
Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles

Balaji Lakshminarayanan, Alexander Pritzel & Charles Blundell
DeepMind, London
{balajiln, apritzel, cblundell}@google.com

Abstract

Deep neural networks are powerful black box predictors that have recently achieved impressive performance on a wide spectrum of tasks. Quantifying predictive uncertainty in neural networks is a challenging and yet unsolved problem. Bayesian neural networks, which learn a distribution over weights, are currently the state-of-the-art for estimating predictive uncertainty; however these require significant modifications to the training procedure and are computationally expensive compared to standard (non-Bayesian) neural networks. We propose an alternative to Bayesian neural networks, that is simple to implement, readily parallelisable and yields high quality predictive uncertainty estimates. Through a series of experiments on classification and regression benchmarks, we demonstrate that our method produces well-calibrated uncertainty estimates which are as good or better than approximate Bayesian neural networks. Finally, we evaluate the predictive uncertainty on test examples from known and unknown classes, and show that our method is able to express higher degree of uncertainty on unknown classes, unlike existing methods which make overconfident predictions even on unknown classes.

1 Introduction

Deep learning methods such as convolutional neural networks and recurrent neural networks have achieved state-of-the-art performance on a wide variety of machine learning tasks and are becoming increasingly popular in domains such as computer vision [19], speech recognition [16], natural language processing [25] and bioinformatics [20]. Despite impressive classification accuracies and mean squared errors in supervised learning problems, neural networks are poor at quantifying predictive uncertainty, and tend to be overconfident in their predictions. Evaluating the quality of predictive uncertainties is challenging as the true conditional probabilities of the data are usually not available. In this work, we shall focus upon two measures of the quality of predictive uncertainty from a neural network. Firstly, we shall examine *calibration* [7, 8]. Formally, calibration is the discrepancy between subjective forecasts and (empirical) long-run frequencies. This is a frequentist notion of uncertainty: if a network claims that 90% of the time a particular label is the correct label, then, during evaluation, 90% of all labels ascribed probability 90% of being correct, should be the correct label. Calibration can be incentivised by *proper scoring rules* [12] such as log predictive probabilities and the Brier score [6]. Interestingly, these two proper scoring rules are commonly used in deep learning, but without reference to their properties for incentivising calibration. Note that calibration is an orthogonal concern to accuracy: a network’s predictions may be accurate and yet miscalibrated. The second notion of quality of predictive uncertainty we consider concerns generalisation of the predictive uncertainty to *unknown unknowns* or model misspecification. For example, if a new class is introduced at test time, then examples with this new class should have high predictive uncertainty. Or, if a network trained on one data set is evaluated on a completely different data set, then the network should have high predictive uncertainty.

Well-calibrated predictions that are robust to model misspecification and domain shift have a number of important practical uses. Calibrated predictions lie at the heart of many forecasting problems

about the physical world (e.g., weather, earthquakes, medical diagnosis, etc) and can be objectively understood and conveyed as a frequentist notion of uncertainty (often being captured, in part, by notions such as mean-time before failure and medical prognosis). Calibrated predictions also permit modularity; if all components of a system are well-calibrated, then the uncertainty of various predictions can easily be integrated as a common currency of uncertainty between different modules. Robustness to misspecification is a common requirement in many real world applications as training data is incomplete or lags behind the most recent data upon which predictions are to be made. Being robust to mild misspecification or domain shift enables a neural network to continue to be useful (or at least known it is not useful) without re-training.

There has been a lot of recent interest in adapting neural networks to encompass uncertainty and probabilistic methods; the majority of this work revolves around a Bayesian formalism [1], whereby a prior distribution is specified upon the weights of a single neural network and then an approximation scheme is derived that infers the posterior distribution upon the weights of the neural network after having incorporated the training data. Approximate Bayesian approaches learn a posterior distribution over the parameters of the neural network and use this approximate posterior distribution to estimate predictive uncertainty. Early work on Bayesian neural networks focused on Markov chain Monte Carlo (MCMC) methods and Laplace approximation, cf. the seminal work of MacKay [23] and Neal [29]. MCMC-based Bayesian neural networks were amongst the best-performing methods in the ‘evaluating predictive uncertainty challenge’ (see [31] for details). While MCMC can be used for small neural networks, it is computationally expensive for large deep neural networks. Hence, recent work on uncertainty in neural networks has focussed mostly on relatively faster approximate Bayesian solutions such as variational Bayesian methods [3, 14, 22], assumed density filtering [15], expectation propagation [21, 35] or stochastic gradient Langevin diffusion methods [18, 36]. These approximations are not guaranteed to provide uncertainty estimates that reflect underlying beliefs (except, possibly, in the limit of infinite data). Indeed, variational Bayes methods typically underestimate posterior uncertainty, whilst expectation propagation methods typically overestimate posterior uncertainty. The quality of predictive distributions obtained using these approaches depends on (1) the degree of approximation (due to computational constraints) and (2) if the prior distribution is ‘correct’ (e.g. the model could be mis-specified); for instance, Rasmussen and Quinero-Candela [32] discuss an example where priors of convenience lead to unreasonable predictive uncertainties. Even the exact Bayesian posterior may not necessarily be robust to misspecification with respect to unknown classes or dataset shift. The enormous size of the parameter space of modern neural networks compounds both of these issues. Bayesian neural networks are often computationally slower to train and harder to implement compared to their non-Bayesian counterparts, which raises the need for a ‘general purpose solution’ that can deliver calibrated uncertainty estimates and yet requires only minor modifications to the standard training pipeline.

Recently, Gal and Ghahramani [10] proposed using *Monte Carlo dropout* (MC-dropout) to estimate predictive uncertainty by using *Dropout* [33] at test time. There has been work on approximate Bayesian interpretation [10, 17, 24] of dropout. However, dropout may also be interpreted as ensemble model combination [33] (as opposed to Bayesian model averaging), where the predictions are averaged over multiple networks. The latter approximation seems more plausible particularly in the scenario where the dropout rates are not tuned based on the training data (since any sensible approximation to the true Bayesian posterior distribution has to depend on the training data). It has long been observed that ensembles of models improve overall performance (see [9] for a review). Bayesian model averaging attempts to find the single best model (or parameters) in a soft manner, assuming the true model lies within the hypothesis class of the prior. In contrast, ensembles perform *model combination*, i.e. they combine the models to obtain a more powerful model; ensembles can be expected to be better when the true model does not lie within the hypothesis class (see [26] for a related discussion). Hence, ensembles potentially provide a complementary source of methods for estimating predictive uncertainty.

Our contribution in this paper is two fold. First, we describe a simple, scalable method for estimating predictive uncertainty estimates from neural networks. We demonstrate that two simple modifications to the training pipeline, namely (i) *ensembles* and (ii) *adversarial training* [13], are sufficient to obtain well-calibrated uncertainty estimates for supervised learning. Secondly, we propose a series of tasks for evaluating the quality of the predictive uncertainty estimates, in terms of calibration and generalisation to unknowns in supervised learning problems. Ensembles of neural networks have long been successfully used to boost predictive performance (e.g. classification accuracy in Imagenet)

and adversarial training has been used to improve robustness to adversarial examples; however, to the best of our knowledge, ours is the first work to investigate their usefulness for predictive uncertainty estimation and compare their performance to current state-of-the-art approximate Bayesian methods on a series of classification and regression benchmark datasets. Compared to Bayesian neural networks (e.g. variational inference or Monte Carlo methods), our method is simpler to implement, well suited for distributed computation, and requires surprisingly few modifications to standard neural networks, thereby making it attractive for large-scale applications.

2 Deep Ensembles: A Recipe For Uncertainty Estimation

2.1 Problem setup and High-level summary

We assume that the training dataset \mathcal{D} consists of N i.i.d. data points $\mathcal{D} = \{\mathbf{x}_n, y_n\}_{n=1}^N$, where $\mathbf{x} \in \mathbb{R}^D$ represents the D -dimensional features. For classification problems, the label is assumed to be one of K classes, that is $y \in \{1, \dots, K\}$. For regression problems, the label is assumed to be real-valued, that is $y \in \mathbb{R}$. Given a test data point \mathbf{x} , the goal is to output a predictive distribution $p_\theta(y|\mathbf{x})$ where θ are the parameters of the neural network.

We suggest a simple recipe: (1) use a proper scoring rule as the training criterion, (2) use *adversarial training* [13] to smooth the predictive distributions, and (3) train an *ensemble*. Let M denote the number of neural networks in the ensemble and $\{\theta_m\}_{m=1}^M$ denote the parameters of the m th net. We first describe how to train a single neural net and then explain how to train an ensemble of networks.

2.2 Proper scoring rules

Scoring rules measure the quality of predictive uncertainty (see [12] for a review). A scoring rule assigns a numerical score to a predictive distribution $p_\theta(y|\mathbf{x})$, rewarding better calibrated predictions over worse. We shall consider scoring rules where a higher numerical score is better. Let a scoring rule be a function $S(p_\theta, (y, \mathbf{x}))$ that evaluates the quality of the predictive distribution $p_\theta(y|\mathbf{x})$ relative to an event $(y, \mathbf{x}) \sim q(y|\mathbf{x})$ where $q(\cdot)$ denotes the true distribution on (y, \mathbf{x}) -tuples. The expected scoring rule is then:

$$S(p_\theta, q) = \int q(y, \mathbf{x}) S(p_\theta, (y, \mathbf{x})) dy d\mathbf{x} \quad (1)$$

A *proper scoring rule* is one where $S(p_\theta, q) \geq S(q, q)$ with equality if and only if $p_\theta(y|\mathbf{x}) = q(y|\mathbf{x})$, for all p_θ and q . A neural network can then be trained according to measure that encourages calibration of predictive uncertainty by minimising the loss $\mathcal{L}(\theta) = -S(p_\theta, q)$.

It turns out many common neural network loss functions are proper scoring rules. For example, when maximising likelihood, the score function is $S(p_\theta, (y, \mathbf{x})) = \log p_\theta(y|\mathbf{x})$, and this is a proper score function due to Gibbs inequality: $S(p_\theta, q) = \mathbb{E}_{q(\mathbf{x})} \text{KL}[q(y|\mathbf{x})||p_\theta(y|\mathbf{x})]$. Interestingly, in the case of K -way classification, $S(p_\theta, (y, \mathbf{x})) = -\sum_{k=1}^K (\delta_{k=y} - p_\theta(y = k|\mathbf{x}))^2$ (i.e., minimising the squared error between the predictive probability of a label and one-hot encoding of the correct label) is also a proper scoring rule, known as the Brier score [6]. This provides justification for this common trick for training neural networks by minimizing the square error between a binary label and its associated probability and shows it is, in fact, a well defined loss with desirable properties.¹

2.3 Adversarial training

Adversarial examples, proposed by Szegedy et al. [34] and extended by Goodfellow et al. [13], are those which are ‘close’ to the original training examples (e.g. an image that is visually indiscernible from the original image), but are misclassified by the neural network. Goodfellow et al. [13] proposed the *fast gradient sign method* as a fast solution to generate adversarial examples. Given an input \mathbf{x} with target y , and loss $\ell(\theta, \mathbf{x}, y)$ (e.g. $-\log p_\theta(y|\mathbf{x})$), the fast gradient sign method generates an adversarial example as

$$\mathbf{x}' = \mathbf{x} + \epsilon \text{sign}\left(\nabla_{\mathbf{x}} \ell(\theta, \mathbf{x}, y)\right) \quad (2)$$

¹Indeed as noted in Gneiting and Raftery [12], it can be shown that asymptotically maximising a proper scoring rule recovers true parameter values.

where ϵ is a small value such that the max-norm of the perturbation is bounded. Intuitively, the adversarial perturbation creates a new training example by adding a perturbation along a direction which the network is likely to increase the loss. Assuming ϵ is small enough, these adversarial examples can be used to augment the original training set by treating (\mathbf{x}', y) as additional training examples. This procedure, referred to as *adversarial training*,² was found to improve the classifier’s robustness.

We add a new perspective to this procedure. Consider the following scoring rule:

$$S(p_\theta, (y, \mathbf{x})) = \alpha \log p_\theta(y|\mathbf{x}) + (1 - \alpha) \log p_\theta(y|\mathbf{x} + \Delta\mathbf{x}) \quad (3)$$

$$\text{where } \Delta\mathbf{x} = \epsilon \text{sign}(-\nabla_{\mathbf{x}} \log p_\theta(y|\mathbf{x})) \quad (4)$$

α trades-off between the usual maximum likelihood score and an adversarial score. If $\alpha = 0$, then (3) is a proper scoring rule since the log likelihood is a proper scoring rule. In our results, we show for $\alpha > 0$, (3) yields good calibration. Intuitively, our goal is to smooth the predictive distributions by increasing the likelihood of the target around an ϵ -neighborhood of the observed training examples. Ideally one would want to smooth the predictive distributions along all 2^D directions in $\{1, -1\}^D$; however this is computationally expensive. A random direction might not necessarily increase the loss; however, adversarial training by definition computes the direction where the loss is high and hence is better for smoothing predictive distributions. Miyato et al. [27] proposed a related idea called *virtual adversarial training* (VAT), where they picked $\Delta\mathbf{x} = \arg \max_{\Delta\mathbf{x}} KL(p(y|\mathbf{x})||p(y|\mathbf{x} + \Delta\mathbf{x}))$; the advantage of VAT is that it does not require knowledge of the true target y and hence can be applied to semi-supervised learning. Miyato et al. [27] showed that distributional smoothing using VAT is beneficial for efficient semi-supervised learning; in contrast, we show that adversarial training is helpful for better predictive uncertainty estimation. Hence, our contributions are complementary; one could use VAT for improving predictive uncertainty in the semi-supervised setting as well.

2.3.1 Training criterion for heteroscedastic regression

For regression problems, neural networks usually output a single value say $\mu(\mathbf{x})$ and the parameters are optimised to minimize the mean squared error (MSE) on the training set, given by $\sum_{n=1}^N (y_n - \mu(\mathbf{x}_n))^2$. However, the MSE does not capture predictive uncertainty. Following [30], we use a network that outputs two values in the final layer, corresponding to the predicted mean $\mu(\mathbf{x})$ and variance³ $\sigma^2(\mathbf{x}) > 0$. By treating the observed value as a sample from a Gaussian distribution with the predicted mean and variance, we minimize the negative log-likelihood criterion:

$$\log p_\theta(y_n|\mathbf{x}_n) = -\frac{1}{2} \log \sigma_\theta^2(\mathbf{x}) - \frac{1}{2\sigma_\theta^2(\mathbf{x})} (y - \mu_\theta(\mathbf{x}))^2 + \text{constant} \quad (5)$$

We found the above to perform satisfactorily in our experiments. However, two simple extensions are worth further investigation: (1) Maximum likelihood estimation over $\mu_\theta(\mathbf{x})$ and $\sigma_\theta^2(\mathbf{x})$ might overfit; one could impose a prior and perform maximum-a-posteriori (MAP) estimation. (2) In cases where the Gaussian is too-restrictive, one could use a complex distribution e.g. mixture density network [2].

It is tempting to use an ensemble of neural networks to obtain multiple predictions and use the empirical variance of the networks’ predictions as an approximate measure of uncertainty. However, this generally does not lead to well-calibrated predictive probabilities. As a motivating example, we report calibration curves (also known as reliability diagrams) on the *Year Prediction MSD* dataset. First, we compute the $z\%$ (e.g. 90%) prediction interval for each test data point based on Gaussian quantiles using predictive mean and variance. Next, we measure what fraction of test observations fall within this prediction interval. For a well-calibrated regressor, the observed fraction should be close to $z\%$. We compute observed fraction for $z = 10\%$ to $z = 90\%$ in increments of 10 and report results in Figure 1. A well-calibrated regressor should lie very close to the diagonal; on the left subplot we observe that the proposed method, which learns the predictive variance, leads to a well-calibrated regressor. However, on the right subplot, we observe that the empirical variance obtained from neural networks which do not learn the predictive variance (specifically, five neural networks trained to minimize MSE) consistently underestimates the true uncertainty. For instance,

²Not to be confused with Generative Adversarial Networks

³We enforce the positivity constraint on the variance by passing the second output through the *softplus* function ($\log(1 + \exp(\cdot))$), and additionally add a minimum variance (e.g. 10^{-6}) for numerical stability.

the 80% prediction interval contains only 20% of the test observations, which means the empirical variance underestimates the true predictive uncertainty.




Figure 1: Calibration results on the Year Prediction MSD dataset: x -axis denotes the expected fraction and y -axis denotes the observed fraction; ideal output is the dashed blue line. Predicted variance (left) is significantly better calibrated than the empirical variance (right). See main text for further details.

2.4 Ensembles: training and prediction

The most popular ensembles use decision trees as the base learners and a wide variety of methods have been explored in the literature on ensembles. Broadly, there are two classes of ensembles: *randomisation*-based approaches such as random forests [5], where the ensemble members can be trained in parallel without any interaction, and *boosting*-based approaches where the ensemble members are fit sequentially. We focus only on the randomisation based approach as it is better suited for distributed computation. Breiman [5] showed that the generalisation error of random forests can be upper bounded by a function of the strength and correlation between individual trees; hence it is desirable to use a *randomisation scheme* that de-correlates the predictions of the individual models as well as ensures that the individual models are strong (e.g. high accuracy). One of the popular strategies is *bagging* or bootstrapping, where ensemble members are trained on different bootstrap samples of the original training set. If the base learner lacks intrinsic randomisation (e.g. it can be trained efficiently by solving a convex optimisation problem), the bootstrap is a good mechanism for inducing diversity. However, if the underlying base learner is unstable (e.g. multiple local optima), the bootstrap is not required and can sometimes hurt performance since a base learner trained on a bootstrap sample sees only 63% unique data points. In the literature on tree ensembles, Breiman [5] proposed to use a combination of *bagging* (a.k.a. bootstrapping) [4] and random subset selection of features at each node. Geurts et al. [11] later showed that the bootstrap is unnecessary if additional randomness can be injected into the random subset selection procedure. Using more data for training the base learners helps reduce their bias and ensembling helps reduce the variance.

We used the entire training dataset to train each net since deep neural networks typically perform better with more data, although it is straightforward to use a random subsample if need be. We found that random initialisation of neural net parameters, along with random shuffling of the data points was sufficient to obtain good performance. The overall training procedure is summarised in Algorithm 1.

Algorithm 1 Pseudocode of the training procedure for our method

- 1: Initialise $\theta_1, \theta_2, \dots, \theta_M$ randomly
 - 2: **for** $m = 1 : M$ **do** \triangleright train networks independently in parallel
 - 3: Sample data point n_m randomly for each net \triangleright single n_m for clarity, minibatch in practice
 - 4: Generate adversarial example using $\mathbf{x}'_{n_m} = \mathbf{x}_{n_m} + \epsilon \text{sign}(\nabla_{\mathbf{x}_{n_m}} \ell(\theta_m, \mathbf{x}_{n_m}, y_{n_m}))$
 - 5: Minimize $\ell(\theta_m, \mathbf{x}_{n_m}, y_{n_m}) + \ell(\theta_m, \mathbf{x}'_{n_m}, y_{n_m})$ w.r.t. θ_m \triangleright adversarial training
-

We treat the ensemble as a uniformly-weighted mixture model and combine the predictions as

$$p(y|\mathbf{x}) = \frac{1}{M} \sum_{m=1}^M p_{\theta_m}(y|\mathbf{x}, \theta_m). \quad (6)$$

For classification, this corresponds to averaging the predicted probabilities. For regression, the prediction is a mixture of Gaussian distributions. For ease of computing quantiles and log predictive probabilities, we further approximate the ensemble prediction as a Gaussian whose mean and variance are respectively the mean and variance of the mixture.⁴

3 Experimental results

3.1 Evaluation metrics and experimental setup

For both classification and regression, we evaluate the negative log likelihood (NLL) which depends on the predictive uncertainty. NLL is a proper scoring rule and a popular metric for evaluating predictive uncertainty [31]. For classification we additionally measure classification accuracy and the Brier score (another proper scoring rule), defined as

$$BS = \frac{1}{K} \sum_{k=1}^K \left(t_k^* - p(y = k | \mathbf{x}^*) \right)^2 \quad \text{where} \quad t_k^* = 1 \text{ if } k = y^*, \text{ and } 0 \text{ otherwise.}$$

Intuitively, our Brier score definition corresponds to creating K binary classification tasks by using a *one-hot* encoding of the true target, and measuring the average MSE between the predicted probability and the true target on the K tasks. For regression problems, we additionally measured the root mean squared error (RMSE). Unless otherwise specified, we used batch size of 100 and ADAM optimiser with fixed learning rate of 0.1 in our experiments. We use the same technique for generating adversarial training examples for regression problems. Goodfellow et al. [13] used a fixed ϵ for all dimensions; this is unsatisfying if the input dimensions have different ranges. Hence, in all of our experiments, we set ϵ to 0.01 times the range of the training data along that particular dimension.

3.2 Regression on toy datasets

First, we qualitatively evaluate the performance of the proposed method on a one-dimensional toy regression dataset. This dataset was used by Hernández-Lobato and Adams [15], and consists of 20 training examples drawn as $y = x^3 + \epsilon$ where $\epsilon \sim \mathcal{N}(0, 3^2)$. The results are shown in Figure 2. The results clearly demonstrate that (i) learning variance leads to improved predictive uncertainty and (ii) ensemble combination improves performance, especially as we move farther from the observed training data.




Figure 2: Results on a toy regression task: x -axis denotes x . On the y -axis, the blue line is the *ground truth* curve, the red dots are observed noisy training data points and the gray lines correspond to the predicted mean along with three standard deviations. Left plot corresponds to empirical variance of 5 networks, middle and right plot show the effect of learning variance using a single net and 5 networks respectively.

In Figure 3, we consider a modified version of the toy regression task; specifically we have added a sinusoidal component to the original curve close to the origin. In absence of any other information about the underlying true function, this sinusoidal component could be either treated as noise or signal. Optimizing for MSE leads to the individual networks predicting the mean close to the sinusoid (as expected). Learning the variance leads to a qualitatively different solution where the network predicts higher uncertainty around the sinusoidal perturbation. In more complicated datasets, the preferred

⁴The mean and variance of a mixture $\frac{1}{M} \sum \mathcal{N}(\mu_{\theta_m}(\mathbf{x}), \sigma_{\theta_m}^2(\mathbf{x}))$ are given by $\mu_*(\mathbf{x}) = \frac{1}{M} \sum_m \mu_{\theta_m}(\mathbf{x})$ and $\sigma_*^2(\mathbf{x}) = \frac{1}{M} \sum_m (\sigma_{\theta_m}^2(\mathbf{x}) + \mu_{\theta_m}^2(\mathbf{x})) - \mu_*^2(\mathbf{x})$ respectively.

solution would of course depend on the number of data points and the network architecture. However this example shows that training using a scoring rule instead of MSE can lead to qualitatively different predictions.




Figure 3: Results on modified toy regression task: see Figure 2 for description of x and y axes. Left plot shows empirical variance of 5 networks optimised using MSE, right plot shows variance of an ensemble of 5 networks optimised using negative log likelihood.

3.3 Regression on real world datasets

In our next experiment, we compare our method to state-of-the-art methods for predictive uncertainty estimation using neural networks on regression tasks. We use the experimental setup proposed by Hernández-Lobato and Adams [15] for evaluating PBP, which was also used by Gal and Ghahramani [10] to evaluate MC-dropout.⁵ Each dataset is split into 20 train-test folds, except for the protein dataset which uses 5 folds and the Year Prediction MSD dataset which uses a single train-test split. We use the same network architecture: 1-hidden layer neural network with ReLU nonlinearity [28], containing 50 hidden units for smaller datasets and 100 hidden units for the larger protein and Year Prediction MSD datasets. We trained for 40 epochs; we refer to [15] for further details about the datasets and the experimental protocol. We used 5 networks in our ensemble. Our results are shown in Table 1, along with the PBP and MC-dropout results reported in their respective papers.

Datasets	RMSE			NLL		
	PBP	MCDropout	our method	PBP	MCDropout	our method
Boston housing	3.01 ± 0.18	2.97 ± 0.85	3.92 ± 1.01	2.57 ± 0.09	2.46 ± 0.25	2.39 ± 0.25
Concrete	5.67 ± 0.09	5.23 ± 0.53	6.03 ± 0.47	3.16 ± 0.02	3.04 ± 0.09	3.05 ± 0.15
Energy	1.80 ± 0.05	1.66 ± 0.19	2.86 ± 0.29	2.04 ± 0.02	1.99 ± 0.09	1.40 ± 0.17
Kin8nm	0.10 ± 0.00	0.10 ± 0.00	0.09 ± 0.00	-0.90 ± 0.01	-0.95 ± 0.03	-1.19 ± 0.02
Naval propulsion plant	0.01 ± 0.00	0.01 ± 0.00	0.00 ± 0.00	-3.73 ± 0.01	-3.80 ± 0.05	-4.17 ± 0.05
Power plant	4.12 ± 0.03	4.02 ± 0.18	4.09 ± 0.16	2.84 ± 0.01	2.80 ± 0.05	2.78 ± 0.04
Protein	4.73 ± 0.01	4.36 ± 0.04	4.71 ± 0.06	2.97 ± 0.00	2.89 ± 0.01	2.84 ± 0.02
Wine	0.64 ± 0.01	0.62 ± 0.04	0.64 ± 0.04	0.97 ± 0.01	0.93 ± 0.06	0.93 ± 0.12
Yacht	1.02 ± 0.05	1.11 ± 0.38	1.58 ± 0.46	1.63 ± 0.02	1.55 ± 0.12	1.18 ± 0.19
Year Prediction MSD	8.88 ± NA	8.85 ± NA	8.94 ± NA	3.60 ± NA	3.59 ± NA	3.37 ± NA

Table 1: Results on regression benchmark datasets comparing RMSE and NLL.

We observe that our method outperforms (or is competitive with) existing methods in terms of NLL. On some datasets, we observe that our method is slightly worse in terms of RMSE. We believe that this might be caused due to the heteroscedastic regression training criterion, which optimises for negative log likelihood instead of mean squared error, as discussed in the toy example in Figure 3.

3.4 Classification on MNIST and SVHN

Next we evaluate the performance on classification tasks using MNIST and SVHN datasets. Our goal is not to achieve the state-of-the-art performance on these problems, but rather to evaluate the effect of adversarial training as well as the number of networks in the ensemble. To verify if adversarial training helps, we also include a baseline which picks a random signed vector. For MNIST, we used an MLP with 3-hidden layers with 200 hidden units per layer and ReLU non-linearities with batch normalisation. For MC-dropout, we added dropout after each non-linearity with 0.1 as the

⁵We do not compare to VI [14] as PBP and MC-dropout outperform VI on these benchmarks.

dropout rate.⁶ Results are shown in Figure 4. We observe that adversarial training and increasing the number of networks in the ensemble significantly improve performance in terms of both classification accuracy as well as NLL and Brier score, illustrating that our method produces well-calibrated uncertainty estimates. Adversarial training leads to better performance than augmenting with random direction. Our method also performs much better than MC-dropout in terms of all the performance measures. Note that augmenting the training dataset with invariances (such as random crop and horizontal flips) is complementary to adversarial training and can potentially improve performance.




Figure 4: Results on MNIST dataset using 3-layer MLP: Both ensembles and adversarial training (AT) significantly improve performance in terms of all 3 metrics. Our method outperforms MC-dropout with the corresponding number of samples.

To measure the sensitivity of the results to the choice of network architecture, we experimented with a two-layer MLP as well as a convolutional neural network; we observed qualitatively similar results; see Section A in the supplementary material for details.

We also report results on the SVHN dataset using an VGG-style convolutional neural network.⁷ The results are in Figure 5. Ensembles outperform MC dropout. However, adversarial training does not seem to significantly for this dataset. If the classes are well-separated, adversarial training might not change the classification boundary significantly. It is not clear if this is the case here, further investigation is required.




Figure 5: Results on SVHN dataset: Ensembles outperform MC-dropout, however adversarial training does not seem to help much on this dataset.

⁶We also tried dropout rate of 0.5, but that produced worse results.

⁷The architecture is described in <http://torch.ch/blog/2015/07/30/cifar.html>

Another advantage of using an ensemble is that it enables us to easily identify training examples where the individual networks disagree or agree the most. This disagreement⁸ provides another useful qualitative way to evaluate predictive uncertainty. Figure 6 reports results on MNIST dataset.




Figure 6: Results on MNIST showing test examples with high or low disagreement between the networks in the ensemble: Top two rows denote the test examples with least disagreement and the bottom two rows denote the test examples with the most disagreement.

3.5 Uncertainty evaluation: test examples from known vs unknown classes

In the final experiment, we evaluate uncertainty on unseen classes. Overconfident predictions on unseen classes pose a challenge for reliable deployment of deep learning models in the wild. We would like the predictions to exhibit higher uncertainty when the test data is very different from the training data. To test if the proposed method possesses this desirable property, we train a MLP on the standard MNIST train/test split using the same architecture as before. However, in addition to the regular test set with known classes, we also evaluate it on a test set containing unknown classes. We used the test split of the NotMNIST⁹ dataset. The images in this dataset have the same size as MNIST, however the labels are alphabets instead of digits. We don't have access to the true conditional probabilities, but we expect the predictions to be closer to uniform on unseen classes compared to the known classes where the predictive probabilities should concentrate on the true targets. We evaluate the entropy of the predictive distribution and use this to evaluate the quality of the uncertainty estimates. The results are shown in Figure 7. For known classes (top row), both our method and MC-dropout have low entropy as expected. For unknown classes (bottom row), as M increases, the entropy of deep ensembles increases much faster than MC-dropout indicating that our method is better suited for handling unseen test examples. In particular, MC-dropout seems to give high confident predictions for some of the test examples, as evidenced by the mode around 0 even for unseen classes. Comparing different variants of our method, the mode for adversarial training increases faster than the mode for vanilla ensembles indicating that adversarial training is beneficial for quantifying uncertainty on unseen classes. As an additional qualitative measure, in Figure 8, we report the examples with lowest disagreement in top two rows and highest disagreement in bottom two rows respectively. From the top two rows, we see that the ensemble agreement is highest for letter 'I' which resembles 1 in the MNIST training dataset. From the bottom two rows, we see that the ensemble disagreement is higher for examples visually different from the MNIST training dataset.

4 Discussion

We have proposed a simple and scalable solution that provides a very strong baseline on evaluation metrics for uncertainty quantification. Our method uses scoring rules as training objectives to encourage the neural network to produce better calibrated predictions and uses a combination of ensembles and adversarial training for robustness to model misspecification and dataset shift. Our method is well suited for large scale distributed computation and can be readily implemented for a wide variety of architectures such as MLPs, CNNs, etc including those which do not use dropout (e.g. residual networks). It is perhaps surprising to the Bayesian deep learning community that a non-Bayesian (yet probabilistic) approach can perform as well as Bayesian neural networks. We hope that this work will encourage community to think about hybrid approaches (e.g. using non-Bayesian approaches such as ensembles) and other interesting metrics for evaluating predictive uncertainty.

⁸More precisely, we define disagreement as $\sum_{m=1}^M KL(p_{\theta_m}(y|\mathbf{x})||p_E(y|\mathbf{x}))$ where KL denotes the Kullback-Leibler divergence and $p_E(y|\mathbf{x}) = M^{-1} \sum_m p_{\theta_m}(y|\mathbf{x})$ is the prediction of the ensemble.

⁹Available at <http://yaroslavvb.blogspot.co.uk/2011/09/notmnist-dataset.html>

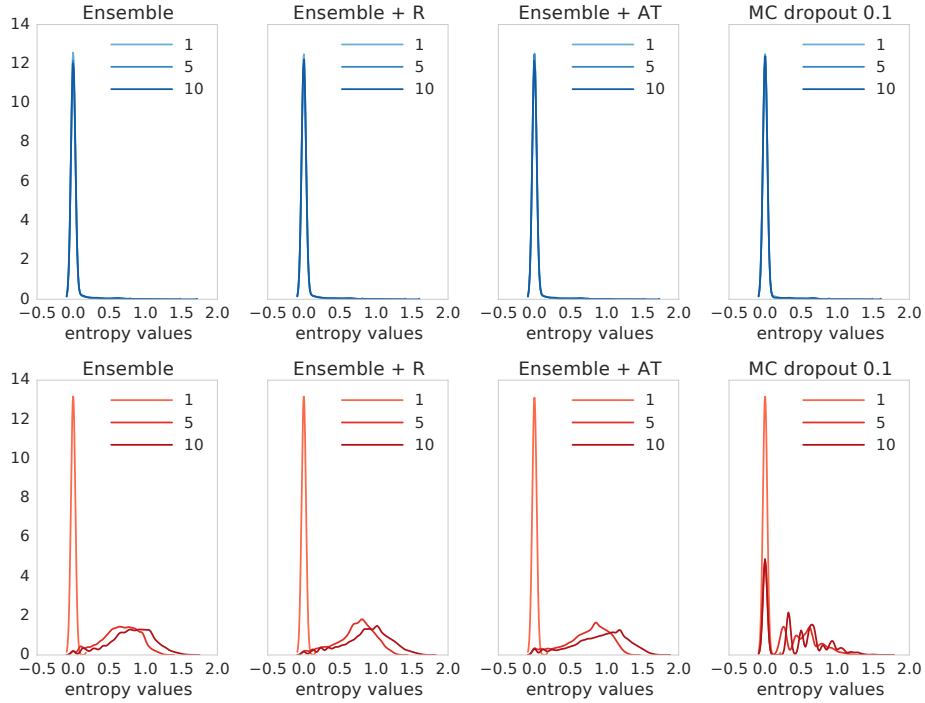


Figure 7: Histogram of the predictive entropy on test examples from known classes (top row) and unknown classes (bottom row), as we vary the number of networks in the ensemble.

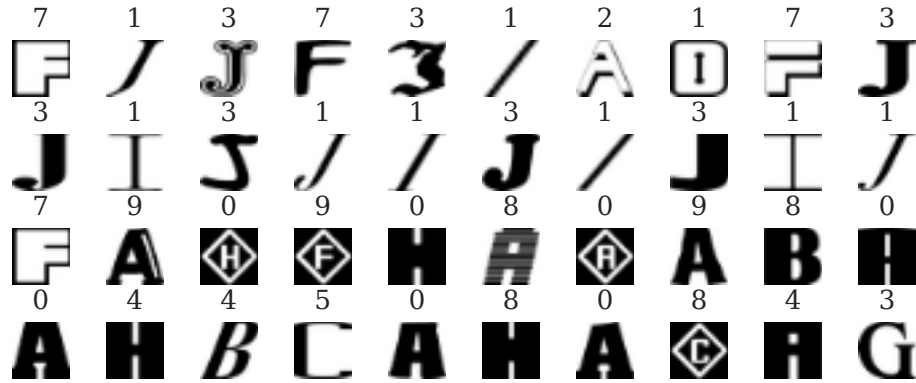


Figure 8: Network trained on MNIST and tested on the NotMNIST dataset containing unseen classes: Top two rows denote the test examples with least disagreement and the bottom two rows denote the test examples with the most disagreement.

Acknowledgments

We would like to thank Daan Wierstra, David Silver, Martin Szummer, Shakir Mohamed, Theophane Weber and Ulrich Paquet for helpful feedback.

References

- [1] J. M. Bernardo and A. F. Smith. *Bayesian Theory*, volume 405. John Wiley & Sons, 2009.
- [2] C. M. Bishop. Mixture density networks. 1994.
- [3] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra. Weight uncertainty in neural networks. In *ICML*, 2015.
- [4] L. Breiman. Bagging predictors. *Mach. Learn.*, 24(2):123–140, 1996.
- [5] L. Breiman. Random forests. *Mach. Learn.*, 45(1):5–32, 2001.
- [6] G. W. Brier. Verification of forecasts expressed in terms of probability. *Monthly weather review*, 1950.
- [7] A. P. Dawid. The well-calibrated Bayesian. *Journal of the American Statistical Association*, 1982.
- [8] M. H. DeGroot and S. E. Fienberg. The comparison and evaluation of forecasters. *The statistician*, 1983.
- [9] T. G. Dietterich. Ensemble methods in machine learning. In *Multiple classifier systems*. 2000.
- [10] Y. Gal and Z. Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, 2016.
- [11] P. Geurts, D. Ernst, and L. Wehenkel. Extremely randomized trees. *Mach. Learn.*, 63(1):3–42, 2006.
- [12] T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- [13] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015.
- [14] A. Graves. Practical variational inference for neural networks. In *NIPS*, 2011.
- [15] J. M. Hernández-Lobato and R. P. Adams. Probabilistic backpropagation for scalable learning of Bayesian neural networks. In *ICML*, 2015.
- [16] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *Signal Processing Magazine, IEEE*, 29(6): 82–97, 2012.
- [17] D. P. Kingma, T. Salimans, and M. Welling. Variational dropout and the local reparameterization trick. In *NIPS*, 2015.
- [18] A. Korattikara, V. Rathod, K. Murphy, and M. Welling. Bayesian dark knowledge. In *NIPS*, 2015.
- [19] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [20] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [21] Y. Li, J. M. Hernández-Lobato, and R. E. Turner. Stochastic expectation propagation. In *NIPS*, 2015.
- [22] C. Louizos and M. Welling. Structured and efficient variational deep learning with matrix gaussian posteriors. *arXiv preprint arXiv:1603.04733*, 2016.
- [23] D. J. MacKay. *Bayesian methods for adaptive models*. PhD thesis, California Institute of Technology, 1992.

- [24] S.-i. Maeda. A Bayesian encourages dropout. *arXiv preprint arXiv:1412.7003*, 2014.
- [25] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [26] T. P. Minka. Bayesian model averaging is not model combination. 2000.
- [27] T. Miyato, S.-i. Maeda, M. Koyama, K. Nakae, and S. Ishii. Distributional smoothing by virtual adversarial examples. In *ICLR*, 2016.
- [28] V. Nair and G. E. Hinton. Rectified linear units improve restricted Boltzmann machines. In *ICML*, 2010.
- [29] R. M. Neal. *Bayesian Learning for Neural Networks*. Springer-Verlag New York, Inc., 1996.
- [30] D. A. Nix and A. S. Weigend. Estimating the mean and variance of the target probability distribution. In *Neural Networks, 1994. IEEE World Congress on Computational Intelligence., 1994 IEEE International Conference on*, volume 1, pages 55–60. IEEE, 1994.
- [31] J. Quinonero-Candela, C. E. Rasmussen, F. Sinz, O. Bousquet, and B. Schölkopf. Evaluating predictive uncertainty challenge. In *Machine Learning Challenges*. Springer, 2006.
- [32] C. E. Rasmussen and J. Quinonero-Candela. Healing the relevance vector machine through augmentation. In *ICML*, 2005.
- [33] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *JMLR*, 2014.
- [34] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. In *ICLR*, 2014.
- [35] Y. W. Teh, L. Hasenclever, T. Lienart, S. Vollmer, S. Webb, B. Lakshminarayanan, and C. Blundell. Distributed Bayesian learning with stochastic natural-gradient expectation propagation and the posterior server. *arXiv preprint arXiv:1512.09327*, 2015.
- [36] M. Welling and Y. W. Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *ICML*, 2011.

Supplementary material

A Additional results on MNIST




Figure 9: Results on MNIST dataset using the same setup as that in Figure 4 except that we use two hidden layers in the MLP instead of three. Ensembles and adversarial training improve performance and our method outperforms MC-dropout.




Figure 10: Results on MNIST dataset using a convolutional network as opposed to the 3-layer MLP in Figure 4. Even on a different architecture, ensembles and adversarial training improve performance and our method outperforms MC-dropout.