
Notes on the Behavior of MC Dropout

Francesco Verdoja¹ Ville Kyrki¹

Abstract

Among the various options to estimate uncertainty in deep neural networks, Monte-Carlo dropout is widely popular for its simplicity and effectiveness. However the quality of the uncertainty estimated through this method varies and choices in architecture design and in training procedures have to be carefully considered and tested to obtain satisfactory results. In this paper we present study offering a different point of view on the behavior of Monte-Carlo dropout, which enable us to observe a few interesting properties of the technique to keep in mind when considering its use for uncertainty estimation.

1. Introduction

The increasing interest in the deployment of deep learning solutions in real safety-critical applications ranging from healthcare to robotics and autonomous vehicles is making apparent the importance to properly estimate the uncertainty of the predictions made by deep neural networks (Kendall et al., 2017; Loquercio et al., 2020).

While most common neural network architectures only provide point estimates, uncertainty can be evaluated with Bayesian neural networks (BNNs) (Denker & LeCun, 1990; MacKay, 1992) where the deterministic weights used in the majority of neural networks are replaced with distributions over the network parameters. Although the formulation of BNNs is relatively easy in theory, their use in practise for most complex problems is often unfeasible due to their need to analytically evaluate the marginal probabilities during training which becomes quickly intractable. Recently, variational inference methods have been proposed as a practical alternative to BNNs, but most of these formulations requires double the number of parameters of a network to

This work was supported by the Strategic Research Council at Academy of Finland, decision 314180.

¹School of Electrical Engineering, Aalto University, Maarintie 8, Espoo, Finland. Correspondence to: Francesco Verdoja <francesco.verdoja@aalto.fi>.

Presented at the ICML 2021 Workshop on Uncertainty and Robustness in Deep Learning., Copyright 2021 by the author(s).

represent its uncertainty which leads to increased computational costs (Blundell et al., 2015; Gal & Ghahramani, 2016).

Another very popular option to model uncertainty in deep neural networks is the use of dropout as a way to approximate Bayesian variational inference (Gal & Ghahramani, 2016). The simplicity of the key idea of this formulation is one the main reasons for its popularity: by enabling dropout not only in training but also during testing, and by doing several forward passes through the network with the same input data, one can use the distribution of the outputs of the different passes to estimate the first two moments (mean and variance) of the predictive distribution. The mean is then used as the estimate, and the variance as a measure of its uncertainty. This technique, called Monte-Carlo dropout (MCD), has proved effective to, *e.g.*, increase visual relocalization accuracy (Kendall & Cipolla, 2016) and semantic segmentation performance on images (Kendall et al., 2017). Despite its success and simplicity however, it has been recognized that the quality of the uncertainty estimates is tied to parameter choices which need to be calibrated to match the uncertainty (Osband et al., 2016; Pearce et al., 2018; Boluki et al., 2020; Caldeira & Nord, 2020). However, when using MCD in practical applications, architectural choices like where to insert the dropout layers, how many to use, and the choice of dropout rate are often either empirically made or set a priori (Kendall et al., 2017; Jungo et al., 2018; Verdoja et al., 2019), leading to possibly suboptimal performance.

In this work, we conduct a study providing some observations both in theory and through experiments over the behavior of MCD. The main contributions of this work are a theoretical analysis over the behavior of MCD on a simple single-layer linear network, extending and correcting the discussion in (Osband, 2016), and an experimental demonstration that the properties found in theory apply to bigger non-linear networks as well. In the discussion, we offer different intuitions over architecture design and training choices for networks using MCD for uncertainty estimation.

2. Behavior of MC Dropout

In this section we expand and correct the intuitions first presented in (Osband, 2016) over the behavior of MCD. In that work, the author conducted a similar analysis but

commented on the results only partially. Here we correct some imprecisions in the way they computed the optimal weights and analyse the results further. To this end, let us consider a single-layer linear network

$$f = \sum_{k=1}^K d_k w_k \quad (1)$$

with weights $w_k \in \mathbb{R}$ and dropouts $d_k \sim \text{Ber}(p)$, *i.e.*, a Bernoulli distribution with success probability p . To note that here p refers to $P(d_k = 1)$, while most dropout implementations require a dropout probability parameter $p_d = P(d_k = 0) = 1 - p$.

Assuming all weights to converge to the same value w , which is to be expected when using dropout, then

$$\begin{aligned} \mathbb{E}[f] &= \mathbb{E}\left[w \sum_{k=1}^K d_k\right] = w \mathbb{E}[\text{B}(K, p)] = \\ &= wKp \\ \text{Var}[f] &= \text{Var}\left[w \sum_{k=1}^K d_k\right] = w^2 \text{Var}[\text{B}(K, p)] = \\ &= w^2 Kp(1-p) \end{aligned} \quad (2)$$

where $\text{B}(K, p)$ is a binomial distribution with K trials and success probability p .

Following (Osband, 2016), given a ground-truth $\{y_1, \dots, y_n\}$ with average $\bar{y} := \sum_{i=1}^n y_i/n$, minimizing the mean squared error (MSE) from the ground-truth means finding the minimum of

$$\begin{aligned} \mathbb{E}[(f - \bar{y})^2] &= \mathbb{E}[f^2 - 2f\bar{y} + \bar{y}^2] \\ &= \mathbb{E}[f^2] - 2\bar{y}\mathbb{E}[f] + \bar{y}^2 \\ &= \text{Var}[f] + \mathbb{E}^2[f] - 2\bar{y}\mathbb{E}[f] + \bar{y}^2 \\ &= w^2 Kp(1-p) + w^2 K^2 p^2 - 2\bar{y}wKp + \bar{y}^2 \\ &= w^2 Kp(Kp - p + 1) - 2\bar{y}wKp + \bar{y}^2 \end{aligned} \quad (3)$$

This can be achieved by solving

$$\begin{aligned} \frac{d}{dw} \mathbb{E}[(f - \bar{y})^2] &= 0 \\ 2wKp(Kp - p + 1) - 2\bar{y}Kp &= 0 \\ w(Kp - p + 1) - \bar{y} &= 0 \end{aligned} \quad (4)$$

which means the optimal weight is

$$w = \frac{\bar{y}}{Kp - p + 1} = \frac{\bar{y}}{K(1 - p_d) + p_d} \quad (5)$$

Consequently, combining Eq. (2) and Eq. (5), we find that

Table 1. Comparison of single-layer network experiment results versus the theoretical expectation from Eqs. (5) and (6)

p_d	Dataset	Theoretical		Experimental	
		w	$\text{Var}[f]$	w	$\text{Var}[f]$
0.2	\mathbf{Y}'	0.025	0.050	0.025	0.058
0.2	\mathbf{Y}''	0.025	0.050	0.026	0.076
0.5	\mathbf{Y}'	0.040	0.199	0.040	0.214
0.5	\mathbf{Y}''	0.040	0.199	0.040	0.249

at convergence:

$$\begin{aligned} \mathbb{E}[f] &= \frac{Kp\bar{y}}{Kp - p + 1} = \frac{K(1 - p_d)\bar{y}}{K(1 - p_d) + p_d} \\ \text{Var}[f] &= \frac{Kp(1-p)\bar{y}^2}{(Kp - p + 1)^2} = \frac{Kp_d(1 - p_d)\bar{y}^2}{(K - Kp_d + p_d)^2} \end{aligned} \quad (6)$$

From Eq. (6), a few observation can be drawn:

1. while one would expect $\mathbb{E}[f] = \bar{y}$, the expected output of the network is actually introducing a bias. For big network however, this bias is negligible, since $Kp \approx Kp - p + 1$;
2. the size of the variance of the posterior distribution generated by MCD on this simple network depends on the interaction between the dropout rate p_d and the model size K ;
3. the posterior distribution has no dependence on the amount of data n , nor the observed variance in the data, which means that it does not concentrate as more data is gathered;
4. finally, the variance of the posterior distribution is proportional to \bar{y}^2 , *i.e.*, the bigger the value to be estimated, the bigger the estimated model uncertainty.

While we demonstrate here in theory that these properties of MCD exist on a very simple network, proving them following a similar theoretical process on bigger networks becomes quickly unfeasible.

For this reason, to try to understand how these interaction work on bigger more realistic networks, we run different experiments, which we will present in the following section.

3. Experiments

To verify that the properties found in Section 2 hold true even for bigger non-linear networks, we created the suite of experiments we report in this section. These experiments have been implemented in *pytorch v0.3.0*¹.

¹Source code for the experiments presented here is available at: github.com/aalto-intelligent-robotics/mc-dropout-notebooks

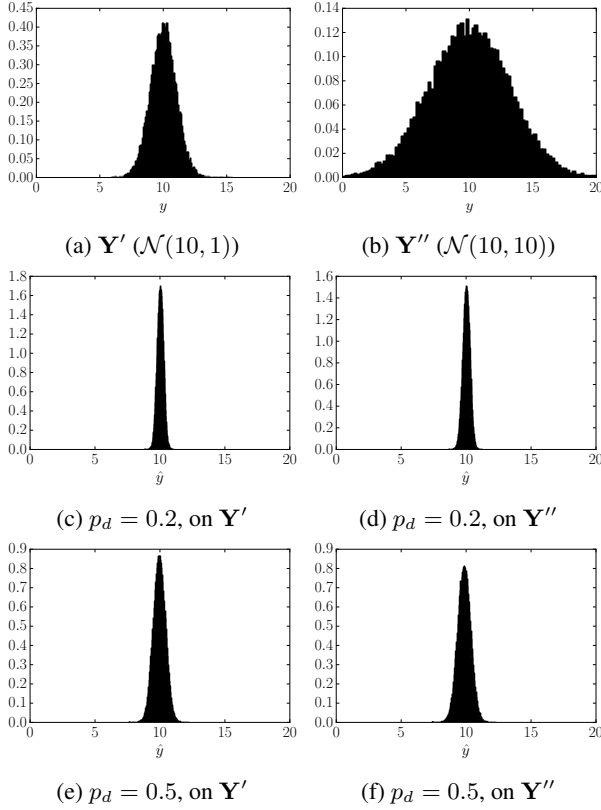


Figure 1. Ground-truth distributions (a, b) and corresponding MCD output distributions for single-layer linear networks with $p_d = 0.2$ (c, d) and $p_d = 0.5$ (e, f) respectively.

3.1. Single-layer linear network

In this first experiment we want to empirically verify the theoretical behavior observed in Section 2. In particular, we implemented two versions of the single-layer linear network in Eq. (1) with $K = 500$, one with $p_d = 0.2$ and one with $p_d = 0.5$ respectively. Once more, note that the dropout probability $p_d = 1 - p$. We created two ground-truth datasets $\mathbf{Y}' = \{y'_1, \dots, y'_n\}$ and $\mathbf{Y}'' = \{y''_1, \dots, y''_n\}$ with $n = 3200$, $y'_i \sim \mathcal{N}(10, 1)$, and $y''_i \sim \mathcal{N}(10, 10)$. We trained both networks on both datasets for 600 epochs with MSE loss and adam optimizer. As a side note, training with MSE holds the same results as training with log-likelihood, since the minimum of the two losses is in the same position in this case. After training, we run multiple forward passes of the networks with dropout active to obtain one million samples to empirically estimate $\text{Var}[f]$.

In Table 1, we can see that the values for w and $\text{Var}[f]$ found experimentally match the theoretical ones computed using Eqs. (5) and (6), aside for noise introduced by the sampling process. Moreover, they confirm how changing the parameter p_d affects the model uncertainty, while changing observed variance in the data (from $\sigma = 1$ in \mathbf{Y}' , to $\sigma = 10$ in \mathbf{Y}'') does not. This phenomenon is also clearly visible

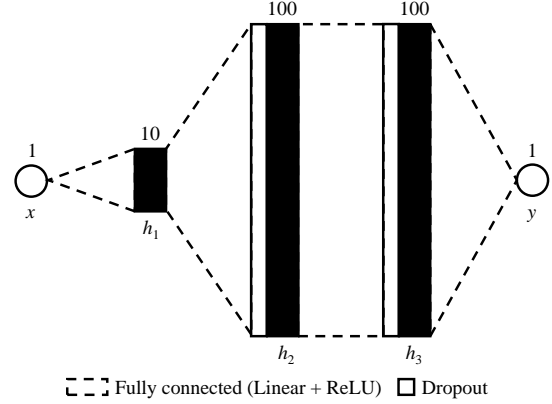


Figure 2. Fully connected network used in the experiments.

in Fig. 1 where it can be noticed how the variance of the output distribution of MCD is impacted by p_d (compare Figs. 1c and 1e, and Figs. 1d and 1f) while being unaffected by the variance in the dataset it trained on (compare Figs. 1c and 1d, and Figs. 1e and 1f).

3.2. Non-linear networks

To verify that effects similar to those presented in the previous section can still be observed after introducing more layers and non-linearities, we performed a similar experiment on the fully connected neural network shown in Fig. 2. While this network is still smaller than most real networks, we decided for this size because it had enough complexity while still being simple enough to analyze. In all the following experiments we trained for 1000 epochs with MSE loss and adam optimizer on a dataset of 32000 samples. We trained two variants of the network with dropout rate $p_d = 0.2$ and $p_d = 0.5$ respectively. As it can be seen in the figure, for these experiments the network has one input x and produces one output y , effectively approximating a function $f : \mathbb{R} \rightarrow \mathbb{R}$.

Fig. 3 shows the results produced by the network on a noisy constant function where for each input $x \in [0, 1]$ multiple possible outputs y are present in the dataset while $\forall x, \mathbb{E}[y] = 0.5$ (Fig. 3a). We first show the behavior of the network when the last linear layer has a bias term in Fig. 3b. It can be seen that the network is nullifying the variance completely. It is indeed setting all weights to 0 and encoding the desired (constant) output in the bias. If we remove the bias from the last linear layer (Figs. 3c and 3d) we obtain constant variance, proportional to the dropout rate.

Fig. 4 shows another experiment, conducted on different functions. In this case the datasets have no noise, *i.e.*, for each x , all samples in the dataset have the same corresponding y . All the functions and network variants we tested are available as additional material, however, due to space limitations, in here we present only two functions (Figs. 4a

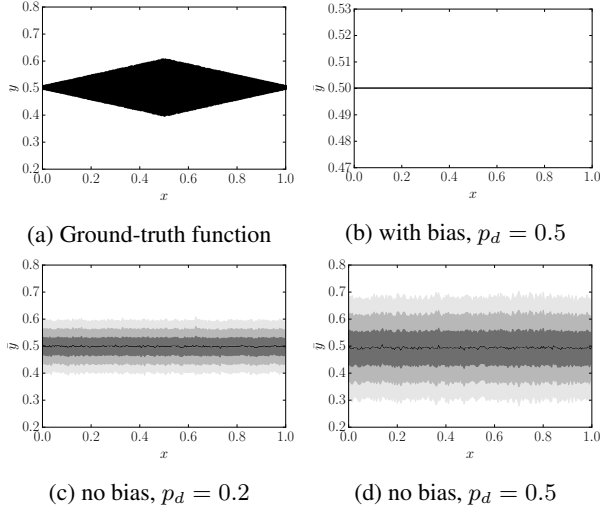


Figure 3. Results obtained by training a non-linear network on the dataset shown in (a); different variants of the network have been trained: one with a bias term in the last linear layer (b) and two without bias in the last linear layer and with different dropout rates (c, d). In (a) each dot represents a datapoint; in (b, c, d) the line represents the average output of 300 forward passes through the network, with the shaded areas representing σ , 2σ , and 3σ respectively.

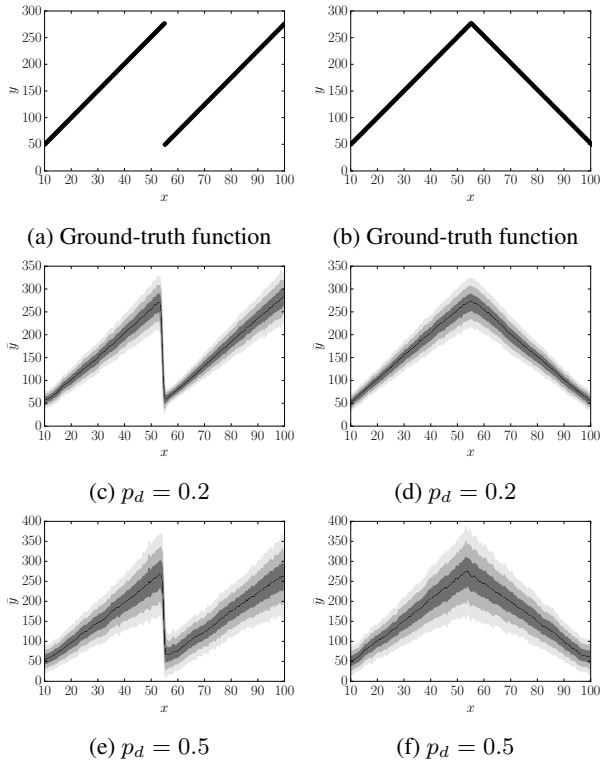


Figure 4. Results obtained by training a non-linear network on the dataset shown in (a, b); two variants of the network have been trained: with dropout rate $p_d = 0.2$ (c,d) and $p_d = 0.5$ (e, f). In (a, b) each dot represents a datapoint; in (c, d, e, f) the line represents the average output of 300 forward passes through the network, with the shaded areas representing σ , 2σ , and 3σ respectively.

and 4b) and only the results for networks where the bias has been disabled on the last linear layer. All results not reported here lead to the same conclusions. From these results, the last of the properties presented in Section 2 can be noticed, namely the fact that the size of the uncertainty scales proportionally to the value of y . These also confirm one more time the dependence of the uncertainty from the dropout parameter p_d .

4. Discussion

The results presented in Section 3 confirm that the behaviors of MCD that we theoretically observed on a single-layer linear network are still present even as the size and complexity of the network grows. When looking at these results, we are able to make a few observations.

First of all, all experiments provided an additional demonstration that the choice of dropout rate p_d is crucial and that in itself the epistemic uncertainty estimate provided by MCD is not affected by the amount of data available during training or its variance, which is a good reminder that the uncertainty estimate produced by MCD is not calibrated and that the dropout rate has to be adjusted to match the epistemic uncertainty. In practise, this is usually done through grid search, although several methods have been proposed to learn the parameter during training, *e.g.*, (Gal et al., 2017; Phan et al., 2018; Boluki et al., 2020).

Another interesting aspect that emerged from our study was the scaling of the uncertainty based on the value of the network output. This effect, if not taken into account, can lead to degraded quality of the uncertainty estimates, which we empirically noticed in the past in a practical robotic application (Verdoja et al., 2019). This effect is particularly visible when dropout is applied before the last linear layer of the network, as is the case in the the experiments presented here. If more non-linear layers are added after the last dropout layer the non-linearities can reduce this effect. This would suggest that in applications with widely varying outputs, it could be useful to limit dropout only to the inner parts of the network, justifying what other studies have found experimentally to helps improve performance (Kendall et al., 2017; Jungo et al., 2018).

5. Conclusions

In this paper we presented a different point of view on the behavior of MCD that enabled us to show both theoretically and experimentally different properties of that approach, in particular the dependency of the uncertainty estimates from the dropout rate and from the value the network is trying to estimate. Our intent was to propose a framework to demonstrate these property in a very understandable fashion, to help in both architecture design and training procedures when using MCD.

References

- Blundell, C., Cornebise, J., Kavukcuoglu, K., and Wierstra, D. Weight Uncertainty in Neural Network. In *International Conference on Machine Learning*, pp. 1613–1622, June 2015.
- Boluki, S., Ardywibowo, R., Dadaneh, S. Z., Zhou, M., and Qian, X. Learnable Bernoulli Dropout for Bayesian Deep Learning. *arXiv:2002.05155 [cs, stat]*, February 2020.
- Caldeira, J. and Nord, B. Deeply Uncertain: Comparing Methods of Uncertainty Quantification in Deep Learning Algorithms. *arXiv:2004.10710 [physics, stat]*, April 2020.
- Denker, J. S. and LeCun, Y. Transforming Neural-net Output Levels to Probability Distributions. In *Proceedings of the 3rd International Conference on Neural Information Processing Systems*, NIPS’90, pp. 853–859, San Francisco, CA, USA, 1990. Morgan Kaufmann Publishers Inc. ISBN 978-1-55860-184-0.
- Gal, Y. and Ghahramani, Z. Dropout As a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning*, volume 48 of *ICML’16*, pp. 1050–1059, New York, NY, USA, 2016.
- Gal, Y., Hron, J., and Kendall, A. Concrete Dropout. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 3581–3590. Curran Associates, Inc., 2017.
- Jungo, A., McKinley, R., Meier, R., Knecht, U., Vera, L., Pérez-Beteta, J., Molina-García, D., Pérez-García, V. M., Wiest, R., and Reyes, M. Towards Uncertainty-Assisted Brain Tumor Segmentation and Survival Prediction. In Crimi, A., Bakas, S., Kuijf, H., Menze, B., and Reyes, M. (eds.), *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*, Lecture Notes in Computer Science, pp. 474–485, Cham, 2018. Springer International Publishing. ISBN 978-3-319-75238-9. doi: 10.1007/978-3-319-75238-9_40.
- Kendall, A. and Cipolla, R. Modelling uncertainty in deep learning for camera relocalization. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4762–4769, May 2016. doi: 10.1109/ICRA.2016.7487679.
- Kendall, A., Badrinarayanan, V., and Cipolla, R. Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. *British Machine Vision Conference 2017, BMVC 2017*, January 2017.
- Loquercio, A., Segù, M., and Scaramuzza, D. A General Framework for Uncertainty Estimation in Deep Learning. *IEEE Robotics and Automation Letters*, 5(2):3153–3160, April 2020. ISSN 2377-3766, 2377-3774. doi: 10.1109/LRA.2020.2974682.
- MacKay, D. J. C. A Practical Bayesian Framework for Backpropagation Networks. *Neural Computation*, 4(3): 448–472, May 1992. ISSN 0899-7667. doi: 10.1162/neco.1992.4.3.448.
- Osband, I. Risk versus Uncertainty in Deep Learning: Bayes, Bootstrap and the Dangers of Dropout. In *NIPS 2016 Workshop on Bayesian Deep Learning*, Barcelona, Spain, October 2016.
- Osband, I., Blundell, C., Pritzel, A., and Van Roy, B. Deep Exploration via Bootstrapped DQN. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29 (NIPS 2016)*, pp. 4026–4034. Curran Associates, Inc., 2016.
- Pearce, T., Anastassacos, N., Zaki, M., and Neely, A. Bayesian Inference with Anchored Ensembles of Neural Networks, and Application to Reinforcement Learning. *arXiv:1805.11324 [cs, stat]*, May 2018.
- Phan, B., Salay, R., Czarnecki, K., Abdelzad, V., Denouden, T., and Vernekar, S. Calibrating Uncertainties in Object Localization Task. *arXiv:1811.11210 [cs, stat]*, November 2018.
- Verdoja, F., Lundell, J., and Kyrki, V. Deep Network Uncertainty Maps for Indoor Navigation. In *2019 IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, Toronto, Canada, October 2019.

Notes on the Behavior of MC Dropout

Additional results obtained by training a non-linear network on different datasets (row a). Different variants of the network have been trained: one with $p_d = 0.2$ and bias term in the last linear layer (rows b, c), one with $p_d = 0.2$ without bias term in the last linear layer (rows d, e), one with $p_d = 0.5$ and bias term in the last linear layer (rows f, g), and one with $p_d = 0.5$ without bias term in the last linear layer (rows h, i). In (rows a, d, h) each dot represents a datapoint; in (rows c, e, g, i) the line represents the average output of 300 forward passes through the network, with the shaded areas representing σ , 2σ , and 3σ respectively.

