

Kalman Bayesian Neural Networks for Closed-Form Online Learning

Philipp Wagner^{1*}, Xinyang Wu^{1*}, Marco F. Huber^{1,2}

¹Department Cyber Cognitive Intelligence (CCI), Fraunhofer Institute for Manufacturing Engineering and Automation IPA, Stuttgart, Germany

²Institute of Industrial Manufacturing and Management IFF, University of Stuttgart, Germany
philipp.wagner@ipa.fraunhofer.de, xinyang.wu@ipa.fraunhofer.de, marco.huber@ieee.org

Abstract

Compared to point estimates calculated by standard neural networks, Bayesian neural networks (BNN) provide probability distributions over the output predictions and model parameters, i.e., the weights. Training the weight distribution of a BNN, however, is more involved due to the intractability of the underlying Bayesian inference problem and thus, requires efficient approximations. In this paper, we propose a novel approach for BNN learning via closed-form Bayesian inference. For this purpose, the calculation of the predictive distribution of the output and the update of the weight distribution are treated as Bayesian filtering and smoothing problems, where the weights are modeled as Gaussian random variables. This allows closed-form expressions for training the network's parameters in a sequential/online fashion without gradient descent. We demonstrate our method on several UCI datasets and compare it to the state of the art.

1 Introduction

Deep Learning has been continuously attracting researchers for its applicability in many fields such as medical diagnostics (Amisha, Pathania, and Rathaur 2019), autonomous control (Zeng et al. 2020), or intelligent mass-productions (El-Shamouty et al. 2019). However, conventional deep Neural Networks (NNs) usually perform maximum likelihood estimation, which results solely in a point estimate without consideration of uncertainty in the data and the learned model. In domains with high safety standards or financial risks this approach is not sufficient and limits the number of possible applications. Bayesian methods offer ways to overcome this issue by quantifying uncertainties using Bayes' rule and probabilistic reasoning, which results in a distribution over network parameters and predictions instead of point estimates. A quantification of the uncertainty indicates whether the predictions are trustworthy and reliable (Begoli, Bhattacharya, and Kusnezov 2019). Popular approaches like Markov Chain Monte Carlo (MCMC) are computationally demanding, whereas variational inference (VI) or ensemble methods rely on noisy gradient computations and need to be trained using batched training data

and several iterations. Due to these characteristics, the mentioned methods are not directly applicable in online learning settings (Parisi et al. 2019), but first approaches using a data memory exist (Nguyen et al. 2018; Kurlle et al. 2019). In addition, gradient-based methods may suffer from poor choices of the optimization hyper-parameters or bad local minima (Bengio 2012). This behavior is mitigated by adaptive learning rates, stochastic gradient descent (SGD), and modern optimizers, but still persists.

In this paper we develop a novel online learning approach for Bayesian Neural Networks (BNN) (MacKay 1992) that is named *Kalman Bayesian Neural Network (KBNN)*. The key idea is to train the BNN via sequential Bayesian filtering without the need of gradient-based optimization. Bayesian filtering is commonly used to estimate probability density functions (PDF) from noisy observations in Markov chains (Särkkä 2013; Huber 2015). Many NN architectures like the one studied in this paper also form a Markov chain (Achille and Soatto 2018). Our approach is motivated by the work of Puskorius and Feldkamp (2001), in which the training procedure is also treated as a filtering problem, which however uses local linearization based on backpropagated gradient information to compute weight updates. While the special case of online training of a single perceptron is discussed by Huber (2020), we aim at constructing a learning algorithm for a universal multilayer perceptron (MLP).

The KBNN consists of a forward pass for inferring the output distribution given an input sample and a backward pass to adjust the network parameters given a corresponding output sample. A part of the backward pass is visualized in Fig. 1. By assuming that the weights and outputs of each network layer are Gaussian distributed, it is sufficient to compute the mean and covariance in a moment matching fashion. In doing so, it is possible to provide closed-form expressions for the forward pass as well as the update equations of the weights in the backward pass for commonly used activation functions like sigmoid or ReLU. The main contributions of this paper are: (a) We introduce an approach that strictly follows Bayesian inference. Thus, learning the network parameters is not reformulated as optimization problem. The use of Bayesian filtering and smoothing techniques instead enables native online training of BNNs, where we show the relationship to the famous Kalman filtering and smoothing equations (Kalman 1960; Rauch, Tung, and Striebel 1965).

*These authors contributed equally.




Figure 1: Probabilistic graphical model for an arbitrary layer $l = 1 \dots L$. During the backward pass, first a_n^l is updated via a Kalman smoothing step and afterwards z^l and w_n^l are updated jointly again via smoothing. For the output layer $l = L$ the ground truth y is known from the data \mathcal{D} and thus, for updating the output z^{l+1} a filtering step instead of a smoothing step is performed. Smoothing I refers to (12) while smoothing II refers to (13) and (14).

(b) We extend the work of Huber (2020) being applicable only for a single neuron to the case of an MLP. In doing so, all calculations necessary are performed in closed form. For ReLU and linear activations they are exact. (c) We compare the performance of our approach on various common benchmarking datasets to inference techniques such as MCMC, VI and expectation propagation (EP).

2 Related Work

Laplace Approximation The idea of Bayesian inference in the NN setting goes back to the work by MacKay (1992), where a probability distribution over the model parameters is learned via Laplace transformation. Here, a Gaussian distribution is fitted to the vicinity of the maximum posterior estimate. Due to that point of time, Bayesian inference was limited to small networks and datasets. This approximation technique recently gained increased interest by (Ritter, Botev, and Barber 2018a; Kristiadi, Hein, and Hennig 2020). Snoek et al. (2015); Kristiadi, Hein, and Hennig (2020) for instance use it to provide uncertainty estimates only for the last layer of an MLP.

MCMC One of the most explored ideas for probabilistic inference in general and for learning BNNs in particular is MCMC (Metropolis et al. 1953), which allows approximating probability integrals with the Monte Carlo method via sampling from a Markov process. Many improvements have been suggested for the initial Metropolis-Hastings algorithm such as Gibbs sampling (Geman and Geman 1984), hybrid Monte Carlo (Neal 1995), or Hamiltonian Monte Carlo (HMC) (Duane et al. 1987). An important extension to HMC is the No-U-Turn Sampler (NUTS) (Hoffman and Gelman 2014), which mostly performs more efficiently. One downside of most MCMC approaches is the high computational cost necessary for performing density estimation. Further, only samples and no closed-form representation of the distribution are available.

Variational Inference The usage of VI for BNN training was introduced by Graves (2011). VI is based on the idea of

approximating the complicated weight posterior by means of a simple distribution like a Gaussian. This approximation is achieved by minimizing the empirical lower bound to the reverse Kullback-Leibler divergence using gradient descent. In Kingma and Welling (2014), the gradient is estimated via sampling leading to high variance gradients and thus, merely a decent performance only in smaller architectures. In Hoffman et al. (2013) a more scalable method called Stochastic Variational Inference (SVI) is proposed, which computes a scaled gradient on randomly sampled subsets of data to update the variational parameters, instead of computing gradients from the full dataset. A deterministic calculation is proposed in Wu et al. (2019), which uses a closed-form forward pass as in our work. Gal and Ghahramani (2016) found that the dropout technique introduced by Srivastava et al. (2014) approximates the variational distribution while being relatively computationally cheap.

Expectation Propagation Minimizing the forward Kullback-Leibler divergence instead of its reverse version leads to EP (Minka 2001). In contrast to VI, EP is not guaranteed to convergence in general. A practical EP version for BNNs named probabilistic backpropagation (PBP) was proposed in Hernández-Lobato and Adams (2015) and extended in Ghosh, Fave, and Yedidia (2016), which share similarities with our work. The forward pass of PBP also utilizes closed-form moment propagation. However, Ghosh, Fave, and Yedidia (2016) and Hernández-Lobato and Adams (2015) employ the mean-field approximation, i.e., fully factorized Gaussian distributions, while we allow correlated weights per neuron. Significant difference are given for the backward pass. Here, the KBNN requires no explicit calculation of the marginal likelihood and its first and second order derivatives.

Kalman Filtering Closest to our work is the usage of Kalman filtering for training BNNs. One of the first approaches was proposed by Watanabe and Tzafesta (1990), where the weights of the networks are assumed to be Gaussian. The mean and variance of the individual weights are updated by means of an extended Kalman filter, which however requires local linearization for updating the hidden neurons. This work was extended by Puskorius and Feldkamp (2001) to allow layer-wise correlated or even network-wide correlated neurons. To avoid linearization, Huber (2020) proposes the so-called Bayesian perceptron. Even though limited to a single neuron, this work shows that closed-form Bayesian inference for calculating the mean and covariance parameters of the weight posterior distribution is possible. In the following, we extend this single neuron approach to an MLP by utilizing Bayesian filtering and smoothing.

Online Learning In the context of online learning Bayesian methods are a popular choice, since uncertainties over the data and the model can be taken into account directly. Opper (1998) use a moment matching approach for online learning which is similar to our work. Kirkpatrick et al. (2017) and Ritter, Botev, and Barber (2018b) deal with the issue of catastrophic forgetting in neural networks for continual learning tasks. There are a few works that include

data memory to improve online learning capabilities (Minka, Xiang, and Qi 2009; Nguyen et al. 2018; Kurle et al. 2019).

3 Problem Formulation

Given a dataset $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$ of N i.i.d. pairs of training instances with inputs $\mathbf{x}_i \in \mathbb{R}^d$ and outputs $\mathbf{y}_i \in \mathbb{R}^e$, we want to train an MLP with L layers in a supervised learning setup. In each layer $l = 1 \dots L$, a nonlinear transformation

$$\mathbf{z}^{l+1} = f(\mathbf{a}^l) \text{ with } \mathbf{a}^l = \mathbf{W}^l \cdot \mathbf{z}^l + \mathbf{w}_0^l \quad (1)$$

is performed with weight matrix $\mathbf{W}^l \triangleq [\mathbf{w}_1 \dots \mathbf{w}_{M_l}]^T \in \mathbb{R}^{M_l \times M_{l-1}}$ with $\mathbf{w}_i \in \mathbb{R}^{M_{l-1}}$ being the i -th neuron's weight vector, bias $\mathbf{w}_0^l \in \mathbb{R}^{M_l}$, and nonlinear activation function $f(\cdot)$, where M_l is the number of neurons of the l -th layer. The output \mathbf{z}^{l+1} of layer l becomes the input of the subsequent layer $l + 1$. For the first layer $\mathbf{z}^1 = \mathbf{x}$ and for the last layer $\mathbf{z}^{L+1} = \mathbf{y}$. To simplify the notation, we avoid the layer superscript l for \mathbf{a} and \mathbf{W} whenever possible.

By redefining the input $\mathbf{z}^l \triangleq [1 \ z_1^l \dots z_{M_{l-1}}^l]^T$ we can conveniently incorporate the bias \mathbf{w}_0 into the weights according to $\mathbf{w}_i \triangleq [w_0^i \ w_1^i \dots w_{M_{l-1}}^i]^T$ where w_0^i is the i -th element of \mathbf{w}_0 and thus, $\mathbf{W} \in \mathbb{R}^{M_l \times (M_{l-1} + 1)}$. Further, $\mathcal{W} \triangleq \{\mathbf{W}^l\}_{l=1}^L$ comprises all weight matrices. Whenever appropriate, we use $\mathbf{w} \triangleq [\mathbf{w}_1^T \dots \mathbf{w}_{M_l}^T]^T = \text{vec}(\mathbf{W})$ to simplify the calculations and notation.

The MLP is treated as a BNN. Accordingly, the weights in \mathcal{W} are random variables with (prior) probability distribution $p(\mathcal{W})$. The task now is two-fold (MacKay 1992; Neal 1995):

(i) Calculating the posterior distribution of the weights

$$p(\mathcal{W}|\mathcal{D}) = \frac{p(\mathbf{Y}|\mathbf{X}, \mathcal{W}) \cdot p(\mathcal{W})}{p(\mathbf{Y}|\mathbf{X})}, \quad (2)$$

with normalization term $p(\mathbf{Y}|\mathbf{X})$ and $\mathbf{X} \triangleq [\mathbf{x}_1 \dots \mathbf{x}_N]$, $\mathbf{Y} \triangleq [\mathbf{y}_1 \dots \mathbf{y}_N]$ being the input and output data from \mathcal{D} , respectively. (ii) Calculating the predictive distribution

$$p(\mathbf{y}|\mathbf{x}, \mathcal{D}) = \int p(\mathbf{y}|\mathbf{x}, \mathcal{W}) \cdot p(\mathcal{W}|\mathcal{D}) \, d\mathcal{W} \quad (3)$$

of the BNN given a new input \mathbf{x} . Unfortunately, both equations cannot be solved exactly in closed form in general. To provide an approximate but closed-form solution we employ techniques from *Bayesian filtering and smoothing*, which is usually applied to Markov processes in order to estimate a state variable over time from noisy observations. Equation (1) forms a continuous-valued Markov chain with random variables \mathbf{a} , \mathbf{W} , and \mathbf{z} as depicted in the graphical model Fig. 1. Hence, (3) can be solved layer-wise by means of consecutive prediction steps of a Bayesian filter and will be addressed in the *forward pass* of the proposed KBNN. Solving (2) to train the weights requires filtering and smoothing steps of a Bayesian filter and is covered by the *backward pass* of the KBNN. To obtain closed-form expressions in both passes, we make two key assumptions.

Assumption 1 For BNNs it is very common to apply the strict version of the mean-field approximation, where all weights are assumed to be independent. In this paper, it is

merely assumed that all neurons are pair-wise independent, so that the weights of individual neurons are dependent. This assumption significantly simplifies the calculations. The implications of dependent neurons are discussed in Puskorius and Feldkamp (2001) and Sec. 6.

Assumption 2 For a single neuron, the corresponding quantities in (1) are assumed to be jointly Gaussian distributed. Due to this assumption, particularly the posterior in (2) and the predictive distribution in (3) are approximated by means of Gaussian distributions and thus, our approach becomes an *assumed density filter* (Maybeck 1979; Oppen 1998). In doing so, it is sufficient to calculate the first two moments (mean and covariance) of the posterior and predictive distribution. For ReLU activations this approach even transforms into *moment matching*.

Based on these assumptions, the posterior weight distribution of a layer is given in factorized form $p(\mathbf{W}|\mathcal{D}) = \prod_{i=1}^{M_l} \mathcal{N}(\mathbf{w}_i | \boldsymbol{\mu}_w^i, \mathbf{C}_w^i)$, where $\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_x, \mathbf{C}_x)$ is a Gaussian PDF with mean vector $\boldsymbol{\mu}_x$ and covariance matrix \mathbf{C}_x . The limitations arising from these assumptions are discussed in greater detail in the Sec. 6.

4 The Kalman Bayesian Neural Network

For deriving the forward and backward pass of the KBNN we process each training data instance $(\mathbf{x}_i, \mathbf{y}_i)$ individually and sequentially. This is possible as the data is assumed to be i.i.d. and thus, the likelihood in (2) can be factorized according to $p(\mathbf{Y}|\mathbf{X}, \mathcal{W}) = \prod_i p(\mathbf{y}_i | \mathbf{x}_i, \mathcal{W})$. Hence, we obtain the posterior $p(\mathcal{W}|\mathcal{D}_i) \propto p(\mathbf{y}_i | \mathbf{x}_i, \mathcal{W}) \cdot p(\mathcal{W}|\mathcal{D}_{i-1})$, with $\mathcal{D}_i = \{(\mathbf{x}_j, \mathbf{y}_j)\}_{j=1}^i \subset \mathcal{D}$ and $p(\mathcal{W}|\mathcal{D}) \equiv p(\mathcal{W}|\mathcal{D}_N)$, by means of recursively processing the data instances, where the recursion commences from the prior $p(\mathcal{W})$. During each recursion step it is not necessary to update all weights of the BNN simultaneously. Instead, we can make use of the Markov chain characteristic of the network (cf. Fig. 1). In doing so, the weight matrix of each layer can be updated one after the other. This updating is essentially performed during the backward pass, but requires intermediate predictive quantities \mathbf{a} and \mathbf{z} that are calculated during the forward pass. Thus, the forward pass is not only necessary to calculate the predictive distribution of the BNN for new inputs, but is also a crucial component of the backward pass. Hence, we start with deriving the forward pass, where we omit the quantity \mathcal{D} in the following to simplify the notation.

4.1 Forward Pass

During the forward pass the predictive distribution $p(\mathbf{y}|\mathbf{x}, \mathcal{D})$ for a given input \mathbf{x} has to be computed. For this purpose, information is propagated forward through the BNN in order to calculate the predictive distributions of all random variables \mathbf{a} and \mathbf{z} along the path from the input to the output. Since these two quantities occur in each layer with the same dependencies, we restrict the derivation to a single layer without loss of generality. For the l -th layer the predictive distribution of the output \mathbf{z}^{l+1} is given by

$$p(\mathbf{z}^{l+1}|\mathbf{x}) = \int p(\mathbf{z}^{l+1}|\mathbf{a}) \cdot p(\mathbf{a}|\mathbf{x}) \, d\mathbf{a} \quad (4)$$

with

$$p(\mathbf{a}|\mathbf{x}) = \int p(\mathbf{a}|z^l, \mathbf{W}) \cdot p(z^l|\mathbf{x}) \cdot p(\mathbf{W}) \mathrm{d}z^l \mathrm{d}\mathbf{W}. \quad (5)$$

All quantities in (5) are related according to (1). Further, $p(z^l|\mathbf{x})$ and $p(\mathbf{W})$ are assumed to be Gaussian. The predictive distribution $p(\mathbf{a}|\mathbf{x})$ however, is not Gaussian due to the multiplicative nonlinearity in (1), but it is approximated by the Gaussian $\mathcal{N}(\mathbf{a}|\boldsymbol{\mu}_a, \mathbf{C}_a)$ with mean vector and covariance matrix matching the moments of the true distribution. The elements of the mean vector $\boldsymbol{\mu}_a$ are given by

$$\mu_a^n = \mathbb{E}[\mathbf{w}_n^T \cdot z^l] = \mathbb{E}[\mathbf{w}_n^T] \cdot \mathbb{E}[z^l] = (\boldsymbol{\mu}_w^n)^T \cdot \boldsymbol{\mu}_z^l, \quad (6)$$

while the covariance matrix is diagonal due to Assumption 1 with elements

$$\begin{aligned} (\sigma_a^n)^2 &= \mathbb{E}[a_n^2] - (\mu_a^n)^2 = \mathbb{E}[(\mathbf{w}_n^T \cdot z^l)^2] - (\mu_a^n)^2 \\ &= (\boldsymbol{\mu}_w^n)^T \mathbf{C}_z^l \boldsymbol{\mu}_w^n + (\boldsymbol{\mu}_z^l)^T \mathbf{C}_w^n \boldsymbol{\mu}_z^l + \mathrm{Tr}(\mathbf{C}_w^n \mathbf{C}_z^l), \end{aligned} \quad (7)$$

where $n = 1 \dots M_l$ is the neuron index, $\mathrm{Tr}(\cdot)$ is the matrix trace, and $\boldsymbol{\mu}_z^l, \mathbf{C}_z^l$ are the mean and covariance of z^l . For the first layer $z^l = \mathbf{x}$ and thus, no random variable. This allows solving (5) exactly as \mathbf{a}^l in (1) becomes a linear function, where $p(\mathbf{a}|\mathbf{x})$ is actually Gaussian. With $p(z^l|\mathbf{x}) = \delta(z^l - \mathbf{x})$ in (5) the means (6) and variances (7) become $\mu_a^n = \mathbf{x}^T \cdot \boldsymbol{\mu}_w^n$ and $(\sigma_a^n)^2 = \mathbf{x}^T \mathbf{C}_w^n \mathbf{x}$, respectively, which corresponds to a *Kalman prediction* step.

The predictive distribution z^{l+1} in (4) is also approximated with a Gaussian $\mathcal{N}(z^{l+1}|\boldsymbol{\mu}_z^{l+1}, \mathbf{C}_z^{l+1})$, where the elements of the mean vector and (diagonal) covariance matrix are given by

$$\boldsymbol{\mu}_z^{l+1,n} = \mathbb{E}[f(a_n)], \quad (8)$$

$$(\sigma_z^{l+1,n})^2 = \mathbb{E}[f(a_n)^2] - (\boldsymbol{\mu}_z^{l+1,n})^2, \quad (9)$$

respectively, and thus depend on the nonlinear activation function. For ReLU activations, the expected values in (8) and (9) can be calculated exactly in closed form and thus, we obtain a moment matching approximation. For sigmoidal activations like sigmoid or hyperbolic tangent, the expected values can be tightly approximated in closed form, except for the special case of a probit activation, where we again obtain a moment matching. Detailed derivations for both activations can be found in Huber (2020).

The distribution $p(z^{l+1}|\mathbf{x})$ is then used for solving the corresponding integrals (4) and (5) of the subsequent layer $l+1$. For $l=L$, we have $z^{l+1} = \mathbf{y}$ and thus $p(z^{l+1}|\mathbf{x})$ coincides with the desired predictive distribution $p(\mathbf{y}|\mathbf{x}, \mathcal{D})$. It is worth mentioning that the calculated moments of each layer must be stored, as they are needed for the weight update procedure during the backward pass.

4.2 Backward Pass

The training of conventional MLPs relies on a problem specific loss function being optimized with SGD, where the entire dataset \mathcal{D} is processed repeatedly. The backward pass of the KBNN updates the weights by means of sequentially processing the data once without gradient-based optimization thanks to strictly following Bayes' rule in (2). Like with

the forward pass, the Markov property of the network allows updating the weights layer-wise. Given any training instance (\mathbf{x}, \mathbf{y}) , updating the l -th layer requires considering joint Gaussian PDFs of the form

$$\begin{aligned} p(\mathbf{s}, \mathbf{t}|\mathbf{x}, \mathbf{y}) &= \mathcal{N}\left(\mathbf{s}, \mathbf{t} \left[\begin{array}{c} \boldsymbol{\mu}_s^+ \\ \boldsymbol{\mu}_t^+ \end{array} \right], \left[\begin{array}{cc} \mathbf{C}_s^+ & \mathbf{C}_{st}^+ \\ \mathbf{C}_{st}^T & \mathbf{C}_t^+ \end{array} \right] \right) \\ &= p(\mathbf{s}|\mathbf{t}, \mathbf{x}) \cdot p(\mathbf{t}|\mathbf{x}, \mathbf{y}) \end{aligned} \quad (10)$$

twice: (I) $\mathbf{s} = a_n, \mathbf{t} = z_n^{l+1}$ and (II) $\mathbf{s} = [\mathbf{w}^T(z^l)^T]^T, \mathbf{t} = \mathbf{a}$ as the graphical model in Fig. 1 indicates. Here, $\mathbf{w} = \mathrm{vec}(\mathbf{W})$ is the vectorized weight matrix as defined in Sec. 3. The Gaussian $p(\mathbf{t}|\mathbf{x}, \mathbf{y}) = \mathcal{N}(\mathbf{t}|\boldsymbol{\mu}_t^+, \mathbf{C}_t^+)$ is already known from the previous step, while the conditional Gaussian $p(\mathbf{s}|\mathbf{t}, \mathbf{x}) = \mathcal{N}(\mathbf{s}|\boldsymbol{\mu}_s^- + \mathbf{J} \cdot (\mathbf{t} - \boldsymbol{\mu}_t^-), \mathbf{C}_s^- - \mathbf{J} \cdot \mathbf{C}_{st}^T)$ with *Kalman gain* $\mathbf{J} = \mathbf{C}_{st} \cdot (\mathbf{C}_t^-)^{-1}$ (Huber 2015; Särkkä 2013). The superscript $-$ indicates quantities $p(\cdot|\mathbf{x})$ of the forward pass, while $+$ is the updated version $p(\cdot|\mathbf{x}, \mathbf{y})$ resulting from the backward pass. Calculating the product of the two Gaussian PDFs in (10) and marginalizing \mathbf{t} yields

$$\begin{aligned} \boldsymbol{\mu}_s^+ &= \boldsymbol{\mu}_s^- + \mathbf{J} \cdot (\boldsymbol{\mu}_t^+ - \boldsymbol{\mu}_t^-), \\ \mathbf{C}_s^+ &= \mathbf{C}_s^- + \mathbf{J} \cdot (\mathbf{C}_t^+ - \mathbf{C}_t^-) \cdot \mathbf{J}^T \end{aligned} \quad (11)$$

being the mean and covariance of $p(\mathbf{s}|\mathbf{x}, \mathbf{y})$, respectively. These equations correspond to the *Kalman* or *Rauch-Tung-Striebel smoother* (Rauch, Tung, and Striebel 1965).

For the smoothing step (I), (11) becomes

$$\begin{aligned} \mu_a^{n,+} &= \mu_a^{n,-} + \mathbf{k}_n^T \cdot (\boldsymbol{\mu}_z^{l+1,+} - \boldsymbol{\mu}_z^{l+1,-}), \\ (\sigma_a^{n,+})^2 &= (\sigma_a^{n,-})^2 + \mathbf{k}_n^T \cdot (\mathbf{C}_z^{l+1,+} - \mathbf{C}_z^{l+1,-}) \cdot \mathbf{k}_n, \end{aligned} \quad (12)$$

for neuron $n = 1 \dots M_l$ with $\mathbf{k}_n = (\mathbf{C}_z^{l+1,-})^{-1} \cdot (\boldsymbol{\sigma}_{az}^n)^2$. All quantities in (12) can be calculated in closed form but the cross-covariance $(\boldsymbol{\sigma}_{az}^n)^2$, which depends on the activation function. As with the forward pass, ReLU allows an analytical solution, while for sigmoidal activations a closed-form approximation exists. For details be referred to Huber (2020).

The result \mathbf{s} of step (I) becomes the quantity \mathbf{t} of step (II), for which the mean and covariance in (11) are given by

$$\left[\begin{array}{c} \boldsymbol{\mu}_w^+ \\ \boldsymbol{\mu}_z^{l,+} \end{array} \right] = \left[\begin{array}{c} \boldsymbol{\mu}_w^- \\ \boldsymbol{\mu}_z^{l,-} \end{array} \right] + \mathbf{L} \cdot (\boldsymbol{\mu}_a^+ - \boldsymbol{\mu}_a^-), \quad (13)$$

$$\left[\begin{array}{cc} \mathbf{C}_w^+ & \mathbf{C}_{wz}^+ \\ \mathbf{C}_{wz}^T & \mathbf{C}_z^{l,+} \end{array} \right] = \mathbf{C} + \mathbf{L} \cdot (\mathbf{C}_a^+ - \mathbf{C}_a^-) \cdot \mathbf{L}^T, \quad (14)$$

with $\mathbf{L} = \mathbf{C}_{wza} \cdot (\mathbf{C}_a^-)^{-1}$, $\mathbf{C} = \mathrm{diag}(\mathbf{C}_w, \mathbf{C}_z^{l,-})$, and $\mathbf{C}_a^+ = \mathrm{diag}((\sigma_a^{1,+})^2, \dots, (\sigma_a^{M_l,+})^2)$. The structure of the covariance can be explained as follows. At the forward pass, \mathbf{W} and z^l are independent as \mathbf{a} is not observed and these three quantities are connected via a v-structure $\mathbf{W} \rightarrow \mathbf{a} \leftarrow z^l$ (cf. Fig. 1). Thus, \mathbf{C} has a block-diagonal structure. At the backward pass, a descendent of z^l , namely \mathbf{y} is observed and thus, \mathbf{W} and z^l are dependent. The mean $\boldsymbol{\mu}_w^+$ and covariance \mathbf{C}_w^+ are the updated weight parameters of $p(\mathbf{w}|\mathbf{x}, \mathbf{y}, \mathcal{D}_{i-1})$, while $\boldsymbol{\mu}_z^{l,+}$ and $\mathbf{C}_z^{l,+}$ are used for the quantity \mathbf{t} of step (I) of layer $l-1$. This update rule differs from Huber (2020) since

Algorithm 1: Backward pass for training on dataset \mathcal{D}

```

1: for each training instance  $(x_i, y_i) \in \mathcal{D}$  do
2:    $(\mu_a^{l,-}, \mathbf{C}_a^{l,-}, \mu_z^{l+1,-}, \mathbf{C}_z^{l+1,-}) \leftarrow \text{ForwardPass}(x_i)$ 
3:    $(\mu_z^{L+1,+}, \mathbf{C}_z^{L+1,+}) \leftarrow (y_i, \mathbf{0})$ 
4:   for  $l = L$  to 1 do
5:     Update  $\mu_a^+, \mathbf{C}_a^+$  via (12)
6:     Update  $\mu_w^+, \mathbf{C}_w^+, \mu_z^{l,+}, \mathbf{C}_z^{l,+}$  via (13) and (14)
7:     Store  $(\mu_w^+, \mathbf{C}_w^+) \leftarrow (\mu_w^+, \mathbf{C}_w^+)$ 
8:   end for
9: end for

```

z^l is not deterministic for any layer but the input layer. All quantities are known except of \mathbf{C}_{wza} , which is given by

$$\begin{aligned} \mathbf{C}_{wza} &= \mathbb{E} \left[\left(\begin{bmatrix} \mathbf{w} \\ \mathbf{z} \end{bmatrix} - \begin{bmatrix} \mu_w \\ \mu_z \end{bmatrix} \right) \cdot (a - \mu_a^-)^T \right] \\ &= \begin{bmatrix} \text{diag}(\mathbf{C}_w^1 \cdot \mu_z^{l,-}, \dots, \mathbf{C}_w^{M_l} \cdot \mu_z^{l,-}) \\ \mathbf{C}_z^{l,-} \cdot \mu_w^1 & \dots & \mathbf{C}_z^{l,-} \cdot \mu_w^{M_l} \end{bmatrix}. \end{aligned} \quad (15)$$

The black-diagonal structure of the upper part of \mathbf{C}_{wza} is due to Assumption 1.

The sequence of smoothing operations is not surprising as updating is not performed with the data directly but with previously estimated quantities. The only exception is layer L , where the recursion starts. Here, in step (I) $\mathbf{t} = \mathbf{z}^{L+1} = \mathbf{y}$ is deterministic, thus $\mu_z^{L+1,+} = \mathbf{y}$ and $\mathbf{C}_z^{L+1,+} = \mathbf{0}$. By substituting these quantities in (12) the Kalman smoother becomes a (nonlinear) *Kalman filter* (Kalman 1960). The backward pass is summarized in Algorithm 1.

5 Experiments

In this section, we validate the proposed KBNN in both classification and regression tasks on benchmark datasets. Four experiments are conducted: (i) Evaluating the KBNN on a synthetic regression task, (ii) binary classification on the well-known Moon dataset, (iii) online learning on the Moon dataset, and (iv) comparison with other approximate inference approaches on nine UCI regression datasets (Dua and Graff 2017). The KBNN implementation merely requires matrix operations and is realized in PyTorch. The performance of the methods is assessed by means of the root mean square error (RMSE) for regression tasks, the accuracy for classification tasks, the negative log-likelihood (NLL) for quantifying the uncertainty, and the training time. A PC with Intel i7-8850H CPU, 16GB RAM but without GPU is used.

Regression on Synthetic Data We generate a synthetic dataset with 800 data instances from the polynomial $y = x^3 + \epsilon_n$, where $\epsilon_n \sim \mathcal{N}(0, 9)$ similar to Hernández-Lobato and Adams (2015), while x is sampled uniformly from the interval $[-4, 4]$. We use a standard MLP with one hidden layer and 100 hidden neurons, and ReLU activation for the hidden layer. The output activation is linear. We compare KBNN with PBP (Hernández-Lobato and Adams 2015) and Monte Carlo (MC) Dropout (Gal and Ghahramani 2016). For both PBP and MC Dropout we use the implementations of the authors. For MC Dropout we use dropout probability 0.1, same as the authors used for regression tasks (Gal

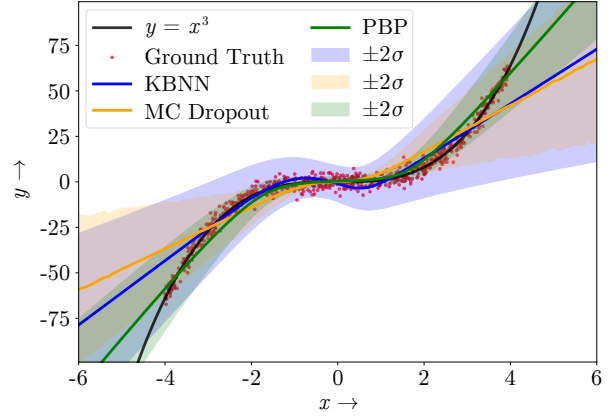


Figure 2: Predictions of KBNN, MC Dropout and PBP trained for one epoch on the regression task $y = x^3 + \epsilon_n$.

# data	Accuracy	NLL	Training Time / s
5	47.53% \pm 0.60%	0.18 \pm 0.02	0.01 \pm 4.00 $\times 10^{-4}$
50	88% \pm 1.63%	0.11 \pm 0.01	0.13 \pm 1.20 $\times 10^{-3}$
500	92.07% \pm 2.28%	0.05 \pm 0.01	1.25 \pm 3.53 $\times 10^{-3}$
1,000	97.87% \pm 2.33%	0.03 \pm 0.004	2.49 \pm 8.13 $\times 10^{-3}$
1,350	99.93% \pm 0.20%	0.03 \pm 0.003	3.40 \pm 5.72 $\times 10^{-3}$

Table 1: Sequential learning on the Moon dataset.

and Ghahramani 2016). All methods merely use one epoch for training in order to simulate an online learning scenario. In Fig. 2 the results of all methods are depicted. KBNN, PBP and MC Dropout are able to approximate the underlying nonlinear function and perform similar, where PBP and MC Dropout tend to underestimate the aleatoric uncertainty.

Binary Classification To validate the capabilities of the KBNN for sequential/online learning in binary classification tasks, we perform classification on the Moon dataset (Pedregosa et al. 2011) with 1,500 instances in total. The data is presented in a continuous stream to the learning algorithm, starting with a single data instance. Training ends, when 90% of the dataset, i.e., 1,350 instances are processed. We measure the performance of the model on the remaining 10% of the instances during training to evaluate the learning progress. To demonstrate learning for multiple layers we use two hidden layers, each with 10 neurons. The hidden activations are ReLU, the output activation is a sigmoid function.

Table 1 lists how the accuracy and NLL of the KBNN on the test set evolve for an increasing number of processed training data. These results are averages over ten random trials. After several seconds of training the proposed model achieves a high accuracy and low NLL on the test set.

Online Learning In order to validate the online learning capability of the KBNN on non-stationary data, it is applied to the classification of a rotating Moon dataset, similar as in (Kurle et al. 2019). We use the same network architecture as in Sec. 5. At the first iteration, we train the KBNN with the standard Moon dataset comprising 1,500 instances. Then, we continue training the KBNN for 18 iterations, where for

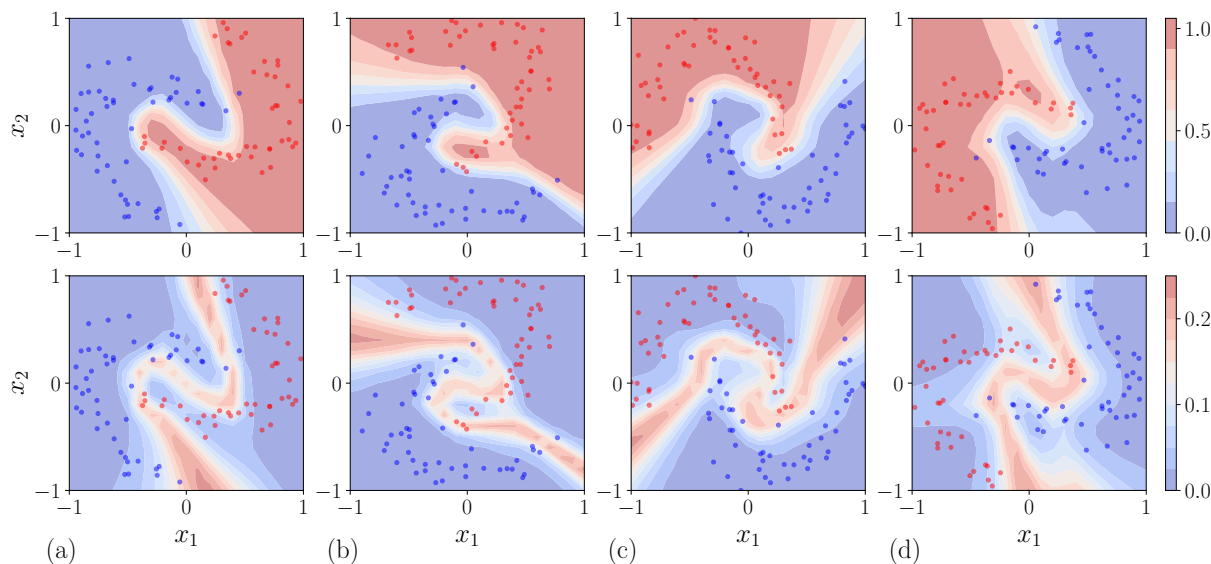


Figure 3: Online learning on the non-stationary Moon dataset. Predictive mean (top row) and variance (bottom row) after rotating by (a) 60, (b) 120, (c) 180, and (d) 240 degrees, respectively.

each iteration the dataset is rotated by 20 degrees and comprises only 100 data instances. Fig. 3 shows the changing decision boundary (predictive mean) after 60, 120, 180, and 240 degrees rotations. After each rotation, the KBNN can always efficiently adapt to the new data distributions.

UCI Regression Datasets In this section we compare the KBNN to SVI, MCMC and PBP for regression tasks on nine UCI datasets: Boston, Concrete, Energy, Wine, Naval, Yacht, Kni8nm, Power and Year. These datasets are commonly used for BNN performance evaluation (cf. Hernández-Lobato and Adams (2015)). Like in the previous sections, the datasets are split into random train and test sets with 90% and 10% of the data, respectively. For SVI and MCMC we use implementations build in the probabilistic programming package Pyro (Bingham et al. 2019). All methods are compared using the same network architecture with one hidden layer comprising 50 units and ReLU activations as proposed in Hernández-Lobato and Adams (2015). We use 40 epochs for PBP as in Hernández-Lobato and Adams (2015) and 5,000 epochs for SVI, after which the trained models converged well mostly. MCMC is conducted with NUTS and we draw 100 samples from the posterior. Although KBNN is designed for online learning, i.e., processing the data only once, we also executed KBNN with 10 epochs of training—denoted KBNN 10 in the following—to improve the performance. We repeat the experiments ten times with random initializations and average the results.

Tables 2 and 3 show the average RMSE and NLL on the test sets as well as the training time. KBNN 1, i.e., the online version, achieves a performance being close to other methods on some datasets while requiring significantly less training time for all datasets. Particularly compared to SVI, the performance gap between KBNN 1 and SVI is narrow. KBNN 10 outperforms SVI in most cases and PBP

and MCMC on some datasets. For the Concrete, Naval and Year datasets, KBNN even outperforms MCMC and PBP in terms of RMSE and PBP also on Boston and Energy. For an increasing number of epochs the NLL value of the KBNN increases in-between for the Boston, Concrete and Wine datasets. This increase is caused by too low variances.

KBNN 1 is clearly faster than the other methods. The training time roughly grows linearly with the number of data instances. Thus, compared to SVI, which is designed for scaling well with large datasets (Zhang et al. 2019), KBNN has a runtime advantage on smaller datasets while this gap closes for larger datasets and more epochs of training. However, it is worth mentioning that as a method with online learning capabilities in contrast to SVI, MCMC and PBP, our method shows great single sample learning efficiency. If SVI or PBP learn for only one epoch, their performance significantly drops and is worse than KBNN 1, especially for the small datasets. Averaged over all datasets, the time of KBNN 1 to process a single input is 1.659 ± 0.041 ms, which is promising for real-time tasks.

We also performed a series of experiments with either a different number of hidden neurons or different number of hidden layers to assess the scalability of the KBNN.

6 Discussion

In this paper we introduced an approach to perform sequential and online learning of BNNs via assumed Gaussian filtering/smoothing. The state of the art in training BNNs are VI-based methods. Although being Bayesian, these methods treat training as an optimization problem. Instead, the proposed KBNN approach is fully Bayesian in the sense that the training strictly aims at (approximately) solving Bayes’ rule (2). Utilizing concepts from Bayesian filtering and smoothing allows updating the mean and covariance of

Dataset	N	d	SVI	MCMC	PBP	KBNN 1	KBNN 10
Boston	506	13	3.434 ± 0.131	2.553 ± 0.027	2.740 ± 0.095	3.893 ± 0.200	2.695 ± 0.155
Concrete	1,030	8	7.597 ± 0.283	6.227 ± 0.108	5.874 ± 0.054	8.396 ± 0.497	5.703 ± 0.183
Energy	768	8	4.025 ± 0.074	0.906 ± 0.049	3.274 ± 0.049	4.155 ± 0.087	2.404 ± 0.259
Wine	4,898	11	0.726 ± 0.007	0.656 ± 0.004	0.667 ± 0.002	0.719 ± 0.011	0.666 ± 0.006
Naval	11,934	16	0.025 ± 0.012	0.008 ± 0.001	$0.006 \pm 6.12 \cdot 10^{-5}$	0.034 ± 0.005	0.004 ± 0.001
Yacht	308	6	1.157 ± 0.222	0.879 ± 0.294	0.867 ± 0.047	3.752 ± 0.240	1.584 ± 0.178
Kin8nm	8,192	8	0.101 ± 0.002	0.081 ± 0.003	0.100 ± 0.003	0.174 ± 0.006	0.110 ± 0.005
Power	9,568	4	4.419 ± 0.046	287.227 ± 200.167	4.060 ± 0.009	4.243 ± 0.011	4.167 ± 0.034
Year	515,345	90	25.163 ± 1.990	NA	8.879 ± 0.004	8.887 ± 0.014	8.874 ± 0.015

Table 2: RMSE on the test set for nine different UCI regression datasets.

Dataset	NLL					training time				
	SVI	MCMC	PBP	KBNN 1	KBNN 10	SVI	MCMC	PBP	KBNN 1	KBNN 10
Boston	2.690 ± 0.04	2.346 ± 0.01	2.421 ± 0.04	3.183 ± 0.05	2.767 ± 0.39	21.4	446.0	8.2	0.8	8.7
Concrete	3.446 ± 0.04	3.236 ± 0.02	3.119 ± 0.01	3.676 ± 0.12	8.428 ± 0.95	22.5	481.7	12.7	1.7	17.5
Energy	2.877 ± 0.03	1.315 ± 0.18	2.680 ± 0.02	3.078 ± 0.02	2.394 ± 0.16	21.6	405.9	10.2	1.2	13.2
Wine	1.107 ± 0.01	1.003 ± 0.01	1.014 ± 0.00	1.529 ± 0.20	1.127 ± 0.12	23.8	520.3	49.1	8.3	86.7
Naval	-1.817 ± 0.18	-3.424 ± 0.09	-3.736 ± 0.02	1.266 ± 0.26	0.128 ± 0.01	42.8	367.0	116.1	20.5	205.3
Yacht	1.435 ± 0.34	1.649 ± 0.46	1.558 ± 0.04	3.033 ± 0.02	2.325 ± 0.06	21.5	357.4	5.9	0.5	5.0
Kin8nm	-0.869 ± 0.03	-1.094 ± 0.03	-0.882 ± 0.03	-0.255 ± 0.14	-0.758 ± 0.04	57.6	1200.0	107.9	20.2	204.0
Power	3.162 ± 0.07	598.460 ± 422.97	2.820 ± 0.00	3.062 ± 0.04	2.922 ± 0.02	56.6	769.9	93.4	20.0	208.6
Year	6.801 ± 0.77	NA	3.588 ± 0.00	4.638 ± 0.22	4.315 ± 0.12	5419.7	NA	5694.9	2021.7	20801.8

Table 3: NLL on the test set and training time on the training set for the UCI regression datasets.

the weight posterior in closed form and in an online fashion, which are two key features compared to the state of the art.

Strengths Given the Assumptions 1 and 2, which do not hinder the learning abilities of the KBNN in practice, our approach performs approximate but fully Bayesian inference for training. For ReLU activations it provides moment matching Gaussian approximations of the predictive and posterior distribution. This is clearly an advantage compared to other methods that rely on stochastic gradient descent. The absence of gradients proves to be data efficient and enables the usage of activation functions that cannot be used in gradient-based learning, e.g., the Heaviside activation or non-differentiable activation schemes.

A second advantage of the proposed method is the ability of learning from sequential data streams without retraining. As shown in the conducted experiments every data instance has to be seen only once during training while still achieving decent performance on the respective test set. This can be especially useful in online learning scenarios or in the context of model-based reinforcement learning where retraining is needed to update the model of the environment.

The update rule of the weights’ means (13) can more abstractly be written as $\text{new} = \text{old} + \mathbf{L} \cdot \Delta$, which is similar to the backpropagation update rule. But instead of a scalar learning rate being a hyper-parameter, KBNN uses the matrix \mathbf{L} , i.e., it uses a matrix-valued, intrinsically calculated learning rate where each weight obtains its individual rate.

Limitations To keep the probability distribution of the network parameters manageable in complexity, independence between the weights of different neurons is assumed (cf. Assumption 1). Abandoning this independence would require the calculation of cross-covariances between neu-

rons. This affects our approach mainly in two ways. First, the memory and computational demand for additionally calculating these terms increases quadratically with the number of neurons per layer. Second, the necessary additional calculation of $E[f(a_i) \cdot f(a_j)]$ to obtain the cross-covariance between the activations of neurons $i, j = 1 \dots M_l$ in the forward pass is challenging. It is questionable if an analytic solution even for ReLU activations exists (Wu et al. 2019).

The considered assumptions significantly simplify the necessary computations and enable closed-form calculations of the quantities of interest. While Assumption 2 is very reasonable for regression tasks, it is not well justified for classification tasks where one would rather want to use for instance a Bernoulli distribution for the output (Kristiadi, Hein, and Hennig 2020). The use of distributions other than a Gaussian as in our case would only be possible if the Kalman filter in the last layer is replaced by more advanced filters such as a particle filter, which uses sampling to approximate the posterior (Särkkä 2013). The Gaussian assumption seems not to impair the performance of the KBNN in classification tasks, at least in the conducted experiments.

Open Issues and Future Work For multi-class classification problems it is common to use a soft-max activation at the output layer. Unfortunately, there generally is no closed-form solution of (8) and (9) if $f(\cdot)$ is a soft-max function. At least (8) can be calculated if the mean-field approximation is applied (Lu, Ie, and Sha 2021). Using a hard-max activation instead of soft-max allows a closed-form solution.

PBP learns its hyper-parameters, which is not the case for the KBNN. To avoid tedious hyper-parameter tuning, adopting a hierarchical Bayesian approach as in (Hernández-Lobato and Adams 2015) is part of future work.

Acknowledgements

This work was partially supported by the Baden-Württemberg Ministry of Economic Affairs, Labor, and Tourism within the KI-Fortschrittszentrum “Lernende Systeme und Kognitive Robotik” under Grant No. 036-140100 and the Fraunhofer-Gesellschaft within ML4Safety.

References

- Achille, A.; and Soatto, S. 2018. Emergence of Invariance and Disentanglement in Deep Representations. *Journal of Machine Learning Research*, 19(1): 1947–1980.
- Amisha, P. M.; Pathania, M.; and Rathaur, V. K. 2019. Overview of artificial intelligence in medicine. *Journal of family medicine and primary care*, 8(7): 2328–2331.
- Begoli, E.; Bhattacharya, T.; and Kusnezov, D. 2019. The need for uncertainty quantification in machine-assisted medical decision making. *Nature Machine Intelligence*, 1(1): 20–23.
- Bengio, Y. 2012. *Practical Recommendations for Gradient-Based Training of Deep Architectures*, 437–478. Springer Berlin Heidelberg.
- Bingham, E.; Chen, J. P.; Jankowiak, M.; Obermeyer, F.; Pradhan, N.; Karaletsos, T.; Singh, R.; Szerlip, P.; Horsfall, P.; and Goodman, N. D. 2019. Pyro: Deep Universal Probabilistic Programming. *Journal of Machine Learning Research*, 20(28): 1–6.
- Dua, D.; and Graff, C. 2017. UCI Machine Learning Repository. *University of California, Irvine, School of Information and Computer Sciences*, <http://archive.ics.uci.edu/ml>.
- Duane, S.; Kennedy, A. D.; Pendleton, B. J.; and Roweth, D. 1987. Hybrid Monte Carlo. *Physics letters B*, 195(2): 216–222.
- El-Shamouty, M.; Kleeberger, K.; Lämmle, A.; and Huber, M. 2019. Simulation-driven machine learning for robotics and automation. *tm-Technisches Messen*, 86(11): 673–684.
- Gal, Y.; and Ghahramani, Z. 2016. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In *Proceedings of The 33rd International Conference on Machine Learning*, Proceedings of Machine Learning Research, 1050–1059.
- Geman, S.; and Geman, D. 1984. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence*, (6): 721–741.
- Ghosh, S.; Fave, F. M. D.; and Yedidia, J. 2016. Assumed Density Filtering Methods for Learning Bayesian Neural Networks. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, 1589–1595.
- Graves, A. 2011. Practical variational inference for neural networks. In *Advances in neural information processing systems*, 2348–2356.
- Hernández-Lobato, J. M.; and Adams, R. 2015. Probabilistic backpropagation for scalable learning of bayesian neural networks. In *International Conference on Machine Learning*, 1861–1869. PMLR.
- Hoffman, M. D.; Blei, D. M.; Wang, C.; and Paisley, J. 2013. Stochastic Variational Inference. *Journal of Machine Learning Research*, 14(4): 1303–1347.
- Hoffman, M. D.; and Gelman, A. 2014. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1): 1593–1623.
- Huber, M. F. 2015. *Nonlinear Gaussian Filtering: Theory, Algorithms, and Applications*, volume 19. KIT Scientific Publishing.
- Huber, M. F. 2020. Bayesian Perceptron: Towards fully Bayesian Neural Networks. In *59th IEEE Conference on Decision and Control*, 3179–3186.
- Kalman, R. E. 1960. A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82(1): 35–45.
- Kingma, D. P.; and Welling, M. 2014. Auto-Encoding Variational Bayes. *CoRR*, abs/1312.6114.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; Hassabis, D.; Clopath, C.; Kumaran, D.; and Hadsell, R. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13): 3521–3526.
- Kristiadi, A.; Hein, M.; and Hennig, P. 2020. Being Bayesian, Even Just a Bit, Fixes Overconfidence in ReLU Networks. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, 5436–5446.
- Kurle, R.; Cseke, B.; Klushyn, A.; van der Smagt, P.; and Günnemann, S. 2019. Continual learning with bayesian neural networks for non-stationary data. In *International Conference on Learning Representations*.
- Lu, Z.; Ie, E.; and Sha, F. 2021. Mean-Field Approximation to Gaussian-Softmax Integral with Application to Uncertainty Estimation. *arXiv preprint arXiv:2006.07584v2*.
- MacKay, D. J. 1992. A practical Bayesian framework for backpropagation networks. *Neural computation*, 4(3): 448–472.
- Maybeck, P. S. 1979. *Stochastic Models, Estimation, and Control*, volume 141 of *Mathematics in Science and Engineering*. Academic Press, Inc.
- Metropolis, N.; Rosenbluth, A. W.; Rosenbluth, M.; Teller, A. H.; and Teller, E. 1953. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21: 1087–1092.
- Minka, T. 2001. *A family of algorithms for approximate Bayesian inference*. Ph.D. thesis, USA.
- Minka, T. P.; Xiang, R.; and Qi, Y. A. 2009. Virtual Vector Machine for Bayesian Online Classification. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, UAI ’09, 411–418. Arlington, Virginia, USA: AUAI Press. ISBN 9780974903958.
- Neal, R. M. 1995. *Bayesian Learning for Neural Networks*. Ph.D. thesis, CAN. AAINN02676.

Nguyen, C. V.; Li, Y.; Bui, T. D.; and Turner, R. E. 2018. Variational Continual Learning. In *International Conference on Learning Representations*.

Opper, M. 1998. A Bayesian Approach to Online Learning. In *Online Learning in Neural Networks*, 363–378.

Parisi, G. I.; Kemker, R.; Part, J. L.; Kanan, C.; and Wermter, S. 2019. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113: 54 – 71.

Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12: 2825–2830.

Puskorius, G. V.; and Feldkamp, L. A. 2001. Parameter-based Kalman Filter Training: Theory and Implementation. *Kalman filtering and neural networks*, 23.

Rauch, H. E.; Tung, F.; and Striebel, C. T. 1965. Maximum likelihood estimates of linear dynamic systems. *AIAA Journal*, 3(8): 1445–1450.

Ritter, H.; Botev, A.; and Barber, D. 2018a. A Scalable Laplace Approximation for Neural Networks. In *Proceedings of the 6th International Conference on Learning Representations*.

Ritter, H.; Botev, A.; and Barber, D. 2018b. Online Structured Laplace Approximations for Overcoming Catastrophic Forgetting. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, 3742–3752. Red Hook, NY, USA: Curran Associates Inc.

Snoek, J.; Rippel, O.; Swersky, K.; Kiros, R.; Satish, N.; Sundaram, N.; Patwary, M. M. A.; Prabhakar, P.; and Adams, R. P. 2015. Scalable Bayesian Optimization Using Deep Neural Networks. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning*, 2171–2180.

Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15(56): 1929–1958.

Särkkä, S. 2013. *Bayesian Filtering and Smoothing*. Institute of Mathematical Statistics Textbooks. Cambridge University Press.

Watanabe, K.; and Tzafesta, S. G. 1990. Learning Algorithms for Neural Networks with the Kalman Filter. *Journal of Intelligent and Robotic Systems*, 3: 305–319.

Wu, A.; Nowozin, S.; Meeds, E.; Turner, R. E.; Hernández-Lobato, J. M.; and Gaunt, A. L. 2019. Deterministic Variational Inference for Robust Bayesian Neural Networks. In *Proceedings of the 7th International Conference on Learning Representations*.

Zeng, A.; Song, S.; Lee, J.; Rodriguez, A.; and Funkhouser, T. 2020. Tossingbot: Learning to Throw Arbitrary Objects with Residual Physics. *IEEE Transactions on Robotics*, 36(4): 1307–1319.

Zhang, C.; Bütepage, J.; Kjellström, H.; and Mandt, S. 2019. Advances in Variational Inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8): 2008–2026.