
Sampling-Free Variational Inference of Bayesian Neural Networks by Variance Backpropagation

Manuel Haußmann¹ Fred A. Hamprecht¹ Melih Kandemir^{2,3}

¹HCI/IWR, Heidelberg University, Germany

²Özyeğin University, Istanbul, Turkey

³Bosch Center for Artificial Intelligence, Renningen, Germany

Abstract

We propose a new Bayesian Neural Net formulation that affords variational inference for which the evidence lower bound is analytically tractable subject to a tight approximation. We achieve this tractability by (i) decomposing ReLU nonlinearities into the product of an identity and a Heaviside step function, (ii) introducing a separate path that decomposes the neural net expectation from its variance. We demonstrate formally that introducing separate latent binary variables to the activations allows representing the neural network likelihood as a chain of linear operations. Performing variational inference on this construction enables a sampling-free computation of the evidence lower bound which is a more effective approximation than the widely applied Monte Carlo sampling and CLT related techniques. We evaluate the model on a range of regression and classification tasks against BNN inference alternatives, showing competitive or improved performance over the current state-of-the-art.

1 INTRODUCTION

The advent of deep learning libraries (Abadi et al., 2015; Theano Development Team, 2016; Paszke et al., 2017) has made fast prototyping of novel neural net architectures possible by writing short and simple high-level code. Their availability triggered an explosion of research output on application-specific neural net design, which in turn allowed for fast improvement of predictive performance in almost all fields where machine learning is used. The next grand challenge is to solve mainstream machine learning tasks with more time-efficient, energy-efficient, and interpretable models that make predictions with attached uncertainty estimates. Industry-scale applications

also require models that are robust to adversarial perturbations (Szegedy et al., 2014; Goodfellow et al., 2015).

The Bayesian modeling approach provides principled solutions to all of the challenges mentioned above of machine learning. Bayesian Neural Networks (BNNs) (MacKay, 1992) lie at the intersection of deep learning and the Bayesian approach that learns the parameters of a machine learning model via posterior inference (MacKay, 1995; Neal, 1995). A deterministic net with an arbitrary architecture and loss function can be upgraded to a BNN simply by placing a prior distribution over its parameters turning them into random variables.

Unfortunately, the non-linear activation functions at the layer outputs render direct methods to estimate the posterior distribution of BNN weights analytically intractable. A recently established technique for approximating this posterior is Stochastic Gradient Variational Bayes (SGVB) (Kingma and Welling, 2014), which suggests reparameterizing the variational distribution and then Monte Carlo integrating the intractable expected data fit part of the ELBO. Sample noise for a cascade of random variables, however, distorts the gradient signal, leading to unstable training. Improving the sampling procedure to reduce the variance of the gradient estimate is an active research topic. Recent advances in this vein include the local reparameterization trick (Kingma et al., 2015) and variance reparameterization (Molchanov et al., 2017; Neklyudov et al., 2017).

We here follow a second research direction (Hernández-Lobato and Adams, 2015; Kandemir, 2018; Wu et al., 2019) of deriving approaches that avoid Monte Carlo sampling and the associated precautions required for variance reduction, and present a novel BNN construction that makes variational inference possible with a closed-form ELBO. Without a substantial loss of generality, we restrict the activation functions of all neurons of a net to the Rectified Linear Unit (ReLU). We build our formulation on the fact that the ReLU function can be ex-

pressed as the product of the identity function and the Heaviside step function: $\max(0, x) = x \cdot \mathbf{1}(x)$. Following [Kandemir \(2018\)](#) we exploit the fact that we are devising a probabilistic learner and introduce latent variables z to mimic the deterministic the Heaviside step functions. This can then be relaxed to a Bernoulli distribution $z \sim \delta_{x>0} \approx \text{Ber}(\sigma(Cx))$ with some $C \gg 0$ and the logistic sigmoid function $\sigma(\cdot)$. The idea is illustrated in [Figure 1](#). We show how the asymptotic account of this relaxation converts the likelihood calculation into a chain of linear matrix operations, giving way to a closed-form computation of the data fit term of the Evidence Lower Bound in mean-field variational BNN inference. In our construction, the data fit term lends itself as the sum of a standard neural net loss (e.g. mean-squared error) on the expected prediction output and the predictor variance. This predictor variance term has a recursive form, describing how the predictor variance back-propagates through the layers of a BNN. We refer to our model as *Variance Back-Propagation* (VBP).

Experiments on several regression and classification tasks show that VBP can perform competitive to and improve upon other recent sampling-free or sampling-based approaches to BNN inference. Last but not least, VBP presents a generic formulation that is directly applicable to all weight prior selections as long as their Kullback-Leibler (KL) divergence with respect to the variational distribution is available in closed form, including the common log-Uniform ([Kingma et al., 2015](#); [Molchanov et al., 2017](#)), Normal ([Hernández-Lobato and Adams, 2015](#)), and horseshoe ([Louizos et al., 2017](#)) priors.

2 BAYESIAN NEURAL NETS WITH DECOMPOSED FEATURE MAPS

Given a data set $\mathcal{D} = \{(\mathbf{x}_n, y_n)_{n=1}^N\}$ consisting of N pairs of d -dimensional feature vectors \mathbf{x}_n and targets y_n , the task is to learn in a regression setting¹ with a normal likelihood

$$\begin{aligned} \mathbf{w} &\sim p(\mathbf{w}), \\ \mathbf{y}|\mathbf{X}, \mathbf{w} &\sim \mathcal{N}(\mathbf{y}|f(\mathbf{X}; \mathbf{w}), \beta^{-1}\mathbf{1}), \end{aligned}$$

for $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ and $\mathbf{y} = \{y_1, \dots, y_N\}$, with β as the observation precision, and $\mathbf{1}$ an identity matrix of suitable size. The function $f(\cdot; \mathbf{w})$ is a feed-forward multi-layer neural net parameterized by weights \mathbf{w} , and ReLU activations between the hidden layers. $p(\mathbf{w})$ is an arbitrary prior over these weights.




Figure 1: **ReLU decomposition.** We decompose the ReLU function into an identity function and a Heaviside step function, which is in turn approximated with a Bernoulli distribution.

2.1 THE IDENTITY-HEAVISIDE DECOMPOSITION

A ReLU function can be decomposed as $\max(0, u) = u \cdot \mathbf{1}(u)$, with $\mathbf{1}$ being the Heaviside step function

$$\mathbf{1}(x) = \begin{cases} 1, & x > 0 \\ 0, & x \leq 0 \end{cases}.$$

This allows us to express the feature map vector \mathbf{h}^l (post-activation) of a data point at layer $l + 1$ as

$$\begin{aligned} \mathbf{h}^l &= \mathbf{f}^l \circ \mathbf{z}^l \\ \text{with } \mathbf{f}^l &= \mathbf{W}_l \mathbf{h}^{l-1} \text{ and } \mathbf{z}^l = \mathbf{1}(\mathbf{W}_l \mathbf{h}^{l-1}), \end{aligned}$$

where \mathbf{h}^{l-1} is the feature map vector of the same data point at layer $l - 1$, the matrix \mathbf{W}_l contains the weights to map from layers $l - 1$ to l and \circ denotes the element-wise Hadamard product of equal sized matrices or vectors. For $l = 0$, i.e. the input, we set $\mathbf{h}^0 = \mathbf{x}$. \mathbf{f}^l is the linear pre-activation output vector of layer l . We refer to \mathbf{z}^l as the activation vector. When the argument of the $\mathbf{1}(\cdot)$ function takes a vector as its input, we mean its elementwise application to all inputs. We denote the above factorized formulation of the feature map \mathbf{h}^l as the *Identity-Heaviside Decomposition*.

Applying this decomposed expression for example on a feed-forward neural net with two hidden layers, we get²

$$y = f(\mathbf{x}; \mathbf{w}) = \mathbf{w}_3^T (\mathbf{z}^2 \circ \mathbf{W}_2 (\mathbf{z}^1 \circ \mathbf{W}_1 \mathbf{x})) + \varepsilon,$$

with $\varepsilon \sim \mathcal{N}(0, \beta^{-1})$, for a data point consisting of the input-output pair (\mathbf{x}, y) . Note that given the binary activations \mathbf{z}^1 and \mathbf{z}^2 of the step function, the predictor output can be computed following a chain of *linear* operations.

¹See [Section 2.7](#) for the extension to classification.

²We suppress the bias from the notation as it can directly be incorporated by extending the respective weights matrices.

2.2 THE PROBABILISTIC MODEL

For z_{nj}^l —the j th activation at the l th layer for data point n —we can approximate the Heaviside step function with a Bernoulli distribution as follows

$$z_{nj}^l \sim \text{Ber}(z_{nj}^l | \sigma(C \cdot \mathbf{w}_{lj}^T \mathbf{h}_n^l)), \quad (1)$$

where $\sigma(\cdot)$ is the logistic sigmoid function.³ The approximation becomes precise as $C \rightarrow \infty$.

Applying the Identity-Heaviside decomposition and Bernoulli relaxation to an L -layer BNN we obtain

$$\begin{aligned} \mathbf{w}_{lj} &\sim p(\mathbf{w}_{lj}), \quad \forall l, j \\ z_{nj}^l | \mathbf{w}_{lj}, \mathbf{h}_n^l &\sim \text{Ber}(z_{nj}^l | \sigma(C \cdot \mathbf{w}_{lj}^T \mathbf{h}_n^l)), \quad \forall l, n, j \\ \mathbf{y} | \mathbf{w}, \mathbf{X}, \mathbf{Z} &\sim \mathcal{N}(\mathbf{y} | f(\mathbf{X}; \mathbf{w}), \beta^{-1} \mathbf{1}), \end{aligned}$$

where \mathbf{Z} is the overall collection of activation variables.

2.2.1 Relation to V-ReLU-Net

The decomposition of the ReLU activation has been proposed before in the context of learning BNNs as the Variational ReLU Network (V-ReLU-Net) (Kandemir, 2018). The goal of that work is to derive a gradient-free closed-form variational inference scheme updating each variable to the local optimum conditioned on the other variables. This necessitates not only the decomposition we described above but also a further mean-field decoupling of the BNN layers. Translated to our notation, the V-ReLU-Net has for h_{nj}^l —the post-activation unit j of layer l and data point n —the following model

$$\begin{aligned} \mathbf{w}_{lj} &\sim \mathcal{N}(\mathbf{w}_{lj} | 0, \alpha^{-1} \mathbf{1}), \\ f_{nj}^l | \mathbf{w}_{lj}, \mathbf{h}_n^{l-1} &\sim \mathcal{N}(f_{nj}^l | \mathbf{w}_{lj}^T \mathbf{h}_n^{l-1}, \beta^{-1}), \\ z_{nj}^l | f_{nj}^l &\sim \text{Ber}(z_{nj}^l | \sigma(C f_{nj}^l)), \\ h_{nj}^l | z_{nj}^l, f_{nj}^l &\sim \mathcal{N}(h_{nj}^l | z_{nj}^l f_{nj}^l, \gamma^{-1}), \end{aligned}$$

where α, β, γ are fixed hyperparameters. This means introducing additional distributions over each pre- & post-activation unit of each layer in the network. While such factorization across layers enjoys gradient-free variational update rules, it suffers from poor local maxima due to lack of direct feedback across non-neighboring layers. Our formulation loosens this update scheme in the way shown in the next sections. Our experiments show that the benefit of gradient-free closed-form updates tends to not be worth the added constraints placed on the BNN for our kind of problems.

³ $\sigma(x) = (1 + \exp(-x))^{-1}$

2.3 VARIATIONAL INFERENCE OF THE POSTERIOR

In the Bayesian context, learning consists of inferring the posterior distribution over the free parameters of the model $p(\theta | \mathcal{D}) = p(\mathbf{y} | \theta, \mathbf{X})p(\theta) / \int p(\mathbf{y} | \theta, \mathbf{X})p(\theta)d\theta$ which is intractable for neural nets due to the integral in the denominator. Hence we need to resort to approximations. Our study focuses on variational inference due to its computational efficiency. It approximates the true posterior by a proxy distribution $q_\phi(\theta)$ with a known functional form parameterized by ϕ and minimizes the Kullback-Leibler divergence between $q_\phi(\theta)$ and $p(\theta | \mathcal{D})$

$$\text{KL}[q_\phi(\theta) || p(\theta | \mathcal{D})].$$

After a few algebraic manipulations, minimizing this KL divergence and maximizing the functional below turn out to be equivalent problems

$$\mathcal{L}_{\text{elbo}} = \underbrace{\mathbb{E}_{q_\phi(\theta)}[\log p(\mathbf{y} | \theta, \mathbf{X})]}_{\mathcal{L}_{\text{data}}} - \underbrace{\text{KL}[q_\phi(\theta) || p(\theta)]}_{\mathcal{L}_{\text{reg}}}. \quad (2)$$

This functional is often referred to as the *Evidence Lower Bound (ELBO)*, as it is a lower bound to the log marginal distribution $\log p(\mathbf{y} | \mathbf{X})$ known as the evidence. The ELBO has the intuitive interpretation that $\mathcal{L}_{\text{data}}$ is responsible for the data fit, as it maximizes the expected log-likelihood of the data, and \mathcal{L}_{reg} serves as a complexity regularizer by punishing unnecessary divergence of the approximate posterior from the prior.

In our setup, θ consists of the global weights \mathbf{w} and the local activations \mathbf{Z} . For the weights we follow prior art (Kingma et al., 2015; Hernández-Lobato and Adams, 2015; Molchanov et al., 2017) and adopt the mean-field assumption that the variational distribution factorizes across the weights. We assume each variational weight distribution to follow $q(w_{ij}^l) = \mathcal{N}(w_{ij}^l | \mu_{ij}^l, (\sigma_{ij}^l)^2)$ and parameterize the individual variances via their logarithms to avoid the positivity constraint giving us $\phi = \{(\mu_{ij}^l, \log \sigma_{ij}^l)_{ijl}\}$. We also assign an individual factor to each z_{nj}^l local variable. Rather than handcrafting the functional form of this factor, we calculate its ideal form having other factors fixed, as detailed in Section 2.5. The final variational distribution is given as

$$q(\mathbf{Z})q_\phi(\mathbf{W}) = \prod_{n=1}^N \prod_{l=1}^L \prod_{i=1}^{J_{l-1}} \prod_{j=1}^{J_l} q(z_{nj}^l)q_\phi(w_{ij}^l), \quad (3)$$

where J_l denotes the number of units at layer l .

This allows us to rewrite $\mathcal{L}_{\text{elbo}}$ as

$$\begin{aligned} \mathcal{L}_{\text{elbo}} &= \mathbb{E}_{q(\mathbf{Z})q_\phi(\mathbf{W})} [\log p(\mathbf{y} | \mathbf{W}, \mathbf{Z}, \mathbf{X})] \\ &\quad - \mathbb{E}_{q_\phi(\mathbf{W})} [\text{KL}(q(\mathbf{Z}) || p(\mathbf{Z} | \mathbf{W}, \mathbf{X}))] \\ &\quad - \text{KL}(q(\mathbf{W}) || p(\mathbf{W})), \end{aligned}$$

splitting it into three terms. As our ultimate goal is to obtain the ELBO in closed form, we have that for the third term any prior on weights that lends itself to an analytical solution of $\text{KL}(q(\mathbf{W})||p(\mathbf{W}))$ is acceptable. We have a list of attractive and well-settled possibilities to choose from, including: i) the Normal prior (Blundell et al., 2015) for mere model selection, ii) the log-Uniform prior (Kingma et al., 2015; Molchanov et al., 2017) for atomic sparsity induction and aggressive synaptic connection pruning, and iii) the horseshoe prior (Louizos and Welling, 2017) for group sparsity induction and neuron-level pruning. In this work, we stick to a simple Normal prior to be maximally comparable to our baselines.

We will discuss the second term of $\mathcal{L}_{\text{elbo}}$, $\mathbb{E}_{q_\phi(\mathbf{w})} [\text{KL}(q(\mathbf{Z})||p(\mathbf{Z}|\mathbf{W}, \mathbf{X}))]$ in greater detail in Proposition 2, which leaves the first term that is responsible for the data fit. For our regression likelihood, we can decompose $\mathcal{L}_{\text{data}}$ as

$$\begin{aligned} \mathcal{L}_{\text{data}} &= \mathbb{E}_{q(\mathbf{Z})q_\phi(\mathbf{w})} [\log p(\mathbf{y}|\mathbf{W}, \mathbf{Z}, \mathbf{X})] \\ &\stackrel{c}{=} -\frac{\beta}{2} \sum_{n=1}^N \mathbb{E}_{q(\mathbf{Z})q_\phi(\mathbf{w})} \left[(y_n - f(\mathbf{x}_n; \mathbf{w}))^2 \right] \\ &= -\frac{\beta}{2} \sum_{n=1}^N \left\{ (y_n - \mathbb{E}_{q(\mathbf{Z})q_\phi(\mathbf{w})} [f(\mathbf{x}_n; \mathbf{w})])^2 \right. \\ &\quad \left. + \text{var}_{q(\mathbf{Z})q_\phi(\mathbf{w})} [f(x_n; \mathbf{w})] \right\}, \quad (4) \end{aligned}$$

where $\stackrel{c}{=}$ indicates equality up to an additive constant.

In this form, the first term is the squared error evaluated at the mean of the predictor $f(\cdot)$ and the second term is its variance, which infers the total amount of model variance to account for the epistemic uncertainty in the learning task (Kendall and Gal, 2017). A sampling-free solution to (4) therefore translates to the requirement of an analytical solution to the expectation and variance terms.

2.4 CLOSED-FORM CALCULATION OF THE DATA FIT TERM

2.4.1 The Expectation Term

When all feature maps of the predictor are Identity-Heaviside decomposed and the distributions over the z_{nj}^l 's are approximated by Bernoulli distributions as described in Section 2.2, the expectation term

$$\mathbb{E}_{q(\mathbf{Z})q_\phi(\mathbf{w})} [f(x_n; \mathbf{w})] \quad (5)$$

can be calculated in closed form, as $f(\mathbf{x}_n; \mathbf{w})$ consists only of linear operations over independent variables according to our mean-field variational posterior with which the expectation can commute operation orders. This order interchangeability allows us to compute the expectation

term in a single forward pass where each weight takes its mean value with respect to its related factor in the approximate distribution $q(\mathbf{Z})q_\phi(\mathbf{W})$. For instance, for a Bayesian neural net with two hidden layers, we have

$$\begin{aligned} &\mathbb{E}_{q(\mathbf{Z})q_\phi(\mathbf{w})} [f(\mathbf{x}; \mathbf{w})] \\ &= \mathbb{E}_{q(\mathbf{Z})q_\phi(\mathbf{w})} [\mathbf{w}_3^T (\mathbf{z}_n^2 \circ \mathbf{W}_2 (\mathbf{z}_n^1 \circ \mathbf{W}_1 \mathbf{x}_n))] \\ &= \mathbb{E}[\mathbf{w}_3]^T \left(\mathbb{E}[\mathbf{z}_n^2] \circ \mathbb{E}[\mathbf{W}_2] \left(\mathbb{E}[\mathbf{z}_n^1] \circ \mathbb{E}[\mathbf{W}_1] \mathbf{x}_n \right) \right). \end{aligned}$$

Consequently, the squared error part of the data fit term can be calculated in closed form. This interchangeability property of linear operations against expectations holds as long as we keep independence between the layers, hence it could also be extended to a non-mean-field case.

2.4.2 The Variance Term Calculated via Recursion

The second term in (4) that requires an analytical solution is the variance

$$\text{var}_{q(\mathbf{Z})q_\phi(\mathbf{w})} [f(x_n; \mathbf{w})]. \quad (6)$$

Its derivation comes after using the following two identities on the relationship between the variances of two independent random variables a and b :

$$\begin{aligned} \text{var}[a + b] &= \text{var}[a] + \text{var}[b], \quad (7) \\ \text{var}[a \cdot b] &= \text{var}[a]\text{var}[b] + \mathbb{E}[a]^2\text{var}[b] \\ &\quad + \text{var}[a]\mathbb{E}[b]^2, \\ &= \mathbb{E}[a^2]\text{var}[b] + \text{var}[a]\mathbb{E}[b]^2. \quad (8) \end{aligned}$$

Applying these well-known identities to the linear output layer activations \mathbf{f}^L of the n th data point⁴ we have

$$\begin{aligned} \text{var}[\mathbf{f}^L] &= \text{var}[\mathbf{w}_L^T \mathbf{h}^{L-1}] = \text{var} \left[\sum_{j=1}^{J_L} w_{Lj} h_j^{L-1} \right] \\ &= \sum_{j=1}^{J_L} \mathbb{E}[w_{Lj}^2] \text{var}[h_j^{L-1}] + \text{var}[w_{Lj}] \mathbb{E}[h_j^{L-1}]^2. \end{aligned}$$

Given the normal variational posterior over the weights, we directly get

$$\begin{aligned} \mathbb{E}[w_{Lj}^2] &= \mu_{Lj}^2 + \sigma_{Lj}^2 \\ \text{var}[w_{Lj}] &= \sigma_{Lj}^2, \end{aligned}$$

while $\mathbb{E}[h_j^{L-1}]$ can be computed as described in Section 2.4.1. For $\text{var}[h_j^{L-1}]$ finally we can use the second variance identity again and arrive at

$$\begin{aligned} \text{var}[h_j^{L-1}] &= \text{var}[z_j^{L-1} \cdot f_j^{L-1}] \\ &= \mathbb{E}[(z_j^{L-1})^2] \text{var}[f_j^{L-1}] + \text{var}[z_j^{L-1}] \mathbb{E}[f_j^{L-1}]^2. \end{aligned}$$

⁴We suppress the n index throughout following derivations

Combining these results we have for the output j of an arbitrary hidden layer l that

$$\begin{aligned}
\text{var}[h_j^l] &= \text{var}\left[\sum_{i=1}^{J_l} z_i^l w_{ij}^l h_i^{l-1}\right] = \sum_{i=1}^{J_l} \text{var}\left[z_i^l w_{ij}^l h_i^{l-1}\right] \\
&= \sum_{i=1}^{J_l} \mathbb{E}[(z_i^l)^2] \text{var}[w_{ij}^l h_i^{l-1}] \\
&\quad + \text{var}[z_i^l] (\mathbb{E}[w_{ij}^l] \mathbb{E}[h_i^{l-1}])^2 \\
&= \sum_{i=1}^{J_l} \mathbb{E}[(z_i^l)^2] \left\{ \mathbb{E}[(w_{ij}^l)^2] \text{var}[h_i^{l-1}] \right. \\
&\quad \left. + \text{var}[w_{ij}^l] \mathbb{E}[h_i^{l-1}]^2 \right\} \\
&\quad + \text{var}[z_i^l] (\mathbb{E}[w_{ij}^l] \mathbb{E}[h_i^{l-1}])^2. \quad (9)
\end{aligned}$$

Assuming that we can evaluate $\mathbb{E}[(z_i^l)^2]$ and $\text{var}[z_i^l]$, which we discuss in the next section, the only term left to evaluate is the variance of the activations at the previous layer $\text{var}[h_i^{l-1}]$. Hence, we arrive at a *recursive* description of the model variance. Following this formula, we can express $\text{var}[f(x_n; \mathbf{w})]$ as a function of $\text{var}[\mathbf{h}_n^{L-1}]$, then $\text{var}[\mathbf{h}_n^{L-1}]$ as a function of $\text{var}[\mathbf{h}_n^{L-2}]$, and repeat this procedure until the observed input layer, where variance is zero $\text{var}[\mathbf{h}_n^0] = \text{var}[\mathbf{x}_n] = 0$. For noisy inputs, the desired variance model of the input data can be directly injected to the input layer, which would still not break the recursion and keep the formula valid. Computing this variance term thus only requires a second pass through the network in addition to the one required when the expectation term is computed, sharing many of the required calculations. As this formula reveals how the analytical variance computation recursively back-propagates through the layers, we refer to our construction as *Variance Back-Propagation* (VBP).

Learning the parameters ϕ of the variational posterior $q_\phi(\mathbf{W})$ to maximize this analytical form of the ELBO then follows via mini-batch stochastic gradient descent.

2.5 UPDATING THE BINARY ACTIVATIONS

The results so far are analytical contingent upon having the existence of a tractable expression for each of the $q(z_{nj}^l)$ factors in the variational posterior in Equation (3). While we update the variational parameters of the weight factors via gradient descent of the ELBO, for the binary activation distributions $q(z_{nj}^l)$, we choose to perform the update at the function level. Benefiting from variational calculus, we fix all other factors in $q(\mathbf{Z})q_\phi(\mathbf{W})$ except for a single $q(z_{nj}^l)$ and find the optimal functional form for this remaining factor. We first devise in Proposition 1 a generic approach for calculating variational update rules

of this sort. The proofs of all propositions can be found in the Appendix.

Proposition 1. *Consider a Bayesian model including the generative process excerpt below*

$$\begin{aligned}
&\dots \\
&a \sim p(a) \\
&z|a \sim \delta_{a>0} \approx \text{Ber}(z|\sigma(Ca)) \\
&b|z, a \sim p(b|g(z, a)) \\
&\dots
\end{aligned}$$

for some arbitrary function $g(z, a)$ and $C \gg 0$. If the variational inference of this model is to be performed with an approximate distribution⁵ $Q = \dots q(a)q(z) \dots$, the optimal closed-form update for z is

$$q(z) \leftarrow \text{Ber}(z|\sigma(C\mathbb{E}_{q(a)}[a])),$$

and for $C \rightarrow \infty$

$$q(z) \leftarrow \delta_{\mathbb{E}_{q(a)}[a]>0}.$$

For our specific case this translates for a finite C to

$$q(z_{nj}^l) \leftarrow \text{Ber}\left(z_{nj}^l | \sigma\left(C \cdot \sum_i \mathbb{E}[w_{ij}^l] \mathbb{E}[h_{ni}^{l-1}]\right)\right),$$

and in the limit⁶

$$q(z_{nj}^l) \leftarrow \delta_{\mathbb{E}[\mathbf{w}_j^l]^T \mathbb{E}[\mathbf{h}_n^{l-1}]>0},$$

involving only terms that are already computed during the forward pass expectation term computation and can be done concurrently to that forward pass. The Bernoulli distribution also provides us analytical expressions for the remaining expectation and variance terms in the computation of Equation (9).

2.5.1 The Expected KL Term on Z

A side benefit of the resultant $q(z_{nj}^l)$ distributions is that the complicated $\mathbb{E}_{q_\phi(\mathbf{W})}[\text{KL}(q(\mathbf{Z})||p(\mathbf{Z}|\mathbf{W}, \mathbf{X}))]$ term can be calculated analytically subject to a controllable degree of relaxation as we devise in Proposition 2.

Proposition 2. *For the model and the inference scheme in Proposition 1 with $q(a) = \mathcal{N}(a|\mu, \sigma^2)$, in the relaxed delta function formulation $\delta_{a>0} \approx \text{Ber}(a|\sigma(Ca))$ with some finite $C > 0$, the expression $\mathbb{E}_{q(a)}[\text{KL}[q(z)||p(z|a)]]$ is (i) approximately analytically tractable and (ii) its magnitude goes to 0 quickly as $|\mu|$ increases, with σ controlling how fast it drops towards 0.*

⁵Note that $q(b)$ might or might not exist depending on whether b is latent or observed.

⁶Note that the expectation of a delta function is the binary outcome of the condition it tests and the variance is zero.




Figure 2: **Visualization of** $\mathbb{E}_{q(a)}[\text{KL}[q(z)||p(z|a)]]$. While the second term in Eq (10) can be computed directly, the expected softplus function is approximated with a large sample for each μ, σ pair. (Here $C = 1$).

This proposition deserves several comments. Firstly, the approximation $\delta_{a>0} \approx \text{Ber}(a|\sigma(Ca))$ is tight even for decently small C values (≈ 10), which allows us to keep a close relationship both to the Identity-Heaviside Decomposition, as well as to the theoretical requirements arising within the proofs as well as the numerical ones within the implementation. Secondly, the Proposition relies on the assumption that a follows a normal distribution. In our case we have for z_{nj}^l that

$$a_{nj}^l = \mathbf{w}_{lj}^T \mathbf{h}_n^l = \sum_{i=1}^{J_l} w_{ij}^l h_{ni}^l.$$

This sum allows us to use a central limit theorem (CLT) argument (Wang and Manning, 2013; Wu et al., 2019), to fulfill this assumption. The relevant $\mu_{nj}^l, \sigma_{nj}^l$ parameters of the can be computed via a moment-matching approach, analogously to our general derivations above. As we show in the appendix, with the Bernoulli relaxation we have that the expression can be simplified to

$$\begin{aligned} \mathbb{E}_{q(a)}[\text{KL}[q(z)||p(z|a)]] \\ = \mathbb{E}_{q(a)}[\log(1 + \exp(Ca))] \\ - \log(1 + \exp(C\mathbb{E}[a])). \end{aligned} \quad (10)$$

This expression drops quickly to zero as $|\mu|$ increases with the size of σ controlling the width of this spread around zero, as we visualize in Figure 2.

In order to arrive at an analytical expression we can further approximate each of the two softplus terms $\text{soft}(\cdot) = \log(1 + \exp(\cdot))$ with their ReLU counterpart, an approxi-

mation that gets tighter as we increase C .

$$\begin{aligned} \mathbb{E}_{q(a)}[\text{soft}(Ca)] - \text{soft}(C\mathbb{E}_{q(a)}[a]) \\ \approx C (\mathbb{E}_{q(a)}[\max(0, a)] - \max(0, \mathbb{E}_{q(a)}[a])). \end{aligned}$$

For this expression in turn we can get the analytical

$$C \left(\mu \Phi \left(\frac{\mu}{\sigma} \right) + \sigma \phi \left(\frac{\mu}{\sigma} \right) - \max(0, \mu) \right), \quad (11)$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the pdf and cdf of a standard Normal distribution.

In practice we drop this expected KL term from the ELBO as it becomes negligible for a sufficiently constrained variance. This soft constrained on the variance terms of the layers however is already enforced through the $\text{var}_{q(\mathbf{Z})_{q_\phi}(\mathbf{w})}[f(x_n; \mathbf{w})]$ term in Equation (4).

2.6 HANDLING CONVOLUTION AND POOLING

As a linear operation, convolution is directly applicable to the VBP formulation by modifying all the sums between weights and feature maps with sliding windows. Doing the same will suffice also for the calculation of $\text{var}[f(\mathbf{x}_n; \mathbf{w})]$. In VBP, one layer affects the next only via sums and products of variables, which is not the case for max-pooling. Even though convolutions are found to be sufficient for building state-of-the-art architectures (Springenberg et al., 2015), we show in the Appendix with Proposition 3 that max-pooling is also directly applicable to VBP by extending Proposition 1.

2.7 HANDLING CLASSIFICATION

For classification, we cannot directly decompose

$$\mathbb{E}_{q(\mathbf{Z})_{q_\phi}(\mathbf{w})}[\log p(\mathbf{y}|\mathbf{W}, \mathbf{Z}, \mathbf{X})],$$

as we did in Equation (4) for regression.

For binary classification, we treat \mathbf{y} as a vector of latent decision margins and squash it with a binary-output likelihood $p(\mathbf{t}|\mathbf{y})$. From (Hensman et al., 2013), the log-marginal likelihood of the resultant bound is given by

$$\begin{aligned} \log p(\mathbf{t}|\mathbf{x}) &= \log \int p(\mathbf{t}|\mathbf{y})p(\mathbf{y}|\mathbf{x})d\mathbf{y} \\ &\geq \log \int p(\mathbf{t}|\mathbf{y}) \exp(\mathcal{L}_{\text{elbo}})d\mathbf{y} = \mathcal{L}_{\text{clsf}}. \end{aligned}$$

Choosing a probit likelihood $p(\mathbf{t}|\mathbf{y}) = \text{Ber}(\mathbf{t}|\Phi(\mathbf{y}))$, the integral becomes tractable. This can also directly be extended to multi-class, multi-label classification.⁷

⁷where for a C class problem $\mathbf{y}_n \in \{0, 1\}^C$

In the experiments we instead focus on multi-class classification with a unique label,⁸ i.e. a categorical likelihood

$$\mathbf{y}|\mathbf{w}, \mathbf{X}, \mathbf{Z} \sim \text{Cat}(\mathbf{y}|\zeta(\mathbf{f}(\mathbf{X}; \mathbf{w}))),$$

where $\zeta(\cdot)$ is the softmax function.⁹ In this case we have

$$\begin{aligned} & \mathbb{E}_{q(\mathbf{Z})q_\phi(\mathbf{w})} [\log p(\mathbf{y}|\mathbf{W}, \mathbf{Z}, \mathbf{X})] \\ &= \sum_{n=1}^N \mathbb{E}_{q(\mathbf{Z})q_\phi(\mathbf{w})} [\mathbf{y}_n^T \mathbf{f}(\mathbf{x}_n; \mathbf{w}) - \text{lse}(\mathbf{f}(\mathbf{x}_n; \mathbf{w}))] \\ &= \sum_{n=1}^N \mathbf{y}_n^T \mathbb{E}[\mathbf{f}(\mathbf{x}_n; \mathbf{w})] - \mathbb{E}[\text{lse}(\mathbf{f}(\mathbf{x}_n; \mathbf{w}))], \end{aligned}$$

where $\text{lse}(\cdot)$ is the logsumexp function.¹⁰ The expectation in the first term can be computed analytically as detailed in Section 2.4.1, while the second is intractable. We follow Wu et al. (2019) and derive the following approximation to the second term using a Taylor expansion

$$\begin{aligned} \mathbb{E}[\text{lse}(\mathbf{f})] &\approx \text{lse}(\mathbb{E}[\mathbf{f}]) \\ &+ \frac{1}{2} \sum_{c=1}^C \text{var}[f_c] \left(\zeta(\mathbb{E}[\mathbf{f}])_c - \zeta(\mathbb{E}[\mathbf{f}]_c^2) \right), \end{aligned}$$

with $\mathbf{f} = \mathbf{f}(\mathbf{x}_n; \mathbf{w})$. The resulting expectation and variance terms can be computed as in the regression case.

2.8 RELATION TO CLT BASED APPROACHES

We close this section with a short comparison with two other state-of-the-art sampling-free approaches and how VBP differs from them. These are *Deterministic Variational Inference* (DVI) (Wu et al., 2019)¹¹ and *Probabilistic Backpropagation* (PBP) (Hernández-Lobato and Adams, 2015). PBP follows an assumed density filtering approach instead of variational inference, but relies as DVI does on a CLT argument as a major part of the pipeline. It is based on the observation that the pre-activations of each layer

$$f_{nj}^l = \mathbf{w}_{lj}^T \mathbf{h}_n^l = \sum_{i=1}^{J_l} w_{ij}^l h_{ni}^l$$

are approximately normally distributed with mean and variance parameters that can be computed by moment matching from the earlier layers as in our case. The mean and variance of the post-activation feature maps

⁸where $\mathbf{y}_n \in \{0, 1\}^C$ with the constraint $\sum_j y_{nj} = 1$

⁹ $\zeta(\mathbf{x})_j = \exp(x_j) / \sum_i \exp(x_i)$

¹⁰ $\text{lse}(\mathbf{x}) = \log \sum_j \exp(x_j)$

¹¹For simplicity we focus only on the homoscedastic, mean-field variation of DVI here.

$h = \max(0, f)$ are tractable for ReLU activations (Frey and Hinton, 1999), i.e.

$$\begin{aligned} \mathbb{E}_{\mathcal{N}(f|\mu, \sigma^2)} [\max(0, f)] &= \mu \Phi\left(\frac{\mu}{\sigma}\right) + \sigma \phi\left(\frac{\mu}{\sigma}\right) \\ \text{var}[\max(0, f)] &= (\mu^2 + \sigma^2) \Phi\left(\frac{\mu}{\sigma}\right) \\ &+ \mu \sigma \phi\left(\frac{\mu}{\sigma}\right) - \mathbb{E}[h]^2, \end{aligned}$$

and can then be propagated forward to the next layer. Due to the decomposition we do not rely on the CLT and get the corresponding expressions as

$$\begin{aligned} \mathbb{E}[h] &= \mathbb{E}[z] \mathbb{E}[f] = \mathbb{E}[z] \mu \\ \text{var}[h] &= \mathbb{E}[z^2] \text{var}[f] + \text{var}[z] \mathbb{E}[f]^2 \\ &= \mathbb{E}[z^2] \sigma^2 + \text{var}[z] \mu^2 \approx \mathbb{E}[z] \sigma^2, \end{aligned}$$

where the value of $\mathbb{E}[z]$ decides whether the mean and variance of the pre-activation f are propagated or blocked.

3 RELATED WORK

Several approaches have been introduced for approximating the intractable posterior of BNNs. One line is model-based Markov Chain Monte Carlo, such as Hamiltonian Monte Carlo (HMC) (Neal, 2010) and Stochastic Gradient Langevin Dynamics (SGLD) (Welling and Teh, 2011). Chen et al. (2017) adapted HMC to stochastic gradients by quantifying the entropy overhead stemming from the stochasticity of mini-batch selection.

While being actively used for a wide spectrum of models, successful application of variational inference to deep neural nets has taken place only recently. The earliest study to infer a BNN with variational inference (Hinton and Camp, 1993) was applicable for only one hidden layer. This limitation has been overcome only recently (Graves, 2011) by approximating intractable expectations by numerical integration. Further scalability has been achieved after SGVB is made applicable to BNN inference using weight reparameterizations (Kingma and Welling, 2014; Rezende et al., 2014).

Dropout has strong connections to variational inference of BNNs (Srivastava et al., 2016). Gal and Ghahramani (2016) developed a theoretical link between a dropout network and a deep Gaussian process (Damianou and Lawrence, 2013) inferred by variational inference. Kingma et al. (2015) showed that extending the Bayesian model selection interpretation of Gaussian Dropout with a log-uniform prior on weights leads to a BNN inferred by SGVB. The proposed model can also be interpreted as an input-dependent dropout (Ba and Frey, 2013; Lee et al., 2018) applied to a linear net. Yet it differs from them and the standard dropout in that the masking variable always

Table 1: **Regression Results.** Average test log-likelihood \pm standard error over 20 random train/test splits.

| N/d | boston 506/13 | concrete 1030/8 | energy 768/8 | kin8nm 8192/8 | naval 11934/16 | power 9568/4 | protein 45730/9 | wine 1599/11 | yacht 308/6 |
|------------|------------------------------------|------------------------------------|------------------------------------|-----------------------------------|-----------------------------------|------------------------------------|------------------------------------|------------------------------------|------------------------------------|
| DVI | -2.58 ± 0.04 | -3.23 ± 0.01 | -2.09 ± 0.06 | 1.01 ± 0.01 | 5.84 ± 0.06 | -2.82 ± 0.00 | -2.94 ± 0.00 | -0.96 ± 0.01 | -1.41 ± 0.03 |
| PBP | -2.57 ± 0.09 | -3.16 ± 0.02 | -2.04 ± 0.02 | 0.90 ± 0.01 | 3.73 ± 0.01 | -2.84 ± 0.01 | -2.97 ± 0.00 | -0.97 ± 0.01 | -1.63 ± 0.02 |
| VarOut | -2.59 ± 0.03 | -3.18 ± 0.02 | -1.25 ± 0.05 | 1.02 ± 0.01 | 5.52 ± 0.04 | -2.83 ± 0.01 | -2.92 ± 0.00 | -0.96 ± 0.01 | -1.65 ± 0.05 |
| VBP (ours) | -2.59 ± 0.03 | -3.15 ± 0.02 | -1.11 ± 0.07 | 1.04 ± 0.01 | 5.79 ± 0.07 | -2.85 ± 0.01 | -2.92 ± 0.00 | -0.96 ± 0.01 | -1.54 ± 0.06 |

shuts down negative activations, hence does not serve as a regularizer but instead implements the ReLU nonlinearity.

A fundamental step in the reduction of ELBO gradient variance has been taken by Kingma et al. (2015) with local reparameterization, which suggests taking the Monte Carlo integrals by sampling the linear activations rather than the weights. Further variance reduction has been achieved by defining the variances of the variational distribution factors as free parameters and the dropout rate as a function of them (Molchanov et al., 2017). Theoretical treatments of the same problem have also been recently studied (Miller et al., 2017; Roeder et al., 2017).

SGVB has been introduced initially for fully factorized variational distributions, which provides limited support for feasible posteriors that can be inferred. Strategies for improving the approximation quality of variational BNN inference include employment of structured versions of dropout matrix normals (Louizos and Welling, 2016), repetitive invertible transformations of latent variables (Normalizing Flows) (Rezende and Mohamed, 2015) and their application to variational dropout (Louizos and Welling, 2017). Wu et al. (2019) use the CLT argument to move beyond mean-field variational inference and demonstrate how to adapt it to learn layerwise covariance structures for the variational posteriors.

Lastly, there is active research on enriching variational inference using its interpolative connection to expectation propagation (Hernández-Lobato and Adams, 2015; Li and Turner, 2016; Li and Gal, 2017).

4 EXPERIMENTS

We evaluate the proposed model on a wide variety of regression and classification data sets. Details on hyperparameters and architectures not provided in the main text can be found in the appendix.¹²

4.1 REGRESSION

For the regression experiments we follow the experimental setup introduced by Hernández-Lobato and Adams (2015) and evaluate the performance on nine UCI bench-

¹²See <https://github.com/manuelhaussmann/vbp> for a reference implementation and the appendix.

mark data sets. For each data set we train a BNN with one hidden layer of 50 units.¹³ Each data set is randomly split into train and test data, consisting of 90% and 10% of the data respectively. We optimize the model using the Adam optimizer (Kingma and Ba, 2015) with their proposed default parameters and a learning rate of $\lambda = 0.01$.

We compare against the two sampling-free approaches *Probabilistic Back-Propagation* (PBP) (Hernández-Lobato and Adams, 2015) and *Deterministic Variational Inference* (DVI) (Wu et al., 2019) as well as the sampling based *Variational Dropout* (VarOut) (Kingma et al., 2015) in the formulation by Molchanov et al. (2017).¹⁴

In order to learn the observation precision β , we follow a Type II Maximum Likelihood approach and after each training epoch choose it so that the ELBO is maximized, which reduces to choosing β to maximize the data fit, i.e.

$$\beta^* = \arg \max_{\beta} \mathbb{E}_{q(\theta)} [\log p(\mathbf{y}|\theta, \mathbf{x})],$$

which gives us

$$\frac{1}{\beta^*} = \frac{1}{N} \sum_{n=1}^N \mathbb{E}_{q(\theta)} [(y_n - f(\mathbf{x}_n; \theta))^2]. \quad (12)$$

This expression is either evaluated via samples for VarOut or deterministically for VBP as we have shown above. One could also introduce a hyperprior over the prior weight precisions, and learn them via also via a Type-II approach (Wu et al., 2019). Instead we use a fixed normal prior $p(\mathbf{w}) = \mathcal{N}(\mathbf{w}|0, \alpha^{-1}\mathbf{1})$. We set $\alpha = 10$ and use $\beta = 1$ as the initial observation precision, observing quick convergence in general.

We summarize the average test log-likelihood over twenty random splits in Table 1. VBP either outperforms the baselines or performs competitively with them.

4.2 CLASSIFICATION

Our main classification experiment is an evaluation on four image classification data sets of increasing complexity: MNIST (LeCun et al., 1998), FashionMNIST (Xiao

¹³100 units for protein

¹⁴The results for these baselines are taken from the respective papers, while for VarOut we rely on our own implementation, replacing the improper log-uniform prior with a proper Gaussian prior to avoid a possible improper posterior (Hron et al., 2017).

Table 2: **Classification Results.** Average error rate and test log-likelihoods \pm standard deviation over five runs

| | AVERAGE ERROR (in %) | | | AVERAGE TEST LOG-LIKELIHOOD (in %) | | |
|--------------|----------------------|------------------|------------------------------------|------------------------------------|------------------------------------|------------------------------------|
| | VarOut | DVI | VBP (ours) | VarOut | DVI | VBP (ours) |
| MNIST | 1.12 \pm 0.05 | 0.97 \pm 0.06 | 0.85 \pm 0.06 | -0.03 \pm 0.00 | -0.03 \pm 0.00 | -0.03 \pm 0.00 |
| FASHIONMNIST | 10.81 \pm 0.24 | 10.33 \pm 0.07 | 10.10 \pm 0.16 | -0.30 \pm 0.00 | -0.29 \pm 0.00 | -0.28 \pm 0.00 |
| CIFAR-10 | 33.40 \pm 1.06 | 35.15 \pm 1.13 | 31.33 \pm 1.01 | -0.95 \pm 0.03 | -1.00 \pm 0.03 | -0.90 \pm 0.03 |
| CIFAR-100 | 62.85 \pm 1.43 | 66.21 \pm 1.14 | 60.97 \pm 1.61 | -2.44 \pm 0.06 | -2.61 \pm 0.06 | -2.37 \pm 0.08 |

et al., 2017), CIFAR-10, and CIFAR-100 (Krizhevsky and Hinton, 2009), in order to evaluate how the three variational inference based approaches of either taking samples, relying on CLT, or the ReLU decomposition to learn compare as the depth of the BNN increases.

We use a modified LeNet5 sized architecture consisting of two convolutional layers and two fully connected layers, with more filters/units per layer for the two CIFAR data sets. As discussed in Section 2.6, VBP can handle max-pooling layers, but they require a careful tracking of indices between the data fit and variance terms, which comes at some extra run time cost in present deep learning libraries. Instead, we provide a reference implementation on how to do this, but stick in the experiments with strided convolutions following the recent trend of “all-convolutional-nets” (Springenberg et al., 2015; Yu et al., 2017; Redmon and Farhadi, 2018).

We compare VBP against VarOut and our own implementation of DVI.¹⁵ In order to ensure maximal comparability between the three methods all of them share the same normal prior $p(\mathbf{w}) = \mathcal{N}(\mathbf{w}|0, \alpha^{-1} \mathbf{1})$ ($\alpha = 100$), initialization and other hyperparameters. They are optimized with the Adam optimizer (Kingma and Ba, 2015) and a learning rate of $\lambda = 0.001$ over 100 epochs.

We summarize the results in Table 2. We observe that DVI, making efficient use of the CLT based moment-matching approach, improves upon the sampling based VarOut on the two easier data sets, while it struggles on CIFAR. VBP can deal with the increasing width from (Fashion)MNIST to CIFAR- $\{10,100\}$ a lot better, improving upon both VarOut as well as DVI on all four data sets. As the depth increases from the regression to the classification experiment the difference that was small for the shallow network also becomes more and more pronounced.

4.2.1 Online Learning Comparison

As the final experiment we follow the setup of Kandemir (2018). They argue that their focus on closed-form updates instead of having to rely on gradients gives them an advantage in an online learning setup that has the con-

¹⁵We use the mean-field setup of DVI to be comparable to the mean-field variational posteriors learned with VBP and VarOut.

Table 3: **Online Learning Results.** Average test set accuracy (in %) \pm standard deviation over ten runs.

| | MNIST | CIFAR-10 |
|------------|----------------------------------|----------------------------------|
| V-ReLU-Net | 92.8 \pm 0.2 | 47.1 \pm 0.2 |
| VarOut | 96.7 \pm 0.2 | 47.8 \pm 0.3 |
| DVI | 96.6 \pm 0.2 | 48.7 \pm 0.4 |
| VBP (ours) | 96.6 \pm 0.2 | 48.3 \pm 0.4 |

straint of allowing only a single pass over the data. They report results on MNIST and CIFAR-10, using a net with a single hidden layer of 500 units for MNIST and a two hidden layer net with 2048/1024 units for CIFAR-10. We summarize the results in Table 3. While the closed-form updates of V-ReLU-Net have the advantage of removing the need for gradients, the required mean-field approximation over the layers substantially constrains it compared to the more flexible VBP structure.

5 CONCLUSION

Our experiments demonstrate that the Identity-Heaviside decomposition and especially the variance back-propagation we propose offer a powerful alternative to other recent deterministic approaches of training deterministic BNNs.

Following the No-Free-Lunch theorem, our closed-form available ELBO comes at the expense of a number of restrictions, such as a fully factorized approximate posterior, sticking to ReLU activations, and inapplicability of Batch Normalization (Ioffe and Szegedy, 2015). An immediate implication of this work is to explore ways to relax the mean-field assumption and incorporate normalizing flows without sacrificing from the closed-form solution. Because Equations (7) and (8) extend easily to dependent variables after adding the covariance of each variable pair as done by Wu et al. (2019), our formulation is applicable to structured variational inference schemes without major theoretical obstacles. Further extensions that are directly applicable to our construction are the inclusion residual He et al. (2016) and skip connections Huang et al. (2017), which is an interesting direction for future work as it will allow this approach to scale to deeper architectures.

References

- M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015.
- Jimmy Ba and Brendan Frey. Adaptive dropout for training deep neural networks. In *NIPS*, 2013.
- Christopher M. Bishop. *Pattern recognition and machine learning*. springer, 2006.
- C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra. Weight uncertainty in neural networks. In *ICML*, 2015.
- T. Chen, E.B. Fox, and C. Guestrin. Stochastic gradient Hamiltonian monte carlo. In *ICML*, 2017.
- A. Damianou and N.D. Lawrence. Deep Gaussian processes. In *AISTATS*, 2013.
- B. Frey and G. Hinton. Variational learning in nonlinear Gaussian belief networks. *Neural Computation*, 1999.
- Y. Gal and Z. Ghahramani. Bayesian convolutional neural networks with bernoulli approximate variational inference. 2015.
- Y. Gal and Z. Ghahramanimani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, 2016.
- I.J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015.
- A. Graves. Practical variational inference for neural networks. In *NIPS*, 2011.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, 2015.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- J. Hensman, N. Fusi, and N.D. Lawrence. Gaussian processes for big data. In *UAI*, 2013.
- J. M. Hernández-Lobato and R. Adams. Probabilistic backpropagation for scalable learning of bayesian neural networks. In *ICML*, 2015.
- G.E. Hinton and D.van Camp. Keeping neural networks simple by minimizing the description length of the weights. In *COLT*, 1993.
- J. Hron, A.G. Matthews, and Z.Ghahramani. Variational Gaussian dropout is not Bayesian. *arXiv preprint arXiv:1711.02989*, 2017.
- Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *CVPR*, 2017.
- S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- M. Kandemir. Variational closed-form deep neural net inference. *Pattern Recognition Letters*, 2018.
- A. Kendall and Y. Gal. What uncertainties do we need in Bayesian deep learning for computer vision? In *NIPS*, 2017.
- D. Kingma and M. Welling. Auto-encoding variational Bayes. In *ICLR*, 2014.
- D. Kingma, T. Salimans, and M. Welling. Variational dropout and the local reparameterization trick. In *NIPS*, 2015.
- D.P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. 2009.
- Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Juho Lee, Saehoon Kim, Jaehong Yoon, Hae Beom Lee, Eunho Yang, and Sungjoo Hwang. Adaptive network sparsification via dependent variational beta-bernoulli dropout. *arXiv preprint arXiv:1805.10896*, 2018.
- Y. Li and Y. Gal. Dropout inference in Bayesian neural networks with alpha-divergences. In *ICML*, 2017.
- Y. Li and R.T. Turner. Rényi divergence variational inference. In *NIPS*, 2016.
- C. Louizos and M. Welling. Structured and efficient variational deep learning with matrix Gaussian posteriors. In *ICML*, 2016.
- C. Louizos and M. Welling. Multiplicative normalizing flows for variational Bayesian neural networks. In *ICML*, 2017.
- C. Louizos, K. Ullrich, and M. Welling. Bayesian compression for deep learning. In *NIPS*, 2017.
- Andrew L Maas, Awni Y Hannun, and Andrew Y Ng. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, 2013.

- D.J. MacKay. A practical Bayesian framework for backpropagation networks. *Neural Computation*, 1992.
- D.J. MacKay. Probable networks and plausible predictions – a review of practical Bayesian methods for supervised neural networks. *Network: Computation in Neural Systems*, 1995.
- A. Miller, N. Foti, A. D’Amour, and R.P. Adams. Reducing reparameterization gradient variance. In *NIPS*, 2017.
- D. Molchanov, A. Ashukha, and D. Vetrov. Variational dropout sparsifies deep neural networks. In *ICML*, 2017.
- R. Neal. *Bayesian learning for neural networks*. PhD Thesis, 1995.
- R. Neal. MCMC using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 54:113–162, 2010.
- K. Neklyudov, D. Molchanov, A. Ashuka, and D. Vetrov. Structured Bayesian pruning via log-normal multiplicative noise. In *NIPS*, 2017.
- A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- J. Redmon and A. Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- D.J. Rezende and S. Mohamed. Variational inference with normalizing flows. In *ICML*, 2015.
- G. Roeder, Y. Wu, and D. Duvenaud. Sticking the landing: Simple, lower-variance gradient estimators for variational inference. In *NIPS*, 2017.
- J.T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller. Striving for simplicity: The all convolutional net. In *ICLR*, 2015.
- N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 2016.
- C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. In *ICLR*, 2014.
- Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, 2016.
- S.I. Wang and C.D. Manning. Fast dropout training. In *ICML*, 2013.
- M. Welling and Y.W. Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *ICML*, 2011.
- 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 *ICLR*, 2019.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.
- F. Yu, V. Koltun, and T. Funkhouser. Dilated residual networks. In *CVPR*, 2017.