and Frank Hutter. Bayesian optimization with robust bayesian neural networks. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [47] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014.
- [48] Supplementary material.
- [49] Tishby, Levin, and Solla. Consistent inference of probabilities in layered networks: predictions and generalizations. In *International 1989 Joint Conference on Neural Networks*, pages 403–409 vol.2, 1989.
- [50] Aleksei Tiulpin and Matthew B Blaschko. Greedy bayesian posterior approximation with deep ensembles. *arXiv preprint arXiv:2105.14275*, 2021.
- [51] Li Wan, Matthew Zeiler, Sixin Zhang, Yann Le Cun, and Rob Fergus. Regularization of neural networks using drop-connect. In *International conference on machine learning*, pages 1058–1066. PMLR, 2013.
- [52] Max Welling and Yee Whye Teh. Bayesian learning via stochastic gradient langevin dynamics. In *ICML*, 2011.
- [53] Yeming Wen, Dustin Tran, and Jimmy Ba. BatchEnsemble: An Alternative Approach to Efficient Ensemble and Lifelong Learning. *arXiv e-prints*, page arXiv:2002.06715, Feb. 2020.
- [54] Florian Wenzel, Jasper Snoek, Dustin Tran, and Rodolphe Jenatton. Hyperparameter ensembles for robustness and uncertainty quantification, 2020.
- [55] Andrew G Wilson and Pavel Izmailov. Bayesian deep learning and a probabilistic perspective of generalization. *Advances in neural information processing systems*, 33:4697–4708, 2020.
- [56] Anqi Wu, Sebastian Nowozin, Edward Meeds, Richard E Turner, Jose Miguel Hernandez-Lobato, and Alexander L Gaunt. Deterministic variational inference for robust bayesian neural networks. *arXiv preprint arXiv:1810.03958*, 2018.
- [57] Anqi Wu, Sebastian Nowozin, Edward Meeds, Richard E. Turner, José Miguel Hernández-Lobato, and Alexander L. Gaunt. Deterministic Variational Inference for Robust Bayesian Neural Networks. *arXiv e-prints*, page arXiv:1810.03958, Oct. 2018.



- [58] Sheheryar Zaidi, Arber Zela, Thomas Elsken, Chris C Holmes, Frank Hutter, and Yee Teh. Neural ensemble search for uncertainty estimation and dataset shift. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 7898–7911. Curran Associates, Inc., 2021.