
Robustly representing uncertainty through sampling in deep neural networks

Patrick McClure

MRC Cognition and Brain Sciences Unit
University of Cambridge
patrick.mcclure@mrc-cbu.cam.ac.uk

Nikolaus Kriegeskorte

Department of Psychology
Columbia University
nk2765@columbia.edu

Abstract

As deep neural networks (DNNs) are applied to increasingly challenging problems, they will need to be able to represent their own uncertainty. Modeling uncertainty is one of the key features of Bayesian methods. Using Bernoulli dropout with sampling at prediction time has recently been proposed as an efficient and well performing variational inference method for DNNs. However, sampling from other multiplicative noise based variational distributions has not been investigated in depth. We evaluated Bayesian DNNs trained with Bernoulli or Gaussian multiplicative masking of either the units (dropout) or the weights (dropconnect). We tested the calibration of the probabilistic predictions of Bayesian convolutional neural networks (CNNs) on MNIST and CIFAR-10. Sampling at prediction time increased the calibration of the DNNs' probabilistic predictions. Sampling weights, whether Gaussian or Bernoulli, led to more robust representation of uncertainty compared to sampling of units. However, using either Gaussian or Bernoulli dropout led to increased test set classification accuracy. Based on these findings we used both Bernoulli dropout and Gaussian dropconnect concurrently, which we show approximates the use of a spike-and-slab variational distribution without increasing the number of learned parameters. We found that spike-and-slab sampling had higher test set performance than Gaussian dropconnect and more robustly represented its uncertainty compared to Bernoulli dropout.

1 Introduction

Deep neural networks (DNNs), particularly convolutional neural networks (CNNs), have recently been used to solve complex perceptual and decision tasks [15; 21; 23]. While these models take into account aleatoric uncertainty via their softmax output (i.e. the uncertainty present in the training data), they do not take into account epistemic uncertainty (i.e. parameter uncertainty) [12]. Bayesian DNNs attempt to learn a distribution over their parameters thereby allowing for the computation of the uncertainty of their outputs given the parameters. However, ideal Bayesian methods do not scale well due to the difficulty in computing the posterior of a network's parameters.

As a result, several approximate Bayesian methods have been proposed for DNNs. Using the Laplace approximation was proposed by [18]. Using Markov chain Monte Carlo (MCMC) has been suggested to estimate the posterior of the networks weights given the training data [22; 26]. Using expectation propagation has also been proposed [11; 8]. However, these methods can be difficult to implement for the very large CNNs commonly used for object recognition. Variational inference methods have also been used to make Bayesian NNs more tractable [9; 1; 7; 2]. Due in large part to the fact that these methods substantially increase the number of parameters in a network, they have not been extensively applied to large DNNs. Gal and Ghahramani [5] and Kingma et al. [13] bypassed this issue by developing Bayesian CNNs using Bernoulli and Gaussian dropout [24], respectively. While independent weight sampling with additive Gaussian noise has been investigated [9; 1; 7; 2], independently sampling weights using multiplicative Bernoulli noise, i.e. dropconnect [25], or independently sampled multiplicative Gaussian noise has not been thoroughly evaluated.

In addition to Bernoulli and Gaussian distributions, spike-and-slab distributions, a combination of the two, have been investigated, particularly for linear models [20; 19; 6; 10]. Interestingly, Bernoulli dropout and dropconnect can be seen as approximations to spike-and-slab distributions for units and weights, respectively [17; 3]. Spike-and-slab variational distributions have been implemented using Bernoulli dropout with additive weight noise sampled from a Gaussian with a learned standard deviation [17]. This approach more than doubled the number of learned parameters, since the mean and the standard deviation of each weight as well as the dropout rate for each unit were learned. However, this method did not consistently outperform standard neural networks. Gal [3] also discussed motivations for spike-and-slab variational distributions, but did not suggest a practical implementation.

We evaluated the performance Bayesian CNNs with different variational distributions on CIFAR-10. We also investigate how adding Gaussian image noise with varying standard deviations to the test set affected each network’s learned uncertainty. We did this to test how networks responded to inputs not drawn from the data distribution used to create the training and test sets. We also propose an approximation of the spike-and-slab variational inference based on Bernoulli dropout and Gaussian dropconnect, which combines the advantages of Gaussian dropconnect and Bernoulli dropout sampling leading to better uncertainty estimates and good test set generalization without increasing the number of learned parameters.

2 Methods

2.1 Bayesian Deep Neural Networks

DNNs are commonly trained by finding the maximum a posteriori (MAP) weights given the training data (D_{train}) and a prior over the weight matrix W ($p(W)$). However, ideal Bayesian learning would involve computing the full posterior. This can be intractable due to both the difficulty in calculating $p(D_{train})$ and in calculating the joint distribution of a large number of parameters. Instead, $p(W|D_{train})$ can be approximated using a variational distribution $q(W)$. This distribution is constructed to allow for easy generation of samples. The objective of variational inference is to optimize the variational parameters V so that the Kullback-Leiber (KL) divergence between $q_V(W)$ and $p(W|D_{train})$ is minimized [9; 11; 7; 2]:

$$V^* = \underset{V}{\operatorname{argmin}} KL[q_V(W)||p(W)] - \int q_V(W) \log p(D_{train}|W) dW \quad (1)$$

Using Monte Carlo (MC) methods to estimate $E_{q_V(w)}[\log p(D_{train}|W)]$, using weight samples $\hat{W}^i \sim q_V(W)$, results in the following loss function:

$$\mathcal{L} := KL(q_V(W)||p(W)) - \frac{1}{n} \sum_i^n \log p(D_{train}|\hat{W}^i) \quad (2)$$

MC sampling can also be used to estimate the probability of test data:

$$p(D_{test}) \approx \frac{1}{n} \sum_i^n p(D_{test}|\hat{W}^i) \quad (3)$$

2.2 Variational Distributions

The number and continuous nature of the parameters in DNNs makes sampling from the entire distribution of possible weight matrices computationally challenging. However, variational distributions can make sampling easier. In deep learning, the most common sampling method is using multiplicative noise masks drawn from some distribution. Several of these methods can be formulated as variational distributions where weights are sampled by element-wise multiplication of the variational




Figure 1: An illustration of sampling network weights using the different variational distributions.

parameters V , the $n \times n$ connection matrix with an element for each connection between the n units in the network, by a mask \hat{M} , which is sampled from some probability distribution:

$$\hat{W} = V \circ \hat{M} \text{ where } \hat{M} \sim p(M) \quad (4)$$

From this perspective, the difference between dropout and dropconnect, as well as Bernoulli and Gaussian methods, is simply the probability distribution used to generate the mask sample (Figure 1).

2.2.1 Bernoulli Dropconnect & Dropout

In Bernoulli dropconnect, each element of the mask is sampled independently, so $\hat{m}_{i,j} \sim \text{Bernoulli}(1 - p)$ where p is the probability of dropping a connection. In Bernoulli dropout, however, the weights are not sampled independently. Instead, one Bernoulli variable is sampled for each row of the weight matrix, so $\hat{m}_{i,*} \sim \text{Bernoulli}(1 - p)$ where p is the probability of dropping a unit.

2.2.2 Gaussian Dropconnect & Dropout

In Gaussian dropconnect and dropout, $\hat{w}_{i,j}$ is sampled from a Gaussian distribution centered at variational parameter $v_{i,j}$. This is accomplished by sampling the multiplicative mask using Gaussian distributions with a mean of 1 and a variance of $\sigma_{dc}^2 = p/(1 - p)$, which matches the mean and variance of Bernoulli dropout when training time scaling is used [24]. In Gaussian dropconnect, each element of the mask is sampled independently, which results in $\hat{m}_{i,j} \sim \mathcal{N}(1, \sigma_{dc}^2)$. In Gaussian dropout, each element in a row has the same random variable, so $\hat{m}_{i,*} \sim \mathcal{N}(1, \sigma_{dc}^2)$. It can be shown that using Gaussian dropconnect or dropout with L2-regularization leads to optimizing a stochastic lower-bound of the variational objective function (See Supplementary Material).

2.2.3 Spike-and-Slab Dropout

A spike-and-slab distribution is the normalized linear combination of a "spike" of probability mass at zero and a "slab" consisting of a Gaussian distribution. This spike-and-slab returns a 0 with probability p_{spike} or a random sample from a Gaussian distribution $\mathcal{N}(\mu_{slab}, \sigma_{slab}^2)$ with probability $1 - p_{spike}$. We propose concurrently using Bernoulli dropout and Gaussian dropconnect to approximate the use of a spike-and-slab variational distribution and spike-and-slab prior by optimizing a lower-bound of the variational objective function (See Supplementary Material). In this formulation, $m_{i,j} \sim b_{i,*} \mathcal{N}(1, \sigma_{dc}^2)$, where $b_{i,*} \sim \text{Bern}(1 - p_{do})$ for each mask row and $\sigma_{dc}^2 = p_{dc}/(1 - p_{dc})$. As for Bernoulli dropout, each row of the mask M is multiplied by 0 with probability p_{do} , otherwise each element in that row is multiplied by a value independently sampled from a Gaussian distribution as in Gaussian dropconnect. During non-sampling inference, spike-and-slab dropout uses the mean weight values and, per Bernoulli dropout, multiplies unit outputs by $1 - p_{do}$.




Figure 2: The probabilistic logistic regression decision boundaries of a linear network for Bernoulli dropconnect (BDC), Gaussian dropconnect (GDC), Bernoulli dropout (BDO), Gaussian dropout (GDO), and spike-and-slab dropout (SSD) with and without MC sampling compared to the decision boundaries for the MAP solution and stochastic gradient Langevin dynamics (SGLD) [26].

3 Experiments

3.1 Logistic Regression

In order to visualize the effects of each variational distribution, we trained linear networks with five hidden units to classify data drawn from two 2D multivariate Gaussian distributions. Multiple linear units were used so that Bernoulli dropout would not dropout the only unit in the network. For the dropout methods, unit sampling was performed on the linear hidden layer. For the dropconnect methods, every weight was sampled. Dropout and dropconnect probabilities of $p = 0.4$ were used for each of these networks, except for the spike-and-slab dropconnect probability which was 0.2. In Figure 2, we show the decision boundaries learned by the various networks. Higher variability in the decision boundary corresponds to higher uncertainty. All of the MC sampling methods predict with higher uncertainty as points become further away from the training data. This is particularly true for the dropconnect and spike-and-slab methods.

3.2 Convolutional Neural Networks

We trained CNNs on the MNIST [16] and CIFAR-10 [14] datasets. For each dataset, a 10,000 image subset of the training set was used for validation. For MNIST, each CNN had two convolutional layers followed by a fully connected layer and a softmax layer. For CIFAR-10, each CNN had 13 convolutional layers followed by a fully connected layer and a softmax layer. (See Supplementary Material for the detailed architectures.) For the dropout networks, dropout was used after each convolutional and fully-connected layer, but before the non-linearity. For the dropconnect networks, all weights were sampled. All p s were treated as network-wide hyperparameters. For L2-regularization, L2-coefficients of $1e-5$ (MNIST) and $4e-5$ (CIFAR-10) were used for all weights. No data augmentation was used for MNIST. Random horizontal flipping was used during CIFAR-10 training. We

Table 1: MNIST and CIFAR-10 mean and standard deviation of test errors for the trained convolutional neural networks (CNNs) with and without Monte-Carlo (MC) across 5 runs, each MC run using 10 samples.

Method	MNIST		CIFAR-10	
	Mean Error (%)	Error Std. Dev.	Mean Error (%)	Error Std. Dev.
MAP	0.76	-	25.86	-
Bernoulli DropConnect	0.56	-	16.46	-
MC Bernoulli DropConnect	0.56	0.03	16.59	0.11
Gaussian DropConnect	0.56	-	16.78	-
MC Gaussian DropConnect	0.58	0.02	16.65	0.11
Bernoulli Dropout	0.49	-	11.23	-
MC Bernoulli Dropout	0.48	0.03	9.95	0.08
Gaussian Dropout	0.42	-	9.07	-
MC Gaussian Dropout	0.36	0.04	9.00	0.10
Spike-and-Slab Dropout	0.48	-	10.64	-
MC Spike-and-Slab Dropout	0.46	0.01	10.05	0.06




Figure 3: The MNIST (a) classification error for additive Gaussian noise using standard deviations St.D of 0, 1, 2, 3, 5, and 5, (b) mean squared error (MSE) between the $x = y$ line and the calibration plot (i.e. the frequency of the true label vs predicted probability of that label) for varying Gaussian image noise StD., and (c) The calibration MSE versus the classification error for predictions across all noise StD. for Bernoulli dropconnect (BDC), Gaussian dropconnect (GDC), Bernoulli dropout (BDO), Gaussian dropout (GDO), and spike-and-slab dropout (SSD) with and without MC sampling using 10 samples.

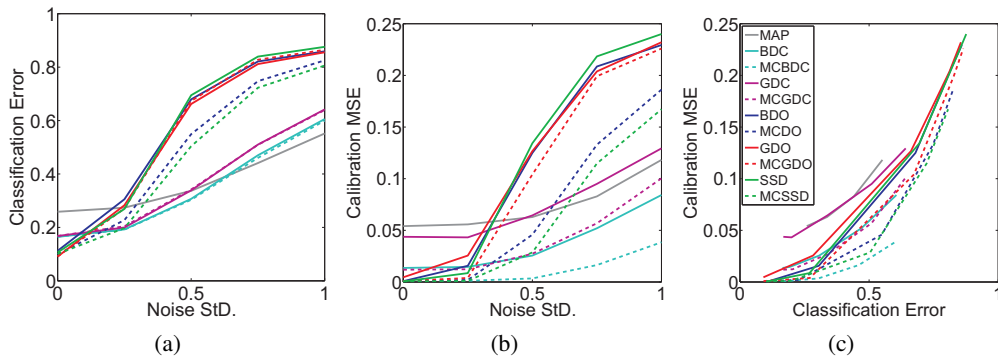


Figure 4: The CIFAR-10 (a) classification error for additive Gaussian noise using standard deviations St.D of 0, 0.25, 0.5, 0.75, and 1, (b) mean squared error (MSE) between the $x = y$ line and the calibration plot (i.e. the frequency of the true label vs predicted probability of that label) for varying Gaussian image noise StD., and (c) The calibration MSE versus the classification error for predictions across all noise StD. for Bernoulli dropconnect (BDC), Gaussian dropconnect (GDC), Bernoulli dropout (BDO), Gaussian dropout (GDO), and spike-and-slab dropout (SSD) with and without MC sampling using 10 samples.

evaluated the trained CNNs using the original testing sets and using the testing images with added random Gaussian noise of increasing variance in order to test each network’s uncertainty for the regions of input space not seen in the training set.

While the dropout-based methods were the most accurate on the test-set (Table 1), as image noise was added they became increasingly worse compared to the dropconnect-based networks (Figure 3a and 4a). Sampling only consistently improved the accuracy for Bernoulli and spike-and-slab dropout. However, sampling did consistently improve the calibration of the networks as the image noise was increased (Figure 3b and 4b). For a given accuracy across each set of noisy test images, sampling also generally lead to better calibration (Figure 3c and 4c). (See Supplementary Material for a table of the calibration plots and training plots for the different CNNs.) Gaussian dropout led to the highest test set accuracy, but it also led to reduced robustness to noise. While slightly less accurate on the test set, Bernoulli dropout and spike-and-slab dropout were much more robust.

Seemingly contradictory results have been reported in the literature regarding CIFAR-10 and MC Bernoulli dropout. Gal and Ghahramani [4] found that standard Bernoulli dropout methods led to relatively inaccurate networks when dropout was used at every layer in a CNN, whereas MC sampling increased the accuracy of these networks. However, Srivastava et al. [24] found that using dropout at




Figure 5: The classification error for the CIFAR-10 test set using different p hyperparameter values. For spike-and-slab dropout, p_{do} was varied, while p_{dc} was fixed.

every layer led to increased generalization performance even without sampling at prediction time. In our CIFAR-10 experiments, but not our MNIST experiments, we have found that using sampling at prediction time makes networks more robust to high variance dropout. Using lower variance dropout results in standard and MC methods having similar accuracies, while using higher variance distributions results in MC inference outperforming standard methods (Figure 5). These results indicate that Bernoulli or Gaussian dropout with MC sampling are less dependent on the exact value of p and can allow higher levels of dropout regularization to be used.

4 Discussion

L2 regularization and Bernoulli dropout are widely used for regularization and routinely lead to increased testing accuracy. However, the uncertainty learned do not generalize well. However, performing approximate Bayesian inference via sampling during training and testing allowed CNNs to better model their uncertainty. Dropconnect-based CNNs performed worse on the unmodified test set, but were much more robust to deviations from the training distribution. On the other hand, dropout-based networks, particularly MC Gaussian dropout, performed well on the unmodified test set, but were not as robust. Using sampling and combining Bernoulli dropout and Gaussian dropconnect to approximate the use of spike-and-slab variational distributions lead to a CNN that performed better near the test set than the dropconnect methods and more robustly represented its uncertainty compared to the dropout methods.

References

- [1] David Barber and Christopher M Bishop. Ensemble learning in bayesian neural networks. *NATO ASI SERIES F COMPUTER AND SYSTEMS SCIENCES*, 168:215–238, 1998.
- [2] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *Proceedings of The 32nd International Conference on Machine Learning*, pages 1613–1622, 2015.
- [3] Yarin Gal. *Uncertainty in Deep Learning*. PhD thesis, University of Cambridge, 2016.
- [4] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Insights and applications. In *Deep Learning Workshop, ICML*, 2015.
- [5] Yarin Gal and Zoubin Ghahramani. Bayesian convolutional neural networks with Bernoulli approximate variational inference. In *4th International Conference on Learning Representations (ICLR) workshop track*, 2016.
- [6] Edward I George and Robert E McCulloch. Approaches for bayesian variable selection. *Statistica sinica*, pages 339–373, 1997.
- [7] Alex Graves. Practical variational inference for neural networks. In *Advances in Neural Information Processing Systems*, pages 2348–2356, 2011.

- [8] José Miguel Hernández-Lobato and Ryan Adams. Probabilistic backpropagation for scalable learning of bayesian neural networks. In *International Conference on Machine Learning*, pages 1861–1869, 2015.
- [9] 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*, pages 5–13. ACM, 1993.
- [10] Hemant Ishwaran and J Sunil Rao. Spike and slab variable selection: frequentist and bayesian strategies. *Annals of Statistics*, pages 730–773, 2005.
- [11] Pasi Jylänki, Aapo Nummenmaa, and Aki Vehtari. Expectation propagation for neural networks with sparsity-promoting priors. *Journal of Machine Learning Research*, 15(1):1849–1901, 2014.
- [12] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *arXiv preprint arXiv:1703.04977*, 2017.
- [13] Diederik P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick. In *Advances in Neural Information Processing Systems*, pages 2575–2583, 2015.
- [14] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images, 2009.
- [15] Alex Krizhevsky, Ilya Sutskever, Geoffrey E Hinton, Krizhevsky Alex, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances In Neural Information Processing Systems*, pages 1–9, 2012. ISBN 9781627480031. URL <http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf>.
- [16] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [17] Christos Louizos. Smart regularization of deep architectures. Master’s thesis, University of Amsterdam, 2015.
- [18] David JC MacKay. A practical bayesian framework for backpropagation networks. *Neural computation*, 4(3):448–472, 1992.
- [19] David Madigan and Adrian E Raftery. Model selection and accounting for model uncertainty in graphical models using occam’s window. *Journal of the American Statistical Association*, 89(428):1535–1546, 1994.
- [20] Toby J Mitchell and John J Beauchamp. Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, 83(404):1023–1032, 1988.
- [21] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei a Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015. ISSN 0028-0836. doi: 10.1038/nature14236. URL <http://dx.doi.org/10.1038/nature14236>.
- [22] Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- [23] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

- [24] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [25] Li Wan, Matthew Zeiler, Sixin Zhang, Yann L Cun, and Rob Fergus. Regularization of neural networks using dropconnect. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 1058–1066, 2013.
- [26] Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 681–688, 2011.

Supplementary Material: Robustly representing uncertainty through sampling in deep neural networks

Patrick McClure
MRC Cognition and Brain Sciences Unit
University of Cambridge
patrick.mcclure@mrc-cbu.cam.ac.uk

Nikolaus Kriegeskorte
Department of Psychology
Columbia University
nk2765@columbia.edu

1 Methods

1.1 L2 regularization and the KLD between Gaussians

The Kullback–Leibler divergence (KLD) between $\mathcal{N}(\mu_q, \sigma_q^2)$ and $\mathcal{N}(\mu_p, \sigma_p^2)$ can be calculated using:

$$KL(q(w_{i,j})||p(w_{i,j})) = \frac{(\mu_q - \mu_p)^2}{2\sigma_p^2} + \log \frac{\sigma_p}{\sigma_q} + \frac{\sigma_q^2}{2\sigma_p^2} - \frac{1}{2} \quad (1)$$

In the case where $\mathcal{N}(\mu_p, \sigma_p^2)$ is a pre-defined prior and σ_q is not a function of the learnable parameters V :

$$\operatorname{argmin}_V KL(q(w_{i,j})||p(w_{i,j})) = \operatorname{argmin}_V \frac{(\mu_q - \mu_p)^2}{2\sigma_p^2} \quad (2)$$

For $\mu_p = 0$, this is equivalent to L2 regularization where the L2-coefficient is equal to $1/\sigma_p^2$. However, in the case where σ_q is a function of V , such as for Gaussian dropout/dropconnect, this equivalence does not hold. In [5], Kingma et al. used a log-uniform prior instead of a Gaussian prior in order to bypass this and make the KLD not a function of V . In our derivations, we minimize a lower bound of Equation 1 constructed using the fact that the sum of the terms that include σ_q and the constant term is greater than or equal to 0:

$$KL(q(w_{i,j})||p(w_{i,j})) \geq \frac{(\mu_q - \mu_p)^2}{2\sigma_p^2} \quad (3)$$

1.2 Gaussian "reparameterization trick"

As discussed in [5], for a matrix W of Gaussian random variables can be sampled using the "reparameterization trick":

$$w_{i,j} \sim \mathcal{N}(\mu_{v_{i,j}}, \alpha v_{i,j}^2) \quad (4)$$

$$w_{i,j} = f(v_{i,j}, \epsilon_{i,j}) = v_{i,j} + \sqrt{\alpha} v_{i,j} \epsilon_{i,j} \quad (5)$$

where $\epsilon_{i,j} \sim \mathcal{N}(0, 1)$, $\alpha = p/(1-p)$, and p is the dropout or dropconnect drop probability. Given a deterministic, differentiable, and monotonic mapping $W = f(V, \epsilon)$, $q_V(W)dW = p(\epsilon)d\epsilon$. As a

result:

$$\int q_V(W)l(W)dW = \int p(\epsilon)l(W)d\epsilon = \int p(\epsilon)l(f(V, \epsilon))d\epsilon \quad (6)$$

1.3 MC Gaussian Dropout

For approximate inference, variational distribution $q_V(W)$ is learned by maximizing the log-evidence lower bound over parameters V [1, 2, 3, 4]:

$$\log(p(D_{train})) \geq \int \log p(D_{train}|W)q_V(W)dW - KL(q_V(W)||p(W)) \quad (7)$$

For either Gaussian dropout or dropconnect, each element of W is sampled from a Gaussian distribution, $\mathcal{N}(v_{i,j}, \sigma_{v_{i,j}}^2)$, where $\sigma_{v_{i,j}}^2 = \alpha\mu_{v_{i,j}}^2$. W can then be sampled using the Gaussian "reparameterization trick", which allows Equation 7 to be rewritten as:

$$\log(p(D_{train})) \geq \int_{\epsilon} \log p(D_{train}|W)q(\epsilon)d\epsilon - KL(q_V(W)||p(W)) \quad (8)$$

where ϵ is a vector containing each $\epsilon_{i,j}$.

This results in the following minimization objective function:

$$\mathcal{L}_V := - \int_{\epsilon} \log(p(D_{train}|W))q(\epsilon)d\epsilon + KL(q_V(W)||p(W)) \quad (9)$$

By using L2 regularization, we are optimizing a lower-bound of the KLD between $q_V(W)$ and the prior $p(w_{i,j} = \mathcal{N}(0, \lambda^{-1}))$ as previously shown:

$$\mathcal{L}_V \geq \tilde{\mathcal{L}}_V := - \int_{\epsilon} \log(p(D_{train}|W))q(\epsilon)d\epsilon + \frac{\lambda}{2} \mathbf{v}\mathbf{v}^T \quad (10)$$

where \mathbf{v} is a vector containing each $v_{i,j}$ and ϵ is a vector containing each $\epsilon_{i,j}$.

Approximating using Monte Carlo integration for learning (Eq. 11) and inference (Eq. 12):

$$\tilde{\mathcal{L}}_V \approx -\frac{1}{n} \sum_{\epsilon} \log(p(D_{train}|W)) + \frac{\lambda}{2} \mathbf{v}\mathbf{v}^T \quad (11)$$

$$p(D_{test}) \approx \frac{1}{n} \sum_{\epsilon} p(D_{test}|W) \quad (12)$$

where $\epsilon_{i,j} \sim \mathcal{N}(0, 1)$ for Gaussian dropconnect and $\epsilon_{i,*} \sim \mathcal{N}(0, 1)$ for Gaussian dropout.

1.4 MC spike-and-slab Dropout

For MC spike-and-slab dropout, the weight matrix $W = B \circ G$ where $b_{i,*} \sim \text{Bern}(1 - p_{do})$ and $g_{i,j} \sim \mathcal{N}(v_{i,j}, \sigma_{v_{i,j}}^2)$, similar to the method discussed in [6]. Instead of directly performing variational inference for $p(W|D_{train})$, we find a variational distribution, $q_V(B, G)$ for $p(B, G|D_{train})$ using:

$$\log(p(D_{train})) \geq \sum_B \int_G \log(p(D_{train}|B, G))q_V(B, G)dG - KL(q_V(B, G)||p(B, G)) \quad (13)$$

Assuming independence between the random variables B and G , $q(B, G) = q(B)q(G)$, so:

$$\log(p(D_{train})) \geq \sum_B \int_G \log(p(D_{train}|B, G))q(B)q_V(G)dG - KL(q(B)||p(B)) - KL(q_V(G)||p(G)) \quad (14)$$

For a spike-and-slab distribution, each element of G is independently sampled from a Gaussian distribution, $\mathcal{N}(v_{i,j}, \sigma_{v_{i,j}}^2)$, where $\sigma_{v_{i,j}}^2 = \alpha \mu_{v_{i,j}}^2$. G can be sampled using the Gaussian "reparameterization trick". This allows Equation [14](#) to be rewritten as:

$$\log(p(D_{train})) \geq \sum_B \int_{\epsilon} \log(p(D_{train}|B, G)q(\epsilon)q(B)d\epsilon - KL(q(B)||p(B)) - KL(q_V(G)||p(G)) \quad (15)$$

ϵ is a vector containing each $\epsilon_{i,j}$.

This results in the following minimization objective function:

$$\mathcal{L}_V := - \sum_B \int_{\epsilon} \log(p(D_{train}|B, G)q(\epsilon)q(B)d\epsilon + KL(q(B)||p(B)) + KL(q_V(G)||p(G)) \quad (16)$$

Using $Bern(1 - p_{do})$ as a prior for each element of B leads to a constant KLD of zero for Bernoulli dropout with a drop probability of p_{do} and using a prior of $\mathcal{N}(0, \sigma_p^2)$ for each element of G leads to L2-regularization being a lowerbound of the KLD between $q_V(G)$ and $\mathcal{N}(0, \lambda^{-1})$:

$$\mathcal{L}_V \geq \tilde{\mathcal{L}}_V := - \sum_B \int_{\epsilon} \log(p(D_{train}|B, G)q(\epsilon)q(B)d\epsilon + \frac{\lambda}{2} \mathbf{v}\mathbf{v}^T \quad (17)$$

where \mathbf{v} is a vector containing each $v_{i,j}$ and ϵ is a vector containing each $\epsilon_{i,j}$.

Approximating using Monte Carlo integration for learning (Eq. [18](#)) and inference (Eq. [19](#)):

$$\tilde{\mathcal{L}}_V := -\frac{1}{n} \sum_{(B, \epsilon)} \log(p(D_{train}|B, G)) + \frac{\lambda}{2} \mathbf{v}\mathbf{v}^T \quad (18)$$

$$p(D_{test}) \approx \frac{1}{n} \sum_{(B, \epsilon)} p(D_{test}|B, G) \quad (19)$$

where $b_{i,*} \sim Bern(1 - p_{do})$ and $\epsilon_{i,j} \sim \mathcal{N}(0, 1)$.

2 Experiments

2.1 Architectures

Table 1: The convolutional neural network (CNN) architecture used for MNIST.

Layer	Kernel Size	# Features	Stride	Non-linearity
Conv-1	5x5	32	1	ReLU
MaxPool-1	2x2	32	2	Max
Conv-2	5x5	64	1	ReLU
MaxPool-2	2x2	64	2	Max
FC	1500	500	-	ReLU
FC	500	10	-	Softmax

Table 2: The convolutional neural network (CNN) architecture used for CIFAR-10.

Layer	Kernel Size	# Features	Stride	Non-linearity
Conv-1	3x3	64	1	ReLU
Conv-2	3x3	64	1	ReLU
MaxPool-1	2x2	64	2	Max
Conv-3	3x3	128	1	ReLU
Conv-4	3x3	128	1	ReLU
MaxPool-2	2x2	128	2	Max
Conv-5	3x3	256	1	ReLU
Conv-6	3x3	256	1	ReLU
Conv-7	3x3	256	1	ReLU
MaxPool-3	2x2	256	2	Max
Conv-8	3x3	512	1	ReLU
Conv-9	3x3	512	1	ReLU
Conv-10	3x3	512	1	ReLU
MaxPool-4	2x2	512	2	Max
Conv-11	3x3	512	1	ReLU
Conv-12	3x3	512	1	ReLU
Conv-13	3x3	512	1	ReLU
MaxPool-5	2x2	512	2	Max
FC	512	512	-	ReLU
FC	512	10	-	Softmax

2.2 Additional results

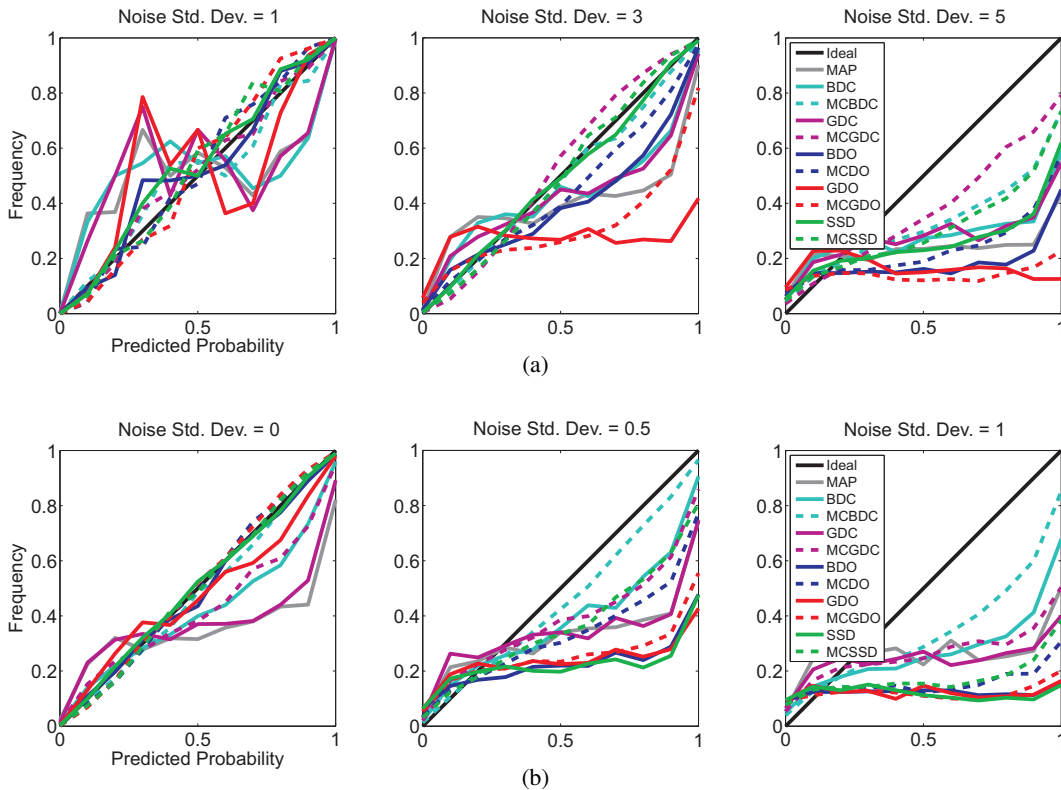


Figure 1: The $x = y$ line (Ideal) and the calibration plot (i.e. the frequency of the true label vs predicted probability of that label) for varying Gaussian image noise StD. for the a) MNIST or (b) CIFAR-10 trained Bernoulli dropconnect (BDC), Gaussian dropconnect (GDC), Bernoulli dropout (BDO), Gaussian dropout (GDO), and spike-and-slab dropout (SSD) networks with and without MC sampling using 10 samples.

References

- [1] David Barber and Christopher M Bishop. Ensemble learning in bayesian neural networks. *NATO ASI SERIES F COMPUTER AND SYSTEMS SCIENCES*, 168:215–238, 1998.
- [2] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *Proceedings of The 32nd International Conference on Machine Learning*, pages 1613–1622, 2015.
- [3] Alex Graves. Practical variational inference for neural networks. In *Advances in Neural Information Processing Systems*, pages 2348–2356, 2011.
- [4] 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*, pages 5–13. ACM, 1993.
- [5] Diederik P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick. In *Advances in Neural Information Processing Systems*, pages 2575–2583, 2015.
- [6] Michalis K Titsias and Miguel Lázaro-Gredilla. Spike and slab variational inference for multi-task and multiple kernel learning. In *Advances in neural information processing systems*, pages 2339–2347, 2011.