
L2M: Practical posterior Laplace approximation with optimization-driven second moment estimation

Christian S. Perone¹ Roberto Pereira Silveira¹ Thomas Paula¹

Abstract

Uncertainty quantification for deep neural networks has recently evolved through many techniques. In this work, we revisit Laplace approximation, a classical approach for posterior approximation that is computationally attractive. However, instead of computing the curvature matrix, we show that, under some regularity conditions, the Laplace approximation can be easily constructed using the gradient second moment. This quantity is already estimated by many exponential moving average variants of Adagrad such as Adam and RMSprop, but is traditionally discarded after training. We show that our method (L2M) does not require changes in models or optimization, can be implemented in a few lines of code to yield reasonable results, and it does not require any extra computational steps besides what is already being computed by optimizers, without introducing any new hyperparameter. We hope our method can open new research directions on using quantities already computed by optimizers for uncertainty estimation in deep neural networks.

1. Introduction

Uncertainty quantification is an important component of learning systems. It is often a required quantity for the deployment of deep learning models in real-world applications, especially when it involves decision making. Uncertainty is also widely used in many domains such as reinforcement learning (Osband et al., 2018), active learning (Gal et al., 2017), out-of-distribution detection (Lakshminarayanan et al., 2017), and planning (Tigas et al., 2019), to name a few.

Although many methods of uncertainty estimation for modern neural networks were recently developed, most of them

¹Machine Learning Porto Alegre. Porto Alegre/RS, Brazil. Correspondence to: Christian S. Perone <christian.perone@gmail.com>.

rely on computationally expensive training and management procedures, sometimes of multiple models, with implementation complexity requiring specialized code for certain layer implementations or post-training tuning/estimation of hyperparameters in some cases.

Our contributions in this work are:

- We show that under some regularity conditions, a diagonal Laplace approximation can be constructed without computing anything besides what is already being computed by widely used optimizers;
- We qualitatively compare this approximation with methods such as deep ensembles (Lakshminarayanan et al., 2017), MC Dropout (Gal & Ghahramani, 2016), Hamiltonian Monte Carlo (HMC) (Cobb & Jalaian, 2020), among others;
- We also show that our approximation is orthogonal to methods such as ensembling (Lakshminarayanan et al., 2017) and does not require changing training procedures, estimating new quantities, or adding new hyperparameters.

To the best of our knowledge, it is the first time that the relationship between the raw second moment estimation in optimizers and uncertainty estimation is established.

1.1. Related Work

The use of the Laplace method for posterior approximation in neural networks can be dated back to Wray L. Buntine & Andreas S. Weigend (1991) and MacKay (1992). Recently, it was also employed by Ritter et al. (2018) in deep neural networks where a Kronecker factorization of the curvature matrix was employed. In Kristiadi et al. (2020), they also used Laplace approximation for the last layer of networks to overcome the confidence of ReLUs. In Kirkpatrick et al. (2017), a diagonal Laplace approximation was used to overcome catastrophic forgetting in neural networks.

Another line of work to estimate uncertainty in neural networks is through the use of ensembling methods as in Lakshminarayanan et al. (2017) where multiple models are trained and then used to estimate the prediction variance.

This method, together with random priors (randomly initialized and non-trainable networks) and bootstrapping was also employed in [Osband et al. \(2018\)](#).

Many variational approximations ([Graves, 2011](#)) were also employed for the posterior approximation, including an efficient and appealing method called MC Dropout by [Gal & Ghahramani \(2016\)](#), where dropout is left turned on during inference time.

Another method that was recently developed is called SWA-Gaussian (SWAG) ([Maddox et al., 2019](#)), which relies on SWA ([Izmailov et al., 2018](#)) running averages for the first and second moment of parameters to build a Gaussian approximation of the posterior. The main difference of SWAG to our method is that it uses SGD and compute the moments of parameters, while our method re-uses optimizer second moment estimation of the gradients. Computationally, they have similar performance when using SWAG with SGD, but that is not the case if SWAG is used with other optimizers such as Adam ([Kingma & Ba, 2015](#)).

2. Laplace approximation introduction

The main idea of the Laplace approximation ([Azevedo-Filho & Shachter, 1994](#)) can be stated as follows: locate the posterior mode and then construct a Gaussian distribution around the mode using as covariance matrix the inverse of the curvature around the mode.

The Laplace approximation can be derived from the second-order Taylor expansion around the mode (which can be found with optimization of a maximum a posterior θ_{MAP} estimator):

$$\log p(\theta|\mathcal{D}) \approx \log p(\theta_{\text{MAP}}|\mathcal{D}) - \frac{1}{2}(\theta - \theta_{\text{MAP}})^\top \tilde{H}(\theta - \theta_{\text{MAP}}) \quad (1)$$

Where \tilde{H} is $\mathbb{E}[H]$, the expected Hessian of the negative log posterior. Given the property that we are expanding around θ_{MAP} where the gradient is zero, the first order term of the expansion becomes zero. By exponentiation we can then reach the functional form of the Gaussian posterior approximation:

$$\theta \sim \mathcal{N}(\theta_{\text{MAP}}, \tilde{H}^{-1}) \quad (2)$$

During inference, we can approximate the mean or any other statistic by taking S Monte Carlo samples $\theta^{(s)}$ from this approximation:

$$p(\mathcal{D}^*|\mathcal{D}) = \int p(\mathcal{D}^*|\theta)p(\theta|\mathcal{D})d\theta \approx \frac{1}{S} \sum_{S=1}^S p(\mathcal{D}^*|\theta^{(s)}) \quad (3)$$

Which gives us the posterior predictive distribution.

3. Laplace approximation in neural networks

The use of Laplace approximation of the posterior in neural networks can be dated back to [Wray L. Buntine & Andreas S. Weigend \(1991\)](#) and [MacKay \(1992\)](#). Only recently, Laplace approximation was used for larger modern deep neural networks ([Ritter et al., 2018](#)), using a Kronecker factorization ([Martens & Grosse, 2015](#)) of the curvature matrix.

Although the Kronecker factorization can yield a better approximation when compared to a diagonal approximation of the curvature matrix, the Kronecker factors still have to be computed on another step after training the network and can require specialized computation depending on the layers ([Grosse & Martens, 2016](#)) used in the network, making it difficult to use in practice. However, it is important to note that our method can be used together with KFAC ([Martens & Grosse, 2015](#)) instead of Adam, as the same approach described in this work can be used with the KFAC approximation instead of the raw diagonal second moment estimated by optimizers such as Adam ([Kingma & Ba, 2015](#)).

4. Diagonal Laplace approximation using optimizer second moment

4.1. Diagonal Empirical Fisher estimation

Many exponential moving average variants of Adagrad ([Duchi et al., 2010](#)) are popular in the deep learning community. RMSprop ([Tieleman & Hinton, 2012](#)), Adam ([Kingma & Ba, 2015](#)), and Adadelta ([Zeiler, 2012](#)) are some examples. In this work, we will focus on Adam ([Kingma & Ba, 2015](#)), as it is one of the most used optimizers for deep neural networks.

The uncentered second moment of the gradients can be seen as a diagonal approximation of the Fisher Information Matrix (FIM) used in the Natural Gradient Descent ([Amari, 1998](#)):

$$\mathbf{F}_\theta = \mathbb{E}_{\substack{y \sim p_\theta(y|x) \\ x \sim p_{\text{data}}}} [\nabla_\theta \log p_\theta(y|x) \nabla_\theta \log p_\theta(y|x)^\top] \quad (4)$$

However, as noted by [Kunstner et al. \(2019\)](#), Adam optimizer does not use the expectation with $y \sim p_\theta(y|x)$ (with y from the model distribution), but instead it uses the ground-truth labels. This matrix, which is often called ‘‘empirical’’ Fisher (EF)¹, is often confused with the actual FIM. Also, in

¹We use the same commonly nomenclature here as ‘‘empirical’’ Fisher, however, as pointed by others ([Thomas et al., 2019](#); [Kunstner et al., 2019](#)), this term can be inaccurate.

the preconditioning with the FIM, Adam employs a square root of the Fisher, another departure from the natural gradient.

4.2. Construction of the Laplace approximation

Under the condition where the model is realizable and there is enough data to recover the true parameters, then at the minimum, the EF converges to the FIM, as shown in (Kunstner et al., 2019), given the consistency of the maximum likelihood estimator. Also, for optimizers using exponential moving averages of the gradient’s second moment, the influence of gradients beyond a fixed window size becomes negligibly small (Reddi et al., 2018).

Under the aforementioned regularity conditions regarding the convergence of the empirical Fisher to the \mathbf{F}_θ matrix, one can easily construct the posterior approximation using a diagonal Laplace approximation with the minima θ_{MAP} and the empirical Fisher diagonal $\text{diag}(\mathbf{F}_\theta)$, quantities already estimated by many modern optimizers (and that are usually discarded after training):

$$\theta \sim \mathcal{N}(\theta_{\text{MAP}}, \text{diag}(\mathbf{F}_\theta)^{-1}) \quad (5)$$

The \mathbf{F}_θ is connected to the Hessian through the Gauss-Newton matrix (Botev et al., 2017), which is an approximation of the Hessian where the second-order term is ignored, as it becomes small for models fitting data well. For exponential family distributions, the Gauss-Newton coincides with the Hessian, and as opposed to the Hessian – that can be degenerate for neural networks – the \mathbf{F}_θ and the Gauss-Newton are guaranteed to be positive semi-definite (p.s.d.).

In practice, we also assume a Gaussian prior, that would be equivalent to the L2 term in regularized loss where $\theta \sim \mathcal{N}(0, \frac{1}{\lambda}\mathbb{I})$, λ is the penalty term, and \mathbb{I} denotes the identity matrix. This leads to the following approximation:

$$\theta \sim \mathcal{N}(\theta_{\text{MAP}}, \left[\text{diag}(\mathbf{F}_\theta) + \frac{1}{\lambda} + \epsilon \right]^{-1}) \quad (6)$$

Where ϵ is a small term to avoid division by zero. This approximation is very attractive in computational terms and simplicity, as it takes leverage of Adam’s gradient second moment buffer (that are nowadays discarded) and uses the minima found by the optimizer as well. We call this approximation L2M (from Laplace 2nd Moment).

It is also important to highlight that our method is orthogonal to methods such as deep ensembles (Lakshminarayanan et al., 2017), since ensembles can be used to expand the mode coverage of the posterior while still using our method

on each model of the ensemble, leveraging the local curvature information around individual modes of the posterior.

4.3. Limitations

Although our approximation is scalable and easy to implement, it also comes with limitations that are important to highlight. The Laplace approximation is a Taylor expansion around the MAP estimate and, in our case, its curvature is built with a diagonal approximation, which means that parameters are assumed to be independent.

Also, the EF converges to the FIM only at the minimum (Kunstner et al., 2019), and although it can have remarkable angle similarity (Thomas et al., 2019), the ratio of traces from the EF and FIM can be quite different, which can lead to potential variance underestimation. However, up to a multiplicative constant, we can expect these two quantities to exhibit high similarity (Thomas et al., 2019).

We note that priors other than Gaussian (which has L2 regularization equivalence) would require changes in our method. Nevertheless, for many applications, this approximation can be an excellent trade-off while providing reasonable uncertainty estimation with a simple, easy to implement, and scalable approach.

4.4. Note regarding the ϵ term

The ϵ term, also introduced in Adam (Kingma & Ba, 2015) to avoid division by zero, can also be seen as a damping term improving the conditioning of the Fisher or setting a trust region radius during optimization, as noted by Choi et al. (2019). In our case, ϵ can have an effect of reducing the posterior variance, as it is added to the Fisher’s diagonal before the inversion.

5. Evaluation

Comparing different uncertainty estimation methods is a complex subject as there is no ground truth for uncertainty. Therefore, many works compare methods on the grounds of a qualitative evaluation on toy regression datasets or on out-of-distribution detection tasks.

In Figure 1, we show how our method compares to other methods on a toy regression dataset where $y \sim x^3 + \mathcal{N}(0, 3^2)$ as in Ritter et al. (2018). For this comparison, implemented in PyTorch (Paszke et al., 2019), we employed a simple MLP with 2 fully-connected layers, 40 units and ReLU activation. We used an Adam’s decoupled weight decay (Loshchilov & Hutter, 2019) of 0.1 with a learning rate of 0.1 and trained it for 5,000 epochs. For the predictions, we sampled 500 posterior samples and then report the mean and 1, 2, and 3 standard deviations above and below the mean.




Figure 1. Qualitative comparison on a toy regression dataset. Red shades denotes 1, 2, and $3\sigma^2$. The black line shows the mean prediction and the crosses represents data points. **First panel** (from left to right): Randomized prior functions (Osband et al., 2018). **Second panel**: MC Dropout (Gal & Ghahramani, 2016). **Third panel**: Deep Ensembles (Lakshminarayanan et al., 2017). **Fourth panel**: SWAG-Diagonal (Maddox et al., 2019). **Fifth panel**: Hamiltonian Monte Carlo (Cobb & Jalaian, 2020). **Sixth panel**: L2M (our method).

As we can clearly see in Figure 1, our method (L2M) performs on par with other methods, i.e., uncertainty is high in regions with no data points and low in regions with several data points. The method here presented does not need lots of hyperparameters to tune and does not need post-processing steps in order to get to these results, while still being fast to get the posterior predictive distribution.

We leave the evaluation of the integration of our method with deep ensembles (Lakshminarayanan et al., 2017) for future work, as our method can be used on each model of the ensemble to approximate the local curvature around the mode, potentially leading to better uncertainty estimates. The application of our method to classification problems is also left for future work.

6. Conclusion

In this paper, we presented L2M, a method for uncertainty quantification for deep neural networks. We revisited Laplace approximation, a classical approach for posterior approximation that is computationally attractive, and showed that we can construct a posterior approximation with the gradient raw second moment which is already computed by many exponential moving average variants of Adagrad (Duchi et al., 2010). To the best of our knowledge, it is the first time that this relationship between the raw second moment estimation in optimizers and uncertainty estimation was established.

We qualitatively compared the approximations produced by our method with other commonly used methods such as MC Dropout (Gal & Ghahramani, 2016), Deep Ensembles (Lakshminarayanan et al., 2017), among others. Our results are encouraging and open new research directions towards scalable approaches leveraging raw second moment estimation in optimizers for uncertainty estimation.

L2M has some limitations though. First, the curvature is

built with a diagonal approximation, which means that parameters are assumed to be independent. Second, we rely on some regularity conditions such as the convergence of the EF (“Empirical” Fisher) to the FIM (Fisher Information Matrix). Third, priors other than Gaussian would require changes in our method, as the prior equivalence with the Gaussian prior wouldn’t hold anymore. And last, we tested our approach only in a toy regression dataset, which means it would need to be further expanded to other larger datasets, network architectures and tasks.

For future work, we plan to conduct experiments on larger datasets. We also plan to expand this method to handle general classification and language modelling tasks with larger networks. Finally, we aim to explore the use of this method for out-of-distribution problems and incorporate it into other techniques such as deep ensembles (Lakshminarayanan et al., 2017).

Software and Data

A PyTorch (Paszke et al., 2019) implementation of our method can be found at the following repository below:

<https://github.com/anonymous/L2M>

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software

available from tensorflow.org.

- Amari, S. I. Natural Gradient Works Efficiently in Learning. *Neural Computation*, 10(2):251–276, feb 1998. ISSN 08997667. doi: 10.1162/089976698300017746.
- Azevedo-Filho, A. and Shachter, R. D. Laplace’s Method Approximations for Probabilistic Inference in Belief Networks with Continuous Variables. In *Uncertainty Proceedings 1994*, pp. 28–36. Elsevier, feb 1994. doi: 10.1016/b978-1-55860-332-5.50009-2.
- Botev, A., Ritter, H., and Barber, D. Practical Gauss-Newton optimisation for deep learning. In *34th International Conference on Machine Learning, ICML 2017*, volume 2, pp. 860–880, 2017. ISBN 9781510855144.
- Choi, D., Shallue, C. J., Nado, Z., Lee, J., Maddison, C. J., and Dahl, G. E. On Empirical Comparisons of Optimizers for Deep Learning. 2019.
- Cobb, A. D. and Jalaian, B. Scaling hamiltonian monte carlo inference for bayesian neural networks with symmetric splitting. *arXiv preprint arXiv:2010.06772*, 2020.
- Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. In *COLT 2010 - The 23rd Conference on Learning Theory*, volume 12, pp. 257–269, 2010. ISBN 9780982252925.
- 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*, ICML’16, pp. 1050–1059. JMLR.org, 2016.
- Gal, Y., Islam, R., and Ghahramani, Z. Deep Bayesian active learning with image data. Technical report, 2017.
- Graves, A. Practical variational inference for neural networks. In *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011, NIPS 2011*, volume 24, 2011. ISBN 9781618395993.
- Grosse, R. and Martens, J. A kronecker-factored approximate fisher matrix for convolution layers. In *33rd International Conference on Machine Learning, ICML 2016*, volume 2, pp. 851–874. International Machine Learning Society (IMLS), feb 2016. ISBN 9781510829008.
- Izmailov, P., Podoprikin, D., Garipov, T., Vetrov, D., and Wilson, A. G. Averaging weights leads to wider optima and better generalization. In *34th Conference on Uncertainty in Artificial Intelligence 2018, UAI 2018*, volume 2, pp. 876–885. Association For Uncertainty in Artificial Intelligence (AUAI), mar 2018. ISBN 9781510871601.
- Kingma, D. P. and Ba, J. L. Adam: a Method for Stochastic Optimization. *International Conference on Learning Representations 2015*, pp. 1–15, 2015. ISSN 09252312. doi: http://doi.acm.org.ezproxy.lib.ucf.edu/10.1145/1830483.1830503.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., and Hadsell, R. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114(13):3521–3526, dec 2017. ISSN 10916490. doi: 10.1073/pnas.1611835114.
- Kristiadi, A., Hein, M., and Henni, P. Being Bayesian, even Just a Bit, Fixes Overconfidence in ReLU Networks. In *37th International Conference on Machine Learning, ICML 2020*, volume PartF16814, pp. 5392–5402, 2020. ISBN 9781713821120.
- Kunstner, F., Balles, L., and Hennig, P. Limitations of the empirical fisher approximation for natural gradient descent. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. Technical report, 2017.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In *7th International Conference on Learning Representations, ICLR 2019*. International Conference on Learning Representations, ICLR, nov 2019.
- MacKay, D. J. C. A Practical Bayesian Framework for Backpropagation Networks. *Neural Computation*, 4(3): 448–472, 1992. ISSN 0899-7667. doi: 10.1162/neco.1992.4.3.448.
- Maddox, W. J., Garipov, T., Izmailov, Vetrov, D., and Wilson, A. G. A simple baseline for Bayesian uncertainty in deep learning. Technical report, 2019.
- Martens, J. and Grosse, R. Optimizing neural networks with Kronecker-factored approximate curvature. In *32nd International Conference on Machine Learning, ICML 2015*, volume 3, pp. 2398–2407. International Machine Learning Society (IMLS), mar 2015. ISBN 9781510810587.
- Osband, I., Aslanides, J., and Cassirer, A. Randomized prior functions for deep reinforcement learning. Technical report, 2018.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang,

L., Bai, J., and Chintala, S. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, volume 32. Neural information processing systems foundation, dec 2019.

Reddi, S. J., Kale, S., and Kumar, S. On the convergence of Adam and beyond. In *6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings*, 2018.

Ritter, H., Botev, A., and Barber, D. A scalable laplace approximation for neural networks. In *6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings*, 2018.

Thomas, V., Pedregosa, F., van Merriënboer, B., Mangazol, P.-A., Bengio, Y., and Roux, N. L. On the interplay between noise and curvature and its effect on optimization and generalization. 2019.

Tieleman, T. and Hinton, G. Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude. COURSERA: Neural Networks for Machine Learning, 2012.

Tigas, P., Filos, A., Mcallister, R., Rhinehart, N., Levine, S., and Gal, Y. Robust Imitative Planning: Planning from Demonstrations Under Uncertainty. Technical Report NeurIPS, 2019.

Wray L. Buntine and Andreas S. Weigend. Bayesian Back-Propagation by Wray L. Buntine and Andreas S. Weigend. *Complex Systems*, 5(6):603–643, 1991.

Zeiler, M. D. ADADELTA: An Adaptive Learning Rate Method. 2012.

Appendix

Appendix A - L2M approximation pseudo-code

This appendix shows a pseudo-code implementation of our method:

```
# Trained model and optimizer
model = (...)
optimizer = (...)

# Build L2M approximation
grad_second_moment = optimizer.second_moment()
prior_variance = 1.0/weight_decay
fisher_diag = grad_second_moment + \
prior_variance + eps
inv_fisher_diag = 1.0/fisher_diag
l2m_approximation = Normal(model.weights,
sqrt(inv_fisher_diag))

# Multiple posterior samples for inference
outputs = [] # predictive distribution
for s in S:
model.weights = l2m_approximation.sample()
output = model(input)
outputs.append(output)
```
