
Variational Dropout via Empirical Bayes

Valery Kharitonov¹
kharvd@gmail.com

Dmitry Molchanov^{1,2}
dmolch111@gmail.com

Dmitry Vetrov^{1,2}
vetrovd@yandex.ru

¹National Research University Higher School of Economics, Joint Samsung-HSE Lab

²Samsung AI Center in Moscow

Abstract

We study the Automatic Relevance Determination procedure applied to deep neural networks. We show that ARD applied to Bayesian DNNs with Gaussian approximate posterior distributions leads to a variational bound similar to that of variational dropout, and in the case of a fixed dropout rate, objectives are exactly the same. Experimental results show that the two approaches yield comparable results in practice even when the dropout rates are trained. This leads to an alternative Bayesian interpretation of dropout and mitigates some of the theoretical issues that arise with the use of improper priors in the variational dropout model. Additionally, we explore the use of the hierarchical priors in ARD and show that it helps achieve higher sparsity for the same accuracy.

1 Introduction

Dropout [1] is a popular regularization method for neural networks that can be interpreted as a form of approximate Bayesian inference [2, 3]. Sparse variational dropout (Sparse VD) [6] further extends this approach and shows that it can be used to significantly prune neural networks. However, it was recently argued [7] that the use of an improper prior distribution in current formulations of variational dropout leads to an improper posterior, and as such, this model cannot be used to provide a principled interpretation of the empirical behavior of the dropout procedure. In this paper, we use variational inference to perform automatic relevance determination procedure [4, 5] to arrive at an alternative interpretation of dropout that does not have such a drawback. The derived objective is remarkably similar to the approximation in Sparse VD and empirical observations confirm that the two models are effectively equivalent.

2 Variational Automatic Relevance Determination for neural networks

Let $\mathcal{D} = (x_i, y_i)_{i=1}^N$ be a dataset of N samples where x_i are observable variables and y_i are the targets. Suppose that we have some parametric model $p(\mathcal{D}|w) = \prod_{i=1}^N p(y_i|x_i, w)$, $w \in \mathbb{R}^D$ (e.g. a deep neural network) mapping x to the corresponding y using parameters w . The parameters w have a prior distribution $p(w|\tau)$ which is itself parameterized by hyperparameters $\tau \in \mathbb{R}^H$. Following the Bayesian approach, we wish to find the posterior distribution $p(w|\mathcal{D}, \tau) = p(\mathcal{D}|w)p(w|\tau)/p(\mathcal{D}|\tau)$. We choose hyperparameters τ such that the marginal likelihood (evidence) of the dataset is maximized:

$$\tau^* = \arg \max_{\tau} p(\mathcal{D}|\tau) = \arg \max_{\tau} \int p(\mathcal{D}|w)p(w|\tau) dw.$$

This is the *empirical Bayes* (EB) approach to hyperparameter selection. In particular, when $p(w|\tau) = \prod_{i=1}^D \mathcal{N}(w_i | 0, \tau_i^{-1})$, this procedure is called *automatic relevance determination* (ARD) [4, 5].

Since in the case of deep neural networks the marginal likelihood is intractable, we use doubly stochastic variational inference (DSVI) [9] to find an approximate posterior $q(w|\phi)$ from some

parametric family. This is achieved by maximizing the *evidence lower bound (ELBO)* w.r.t. variational parameters ϕ and hyperparameters τ :

$$\log p(\mathcal{D}|\tau) \geq \mathcal{L}(\phi, \tau) = \mathbb{E}_{q(w|\phi)} [\log p(\mathcal{D}|w)] - D_{KL}(q(w|\phi) || p(w|\tau)) \rightarrow \max_{\phi, \tau}.$$

Now suppose that $p(w|\tau) = \prod_{i=1}^D \mathcal{N}(w_i|0, \tau_i^{-1})$ (ARD prior) and $q(w|\mu, \sigma) = \prod_{i=1}^D \mathcal{N}(w_i|\mu_i, \sigma_i^2)$. The optimal value for hyperparameters τ in this case can be found analytically [9]: $\tau_i^* = (\mu_i^2 + \sigma_i^2)^{-1}$. The ELBO then takes the following form [9]:

$$\begin{aligned} \mathcal{L}_{ARD}(\mu, \sigma) &= \sum_{i=1}^N \mathbb{E}_{q(w|\mu, \sigma)} [\log p(y_i|x_i, w)] - \frac{1}{2} \sum_{j=1}^D \log \left(1 + \frac{\mu_j^2}{\sigma_j^2} \right) \\ &= \mathcal{L}_{\mathcal{D}}(\mu, \sigma) + \mathcal{R}_{ARD}(\mu, \sigma) \rightarrow \max_{\mu, \sigma}. \end{aligned} \quad (1)$$

In practice, we estimate the gradients of the ELBO w.r.t. the variational parameters using the local reparameterization trick [2]. The objective (1) was first derived in the context of linear models [9].

3 Connection with variational dropout

Variational dropout (VD) [2] is a generalization of Gaussian dropout which interprets it as an approximate Bayesian inference procedure. It puts an improper scale-invariant log-uniform prior on the weights of a neural network with fully-connected and convolutional layers and uses a factorized Gaussian approximation to the true posterior.

Let us now consider a restricted variational approximation $q(w|\mu) = \prod_{i=1}^D \mathcal{N}(w_i | \mu_i, \alpha \mu_i^2)$ for some constant $\alpha > 0$ in the ARD objective (1). The regularizer term is now constant and does not affect optimization, so the objective takes the form:

$$\tilde{\mathcal{L}}_{ARD}(\mu) = \mathcal{L}_{\mathcal{D}}(\mu) = \sum_{i=1}^N \mathbb{E}_{q(w|\mu)} [\log p(y_i|x_i, w)] \rightarrow \max_{\mu}. \quad (2)$$

It can be shown [2] that optimizing such functional is equivalent to training a neural network with Gaussian dropout which puts multiplicative normal noise $\mathcal{N}(1, \alpha)$ on the input of each dense and convolutional layer in the net, while ignoring the dependencies between the output units of the layer. Note that unlike variational dropout [2], we did not have to use an improper prior to arrive at this objective.

We can assign individual dropout rates to each weight in the network and optimize with respect to them. Such an approach [6] leads to very sparse solutions where most of the weights of a network are assigned high dropout rates, thus effectively being pruned from the network. To tune individual dropout rates in the ARD model, we set $q(w|\mu, \alpha) = \prod_{i=1}^D \mathcal{N}(w_i | \mu_i, \alpha_i \mu_i^2)$ which yields the following objective (ARD Dropout):

$$\mathcal{L}_{ARD}(\mu, \alpha) = \mathcal{L}_{\mathcal{D}}(\mu, \alpha) - \frac{1}{2} \sum_{j=1}^D \log(1 + \alpha_j^{-1}) \rightarrow \max_{\mu, \alpha}. \quad (3)$$

Compare this to the approximation for the Sparse VD objective [6]:

$$\mathcal{L}_{SVDO}(\mu, \alpha) \approx \mathcal{L}_{\mathcal{D}}(\mu, \alpha) + \sum_{j=1}^D \left[k_1 \sigma(k_2 + k_3 \log \alpha_j) - \frac{1}{2} \log(1 + \alpha_j^{-1}) + C \right] \rightarrow \max_{\mu, \alpha}, \quad (4)$$

where $k_1 = 0.63576$, $k_2 = 1.87320$, $k_3 = 1.48695$, $C = -k_1$. As we can see, both (3) and (4) have the same log terms in the regularizer and the sigmoid term in $\mathcal{L}_{SVDO}(\mu, \alpha)$ is bounded. It is also easy to show that Sparse VD objective is actually a lower bound on the ARD objective. Furthermore, when $\alpha_j \rightarrow \infty$, $\mathcal{L}_{SVDO}(\mu, \alpha)$ approaches $\mathcal{L}_{ARD}(\mu, \alpha)$ from below. The common term in the objectives is the one that encourages sparsity, so we expect that ARD would perform similarly to Sparse VD in terms of compression rate. In Table 1, we report accuracy and compression

Objective	Error %	Compression
MNIST (LeNet-5)		
Sparse VD	0.81 (± 0.04)	136 (± 8)
ARD Dropout	0.76 (± 0.09)	132 (± 13)
CIFAR-10 (VGG-like)		
Sparse VD	8.08 (± 0.27)	38 (± 1)
ARD Dropout	7.75 (± 0.28)	34 (± 1)

Table 1: Empirical comparison of ARD and sparse variational dropout. Compression is defined as the total number of weights in the network divided by the number of non-zero weights after trimming.

a	Error %	Compression
CIFAR-10 (VGG-like)		
0.505	7.46 (± 0.24)	52 (± 1)
0.510	7.65 (± 0.25)	85 (± 1)
0.515	7.88 (± 0.13)	114 (± 1)
0.520	7.94 (± 0.18)	136 (± 2)

Table 2: Classification error and compression for Gamma hyperprior MAP-II model, where $b = 10^{-8}$.

for networks trained with ARD and Sparse VD objectives (see Appendix A for details). As we can see, both models show comparable sparsity while maintaining low classification error rate.

The proposed ARD dropout interpretation mitigates one of the issues described in Hron et al. (2018) [7]. The paper identifies two problems with variational dropout: (a) the use of improper prior distribution and (b) singularity of the approximate posterior distribution. Now that we do not use the log-uniform prior anymore, our ARD dropout model is defined correctly, and both the prior and the true posterior distributions for each value of τ are now proper, which fixes (a), and (b) is only present in the model with correlated weight noise which is not the case for the ARD model.

4 Accuracy-compression trade-off using Gamma hyperprior

We can go further and introduce a hyperprior over hyperparameters τ and then perform *maximum a posteriori* estimation for them. Suppose that $p(\tau_i|a, b) = \text{Gamma}(\tau_i|a, b)$. If we maximize the evidence w.r.t. τ in such model (so-called MAP-II estimation), it leads to the following objective when $a > 1/2, b > 0$ (see Appendix B for details):

$$\mathcal{L}_\Gamma(\mu, \sigma) = \mathcal{L}_\mathcal{D}(\mu, \sigma) + \sum_{j=1}^D \left[\frac{1}{2} \log \frac{\sigma_j^2}{(\sigma_j^2 + \mu_j^2 + 2b)^{2a-1}} + C \right] \rightarrow \max_{\mu, \sigma},$$

where $C = 1 - a + a \log b - \log \Gamma(a) + \frac{1}{2}(2a - 1) \log(2a - 1)$, and Γ is the gamma function. Unlike ARD or Sparse VD models, we now have tunable parameters a, b of the hyperprior. With $a = 1, b = 0$ (which gives an improper prior) and ignoring the constant term, we get the usual ARD objective. Smaller values of a (around 0.5) reduce pruning without rescaling the regularizer term, which is a common way to prevent underfitting [11, 12]. Moreover, this approach helps achieve higher values of sparsity for the same accuracy level. See Table 2 for empirical results for different values of a .

Interestingly, if we approach this model in a fully Bayesian way and marginalize the hyperparameters out, we obtain the marginal prior $p(w)$ distributed according to the generalized Student t-distribution $p(w_i) = \text{Student}(\nu = 2a, \mu = 0, \lambda = a/b)$. In particular, the log-uniform prior used in Sparse VD can be considered a limiting case of this model when $a = b \rightarrow 0$. See Appendix C for details.

5 Conclusion

We have shown that it is possible to overcome the theoretical difficulties with variational dropout using a variational approximation to the ARD procedure with fully factorized Gaussian variational posterior distributions. It does not require the use of improper priors and leads to a similar variational lower bound, and thus can be seen as an alternative to variational dropout. Our theoretical study is supported by experimental results that show that in practice, both methods achieve comparably high values of sparsity without a significant drop in classification accuracy. Additionally, such an approach allows applying a hierarchical prior to the hyperparameters, which gives an option to trade off between accuracy and sparsity.

Acknowledgements

We would like to thank Dmitry Kropotov, Arsenii Ashukha, Kirill Neklyudov, and Dmitrii Podoprikin for valuable discussions and feedback.

References

- [1] Srivastava, Nitish, et al. "Dropout: a simple way to prevent neural networks from overfitting." *The Journal of Machine Learning Research* 15.1 (2014): 1929-1958.
- [2] Kingma, Diederik P., Tim Salimans, and Max Welling. "Variational dropout and the local reparameterization trick." *Advances in Neural Information Processing Systems*. 2015.
- [3] Gal, Yarin, and Zoubin Ghahramani. "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning." *International Conference on Machine Learning*. 2016.
- [4] MacKay, David JC. "Probable networks and plausible predictions—a review of practical Bayesian methods for supervised neural networks." *Network: Computation in Neural Systems* 6.3 (1995): 469-505.
- [5] Neal, Radford M. Bayesian learning for neural networks. *Springer*, (1996).
- [6] Molchanov, Dmitry, Arsenii Ashukha, and Dmitry Vetrov. "Variational dropout sparsifies deep neural networks." *arXiv preprint arXiv:1701.05369* (2017).
- [7] Hron, Jiri, Alexander G. de G. Matthews, and Zoubin Ghahramani. "Variational Bayesian dropout: pitfalls and fixes." *arXiv preprint arXiv:1807.01969* (2018).
- [8] Tipping, Michael E. "Sparse Bayesian learning and the relevance vector machine." *Journal of machine learning research* 1.Jun(2001): 211-244.
- [9] Titsias, Michalis, and Miguel Lázaro-Gredilla. "Doubly stochastic variational Bayes for non-conjugate inference." *International Conference on Machine Learning*. 2014.
- [10] Kingma, Diederik P., and Max Welling. "Auto-encoding variational bayes." *arXiv preprint arXiv:1312.6114* (2013).
- [11] Louizos, Christos, Karen Ullrich, and Max Welling. "Bayesian compression for deep learning." In *Advances in Neural Information Processing Systems*, pp. 3288-3298. 2017.
- [12] Ullrich, Karen, Edward Meeds, and Max Welling. "Soft weight-sharing for neural network compression." *arXiv preprint arXiv:1702.04008* (2017).

A Experimental setting

We compare ARD dropout with Sparse VD [6]. We train LeNet-5-Caffe¹ and VGG-like² networks on MNIST and CIFAR-10 correspondingly. To reduce the variance of the gradient, we use the additive parameterization [6], training means and logarithms of standard deviations. We train both networks from the same random initialization (with $\log \sigma$ initialized from $\mathcal{N}(-5, 0.1^2)$) for 200 epochs using Adam optimizer with the initial learning rate 10^{-3} and minibatch size 100. Starting from epoch 100, we linearly reduce the learning rate to zero. To overcome underfitting, for VGG on CIFAR-10, we scale the regularizer terms in both objectives by 0.05, and for both networks, we anneal the regularizer term over 20 epochs. For Gamma hyperprior experiments, posterior variances are clipped so that $\log \sigma < -4$ and the regularizer term is not scaled. After training, we evaluate the networks in "deterministic mode": we set the network weights to the means of their approximate posterior distribution. We then trim the weights of the networks setting those with $|\mu_i| < 10^{-2}$ to zero. We report the mean and the standard deviation for sparsity and error over 5 random seeds.

¹<https://github.com/BVLC/caffe/blob/master/examples/mnist/lenet.prototxt>

²<http://torch.ch/blog/2015/07/30/cifar.html>

B MAP-II Estimation with Gamma Hyperprior

We can introduce a prior distribution over the inverse variance τ :

$$p(y, w, \tau|x) = p(y|x, w)p(w|\tau)p(\tau),$$

where $p(\tau) = \prod_i \text{Gamma}(\tau_i|a, b)$, and $p(w_i|\tau_i) = \mathcal{N}(w_i|0, \tau_i^{-1})$. Then, we perform maximum a posteriori estimation for the hyperparameters of the prior over the weights. The objective takes the form:

$$\mathcal{L}(\mu, \sigma, \tau) = \mathcal{L}_{data}(\mu, \sigma) + \log p(\tau) - D_{KL}(q(w|\mu, \sigma) || p(w|\tau)) \rightarrow \max_{\tau, \mu, \sigma}. \quad (5)$$

The KL term is the same as for the regular ARD:

$$D_{KL}(q(w_i|\mu_i, \sigma_i) || p(w_i|\tau_i)) = -\frac{1}{2} + \frac{\tau_i(\sigma_i^2 + \mu_i^2)}{2} - \frac{\log(\tau_i\sigma_i^2)}{2}.$$

The regularizer:

$$\log p(\tau_i) = a \log b - \log \Gamma(a) + (a - 1) \log \tau_i - b\tau_i.$$

Differentiating (5) with respect to τ_i and equating to zero, we get the optimal value of τ_i which exists when $a > \frac{1}{2}$:

$$\tau_i^* = \frac{2a - 1}{\sigma_i^2 + \mu_i^2 + 2b}.$$

Substituting back, we get the following regularizer for a single weight:

$$R(\mu_i, \sigma_i) = \frac{1}{2} \log \frac{\sigma_i^2}{(\sigma_i^2 + \mu_i^2 + 2b)^{2a-1}} + C,$$

where

$$C = 1 - a + a \log b - \log \Gamma(a) + \frac{1}{2}(2a - 1) \log(2a - 1).$$

C Full Bayes with Gamma Hyperprior

Instead of a MAP estimation, we can go full Bayes and marginalize over hyperparameters:

$$p(w_i) = \int p(w_i | \tau_i) p(\tau_i) d\tau_i = \int \mathcal{N}(w_i | 0, \tau_i^{-1}) \cdot \text{Gamma}(\tau_i | a, b) d\tau_i.$$

It is a well-known fact that such Gaussian scale mixture gives a generalized Student t-distribution:

$$p(w_i) = \text{Student}(\nu = 2a, \mu = 0, \lambda = a/b).$$

Setting $a = b = \xi$ and taking limit as ξ approaches 0 from above, we get a log-uniform distribution:

$$p(\log |w_i|) \propto C,$$

which is exactly the prior for variational dropout.