

REPRESENTATION OF UNCERTAINTY IN DEEP NEURAL NETWORKS THROUGH SAMPLING

Patrick McClure & Nikolaus Kriegeskorte

MRC Cognition and Brain Science Unit

University of Cambridge

Cambridge, UK

{Patrick.McClure,Nikolaus.Kriegeskorte}@mrc-cbu.cam.ac.uk

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. Scalable Bayesian DNNs that use dropout-based variational distributions have recently been proposed. Here we evaluate the ability of Bayesian DNNs trained with Bernoulli or Gaussian distributions over units (dropout) or weights (dropconnect) to represent their own uncertainty at the time of inference through sampling. We tested how well Bayesian fully connected and convolutional DNNs represented their own uncertainty in classifying the MNIST handwritten digits. By adding different levels of Gaussian noise to the test images, we assessed how DNNs represented their uncertainty about regions of input space not covered by the training set. Bayesian DNNs estimated their own uncertainty more accurately than traditional DNNs with a softmax output. These results are important for building better deep learning systems and for investigating the hypothesis that biological neural networks use sampling to represent uncertainty.

1 INTRODUCTION

Deep neural networks (DNNs), particularly convolutional neural networks (CNN), have recently been used to solve complex perceptual and decision tasks (Krizhevsky et al., 2012; Mnih et al., 2015; Silver et al., 2016). However, these models fail to model the uncertainty of their predictions or actions. Although some networks' outputs are probability distributions, these networks deterministically map an input to a probabilistic predictions, but do not model the uncertainty of this mapping. In contrast, Bayesian neural networks (NNs) attempt to learn a distribution over a networks parameters thereby offering uncertainty estimates of their outputs (MacKay, 1992; Neal, 2012). However, these methods do not scale well due to the difficulty in computing the posterior of the network parameters.

Approximate methods, in particular variational inference, have been used to make Bayesian NNs more tractable (Hinton & Van Camp, 1993; Barber & Bishop, 1998; Graves, 2011; Blundell et al., 2015). Due in large part to the fact that these methods substantially increase the number of parameters in a network, they have not been applied to large DNNs, such as CNNs. Gal & Ghahramani (2016) and Kingma et al. (2015) bypassed this issue by developing Bayesian CNNs using dropout (Srivastava et al., 2014). Dropout is a widely used regularization technique that during training drops a unit out of the network with a probability p and during inference multiplies the output of each unit by p . A similar technique is dropconnect (Wan et al., 2013), which drops network connections instead of units. Gal & Ghahramani (2015) detailed how dropping units was equivalent to sampling weights from a Bernoulli-based variational distribution and that in order to make a DNN with dropout Bayesian, sampling should be used during both training and inference. Monte-Carlo (MC) sampling at inference allows a DNN to efficiently model a distribution over its outputs. The uncertainty of a DNN can then be calculated using this probability distribution. By formulating their method as a version of dropout, their Bayesian approach is applicable to most networks that utilize dropout.

The use of sampling to model the uncertainty of a NN has also been investigated in computational neuroscience. The neural sampling hypothesis states that the activity patterns of biological neural networks represent samples from a learned posterior distribution over interpretations given an input (Fiser et al., 2010). Restricted Boltzmann machines (RBMs) (Hinton, 2010) and spiking neural networks with binary stochastic units that perform Bayesian inference have been proposed as models (Buesing et al., 2011; Habenschuss et al., 2013). Like Bayesian DNNs, these models use the variability of unit activations to represent uncertainty. However, these models do not scale as well to complex tasks as other types of NNs, such as CNNs.

In this paper, we investigate how using MC sampling to model the uncertainty of a network affects a network’s predictions. Specifically, we test if using MC sampling improves the calibration of the probabilistic predictions made by Bayesian DNNs with softmax output layers. Unlike previous work, we used variational distributions based on dropout and dropconnect with either Bernoulli or Gaussian sampling during both training and inference. These variational distributions scale extremely well and make the results of this work applicable to a large range of state-of-the-art DNNs.

2 METHODS

2.1 BAYESIAN NEURAL NETWORKS

Artificial neural networks (NNs) can be trained using Bayesian learning by finding the maximum a posteriori (MAP) weights given the training data (D_{train}) and a prior over the weight matrix W ($p(W)$):

$$\max_W p(W|D_{train}) = \max_W p(D_{train}|W)p(W) \quad (1)$$

This is usually done by minimizing the mean squared error (MSE) or cross entropy error for either regression or classification, respectively, while using L2 regularization, which corresponds to a Gaussian prior over weights. At inference, the probability of the test data (D_{test}) is then calculated using only the maximum likelihood estimate (MLE) of the weights (W^*):

$$p(D_{test}|W^*) \quad (2)$$

However, ideally the full posterior distribution over the weights would be learned instead of just the MLE:

$$p(W|D_{train}) = \frac{p(D_{train}|W)p(W)}{p(D_{train})} \quad (3)$$

This can be intractable due to both the difficulty in calculating $p(D_{train})$ and calculating the joint distribution of a large number of parameters, so $p(W|D_{train})$ can be approximated using a variational distribution $q(W)$. This distribution is constructed to be easy to model and to allow for easy generation of samples. Using variational inference, $q(W)$ is learned by minimizing:

$$- \int \log p(D_{train}|W)q(W)dW + KL(q(W)||p(W)) \quad (4)$$

Monte-Carlo (MC) sampling can then be used to estimate the probability of test data using $q(W)$:

$$p(D_{test}) \approx \frac{1}{n} \sum_i^n p(D_{test}|\hat{W}^i) \text{ where } \hat{W}^i \sim q(W) \quad (5)$$

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. As a result, distributions that are easier to sample from can be used. In the deep learning literature, the most common sampling




Figure 1: A visualization of Bernoulli and Gaussian dropconnect and dropout on a simple neural network.

method is dropout with Bernoulli variables. However, dropconnect, having independently sampled variables for each parameter, and Gaussian variables have also been used. (A visualization of the different methods is shown in Figure 1.) All of these methods can be formulated as variational distributions where weights are sampled by element-wise multiplying the variational 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. Mathematically, this can be written as:

$$\hat{W} = V \circ \hat{M} \text{ where } \hat{M} \sim p(M) \tag{6}$$

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, \hat{M} (Figure 2).




Figure 2: An illustration of sampling weights using Bernoulli and Gaussian dropconnect and dropout.

2.2.1 BERNOULLI DROPCONNECT & DROPOUT

Bernoulli distributions are simple distributions which return a 1 with probability p and 0 with a probability $(1 - p)$. In Bernoulli dropconnect, each element of the mask is sampled independently, so $\hat{m}_{i,j} \sim \text{Bernoulli}(p)$. This sets $\hat{w}_{i,j}$ to $v_{i,j}$ with probability p and 0 with a probability $(1 - p)$. In dropout, however, the weights are not sampled independently. Instead, one Bernoulli variable is sampled for each row of the weight matrix, so $\hat{m}_{i,*} \sim \text{Bernoulli}(p)$.

2.2.2 GAUSSIAN DROPCONNECT & DROPOUT

In Gaussian dropconnect and dropout, the mask is sampled from a normal distribution. As proposed by Srivastava et al. (2014), we used a Gaussian distribution with a mean of 1 and a standard deviation of $\sqrt{(1 - p)/p}$, which matches the mean and variance of dropout when training time scaling is used. This corresponds to sampling $\hat{w}_{i,j}$ from a Gaussian distribution centered at variational parameter $v_{i,j}$. In Gaussian dropconnect, each element of the mask is sampled independently, which results in $\hat{m}_{i,j} \sim \mathcal{N}(0, \sqrt{(1 - p)/p})$. In Gaussian dropout, each element in a row has the same random variable, so $\hat{m}_{i,*} \sim \mathcal{N}(0, \sqrt{(1 - p)/p})$.




Figure 3: Examples of noisy MNIST images where Gaussian noise with standard deviations of 0, 1, 2, 3, 4, and 5.

3 RESULTS

In this paper, we investigate the effects of performing probabilistic inference using MC sampling has on how well a DNN models its own uncertainty. To test this, we trained several networks differing only in whether no sampling was performed (baseline NN and NN with L2-regularization), sampling was only performed during learning (dropout and dropconnect), or sampling was performed both during training and inference (MC dropout and MC dropconnect). Additionally, we varied the variational distribution used for each network with sampling. The standard probability of $p = 0.5$ was used for both Bernoulli and Gaussian variational distributions. We trained two groups of DNNs, one with a fully connected (FC) architecture and one with a convolutional architecture, on digit classification using the 28x28 images from the MNIST dataset (LeCun et al., 1998).

We compared the test classification error, the uncertainty of the softmax output, and the calibration of the softmax output for each level of sampling and variational distribution. The test classification error tells us how well the probability distribution learned by each DNN models the data. The uncertainty tells us how the probability distribution learned by each DNN is distributed across classes, a low entropy means that the probability mass is primarily located at a few labels and a high entropy means that the probability mass is distributed across labels. The calibration tells us how well the probability distribution learned by the DNN models its own uncertainty (i.e. when an outcome is predicted with probability p , that outcome should occur with probability p given that prediction). To evaluate how calibrated a prediction was, we calculated the MSE between the observed frequency that a prediction of probability p was the correct label and the value of p . We evaluated these three measures for the trained networks on the MNIST test set with noise sampled from Gaussian distributions with varying standard deviations (Figure 3) to not only test how well modelled each network’s uncertainty was on the MNIST test set, but also in regions of input space not seen in the training set.

Table 1: MNIST test error for the trained fully connected neural networks with and without Monte-Carlo (MC) sampling using 100 samples.

Method	Mean Error (%)	Std. Dev.
NN	1.68	-
NN+L2	1.64	-
Bernoulli DropConnect	1.33	-
MC Bernoulli DropConnect	1.30	0.04
Gaussian DropConnect	1.24	-
MC Gaussian DropConnect	1.27	0.03
Bernoulli Dropout	1.45	-
MC Bernoulli Dropout	1.42	0.03
Gaussian Dropout	1.36	-
MC Gaussian Dropout	1.37	0.03

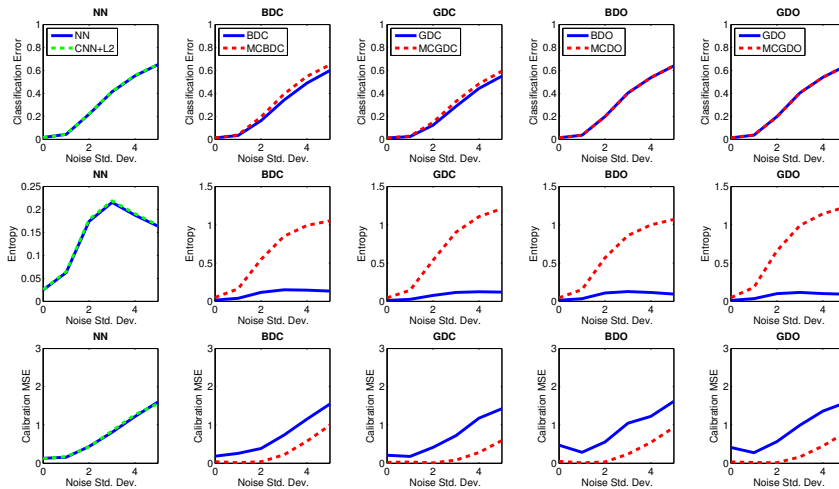


Figure 4: The MNIST test classification error, entropy, and calibration of the predictions of the fully connected networks: NN, NN+L2, Bernoulli DropConnect (BDC) with and without Monte-Carlo (MC) sampling, Gaussian DropConnect (GDC) with and without MC sampling, Bernoulli Dropout (BDO) with and without MC sampling, and Gaussian Dropout with and without MC sampling.

3.1 FULLY CONNECTED NEURAL NETWORKS

First, we trained DNNs with two FC hidden layers, each with 800 units and ReLU non-linearities. For the L2-regularized network, an L2-coefficient of $1e-5$ was used for all weights. For Bernoulli and Gaussian dropout, dropout was performed after each FC layer. For Bernoulli and Gaussian dropconnect, every weight was sampled. The classification error of the networks on the MNIST test set is shown in Table 1. All of the methods with sampling during learning are significantly more accurate than the baseline NN and NN+L2 networks, with the dropconnect-based networks the most accurate. However, MC sampling at inference did not significantly increase the networks’ accuracy.

The classification error, uncertainty, and calibration of the learned probability distribution of each network for varying levels of noise are shown in Figure 4. While not improving accuracy, MC sampling did lead to networks that better represent their own uncertainty. As the noise in the test set was increased, the uncertainty of the networks with MC sampling highly increased, especially when compared to the networks without sampling at inference. This resulted in better calibrated models for all levels of noise.

As can be seen by also taking into account the the calibration curves (Figure 5), sampling only during training, especially when using dropout, led to overconfidence through placing too much probability mass on the most predicted label. In particular, sampling only during training resulted in under-confidence for low predicted probabilities and over-confidence for high predicted probabilities. By distributing probability mass over several labels, the DNNs that sampled at inference better represented the uncertainty of their predictions.

Table 2: The trained convolutional neural network (CNN) architecture.

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
Linear	500	10	-	-

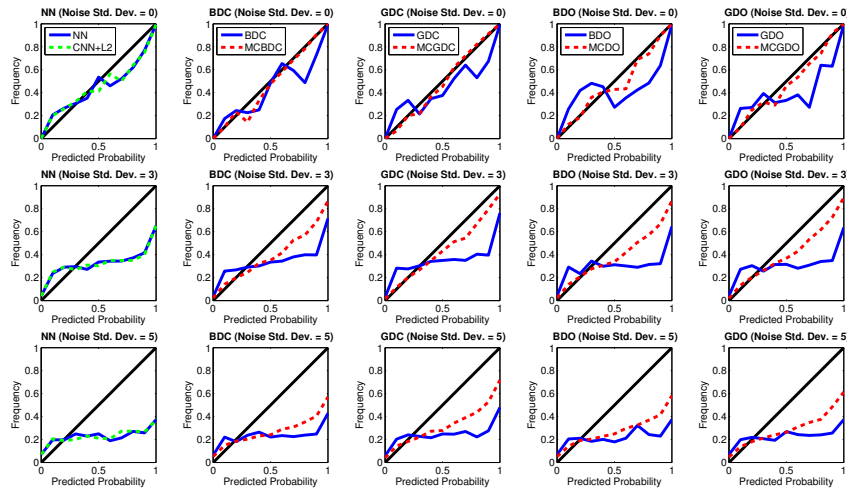


Figure 5: The calibration curves for the MNIST test set with and without Gaussian noise of the softmax outputs of the fully connected networks: NN, NN+L2, Bernoulli DropConnect (BDC) with and without Monte-Carlo (MC) sampling, Gaussian DropConnect (GDC) with and without MC sampling, Bernoulli Dropout (BDO) with and without MC sampling, and Gaussian Dropout with and without MC sampling.

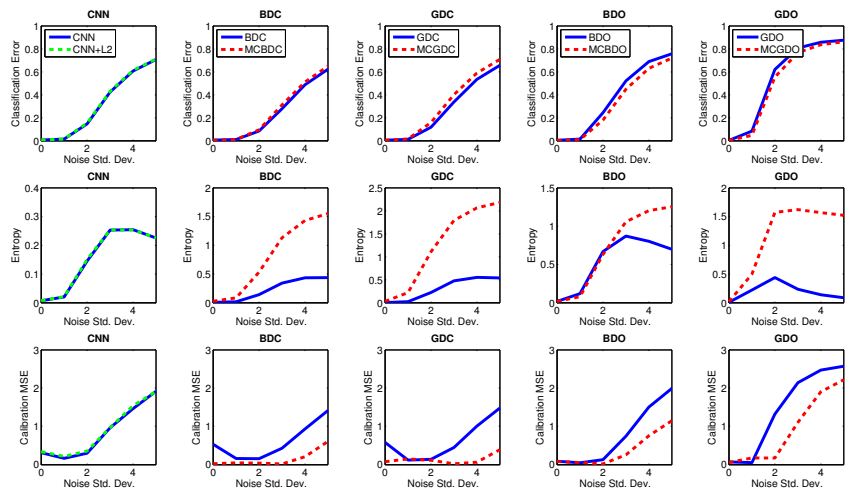


Figure 6: The MNIST test classification error, entropy, and calibration of the predictions of the convolutional networks: CNN, CNN+L2, Bernoulli DropConnect (BDC) with and without Monte-Carlo (MC) sampling, Gaussian DropConnect (GDC) with and without MC sampling, Bernoulli Dropout (BDO) with and without MC sampling, and Gaussian Dropout with and without MC sampling.

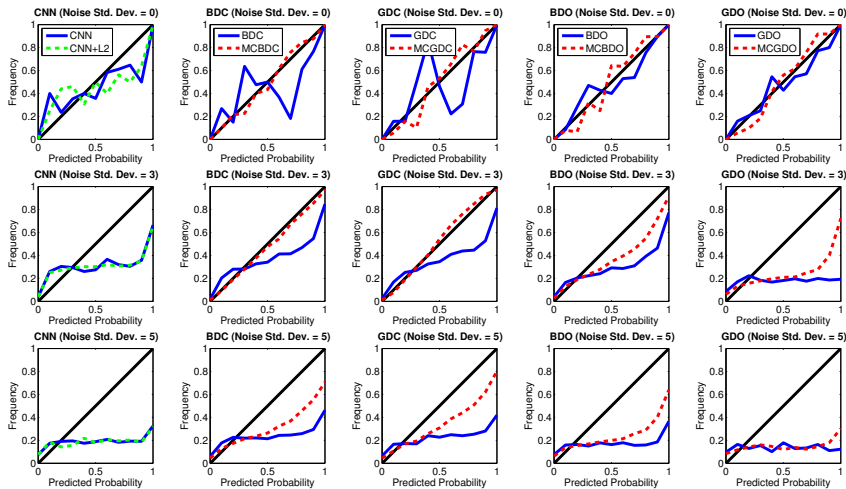


Figure 7: The calibration curves for the MNIST test set with and without Gaussian noise of the softmax outputs of the convolutional networks: CNN, CNN+L2, Bernoulli DropConnect (BDC) with and without Monte-Carlo (MC) sampling, Gaussian DropConnect (GDC) with and without MC sampling, Bernoulli Dropout (BDO) with and without MC sampling, and Gaussian Dropout with and without MC sampling.

3.1.1 CONVOLUTIONAL NEURAL NETWORKS

We also trained CNNs with the architecture shown in Table 2. For the L2-regularized network, an L2-coefficient of $1e-5$ was used for all weights. For Bernoulli and Gaussian dropout, dropout was performed after each convolutional layer and after the FC layer. For Bernoulli and Gaussian dropconnect, every weight was sampled. The classification error of the networks on the MNIST test set is shown in Table 3. Sampling during training significantly increased the accuracy for the all of the networks, but especially for the Gaussian dropout network. However, unlike for the FC networks, the dropout-based methods were more accurate than the dropconnect-based methods. As with the FC networks, MC sampling during inference did not significantly increase the accuracy of the networks.

The classification error, uncertainty, and calibration of the learned probability distribution of each network for varying levels of noise are shown in Figure 6. As with the FC networks, MC sampling at inference greatly increased the CNNs’ ability to estimate their own uncertainty, particularly for inputs that are different from the training set. MC sampling led to increased entropy as inputs became more noisy, which resulted in better calibration. In particular, this is true of both the Bernoulli and Gaussian dropconnect networks, which very accurately represented their uncertainty even for

Table 3: MNIST test error for the trained convolutional neural networks (CNNs) with and without Monte-Carlo (MC) sampling using 100 samples.

Method	Mean Error (%)	Error Std. Dev.
CNN	0.70	-
CNN+L2	0.70	-
Bernoulli DropConnect	0.59	-
MC Bernoulli DropConnect	0.59	0.02
Gaussian DropConnect	0.49	-
MC Gaussian DropConnect	0.49	0.01
Bernoulli Dropout	0.45	-
MC Bernoulli Dropout	0.46	0.01
Gaussian Dropout	0.38	-
MC Gaussian Dropout	0.37	0.01

highly noisy inputs. As shown in the calibration curves (Figure 7), not using MC sampling resulted in networks that were under-confident when making low probability predictions and over-confident when making high probability predictions.

4 DISCUSSION

In this paper, we investigated the ability of MC sampling to improve DNNs’ ability to model their own uncertainty. We did this by training Bayesian DNNs with either dropconnect or dropout and either Bernoulli or Gaussian sampling. Based on the results, we draw the following main conclusions:

1. Sampling during learning improved a network’s ability to represent its own uncertainty

MC sampling at inference improved the calibration of a network’s predictions. This improvement was particularly large for inputs from outside the training set, which traditional models classified with high confidence despite not being trained on similar inputs. This is an interesting finding for machine learning research, but also for computational neuroscience. Showing that sampling at inference can allow a DNN to better represent its own uncertainty by stochastically varying activation patterns supports the neural sampling hypothesis.

2. Sampling weights independently led to the most accurate FC networks, but sampling units led to the most accurate CNNs

For the FC networks, dropconnect sampling, particularly Gaussian dropconnect, resulted in the most accurate networks. However, dropout sampling led to the most accurate CNNs. A potential cause of this is the large correlation in the information contained by the image patches covered by a convolutional kernel. This could mean that sampling the weights of a kernel does not provide as much regularization as the dropout-based methods.

These scalable methods for estimating a network’s uncertainty are widely applicable, since most DNNs already use dropout and getting uncertainty estimates only requires using MC sampling at inference. We plan to further investigate the use of different variational distributions. We also plan to evaluate the use of dropout and dropconnect sampling on large recurrent neural networks. In conclusion, our results demonstrate that sampling at inference allows DNNs to efficiently represent their own uncertainty, an essential part of real-world perception and decision making.

ACKNOWLEDGMENTS

We would like to thank Yarin Gal and Sergii Strelchuk for their helpful discussions regarding the manuscript. This research was funded by the Cambridge Commonwealth, European & International Trust, the UK Medical Research Council (Program MC-A060-5PR20), and a European Research Council Starting Grant (ERC-2010-StG 261352).

REFERENCES

- David Barber and Christopher M Bishop. Ensemble learning in bayesian neural networks. *NATO ASI SERIES F COMPUTER AND SYSTEMS SCIENCES*, 168:215–238, 1998.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *Proceedings of The 32nd International Conference on Machine Learning*, pp. 1613–1622, 2015.
- Lars Buesing, Johannes Bill, Bernhard Nessler, and Wolfgang Maass. Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons. *PLoS Comput Biol*, 7(11):e1002211, 2011.
- József Fiser, Pietro Berkes, Gergő Orbán, and Máté Lengyel. Statistically optimal perception and learning: from behavior to neural representations. *Trends in cognitive sciences*, 14(3):119–130, 2010.
- Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Insights and applications. In *Deep Learning Workshop, ICML*, 2015.

- 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.
- Alex Graves. Practical variational inference for neural networks. In *Advances in Neural Information Processing Systems*, pp. 2348–2356, 2011.
- Stefan Habenschuss, Zeno Jonke, and Wolfgang Maass. Stochastic computations in cortical micro-circuit models. *PLoS Comput Biol*, 9(11):e1003311, 2013.
- Geoffrey Hinton. A practical guide to training restricted boltzmann machines. *Momentum*, 9(1): 926, 2010.
- 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*, pp. 5–13. ACM, 1993.
- Diederik P Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick. In *Advances in Neural Information Processing Systems*, pp. 2575–2583, 2015.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- 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.
- David JC MacKay. A practical bayesian framework for backpropagation networks. *Neural computation*, 4(3):448–472, 1992.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- 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.
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- 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)*, pp. 1058–1066, 2013.