
On the Importance of Strong Baselines in Bayesian Deep Learning

Jishnu Mukhoti

Department of Computer Science
University of Oxford
jishnu.mukhoti@cs.ox.ac.uk

Pontus Stenetorp

Department of Computer Science
University College London
p.stenetorp@cs.ucl.ac.uk

Yarin Gal

Department of Computer Science
University of Oxford
yarin@cs.ox.ac.uk

Abstract

Like all sub-fields of machine learning Bayesian Deep Learning is driven by empirical validation of its theoretical proposals. Given the many aspects of an experiment it is always possible that minor or even major experimental flaws can slip by both authors and reviewers. One of the most popular experiments used to evaluate approximate inference techniques is the regression experiment on UCI datasets. However, in this experiment, models which have been trained to convergence have often been compared with baselines trained only for a fixed number of iterations. We find that a well-established baseline, Monte Carlo dropout, when evaluated under the same experimental settings shows significant improvements. In fact, the baseline outperforms or performs competitively with methods that claimed to be superior to the very same baseline method when they were introduced. Hence, by exposing this flaw in experimental procedure, we highlight the importance of using identical experimental setups to evaluate, compare, and benchmark methods in Bayesian Deep Learning.

1 Introduction

Empiricism is at the very core of machine learning research, where we demand that new methods and approaches compare favorably to previously introduced work. We expect this in terms of performance on either artificially generated data that highlights specific challenges and/or real-world data for specific tasks. In this process, we implicitly rely on fellow scientists and reviewers to note discrepancies (intentional or not, to err is after all human) in the experimental setting – such as for example, any kind of overfitting.¹ Recently, several studies have noted empirical shortcomings in the machine learning literature. For example, Henderson et al. (2017) observed that due to non-determinism, variance, and lack of significance metrics, it is difficult to judge whether claimed advances in reinforcement learning are empirically justified. Also, Melis et al. (2018) established that several years of assumed progress in language modeling did not in fact improve upon a standard stacked LSTM model if the hyperparameters for all models were adequately tuned.

Bayesian Deep Learning applies the ideas of Bayesian inference to deep networks and is an active area of machine learning research. Popular techniques for approximate inference in deep networks

¹<http://hunch.net/?p=22>

include variational inference (VI) (Graves, 2011), probabilistic backpropagation (PBP) (Hernández-Lobato and Adams, 2015) for Bayesian neural nets, dropout as an interpretation of approximate Bayesian inference (Gal and Ghahramani, 2016), Deep Gaussian Processes (DGP) as a multi-layer generalization of Gaussian Processes (Bui et al., 2016), Bayesian neural networks using Variational Matrix Gaussian (VMG) posteriors (Louizos and Welling, 2016), and variants of Stochastic Gradient Hamiltonian Monte Carlo (SGHMC) methods (Springenberg et al., 2016).

A well-defined experimental setup is necessary to compare and benchmark these methods; one of the most popular setups is a regression experiment on a number of curated UCI datasets (Hernández-Lobato and Adams, 2015). The choice of this experiment is appealing for Bayesian neural nets because it provides predictive log-likelihood in addition to RMSE as an evaluation metric. As such, predictive log-likelihood can be used to judge the quality of the uncertainty estimates produced by the model (Gal and Ghahramani, 2016).

In this study, we observe a discrepancy in the setup used for the UCI regression experiment in some recent works, in that they compare their model to baselines obtained under a different experimental setting. Concretely, this applies to VMG (Louizos and Welling, 2016), HS-BNN (Ghosh and Doshi-Velez, 2017), PBP-MV, (Sun et al., 2017) and SGHMC (Springenberg et al., 2016). In order to gauge the impact of this erroneous comparison, we reevaluate the regression experiments of the above-mentioned works compared to MC dropout (Gal and Ghahramani, 2016) in the same setting. The experimental results indicate that the networks with dropout inference – when trained under the same conditions – outperform VMG, HS-BNN, and SGHMC; and they are a close second to PBP-MV. These results suggest that several methods, when introduced, erroneously claimed state of the art performance in their publications.

2 Regression Experiments

We perform the non-linear regression experiments proposed in Hernández-Lobato and Adams (2015), which have been adopted to evaluate approximate inference techniques in numerous subsequent works: Gal and Ghahramani (2016), Bui et al. (2016), Louizos and Welling (2016), Ghosh and Doshi-Velez (2017), Sun et al. (2017), Springenberg et al. (2016) etc. All the datasets from Hernández-Lobato and Adams (2015) are used, apart from the *YearPredictionMSD* dataset. The *YearPredictionMSD* dataset is very large with 515,345 instances, each of which has 90 dimensions. Hence, tuning network hyperparameters over this dataset requires an inordinate amount of time. Our network architecture has a single hidden layer with 50 hidden units for all datasets, except for the *Protein Structure* dataset for which there are 100 hidden units.

There are two ways in which this experiment has been conducted in the past. In the first, the networks are trained for a fixed number of iterations (specifically, 40 epochs), and the average training time of the networks are noted and compared. This setting was used by Hernández-Lobato and Adams (2015), Gal and Ghahramani (2016), and Bui et al. (2016). In this work, we refer to this as the *timed setting* of the experiment. In the second variant of the experiment, the networks are trained to convergence with a higher number of training iterations and was used by Louizos and Welling (2016), Ghosh and Doshi-Velez (2017), Sun et al. (2017), Springenberg et al. (2016), and others. For both variants, the test set RMSE and log-likelihood values are used as the evaluation metrics.

Given these two settings, it is a natural question to ask if models trained under the timed setting and those trained to convergence are comparable. One might argue that networks trained for a fixed number of iterations might not have converged and will thus perform poorly compared to those which have been trained to convergence. To test this hypothesis, we use networks with MC dropout (Gal and Ghahramani, 2016) – one of the de facto standard baselines for approximate inference.

We compare with the following works: VMG (Louizos and Welling, 2016), where models trained to convergence have been benchmarked against MC dropout baselines from the timed setting; Bayesian networks with horseshoe priors (HS-BNN) (Ghosh and Doshi-Velez, 2017) and probabilistic backpropagation with the Matrix Variate Gaussian (MVG) distribution (PBP-MV) (Sun et al., 2017), where the models were only benchmarked against VMG; Stochastic Gradient Hamiltonian Monte Carlo methods (SGHMC) (Springenberg et al., 2016), where the results have been compared with Probabilistic Backpropagation (PBP) (Hernández-Lobato and Adams, 2015) and Variational Inference (VI) (Graves, 2011) baselines obtained under the timed setting.

Table 1: **Average RMSE test performance.** The RMSE values along with corresponding standard errors are presented.

Dataset	Dropout (Timed Setting)	Dropout (Convergence)	Dropout (Hyperparameter tuning)	VMG	HS-BNN	PBP-MV
Boston Housing	2.97 ± 0.19	2.83 ± 0.17	2.90 ± 0.18	2.70 ± 0.13	3.32 ± 0.66	3.11 ± 0.15
Concrete Strength	5.23 ± 0.12	4.93 ± 0.14	4.82 ± 0.16	4.89 ± 0.12	5.66 ± 0.41	5.08 ± 0.14
Energy Efficiency	1.66 ± 0.04	1.08 ± 0.03	0.54 ± 0.06	0.54 ± 0.02	1.99 ± 0.34	0.45 ± 0.01
Kin8nm	0.10 ± 0.00	0.09 ± 0.00	0.08 ± 0.00	0.08 ± 0.00	0.08 ± 0.00	0.07 ± 0.00
Naval Propulsion	0.01 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
Power Plant	4.02 ± 0.04	4.00 ± 0.04	4.01 ± 0.04	4.04 ± 0.04	4.03 ± 0.15	3.91 ± 0.04
Protein Structure	4.36 ± 0.01	4.27 ± 0.01	4.27 ± 0.02	4.13 ± 0.02	4.39 ± 0.04	3.94 ± 0.02
Wine Quality Red	0.62 ± 0.01	0.61 ± 0.01	0.62 ± 0.01	0.63 ± 0.01	0.63 ± 0.04	0.64 ± 0.01
Yacht Hydrodynamics	1.11 ± 0.09	0.70 ± 0.05	0.67 ± 0.05	0.71 ± 0.05	1.58 ± 0.23	0.81 ± 0.06

Table 2: **Average log likelihood test performance.** The log likelihood values along with corresponding standard errors are presented.

Dataset	Dropout (Timed Setting)	Dropout (Convergence)	Dropout (Hyperparameter tuning)	VMG	HS-BNN	PBP-MV	SGHMC (Tuned per dataset)	SGHMC (Scale Adapted)
Boston Housing	-2.46 ± 0.06	-2.40 ± 0.04	-2.40 ± 0.04	-2.46 ± 0.09	-2.54 ± 0.15	-2.54 ± 0.08	-2.49 ± 0.15	-2.54 ± 0.04
Concrete Strength	-3.04 ± 0.02	-2.97 ± 0.02	-2.93 ± 0.02	-3.01 ± 0.03	-3.09 ± 0.06	-3.04 ± 0.03	-4.17 ± 0.72	-3.38 ± 0.24
Energy Efficiency	-1.99 ± 0.02	-1.72 ± 0.01	-1.21 ± 0.01	-1.06 ± 0.03	-2.66 ± 0.13	-1.01 ± 0.01	---	---
Kin8nm	0.95 ± 0.01	0.97 ± 0.00	1.14 ± 0.01	1.10 ± 0.01	1.12 ± 0.03	1.28 ± 0.01	---	---
Naval Propulsion	3.80 ± 0.01	3.91 ± 0.01	4.45 ± 0.00	2.46 ± 0.00	5.52 ± 0.10	4.85 ± 0.06	---	---
Power Plant	-2.80 ± 0.01	-2.79 ± 0.01	-2.80 ± 0.01	-2.82 ± 0.01	-2.81 ± 0.03	-2.78 ± 0.01	---	---
Protein Structure	-2.89 ± 0.00	-2.87 ± 0.00	-2.87 ± 0.00	-2.84 ± 0.00	-2.89 ± 0.00	-2.77 ± 0.01	---	---
Wine Quality Red	-0.93 ± 0.01	-0.92 ± 0.01	-0.93 ± 0.01	-0.95 ± 0.01	-0.95 ± 0.05	-0.97 ± 0.01	-1.29 ± 0.28	-1.04 ± 0.17
Yacht Hydrodynamics	-1.55 ± 0.03	-1.38 ± 0.01	-1.25 ± 0.01	-1.30 ± 0.02	-2.33 ± 0.01	-1.64 ± 0.02	-1.75 ± 0.19	-1.10 ± 0.08

There are two hyperparameters: i) the model precision parameter τ to evaluate the log-likelihood and ii) the dropout rate d . We perform the following two variations of the regression experiment:

1. **Convergence:** The networks are trained to convergence for 4,000 epochs and the hyperparameter values are obtained by Bayesian Optimization (BO) (as described in Gal and Ghahramani (2016)).
2. **Hyperparameter tuning:** Just as in the previous variant the networks are trained for 4,000 epochs, but we also obtain optimal hyperparameter values by performing grid search over a range of (τ, d) pairs and choose the best pair based on performance over a validation set. The validation set is created by randomly choosing 20% of the data points in the training set.

The RMSE and log likelihood values obtained from the above experiments are compared in Table 1 and 2 respectively. It should be noted that the RMSE and log likelihood values of VMG, HS-BNN, PBP-MV, and SGHMC have been taken from their respective papers. The experimental results indicate that the *Convergence* and *Hyperparameter tuning* baseline experiments show a substantial improvement in performance compared to the results in the *timed setting*. We also observe that the baseline outperform the other methods on the *Concrete Strength*, *Naval Propulsion Plants*, *Wine Quality Red*, and *Yacht Hydrodynamics* datasets in terms of RMSE. With respect to log likelihood, the baseline performs best on the *Boston Housing*, *Concrete Strength*, and *Wine Quality Red* datasets. Finally, we also observe that on the other datasets the baseline is competitive – second only to PBP-MV (Sun et al., 2017).

3 Conclusion

The RMSE and log-likelihood values obtained from the *Convergence* and *Hyperparameter tuning* experiments (as given in Tables 1 and 2) provide substantially better results for MC dropout. We conclude that previous comparisons with baselines obtained from the timed setting are unrepresentative, as the models in one setting had not reached convergence. In summary, when benchmarking a method, its performance should always be evaluated using an experimental setup identical to the one used to evaluate its peers.

The source code for our experiments can be found at: <https://github.com/yaringal/DropoutUncertaintyExps>

References

- Bui, T., Hernández-Lobato, D., Hernandez-Lobato, J., Li, Y., and Turner, R. (2016). Deep gaussian processes for regression using approximate expectation propagation. In *International Conference on Machine Learning*, pages 1472–1481.
- Gal, Y. and Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of the 33rd International Conference on Machine Learning (ICML-16)*.
- Ghosh, S. and Doshi-Velez, F. (2017). Model selection in bayesian neural networks via horseshoe priors. *arXiv preprint arXiv:1705.10388*.
- Graves, A. (2011). Practical variational inference for neural networks. In *Advances in neural information processing systems*, pages 2348–2356.
- Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., and Meger, D. (2017). Deep reinforcement learning that matters. *arXiv preprint arXiv:1709.06560*.
- Hernández-Lobato, J. M. and Adams, R. (2015). Probabilistic backpropagation for scalable learning of bayesian neural networks. In *International Conference on Machine Learning*, pages 1861–1869.
- Louizos, C. and Welling, M. (2016). Structured and efficient variational deep learning with matrix gaussian posteriors. In *International Conference on Machine Learning*, pages 1708–1716.
- Melis, G., Dyer, C., and Blunsom, P. (2018). On the state of the art of evaluation in neural language models. In *International Conference on Learning Representations*.
- Springenberg, J. T., Klein, A., Falkner, S., and Hutter, F. (2016). Bayesian optimization with robust bayesian neural networks. In *Advances in Neural Information Processing Systems*, pages 4134–4142.
- Sun, S., Chen, C., and Carin, L. (2017). Learning structured weight uncertainty in bayesian neural networks. In *Artificial Intelligence and Statistics*, pages 1283–1292.