
Stabilising priors for robust Bayesian deep learning

Felix McGregor*

Department of E&E Engineering
Stellenbosch University

Arnu Pretorius

Computer Science Division
Stellenbosch University

Johan du Preez

Department of E&E Engineering
Stellenbosch University

Steve Kroon

Computer Science Division
Stellenbosch University

Abstract

Bayesian neural networks (BNNs) have developed into useful tools for probabilistic modelling due to recent advances in variational inference enabling large scale BNNs. However, BNNs remain brittle and hard to train, especially: (1) when using deep architectures consisting of many hidden layers and (2) in situations with large weight variances. We use signal propagation theory to quantify these challenges and propose *self-stabilising priors*. This is achieved by a reformulation of the ELBO to allow the prior to influence network signal propagation. Then, we develop a stabilising prior, where the distributional parameters of the prior are adjusted before each forward pass to ensure stability of the propagating signal. This stabilised signal propagation leads to improved convergence and robustness making it possible to train deeper networks and in more noisy settings.

1 Introduction

Bayesian neural networks (BNNs) offer a way of combining uncertainty estimation with the expressive power of deep learning. Advances in BNNs (Graves, 2011; Blundell et al., 2015; Kingma and Welling, 2013; Rezende et al., 2014) have made it possible to scale these models but not yet to the level of success of modern deep learning. The reason in part, is that BNNs are difficult to train. Stochastic optimisation methods for these models tend to exhibit high variance and BNNs can be very sensitive to small changes in hyperparameters, architecture and the choice of prior. This work builds on Pretorius et al. (2018), which extended the analysis of signal propagation for deterministic ReLU networks to ReLU networks with stochastic regularisation. In light of recent links between stochastic regularisation techniques and variational inference (Gal and Ghahramani, 2016; Kingma et al., 2015), we relate the initialisation techniques derived in Pretorius et al. (2018), to an iteratively updating prior to stabilise the flow of information through ReLU BNNs throughout training.

Signal propagation analysis of BNNs in the infinite width limit leads us to propose *self-stabilising priors* for robust training. We design an iterative prior with distributional parameters derived to preserve the variance of signals propagating forward through the network. This is a heuristic similar in nature to an iterated application of Empirical Bayes (EB) for setting prior hyperparameters, as is common in BNN training (Graves (2011) and Titsias and Lázaro-Gredilla (2014) propose closed form updates for prior hyperparameters). While EB chooses hyperparameters that optimize the likelihood, our approach chooses prior hyperparameters for each forward pass that attempt to optimize signal propagation behaviour in the network. In order for this prior to influence the signal propagation dynamics, we reformulate the *evidence lower bound* (ELBO) objective to allow this prior to exercise its stabilising effect during the forward pass.

*Correspondence: flx.mcgregor@gmail.com

2 Self-stabilising priors

The prior $p_\alpha(W)$ usually impacts the ELBO through an additive KL term, only affecting the weights after the forward pass, having no effect on the signal propagation dynamics of the network. We instead propose priors that also exert their influence *during* the forward pass, so as to promote stable signal propagation, and improve robustness in deep BNNs. We reformulate the ELBO in terms of a combination of the prior and approximating posterior as $\tilde{q}_{\{\alpha,\phi\}}(W) = p_\alpha(W)q_\phi(W)/Z$. This formulation allows us to estimate an expectation using a Monte Carlo estimator with samples drawn from $\tilde{q}_{\{\alpha,\phi\}}(W)$ instead of $q_\phi(W)$. This ensures that the sampled weights of the network are being influenced by the current prior $p_\alpha(W)$ during the forward pass. From Appendix B we define the ELBO as

$$\mathcal{L}_{\tilde{q}} := \mathbb{E}_{p(\epsilon)} [\log p(\mathbf{y}|\mathbf{x}, \mathbf{b}, W = \xi(\epsilon, \alpha, \phi))] - \text{KL}(\tilde{q}_{\{\alpha,\phi\}}(W)||p_\alpha(W)), \quad (1)$$

where $\epsilon \sim \mathcal{N}(0, I)$. The prior is adjusted after every gradient update to adapt to the updated posterior $q_\phi(W)$. These prior parameters are optimal in the sense that together with the reformulated ELBO in (1), the sampled weights from $\tilde{q}_{\{\alpha,\phi\}}(W)$ have a stabilising effect on the signal propagation dynamics of a Bayesian deep neural network.

To quantify the signal propagation dynamics we make use of the *mean field* assumption (Saxe et al., 2014; Poole et al., 2016) which allows for the components of the pre-activation vectors \mathbf{h}^l to be treated as independent Gaussian random variables. Then, according to the central limit theorem, these pre-activations are Gaussian distributed. We consider independent Gaussian priors $p_\alpha(w_{i,j}^l) = \mathcal{N}(\mu_{p_{i,j}}^l, (\sigma_{p_{i,j}}^l)^2)$ and posteriors $q_\phi(w_{i,j}^l) = \mathcal{N}(\mu_{q_{i,j}}^l, (\sigma_{q_{i,j}}^l)^2)$. Focusing on ReLU activations, i.e. $g(a) = \max(0, a)$, we consider a single scalar hidden unit pre-activation h_j^l at an arbitrary layer l of the network. In Appendix A, we show that in expectation over the weights and biases, under the assumption of zero mean inputs at each layer (true at initialisation but a somewhat unrealistic assumption during training further discussed in Appendix A), the variance of h_j^l , governing signal propagation in the forward pass is given by

$$\text{Var}[h_j^l] = \left[\left(1 - \frac{1}{\pi}\right) (\mu_{q_j}^l)^2 + (\sigma_{q_j}^l)^2 \right] \frac{\text{Var}[h_{j'}^{l-1}]}{2}. \quad (2)$$

where $j' \in \{1, \dots, D_{l-1}\}$. Then, to stabilise the signal, we want to find prior parameters $\alpha = \{\mu_{p_{i,j}}^l, \sigma_{p_{i,j}}^l\}$ that can preserve the variance during the forward pass. Specifically, we want α to ensure $\text{Var}[h_j^l] = \text{Var}[h_{j'}^{l-1}]$ in (2). To achieve this, we require $\mu_{p_{i,j}}^l = \mu_{q_{i,j}}^l$ giving the stabilising prior parameters α as

$$\mu_{p_{i,j}}^l = \mu_{q_{i,j}}^l, \quad \sigma_{p_j}^l = \sqrt{\frac{(\sigma_{q_j}^l)^2 \gamma}{(\sigma_{q_j}^l)^2 - \gamma}}, \quad (3)$$

where $\gamma = |2 - (1 - 1/\pi)(\mu_{q_j}^l)^2|$ and we take the absolute value to ensure positive variances. Note that we can apply the result in (2) recursively for all layers $l = 1, \dots, L$, with base case $\text{Var}[x^0] = \frac{1}{D_0} \mathbf{x}^0 \cdot \mathbf{x}^0$. Therefore, sampling the weights $W \sim \tilde{q}_{\{\alpha,\phi\}}(W)$ at each forward pass, while setting the prior $p_\alpha(w) = \mathcal{N}(\mu_q^l, |(\sigma_q^l)^2 \gamma / ((\sigma_q^l)^2 - \gamma)| / D_{l-1})$, promoting stable signal propagation.

The effect of the stabilising prior is that the larger the mean and variance of the incoming weights, the more likely it is to destroy the signal, whereas if the second moment is small, it is not likely to add noise to the signal. The prior encourages the weight distribution to sample closer to its means when the second moment of the distribution is large.

3 Experiments

Limits of trainability. We restrict ourselves to BNNs with fully connected layers with a specified number of hidden layers all of constant width.² We investigate performance at extremities by training a series of networks with varying depths and initial variances. We compare a series of networks with our proposed stabilising prior incorporated on the forward pass with a standard non-conjugate Gaussian prior Titsias and Lázaro-Gredilla (2014) which we report in Figures 1 (a) and (b). We observe our stabilising prior makes it possible to train deeper BNNs and in more noisy conditions.

²Code available <https://git.io/JeLkH>

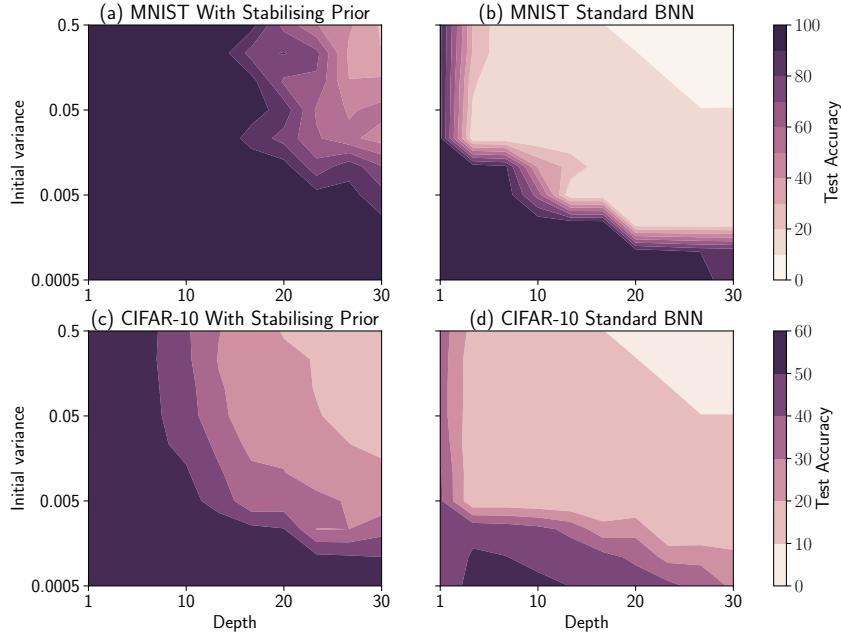


Figure 1: MNIST and CIFAR-10 large scale experiments. Classification accuracy grid of ReLU networks trained with varying depths and initial variance conditions trained for 50 epochs.

Accelerated training. In general, we also observe that our prior improves convergence as demonstrated in Figure 2. We compare with Empirical Bayes (EB) as in Titsias and Lázaro-Gredilla (2014) which uses the gradient to find optimal hyperparameters for the prior. We further compare these priors with a Gaussian prior and report their results for both the reparametrisation trick (RT) Kingma and Welling (2013) and local reparametrisation trick (LRT) Kingma et al. (2015).

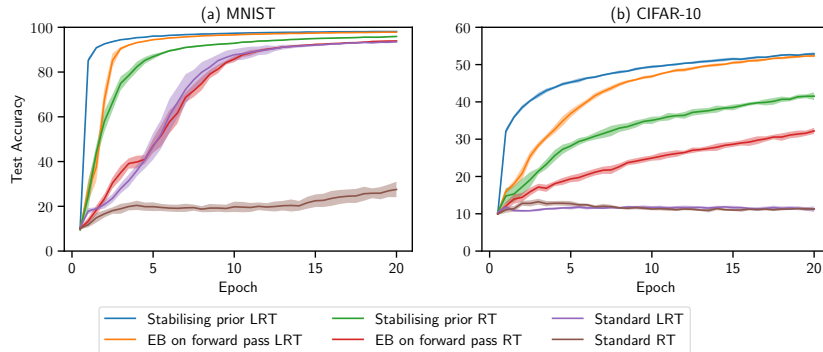


Figure 2: Progression of test accuracy for various networks through training averaged over 10 runs for a 5 layer deep 512 wide network with an initial variance of 0.001.

4 Discussion

We used signal propagation theory to derive priors for BNNs that promote stable signal propagation. The prior incorporates knowledge of model architecture and activation function derived from how signals propagate in the network in the infinite width limit. We showed that these priors, when their effect is exerted in the forward pass, makes it possible to train deeper networks and in more noisy conditions. This alleviates the need to tune hyper-parameters and extends BNNs to deeper architectures. We also observe that stable signal propagation accelerates training, which we attribute to cleaner signals and gradients being propagated through the network and more efficient expectations.

Acknowledgements

We would like to thank the Technology and Human Resource for Industry Programme (THRIP) as well as Centre for Artificial Intelligence Research (CAIR) for financial support. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research.

References

- A. Graves, “Practical variational inference for neural networks,” in *Advances in neural information processing systems*, 2011, pp. 2348–2356.
- C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, “Weight uncertainty in neural networks,” *arXiv preprint arXiv:1505.05424*, 2015.
- D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- D. J. Rezende, S. Mohamed, and D. Wierstra, “Stochastic backpropagation and approximate inference in deep generative models,” *arXiv preprint arXiv:1401.4082*, 2014.
- A. Pretorius, E. Van Biljon, S. Kroon, and H. Kamper, “Critical initialisation for deep signal propagation in noisy rectifier neural networks,” in *Advances in Neural Information Processing Systems*, 2018, pp. 5722–5731.
- Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” in *international conference on machine learning*, 2016, pp. 1050–1059.
- D. P. Kingma, T. Salimans, and M. Welling, “Variational dropout and the local reparameterization trick,” in *Advances in Neural Information Processing Systems*, 2015, pp. 2575–2583.
- M. Titsias and M. Lázaro-Gredilla, “Doubly stochastic variational bayes for non-conjugate inference,” in *International conference on machine learning*, 2014, pp. 1971–1979.
- A. M. Saxe, J. L. McClelland, and S. Ganguli, “Exact solutions to the nonlinear dynamics of learning in deep linear neural networks,” *Proceedings of the International Conference on Learning Representations*, 2014.
- B. Poole, S. Lahiri, M. Raghu, J. Sohl-Dickstein, and S. Ganguli, “Exponential expressivity in deep neural networks through transient chaos,” in *Advances in Neural Information Processing Systems*, 2016, pp. 3360–3368.
- A. Wu, S. Nowozin, E. Meeds, R. E. Turner, J. M. Hernandez-Lobato, and A. L. Gaunt, “Deterministic variational inference for robust bayesian neural networks,” in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=B1108oAct7>
- S. S. Schoenholz, J. Gilmer, S. Ganguli, and J. Sohl-Dickstein, “Deep information propagation,” *Proceedings of the International Conference on Learning Representations*, 2017.
- B. Lakshminarayanan, A. Pritzel, and C. Blundell, “Simple and scalable predictive uncertainty estimation using deep ensembles,” in *Advances in Neural Information Processing Systems*, 2017, pp. 6402–6413.

Appendices

A Signal propagation in BNNs

We consider the quantity $\mathbb{E}[(h_j^l)^2]$ during a forward pass. This quantity is the second moment of a single hidden unit h_j^l in layer l , consisting of D_{l-1} incoming connections, where the expectation taken is over the network parameters and is given by

$$\begin{aligned}\mathbb{E}[(h_j^l)^2] &= \mathbb{E} \left[\left(\sum_{i=1}^{D_{l-1}} w_{ij}^l x_i^l + b_j^l \right)^2 \right] \\ &= \sum_{i=1}^{D_{l-1}} \mathbb{E}[(w_{ij}^l)^2] \mathbb{E}[(x_i^l)^2] + \mathbb{E}[(b_j^l)^2].\end{aligned}\quad (4)$$

We define

$$\sigma_{\tilde{q}_j}^l = \sum_{i=1}^{D_{l-1}} (\sigma_{ij}^l)^2 \quad \text{and} \quad \mu_{\tilde{q}_j}^l = \sum_{i=1}^{D_{l-1}} \mu_{ij}^l \quad (5)$$

to use as the statistics to describe $\sum_{i=1}^{D_{l-1}} w_{ij}$. Note that while it is true that we can write

$$\sum_{i=1}^{D_{l-1}} w_{ij}^l \sim \mathcal{N} \left(\mu_{\tilde{q}_j}^l, (\sigma_{\tilde{q}_j}^l)^2 \right) \quad (6)$$

at initialisation, in our analysis of the network at an arbitrary stage of the stage of training the i.i.d. assumption of the central limit theorem (CLT) does not strictly hold. As in Wu et al. (2019), we empirically find that some form of the CLT holds for the hidden units during training. We thus continue to approximate the expectation with a Gaussian according to the CLT.

Next, in designing $\tilde{q}_\beta(W)$ we scale the variance and impose $\mathbb{E}[(w_{ij}^l)^2] = \frac{(\mu_{\tilde{q}}^l)^2 + (\sigma_{\tilde{q}}^l)^2}{D_{l-1}}$, $\forall i, j$ to ensure that the variance is bounded in the infinite width limit Schoenholz et al. (2017). This also allows the variance propagated forward to be independent of the layer width. We now have

$$\mathbb{E}[(h_j^l)^2] = ((\mu_{\tilde{q}_j}^l)^2 + (\sigma_{\tilde{q}_j}^l)^2) \frac{1}{D_{l-1}} \sum_{i=1}^{D_{l-1}} g(h_i^{l-1})^2 + ((\mu_b^l)^2 + (\sigma_b^l)^2). \quad (7)$$

As $D_{l-1} \rightarrow \infty$, h_j^l becomes an infinite sum of i.i.d. random variables and becomes Gaussian distributed according to the CLT. We can thus write

$$\mathbb{E}[(h_j^l)^2] = ((\mu_{\tilde{q}}^l)^2 + (\sigma_{\tilde{q}}^l)^2) \mathbb{E}_z[\phi(\tau^{l-1} + \sqrt{\nu^{l-1}}z)^2] + (\mu_b^l)^2 + (\sigma_b^l)^2 \quad (8)$$

where $z \sim \mathcal{N}(0, 1)$, and τ^{l-1} and ν^{l-1} are the incoming signal to layer l 's mean and variance respectively. If we use ReLU as activation, i.e. $g(a) = \max(0, a)$, then

$$\begin{aligned}\mathbb{E}[(h_j^l)^2] &= ((\mu_{\tilde{q}}^l)^2 + (\sigma_{\tilde{q}}^l)^2) \left\{ \int_{-\infty}^{\infty} \Phi(z) \phi(\tau^{l-1} + \sqrt{\nu^{l-1}}z)^2 dz \right\} + (\mu_b^l)^2 + (\sigma_b^l)^2 \\ &= ((\mu_{\tilde{q}}^l)^2 + (\sigma_{\tilde{q}}^l)^2) \left\{ \int_0^{\infty} \Phi(z) \left((\tau^{l-1})^2 + 2\tau^{l-1}\sqrt{\nu^{l-1}}z + \nu^{l-1}z^2 \right) dz \right\} \\ &\quad + (\mu_b^l)^2 + (\sigma_b^l)^2 \\ &= ((\mu_{\tilde{q}}^l)^2 + (\sigma_{\tilde{q}}^l)^2) \left[\frac{(\tau^{l-1})^2}{2} + \frac{2\tau^{l-1}\sqrt{\nu^{l-1}}}{\sqrt{2\pi}} + \frac{\nu^{l-1}}{2} \right] + (\mu_b^l)^2 + (\sigma_b^l)^2\end{aligned}\quad (9)$$

where $\Phi(z) = \frac{e^{-z^2/2}}{\sqrt{2\pi}}$.

Similarly, we can show that the relevant statistics governing signal propagation in the forward pass are given by

$$\begin{aligned}\mathbb{E}[h_j^l] &= \mathbb{E}\left[\sum_{j=1}^{D_{l-1}} w_{ij}^l g(h_j^{l-1} + b_i)\right] \\ &= \mu_{\tilde{q}_j}^l \left(\frac{\tau^{l-1}}{2} + \sqrt{\frac{\nu^{l-1}}{2\pi}}\right) + \mu_b^l\end{aligned}\quad (10)$$

and

$$\begin{aligned}\nu_j^l &= \mathbb{E}\left[\left(\sum_{j=1}^{D_{l-1}} w_{ij}^l g(h_j^{l-1})\right)^2\right] - \mathbb{E}\left[\sum_{j=1}^{D_{l-1}} w_{ij}^l g(h_j^{l-1})\right]^2 \\ &= (\mu_{\tilde{q}_j}^l)^2 \left[\frac{(\tau^{l-1})^2}{4} + \tau^{l-1} \sqrt{\frac{\nu^{l-1}}{2\pi}} + \left(1 - \frac{1}{\pi}\right) \frac{\nu^{l-1}}{2}\right] \\ &\quad + (\sigma_{\tilde{q}_j}^l)^2 \left[\frac{(\tau^{l-1})^2}{2} + 2\tau^{l-1} \sqrt{\frac{\nu^{l-1}}{2\pi}} + \frac{\nu^{l-1}}{2}\right] + (\sigma_b^l)^2.\end{aligned}\quad (11)$$

Described above is the signal propagation dynamics in general. With a mean preserving prior we can only control variance by multiplicatively expanding or squeezing $\sigma_{\tilde{q}_j}^l$. We can only design for conditions that set $\tau^{l-1} = 0$ and $\sigma_b = 0$ which is true at initialisation but starts to break down during training. In order to continue we thus implicitly assume that during training: (1) the mean of the summed weights' means across a hidden layer's pre-activation remain mean zero i.e. $\mathbb{E}[(\sum_{j=1}^{D_{l-1}} \mu_{q_j}^l)] = 0$; (2) biases are zero (note, it is possible to absorb the biases by augmenting the input at each layer with an additional column of ones, this yields more stable signal propagation. We find that treating biases as deterministic parameters aids in training and outweighs the minor gain in stabilising propagation). This allows us to write the variance ν_j^l as

$$\nu_j^l = \left[\left(1 - \frac{1}{\pi}\right) (\mu_{\tilde{q}_j}^l)^2 + (\sigma_{\tilde{q}_j}^l)^2\right] \frac{\nu^{l-1}}{2}.\quad (12)$$

From here we can design $\tilde{q}_{\{\alpha, \phi\}}(W)$ to preserve variances through the network. We find actually forcing our assumptions and setting parameter means and biases to zero does not train.

The parameters for the joint distribution $\tilde{q}_{\{\alpha, \phi\}}(W)$ are determined by α and ϕ as the product of two Gaussian pdfs as

$$\mu_{\tilde{q}_{ij}}^l = \frac{\mu_{p_{ij}}^l (\sigma_{p_{ij}}^l)^2 + \mu_{q_{ij}}^l (\sigma_{q_{ij}}^l)^2}{(\sigma_{p_j}^l)^2 + (\sigma_{q_{ij}}^l)^2}, \quad \sigma_{\tilde{q}_{ij}}^l = \sqrt{\frac{(\sigma_{p_j}^l)^2 (\sigma_{q_{ij}}^l)^2}{(\sigma_{p_j}^l)^2 + (\sigma_{q_{ij}}^l)^2}}.\quad (13)$$

To stabilise the signal, we want to find prior parameters $\alpha \in \{\mu_{p_{ij}}^l, (\sigma_{p_j}^l)^2\}$ that can preserve the variance during the forward pass. Specifically, we want α to ensure $\nu_j^l = \nu_j^{l-1}$ or $\text{Var}[h_j^l] = \text{Var}[h_j^{l-1}]$ in (2) where $j' \in \{1, \dots, D_{l-1}\}$. To achieve this, we require $\mu_{p_{ij}}^l = \mu_{q_{ij}}^l$. Secondly, we find variance parameters of $(\sigma_{p_j}^l)^2$ using (13), to satisfy the condition $(1 - 1/\pi)(\mu_{\tilde{q}_j}^l)^2 + (\sigma_{\tilde{q}_j}^l)^2 = 2$ found by setting $\nu_j^l = \nu_j^{l-1}$ in 12.

B Reformulating the ELBO

We reformulate the ELBO by lower bounding the log marginal likelihood of the data as follows

$$\begin{aligned}\log p(\mathbf{y}|\mathbf{x}) &= \log \int p(\mathbf{y}|\mathbf{x}, W) p_\alpha(W) dW \\ &= \log \int p(\mathbf{y}|\mathbf{x}, W) p_\alpha(W) \frac{q_\phi(W)}{q_\phi(W)} dW,\end{aligned}\quad (14)$$

where we combine the prior and approximating posterior as $\tilde{q}_{\{\alpha, \phi\}}(W) = p_\alpha(W)q_\phi(W)/Z$, where Z is a normalisation constant. Then, we can construct a lower bound making use of Jensen’s inequality which gives

$$\begin{aligned} \log p(\mathbf{y}|\mathbf{x}) &= \log \int \tilde{q}_{\{\alpha, \phi\}}(W) \frac{p(\mathbf{y}|\mathbf{x}, W)}{q_\phi(W)} Z dW \\ &\geq \int \tilde{q}_{\{\alpha, \phi\}}(W) \log \frac{p(\mathbf{y}|\mathbf{x}, W)}{q_\phi(W)} Z dW \\ &= \mathbb{E}_{\tilde{q}_{\{\alpha, \phi\}}(W)} [\log p(\mathbf{y}|\mathbf{x}, W)] - \mathbb{E}_{\tilde{q}_{\{\alpha, \phi\}}(W)} [\log q_\phi(W)] + \log Z. \end{aligned} \quad (15)$$

By reformulating the ELBO in this way, we can estimate the above expectations using a Monte Carlo estimator with samples drawn from $\tilde{q}_{\{\alpha, \phi\}}(W)$ instead of $q_\phi(W)$. This ensures that the sampled weights of the network are being influenced by the current prior $p_\alpha(W)$ during the forward pass. Finally, we ignore the constant term in (15) and replace the cross-entropy term, which is the second term in the equation, with a KL term by introducing the identity ratio $p_\alpha(W)/p_\alpha(W)$ in (14). This gives the following objective when reparametrised using the reparametrisation trick:

$$\mathcal{L}_{\tilde{q}} := \mathbb{E}_{p(\epsilon)} [\log p(\mathbf{y}|\mathbf{x}, \mathbf{b}, W = \xi(\epsilon, \alpha, \phi))] - \text{KL}(\tilde{q}_{\{\alpha, \phi\}}(W) || p_\alpha(W)), \quad (16)$$

where $\epsilon \sim \mathcal{N}(0, I)$. The prior $p_\alpha(W)$ usually impacts the ELBO through an additive KL term, only affecting the weights after the forward pass, having no effect on the signal propagation dynamics of the network. We instead propose priors that also exert their influence *during* the forward pass, so as to promote stable signal propagation, and improve robustness in deep BNNs.

C Use of $\tilde{q}(W)$ at test time

We opt to always include the prior during test time: we sample from $\tilde{q}(W)$ instead of $q(W)$, defining a new posterior of interest, because we find that the accuracy of the model progresses faster. Since we require adjustment during the forward pass throughout training to ensure stable signal propagation (it is often necessary for any training to occur), it seems reasonable to include it in the forward pass at test time. The use of $q(W)$ exhibits performance similar to networks with unstable signal propagation. Note that the adjustment to the variational posterior for signal propagation is unlike the use of EB to choose the hyperparameters maximizing the likelihood, in that case the prior should not be factored in and $q(W)$ would be appropriate. We investigate what effect including the prior has on the quality of prediction uncertainty in the experiments in Appendix D.

D Appendix Experiments: Signal Propagation and Quality of Uncertainty

Signal propagation. We examine signals in a ReLU network in Figure 3 to analyse the effect of the prior on the network signal propagation. We monitor the variance dynamics of the same data point throughout training by calculating the empirical variance of the vector of pre-activations at each layer during a forward pass. In Figures 3 (a) and (b) we show a controlled example where we force our assumptions, setting biases and parameter means to zero and see that the prior preserves the signal, whereas in a standard BNN it explodes. Furthermore, we show a typical training scenario in Figures 3 (c) and (d), where we see that our assumptions hold in the early stages of training and start to break down later in training, yielding less stable signal propagation.

Quality of uncertainty. Finally, we turn to the issue of what effect this prior has on uncertainty and calibration. We measure calibration with the Brier score and, similar to Lakshminarayanan et al. (2017), the accuracy of predictions above 50% and 90% confidence to see whether our models tend towards overconfidence. We monitor the progression of these metrics of models with different priors through 100 epochs reported in Figure 4. As with any iteratively updating prior, we expect that it may adapt to the dataset and overfit, as is shown to be true of our stabilising prior and EB in Figure 4. As an answer to this we explore combining a regularising and stabilising prior which trains faster and results in a well calibrated model with better Brier scores than any solitary prior.

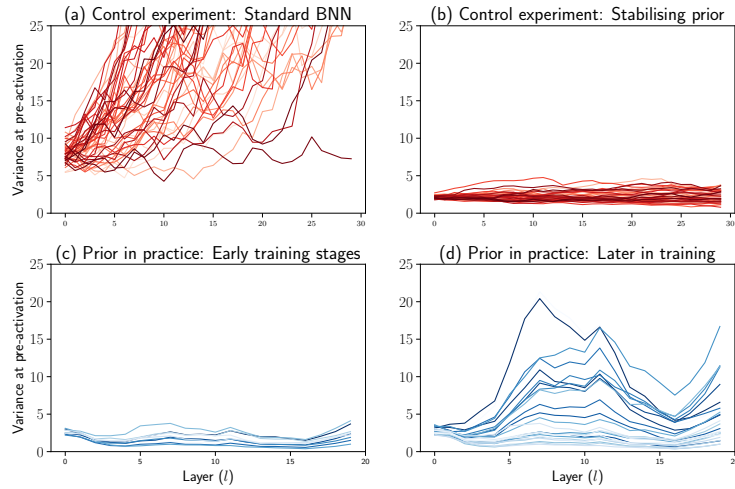


Figure 3: Signal propagation dynamics of the same signal propagated through different networks. We track the variance of the same data point throughout training by calculating the empirical variance of the vector at the pre-activation at each layer. In a controlled setting we can achieve perfectly stable signal propagation. In practice our assumptions hold for the early stages of training.

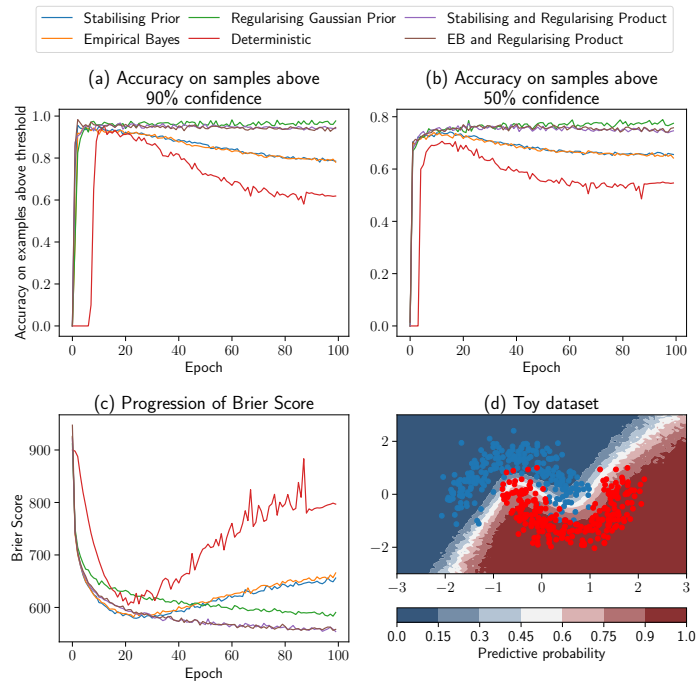


Figure 4: Uncertainty and calibration experiments on CIFAR-10. Iteratively updating priors overfit, however, we can combine regularising and optimal priors to maintain calibrated confidence and better Brier scores. In (d) we also see we are able to get reasonable uncertainty estimates with a deep neural network on a toy dataset.