



Bernstein flows for flexible posteriors in variational Bayes

Oliver Dürr¹ · Stefan Hörting¹ · Danil Dold¹ · Ivonne Kovylov¹ · Beate Sick^{2,3}

Received: 18 November 2022 / Accepted: 6 February 2024

© The Author(s) 2024

Abstract

Black-box variational inference (BBVI) is a technique to approximate the posterior of Bayesian models by optimization. Similar to MCMC, the user only needs to specify the model; then, the inference procedure is done automatically. In contrast to MCMC, BBVI scales to many observations, is faster for some applications, and can take advantage of highly optimized deep learning frameworks since it can be formulated as a minimization task. In the case of complex posteriors, however, other state-of-the-art BBVI approaches often yield unsatisfactory posterior approximations. This paper presents Bernstein flow variational inference (BF-VI), a robust and easy-to-use method flexible enough to approximate complex multivariate posteriors. BF-VI combines ideas from normalizing flows and Bernstein polynomial-based transformation models. In benchmark experiments, we compare BF-VI solutions with exact posteriors, MCMC solutions, and state-of-the-art BBVI methods, including normalizing flow-based BBVI. We show for low-dimensional models that BF-VI accurately approximates the true posterior; in higher-dimensional models, BF-VI compares favorably against other BBVI methods. Further, using BF-VI, we develop a Bayesian model for the semi-structured melanoma challenge data, combining a CNN model part for image data with an interpretable model part for tabular data, and demonstrate, for the first time, the use of BBVI in semi-structured models.

Keywords Variational inference · Deep learning · Transformation models · Bayesian neural network

1 Introduction

Uncertainty quantification is essential, especially if model predictions are used to support high-stakes decision-making. Quantifying uncertainty in statistical or machine learning models is often achieved by Bayesian approaches, where posterior distributions represent the uncertainty of the estimated model parameters.

Oliver Dürr and Beate Sick have contributed equally to this work.

Extended author information available on the last page of the article

Determining the exact posterior distributions is often impossible when the posterior takes a complex shape and the model has many parameters. This is especially true for complex models such as Bayesian neural networks (NNs) or semi-structured models that combine an interpretable model part with deep NNs. Variational inference (VI) is a commonly used approach to approximate complex distributions through optimization (Jordan et al. 1999; Blei et al. 2017). In VI, the complex posterior is approximated by a variational distribution by minimizing a divergence measure between the variational and the true posterior distribution. VI is currently a very active research field tackling different challenges, which can be categorized into the following groups: (1) constructing variational distributions that are flexible enough to match the true posterior distribution, (2) defining optimal variational objective for tuning the variational distribution, which boils down to finding the most suited divergence measure quantifying the difference between a variational distribution and posterior, and (3) developing robust and accurate stochastic optimization frameworks for the variational objective (Dhaka et al. 2020; Blei et al. 2016; Welandawe et al. 2022). Here, we focus on challenge (1) and propose a method to construct a variational distribution that is flexible enough to accurately and robustly approximate complex multidimensional posteriors.

To avoid model-specific calculations, we design our method as a black-box VI (BBVI) approach (Ranganath et al. 2014). In BBVI, the approximative posterior is determined by stochastic gradient descent. The user simply defines the Bayesian model by specifying the likelihood and the prior, after which all subsequent calculations are carried out automatically. Due to its simplicity, BBVI is implemented in many packages for Bayesian modeling, like Stan (Carpenter et al. 2017) and Pyro (Bingham et al. 2019) as an alternative to MCMC. Given BBVI's scalability to large datasets and its widespread applicability, it has emerged as the preferred technique in the field of machine learning (Welandawe et al. 2022).

Our approach uses transformation models (TMs) to construct complex posteriors. Transformation models (TMs) have been introduced for fitting potentially complex outcome distributions for probabilistic regression models (Hothorn et al. 2014). Since then, they have been mainly used to model different outcome types, such as ordinal (Kook et al. 2022; Buri et al. 2020), count (Siegfried and Hothorn 2020), continuous (Lohse et al. 2017), or time-to-event outcomes (Campanella et al. 2022) based on tabular predictors. Moreover, TMs have been used to model multidimensional distributions (Klein et al. 2019). Neural networks can be used to extend TMs to model outcomes for unstructured predictors (e.g., images or text) or a combination of tabular and unstructured predictors (Sick et al. 2021; Baumann et al. 2021; Kook et al. 2022; Rügamer et al. 2021).

The basic idea of TMs is to learn a flexible and monotone transformation function that transforms between a simple latent distribution and a potentially complex conditional outcome distribution. In TMs, the transformation function is parameterized as an expansion of basis functions. In the case of continuous target distributions, most often, Bernstein polynomials (Bernštein 1912) are used because they can easily be constrained to be strictly monotone, and their flexibility can be tuned via the order M . A large order M ensures an accurate approximation of the distribution, which is robust against a further increase of M (Hothorn et al. 2018;

Ramasinghe et al. 2021); this is also demonstrated in our experiments for the BBVI setting.

Independently of TMs, normalizing flows (NFs) have been developed in the deep learning community. NFs and TMs rely on the same idea, but NFs usually construct the transformation by chaining many simple functions, while TMs construct one rather complex transformation function. In NFs, each simple function, such as shifting and scaling, incrementally adds to the complexity of the final transformation. Among the prominent NF implementations are RealNVP (Dinh et al. 2016) and Masked Autoregressive Flow (MAF) (Papamakarios et al. 2017). RealNVP stands out for its efficient, invertible transformations facilitated by a specialized neural network architecture. Its key advantage lies in the efficient computation of the Jacobian matrix's determinant, essential for direct density estimation in the change of variable function (see 6). This efficiency is achieved by iteratively splitting the components of the data into two parts. In each step, the first part of the components is used to train a neural network computing the scale and shift parameters of the transformation. This transformation is then applied to the other components, while the first part remains unaltered. This procedure is repeated multiple times with different partitioning, leading to a triangular Jacobian, thus enabling efficient and invertible transformations. In contrast, MAF adopts a fundamentally different approach to construct transformations (Papamakarios et al. 2017). It utilizes a sequential (autoregressive) framework, facilitated by neural networks. In MAF each output component relies exclusively on its preceding components, a concept often referred to as causality in this context. This design also leads to a triangular Jacobian matrix and thus a fast computation of the change of variable equation. The MAF ensures that the n th output of NN is solely dependent on the first $n - 1$ inputs, yielding an autoregressive model. However, some NF approaches use a single flexible transformation, such as sum-of-squares polynomials (Jaini et al. 2019) or splines (Durkan et al. 2019). Recently, also Bernstein-based polynomials have been used for modeling unconditional multivariate density distributions (Ramasinghe et al. 2021).

NFs were initially introduced for variational inference to approximate potentially complex distributions of latent variables in models such as variational autoencoders (Rezende and Mohamed 2015; Van Den Berg et al. 2018). In the past, often members from simple distribution families have been used to approximate the posterior in BBVI. In the "Bayes by Backprop" method, Blundell et al. used independent Gaussians to approximate the posterior of the weights in a Bayesian Neural Network (BNN). They determined the parameters of these Gaussians using BBVI (Blundell et al. 2015). This approach was made more flexible by using a multivariate Gaussians (Louizos and Welling 2017) as variational distribution. While it is clear that TMs or NFs have the potential to construct flexible variational distributions, the first attempts to use NF-based BBVI were proposed only recently (Agrawal et al. 2020). These NF-based BBVI approaches compare favorably against existing BBVI methods but require a complex training scheme and sometimes exhibit pathological behavior (Dhaka et al. 2021).

Here, we introduce Bernstein flow variational inference (BF-VI), which, for the first time, uses TMs in BBVI. We use TMs based on Bernstein polynomials to construct a

variational distribution that closely approximates a potentially complex posterior in Bayesian models. The proposed method is computationally efficient and applicable to typical statistical models. The proposed method yields superior results in our experiments compared to existing NF approaches (Dhaka et al. 2021). Using BF-VI we further demonstrate, for the first time, how VI can be used to fit Bayesian semi-structured models where interpretable statistical model parts (based on tabular data) and deep NN model parts (based on images) are jointly fitted. We define our method in Sect. 2.1 for one-dimensional examples and generalize it to Bayesian models with multivariate posteriors in Sect. 2.2. In Sect. 3, we benchmark our BF-VI approach against exact Bayesian models, MCMC-Simulations, Gaussian-VI, and NF-based BBVI, showing accurate posterior approximations in low dimensions and superior approximations in higher dimensions when compared to NF-based BBVI and summarize in Sect. 4.

2 Bernstein flow variational inference

In the following, we describe the Bernstein Flow-VI (BF-VI) approach, which we propose for accurately and robustly approximating potentially complex posteriors in Bayesian models. The main idea is to enable the VI procedure to approximate the joined posterior of the p model parameters by a flexible variational distribution. This is done by modeling the transformation function from a predefined simple latent distribution to a potentially complex variational distribution. The number of parameters p in the Bayesian model determines the dimension of both the latent and the variational distribution.

We first explain BF-VI for Bayesian models with a single parameter and hence a one-dimensional posterior and then generalize to models with multivariate posteriors. The code is publicly available on GitHub.¹

2.1 One-dimensional Bernstein flows

BF-VI approximates the bijective transformation $g : Z \rightarrow \theta$ between a latent variable $Z \in \mathbb{R}$ with predefined distribution $F_Z : \mathbb{R} \rightarrow [0, 1]$ with log-concave and continuous density f_Z , and the model parameter $\theta \in \mathbb{R}$ with a potentially complex distribution $F_\theta : \mathbb{R} \rightarrow [0, 1]$ so that $F_Z(z) = F_\theta(g(z))$. Figure 1 visualizes this transformation on the scale of the densities, where $f_Z(z) = f_\theta(g(z)) \left| \frac{\partial g(z)}{\partial z} \right|$ according to the change-of-variable formula.

Hothorn et al. (2018) give theoretical guarantees for the existence and uniqueness of $g = F_\theta^{-1} \circ F_Z$. However, g cannot be computed directly if F_θ is not known (in our application F_θ is the unknown distribution of the posterior). The core of BF-VI is to approximate g , shown in Fig. 1, by f_{BP} Bernstein polynomials (BP)² as

¹ https://github.com/tensorchiefs/bfvi_paper.

² Some authors make a distinction between Bernstein polynomials in which ϑ_i is fixed by the values of the function to be approximated and call expressions like expressions like in (1), where ϑ_i is a fitting parameter, polynomials of Bernstein type.

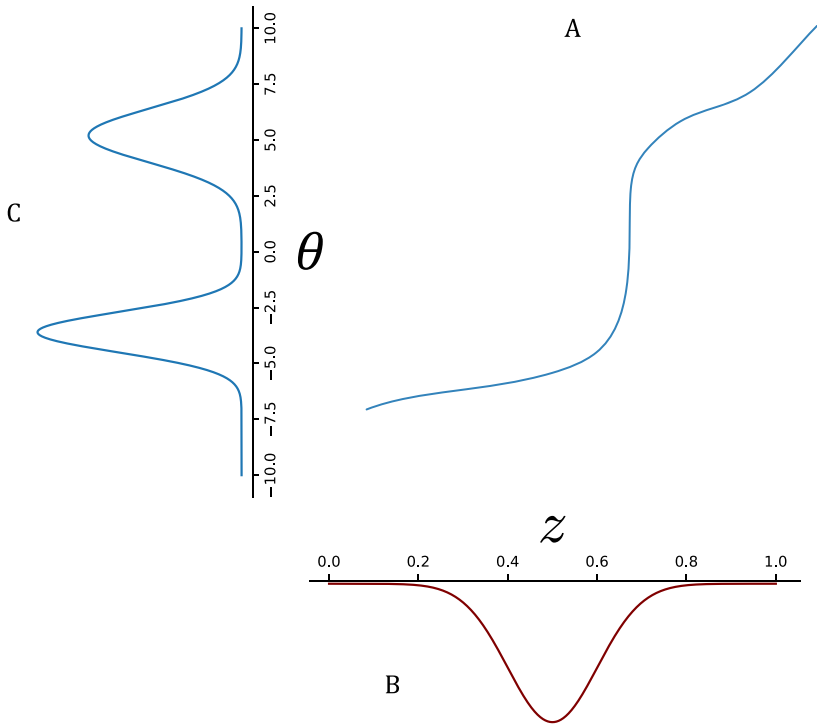


Fig. 1 Overview of the transformation model. **A** Shows the bijective transformation function $g : Z \rightarrow \theta$ (or its approximation f_{BP}) mapping between, **B** a predefined latent density f_Z , and **C** a potentially complex posterior (or its variational distribution)

$$f_{BP}(z) = \sum_{i=0}^M \text{Be}_i(z) \frac{\vartheta_i}{M+1} \tag{1}$$

with $\text{Be}_i(z) = \text{Be}_{(i+1, M-i+1)}(z)$ being the density of a Beta distribution with parameters $i + 1$ and $M - i + 1$. To preserve the bijectivity of g , we use in BF-VI w.l.o.g. a strict monotone increasing BP to approximate g . With the approximation of the transformation function g by f_{BP} , it holds that $f_Z(z)$ can be approximated by $f_\theta(f_{BP}(z)) \left| \frac{\partial f_{BP}(z)}{\partial z} \right|$.

Using a BP for approximating g gives the following theoretical guarantees (Farouki 2012) (1) With increasing order M of the BP, the approximation f_{BP} to g gets arbitrarily close (the BP have been introduced for this very purpose in the constructive proof of the Weierstrass theorem by Bernštein (1912)); (2) the required

strict monotonicity of the approximation f_{BP} can be easily achieved by constraining the coefficients ϑ_i of the BP to be increasing; (3) BPs are robust versus perturbations of the coefficients ϑ_i ; (4) the approximation error decreases with $1/M$ (Voronovskaya's theorem). See Bernštein (1912); Farouki (2012) for more detailed discussions of the beneficial properties of BPs in general and Hothorn et al. (2018); Ramasinghe et al. (2021) for transformation models.

While the output of $f_{\text{BP}}(z)$ is unrestricted, a BP requires a input z within $[0, 1]$. We experimented with several approaches to ensure the restriction $z \in [0, 1]$ resulting in slightly different behavior during the training (see Appendix B.2.2). Based on these experiments, we decided to obtain $z \in [0, 1]$ by sampling values from a standard normal distribution, $z'' \sim N(0, 1)$, then apply the affine transformation $l(z'') = \alpha \cdot z'' + \beta$, followed by a sigmoid $\sigma(z') = 1/(1 + e^{-z'})$. Altogether, we approximate the transformation g by $f : Z \rightarrow \theta$ via $f = f_{\text{BP}} \circ \sigma \circ l$, which we call Bernstein flow.

To allow the application of unconstrained stochastic gradient descent optimization, which is typically used in the deep learning domain, we enforce the strict monotonicity of the flow f as follows: We optimize unrestricted parameters of f , i.e., $\vartheta'_0, \dots, \vartheta'_M, \alpha', \beta'$, and apply the following transformations to determine the parameters of the bijective flow: $\vartheta_0 = \vartheta'_0$, and $\vartheta_i = \vartheta_{i-1} + \text{softplus}(\vartheta'_i)$ for $i = 1, \dots, M$ for getting a strictly increasing BP and $\alpha = \text{softplus}(\alpha')$, $\beta = \beta'$ for getting an increasing affine transformation.

In Appendix A, we show that the resulting variational distribution is a tight approximation to the posterior in the sense that the KL divergence between $q_\lambda(\theta)$ and $p(\theta | D)$ decreases with the order of the BP via $1/M$.

2.2 Multivariate generalization

In the case of a Bayesian model with p parameters, $\theta_1, \theta_2, \dots, \theta_p$, the Bernstein flow bijectively maps a p -dimensional \mathbf{Z}' to a p -dimensional θ . We realize this flow by choosing p independent standard normal Gaussians as simple latent distribution for the p -dimensional \mathbf{Z}' and apply on each component an affine transformation followed by a sigmoid function to achieve a $[0,1]$ restricted \mathbf{Z} . The possible complex dependencies in θ are modeled in the multivariate generalization \mathbf{f}_{BP} of the one-dimensional Bernstein polynomial (see Eq. 2 for the definition of the j th component of \mathbf{f}_{BP}).

$$\theta_j = f_{\text{BP}_j}(z_{1:j}) = \frac{1}{M+1} \sum_{i=0}^M \vartheta_i^j(z_1, \dots, z_{j-1}) \text{Be}_i(z_j) \quad (2)$$

To achieve an efficient computation, we use a triangular map for constructing coefficients ϑ_i^j , $j = 2, \dots, p$, $i = 0, \dots, M$ from \mathbf{Z} . This ensures that the j th BP determining θ_j only depends on the first $j-1$ components of \mathbf{Z} (see Eq. 2). It is known that bijective triangular maps with sufficient flexibility can map a simple p -dimensional distribution into arbitrary complex p -dimensional target distributions (Bogachev et al.

2005). We use a masked autoregressive flow (MAF) (Papamakarios et al. 2017) to map \mathbf{Z} to the BP coefficients $\vartheta_i^j, j = 2, \dots, p, i = 0, \dots, M$ from \mathbf{Z} . The MAF architecture ensures that

that ϑ_i^j depend only on those components of the latent variables $z_{j'}$ with $j' \leq j$ (as required in Eq. 2). Note that the first coefficients in all BPs ϑ^1 do not depend on z and are therefore not modeled via the MAF. Therefore, the Jacobian $\nabla \mathbf{f}_{\text{BP}}$ w.r.t. \mathbf{z} is a triangular matrix, and hence $\det \nabla \mathbf{f}_{\text{BP}}$ is given by the product of the diagonal elements of the Jacobian allowing for efficient computation of the resulting p -dimensional variational distribution $q_\lambda(\boldsymbol{\theta})$ via the multivariate version of the change of variable formula (see Eq. 6). The flexibility of such a p -dimensional bijective Bernstein flow is only limited by the order M of the Bernstein polynomial and the complexity of the MAF. In our experiments, we use an MAF with two hidden layers, each with 10 neurons. The weights \mathbf{w} of the MAF are part of the variational parameters for \mathbf{f}_{BP} . In total, we have $\lambda = (\vartheta^1, \mathbf{w}, \boldsymbol{\alpha}, \boldsymbol{\beta})$ variational parameters.

2.3 Variational inference procedure

In VI the variational parameters λ are tuned such that the resulting variational distribution $q_\lambda(\boldsymbol{\theta})$ is as close to the posterior $p(\boldsymbol{\theta} | D)$ as possible. Here, we do this by minimizing the KL divergence between the variational distribution and the (unknown) posterior:

$$\begin{aligned} \text{KL}(q_\lambda(\boldsymbol{\theta}) \| p(\boldsymbol{\theta} | D)) &= \int q_\lambda(\boldsymbol{\theta}) \log \left(\frac{q_\lambda(\boldsymbol{\theta})}{p(\boldsymbol{\theta} | D)} \right) d\boldsymbol{\theta} \\ &= \log(p(D)) - \underbrace{(\mathbb{E}_{\boldsymbol{\theta} \sim q_\lambda}(\log(p(D | \boldsymbol{\theta}))) - \text{KL}(q_\lambda(\boldsymbol{\theta}) \| p(\boldsymbol{\theta})))}_{\text{ELBO}(\lambda)} \end{aligned} \quad (3)$$

The KL divergence is commonly used in VI, and a recent study showed that it is easier to train than other divergences and applicable to higher-dimensional distributions (Dhaka et al. 2021).

Instead of minimizing (3) usually only the evidence lower bound (ELBO) is maximized (Blundell et al. 2015) which consists of the expected value of the log-likelihood, $\mathbb{E}_{\boldsymbol{\theta} \sim q_\lambda}(\log(p(D | \boldsymbol{\theta})))$, minus the KL divergence between the variational distribution $q_\lambda(\boldsymbol{\theta})$ and the prior $p(\boldsymbol{\theta})$. Note that the ELBO does not explicitly contain the unknown posterior. In practice, we minimize the negative ELBO using stochastic gradient descent facilitated by automatic differentiation. For consistency with Dhaka et al. (2021), we use TensorFlow's RMSprop in all our experiments, configured with the default settings. We follow the BBVI approach and approximate the expected log-likelihood by averaging over S samples $\boldsymbol{\theta}_s \sim q_\lambda(\boldsymbol{\theta})$ via

$$\mathbb{E}_{\boldsymbol{\theta} \sim q_\lambda}(\log(p(D_i | \boldsymbol{\theta}))) \approx \frac{1}{S} \sum_{s,i} \log(p(D_i | \boldsymbol{\theta}_s)). \quad (4)$$

Hereby, we also assume the usual independence of the $i = 1, \dots, N$ training data points D_i . To get these samples θ_s , we use S samples \mathbf{z}'_s from the latent distribution, apply the transformation $f = f_{\text{BP}} \circ \sigma \circ l$, and then compute the corresponding parameter samples via $\theta_s = \mathbf{f}(\mathbf{z}'_s)$. We use the same samples $\theta_s \sim q_\lambda(\theta)$ to approximate the Kullback–Leibler divergence between the variational distribution $q_\lambda(\theta)$ and the prior $p(\theta)$ via:

$$\text{KL}(q_\lambda(\theta) \| p(\theta)) \approx \frac{1}{S} \sum_s \log \left(\frac{q_\lambda(\theta_s)}{p(\theta_s)} \right) \quad (5)$$

where the probability density $q_\lambda(\theta_s)$ can be calculated, from the samples θ_s using the change of variable function as:

$$q_\lambda(\theta_s) = p(\mathbf{z}'_s) \cdot |\det \nabla_{\mathbf{z}'} \mathbf{f}_{\text{BP}}(\sigma(l(\mathbf{z}'_s)))|^{-1} \quad (6)$$

2.4 Evaluation

Evaluating the quality of the fitted variational distributions requires a comparison to the true posterior. In the case of low-dimensional problems, the two distributions can be compared visually. In the case of higher-dimensional problems, this is not possible anymore. While the evidence lower bound (ELBO) is a valuable metric for optimizing the parameters in VI, it is less helpful in comparing different approximations because it depends on the specific parametrization of the model (Yao et al. 2018). Therefore, Yao et al. (2018) introduced \hat{k} as a more suited approach for comparison, which since then has been used in other studies like (Dhaka et al. 2021) to which we compare. The computation of \hat{k} is based on the importance ratios which are defined as

$$r_s = \frac{p(\theta_s, D)}{q_\lambda(\theta_s)} = \frac{p(D | \theta_s)p(\theta_s)}{q_\lambda(\theta_s)} \quad (7)$$

If the variational distribution $q_\lambda(\theta)$ would be a perfect approximation of the posterior $p(\theta | D) \propto p(D | \theta)p(\theta)$, then important ratios r_s would be constant. However, because of the asymmetry of the KL divergence used in the optimization objective (see Eq. 3), the fitted $q_\lambda(\theta)$ tends to have lighter tails than $p(\theta | D)$, with the effect that the distribution of r_s is heavily right-tailed. To quantify the severity of the underestimated tails, a generalized Pareto distribution is fitted to the right tail of the r_s . The estimated shape parameter \hat{k} of the Pareto distribution can be used as a diagnostic tool. A large \hat{k} indicates a pronounced tail in the r_s distribution and, hence, a bad posterior approximation. According to Yao et al. (2018) values of $\hat{k} < 0.5$ indicate that the variational approximation q_λ is good. Values of $0.5 < \hat{k} < 0.7$ indicate the variational approximation q_λ is not perfect but still useful.

3 Experiments

We performed several experiments to benchmark our BF-VI approach versus exact Bayesian solutions, Gaussian-VI, and recent NF-VI approaches. All experiments were conducted using five repetitions, with the observed stability of evidence lower bound (ELBO) optimization generally appearing independent of the randomly chosen starting values. Table 1 shows an overview of the fitted models. The complete model definitions in Stan, along with the code for all experiments, can be found on GitHub.

3.1 Models with a single parameter

First, we demonstrate with two single-parameter experiments that BF-VI can accurately approximate a skewed or bimodal posterior, which is impossible with Gaussian-VI. To obtain complex posterior shapes, we work with small datasets.

Bernoulli experiment

We first look at an unconditional Bayesian model for a random variable Y following a Bernoulli distribution $Y \sim \text{Ber}(\pi)$ which we fit based on data D consisting only of two samples ($y_1 = 1, y_2 = 1$). In this simple Bernoulli model, it is possible to determine the solution for the posterior analytically when using a beta distribution as prior. We choose $p(\pi) = \text{Be}_{(1,1,1,1)}(\pi)$ which leads to the conjugated posterior $p(\pi | D) = \text{Be}_{(\alpha + \sum y_i, \beta + n - \sum y_i)}(\pi)$ (see analytical posterior in Fig. 2).

We now use BF-VI to approximate the posterior. To ensure that the modeled variational distribution for π is restricted to the support of $\pi \in [0, 1]$, we pipe the result of the flow through an additional sigmoid transformation. Figure 2 shows the achieved variational distributions after minimizing the negative ELBO and demonstrates the robustness of BF-VI: When increasing the order M of the BP, the resulting variational distribution gets closer to the posterior up to a certain value of M and then does not deteriorate when further increasing M . The left part of Fig. 2 indicates a convergence order of M , which can also be proven for the one-dimensional case (see Appendix A). As expected, the Gaussian-VI does not have enough flexibility to approximate the analytical posterior nicely (see Fig. 2).

Cauchy experiment

Here, we follow an example from Yao et al. (2022) and fit an unconditional Cauchy model $Y \sim \text{Cauchy}(\xi, \gamma)$ to six samples which we have drawn from a mixture of two Cauchy distributions $Y \sim \text{Cauchy}(\xi_1 = -2.5, \gamma = 0.5) + \text{Cauchy}(\xi_2 = 2.5, \gamma = 0.5)$. Due to the misspecification of the model, the true posterior of the parameter ξ has a bimodal shape which we have determined via MCMC (see Fig. 3). We use BF-VI and Gaussian-VI to approximate the posterior of the Cauchy parameter ξ by a variational distribution. As in the Bernoulli experiment, also here, BF-VI has enough flexibility to accurately approximate the complex shape of the posterior when M is

Table 1 Overview of the fitted Bayesian models in the benchmark experiments and the methods used to get the posteriors

Experiment	Model	Method (ground truth; VI approximations—best is bold)
Bernoulli	$Y \sim \text{Ber}(\pi)$	analytically; BF-VI , MF-Gaussian-VI
Cauchy	$Y \sim \text{Cauchy}(\xi, \gamma)$	MCMC; BF-VI , MF-Gaussian-VI
Toy linear regression	$(Y x_1, x_2) \sim N(\mu_k = \mu_0 + \beta_1 x_1 + \beta_2 x_2, \sigma)$	MCMC; BF-VI , MF-Gaussian-VI
Diamond	$(Y \mathbf{x}) \sim N(\mu_{\mathbf{x}} = \mu_0 + \mathbf{x}^T \boldsymbol{\beta}, \sigma)$	MCMC; BF-VI, MF-Gaussian-VI , NVP-NF-VI, PL-NF-VI
NN nonlinear regression	$(Y x) \sim N(\mu(x), \sigma)$	MCMC; BF-VI, MF-Gaussian-VI
8schools		
CP parametrization	$(Y \theta_n, \sigma_n) \sim N(\theta_n, \sigma_n)$, with $(\theta_n \mu, \tau) \sim N(\mu, \tau)$ $\mu \sim N(0, 5)$, $\tau \sim \text{half-Cauchy}(0, 5)$	MCMC; BF-VI , MF-Gaussian-VI, NVP-NF-VI, PL-NF-VI
NCP parametrization	$(Y \mu, \tau, \tilde{\theta}_n, \sigma_n) \sim N(\mu + \tau \cdot \tilde{\theta}_n, \sigma_n)$, with $\tilde{\theta}_n \sim N(0, 1)$ $\mu \sim N(0, 5)$, $\tau \sim \text{half-Cauchy}(0, 5)$	MCMC; BF-VI , MF-Gaussian-VI, NVP-NF-VI, PL-NF-VI
Melanoma		
M1 (based on image B)	$(Y B) \sim \text{Ber}(\pi_B = \sigma(\mu(B)))$	–; (Ensembling)
M2 (based on tabular feature x)	$(Y x) \sim \text{Ber}(\pi_x = \sigma(\mu_0 + \beta_1 \cdot x))$	MCMC; BF-VI
M3 (semi-structured using x, B)	$(Y B, x) \sim \text{Ber}(\pi_{(B,x)} = \sigma(\mu_0(B) + \beta_1 \cdot x))$	–; BF-VI for β_1 , non-Bayesian CNN for $\mu_0(B)$

MCMC or analytical solutions are used as ground truth, against which the quality of different VI approximations is compared (bold-faced methods outperform all others). The different VI approximations are either compared by the achieved k (see Sect. 2.4) or visually by inspecting their deviation from the ground truth




Fig. 2 Bernoulli experiment. Left panel: comparison of the analytical posterior for the parameter π in the Bernoulli model $Y \sim \text{Ber}(\pi)$ with variational distributions achieved via Gaussian-VI and BF-VI with BP order $M = 1, 10, 50$. Right panel: the dependence of the divergence $KL(q_i(w) \| p(w | D))$ on M for 20 runs




Fig. 3 Cauchy experiment: comparison of MCMC posterior distribution of the parameter ξ in the Cauchy model $Y \sim \text{Cauchy}(\xi, \gamma)$ and the variational distributions estimated via Gaussian-VI or BF-VI ($M = 2, 6, 10, 30, 50$). For the BF-VI method, the curves are overlays of 10 independent runs

chosen large enough. Further increasing M does not deteriorate the approximation. Gauss-VI fails as expected.

3.2 Models with multiple parameters

The following experiments use BF-VI in multi-parameter Bayesian models and benchmark the achieved solutions versus MCMC or published state-of-the-art VI approximations (see Table 1). In the following experiments, we did not tune the flexibility of our BF-VI approach but allowed it to be relatively high ($M = 50$) since BF-VI does not suffer from being too flexible. Further, in all experiments in this section, we trained for 10^5 epochs with 5 repetitions and set the number of samples for MC estimation to $S = 10$ to be comparable with Dhaka et al. (2021). From the repetitions and the posterior samples, we estimated $\hat{\kappa}$ and the

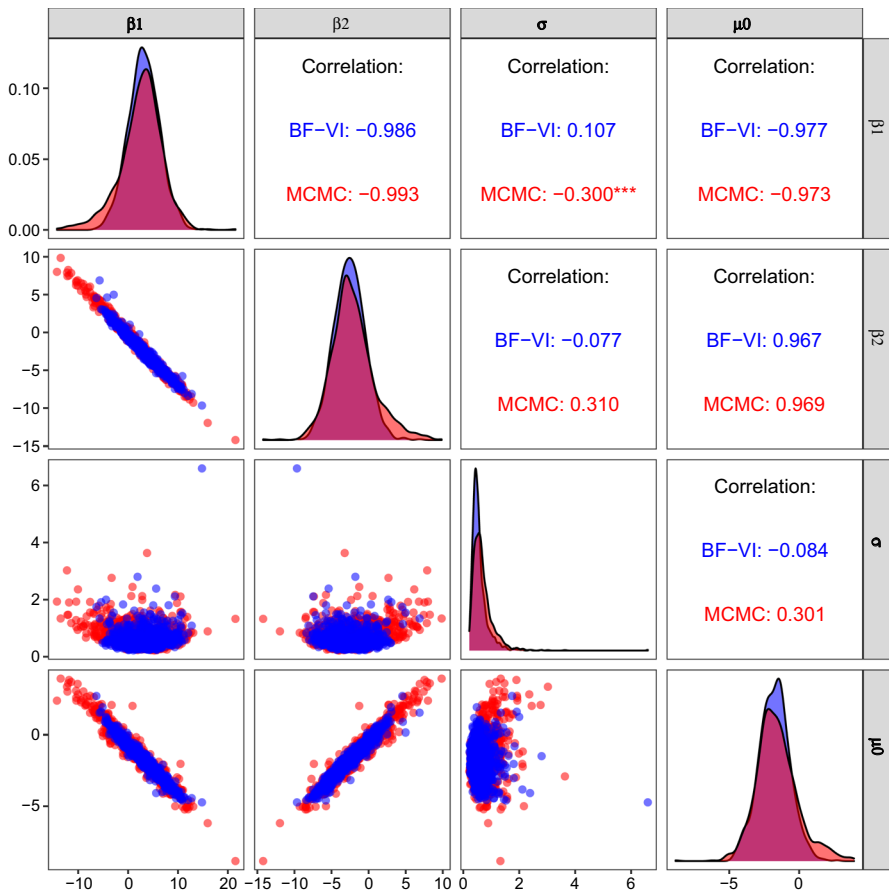


Fig. 4 Toy linear regression example visualization of the posterior. The model has four parameters: two regression coefficients β_1 and β_2 , the intercept μ_0 , and the standard derivation σ . Samples from the true posterior resulting from MCMC (red) are overlaid with samples from the BF-VI approximation (blue)

90 % confidence intervals using the Rubin's rule for the BF-VI and Gaussian-VI methods, with $S = 50'000$ samples. Though runtime is not a consideration in this study, to provide context, a 2023 MacBook Pro's CPU processes approximately 100 epochs per second.

Toy linear regression experiment

To investigate whether dependencies between model parameters are correctly captured, we use a simulated toy dataset with two predictors and six data points to which we fit a Bayesian linear regression modeling the conditional outcome distribution $(Y | x_1, x_2) \sim N(\mu_x = \mu_0 + \beta_1 x_1 + \beta_2 x_2, \sigma)$.

Figure 4 gives a visual impression of the joint true posterior of the four model parameters $(\mu_0, \beta_1, \beta_2, \sigma)$ determined via MCMC samples (red) and its variational

Table 2 Comparison posterior approximation with results from the literature for the Diamond and 8schools dataset in CP and NCP parametrization

	MF-G	MF-T	NVP-VI	P-VI	BF-VI
Diamond	1.2	1.3	∞	∞	5.34(-2.52,13.20)
8Schools (CP)	0.9	0.9	1.3	1.1	0.53(0.11, 0.95)
8Schools (NCP)	0.7	0.6	1.2	0.7	0.36(0.17, 0.55)

Shown is \hat{k} (the lower, the better) for different approximations: mean-Field Gaussian (MF-G) and Student-t (MF-T); NVP (NVP-VI) and Planar (P-VI) flow from Dhaka et al. (2021), and BF-VI

approximation (blue) achieved via BF-VI. The strong correlation between the regression coefficients (β_1, β_2) is nicely captured by the BF-VI approximation. Further, the skewness of posterior marginals involving σ is similar. However, we can see that BF-VI slightly underestimates the long tails of the posterior, confirming the known shortcoming of using the asymmetric KL divergence in the objective function (Blei et al. 2017). BF-VI ($\hat{k} = 0.68(0.51, 0.86)$) is superior to MF-Gaussian-VI ($\hat{k} = 0.90(0.77, 1.01)$), which cannot (by construction) capture the dependencies (MF) or the non-Gaussian shapes (see Fig. S4).

Diamond: linear regression experiment


The *Diamond* linear regression benchmark example ($Y | \mathbf{x}) \sim N(\mu_x = \mu_0 + \mathbf{x}^\top \boldsymbol{\beta}, \sigma)$) has 26 model parameters and 5000 data points (see [37] for reference MCMC samples and Stan code for a complete model definition). Since we have much more data than parameters, the posterior is expected to be a narrow Gaussian around the maximum-likelihood solution, which is indeed seen in the MCMC solution (see Fig. S5). In this setting, BF-VI or NF-VI cannot profit from their ability to fit complex distributions. Still, Dhaka et al. (2021) used this dataset for benchmarking different VI methods, e.g., planar NF (PL-NF-VI), non-volume-preserving NF (NVP-NF-VI), and MF-Gaussian-VI. They achieved the best approximation via the simple Gaussian-VI ($\hat{k} = 1.2$). The posterior approximation via PL-NF-VI and NVP-NF-VI has both been unsatisfactory ($\hat{k} = \infty$). We use the same amount of sampling ($S = 10$) and achieve with BF-VI a better approximation of the posterior $\hat{k} = 5.34(-2.52, 13.20)$ but is still worse than the Gaussian-VI. A large spread in \hat{k} indicates an unstable training procedure. This dataset is also quite challenging for MCMC simulations; we did not get satisfactory MCMC samples and took the reference posterior samples from the posteriorDB.³ See Fig. S5 for a comparison of BF-VI and MCMC.

8schools: hierarchical model experiment

The *8schools* dataset is a benchmark dataset for fitting a Bayesian hierarchical model and is known to be challenging for VI approaches (Yao et al. 2018; Huggins et al. 2020). It has 8 data points, corresponding to eight schools that have conducted

³ <https://github.com/stan-dev/posteriordb>.

Fig. 5 Posterior predictive distribution of the nonlinear regression model ($Y | \mathbf{x}) \sim N(\mu(\mathbf{x}), \sigma = 0.2)$ where the conditional mean is modeled by a BNN using MCMC, MF-Gaussian-VI, or BF-VI



independent coaching programs to enhance the SAT (Scholastic Assessment Test) scores of their students. There are two commonly used parameterizations of the model: centered parameterization (CP) and non-centered parameterization (NCP). NCP uses a transformed parameter to facilitate the MCMC sampling. See Table 1 for more details of these parametrizations and [37] for complete model definitions in Stan. In both parametrizations, the model has 10 parameters. In Dhaka et al. (2021), this benchmark dataset was fitted with two NF-based methods and MF-Gaussian-VI. For the CP version (Dhaka et al. 2021) $\hat{k}_{\text{CP}} = 1.3, 1.1, 0.9$ was reported for PL-NF-VI, NVP-NF-VI, MF-Gaussian-VI, respectively, which are all outperformed by our BF-VI method with $\hat{k}_{\text{CP}} = 0.53(0.11, 0.95)$. For the NCP version (Dhaka et al. 2021) $\hat{k}_{\text{NCP}} = 1.2, 0.7, 0.7$ was reported (same order), and again BF-VI yields a superior $\hat{k}_{\text{NCP}} = 0.36(0.17, 0.55)$. A visual inspection of the true MCMC posterior and its variational approximation again shows the underestimated distribution tails (see Fig. S6). For 8Schools and Diamond, a comparison with state-of-the-art models from the literature is summarized in Table 2.

NN-based nonlinear regression experiment

For this experiment, we use a small Bayesian NN for nonlinear regression fitted on 9 data points. The model for the conditional outcome distribution is $(Y | \mathbf{x}) \sim N(\mu(\mathbf{x}), \sigma = 0.2)$. The small size of the used BNN, with only one hidden layer comprising 3 neurons and one neuron in the output layer giving $\mu(\mathbf{x})$, allows us to determine the posterior via MCMC.

We then use BF-VI and MF-Gaussian-VI to fit this BNN. Because the weights in a BNN with hidden layers are not directly interpretable, they are not of direct interest, and therefore, the fit of a BNN is commonly assessed on the level of the posterior predictive distribution (see Fig. 5). In this example, the more flexible BF-VI shows a slight improvement in approximating the true posterior predictive distribution when compared to the less complex approach with MF-Gaussian-VI, especially inside the regions where there are data (around $x = 0$ in Fig. 5).

3.3 Melanoma: semi-structured NN experiment

In this experiment, we use BF-VI for semi-structured transformation models (Kook et al. 2022) (see Fig. 6), where complex data like images can be modeled by deep

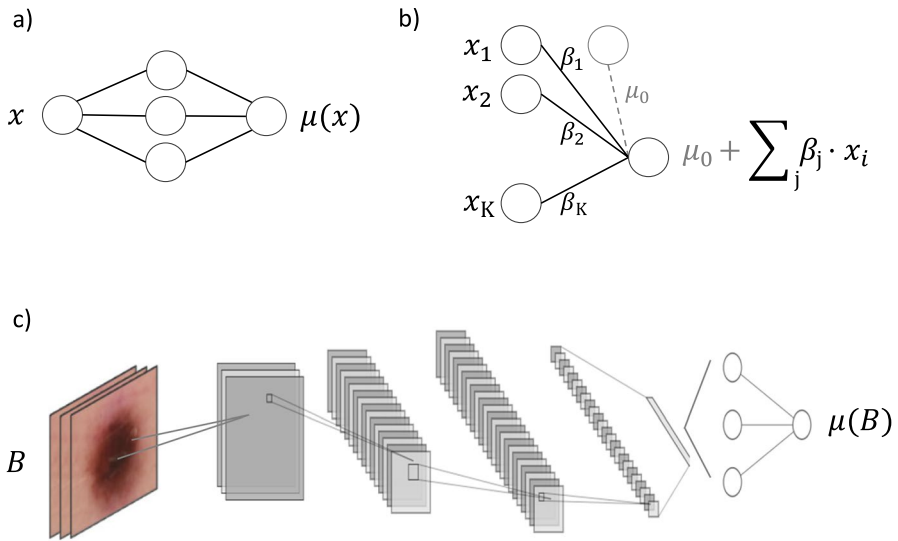


Fig. 6 The architecture of the used NN models or model parts. **a** Dense NN with one hidden layer to model nonlinear dependencies from the input (used in the NN-based nonlinear regression example). **b** Dense NN without hidden layer to model linear dependencies from tabular input data (used in M1 and M3 of the melanoma experiment). **c** CNN to model nonlinear dependencies from the image input (used in M1 and M3)


NNs and tabular data by interpretable model components. Please note that here, both the conditional distribution of the outcome ($y | B, x$) and the unconditional posterior of the parameters are modeled by transformation models. Because of the deep NN model components involved, MCMC is not feasible anymore to determine the posterior. As a dataset, we use the SIIM-ISIC Melanoma Classification Challenge⁴ data. The data come from 33,126 patients (6626 as test set, 21,200 train, and 5300 validation set) with a confirmed diagnosis of their skin lesions, which is in $\approx 98\%$ benign ($y = 0$) and in $\approx 2\%$ malignant ($y = 1$). The provided data $D = (B, x)$ are semi-structured since it comprises (unstructured) image data B from the patient's lesion along with (structured) tabular data x , i.e., the patient's age.

We fit the conditional outcome distribution $(Y | D) \sim \text{Ber}(\pi_D)$ by modeling the probability for a lesion to be malignant $\pi_D = p(y = 1 | D) = \sigma(h)$ applying the sigmoid function $\sigma(\cdot)$ to a fitted transformation function $h : Y \rightarrow Z$. We study three models for h depending on x alone, B alone, and in combination B and x :

M1 (DL-Model) $h = \mu(B)$: As a baseline, we use a deep convolutional neural network (CNN) based on the melanoma image data (see Fig. 6c) with a total of 419,489 weights to take advantage of the predictive power of DL on complex image data. For this DL model, we use deep ensembling (Lakshminarayanan et al. 2017) by fitting three CNN models with different random initializations and averaging the

⁴ <https://challenge2020.isic-archive.com>

Fig. 7 Posteriors for the age-effect parameter β_1 in the melanoma models M2 and M3



predicted probabilities. The achieved test predictive performance and its comparison to other models are discussed in the last paragraph of this section.

M2 (Logistic Regression) $h = \mu_0 + \beta_1 \cdot x$: When using only tabular features x , interpretable models can be built. We consider a Bayesian logistic regression with age as the only explanatory variable x and use a BNN without a hidden layer to set up the model (see Fig. 6b with only one input feature x). In logistic regression, a latent variable is modeled by a linear predictor $h = \mu_0 + \beta_1 \cdot x$, which determines the probability for a lesion to be malignant via $\pi_x = \sigma(\mu_0 + \beta_1 \cdot x)$ allowing to interpret e^{β_1} as the odds ratio, i.e., the factor by which the odds for lesions to be malignant changes when increasing the predictor x by one unit. In Fig. 7, we compare the exact MCMC posterior of β_1 with the BF-VI approximation, demonstrating that BF-VI accurately approximates the posterior.

M3 (semi-structured) $h = \mu(B) + \beta_1 \cdot x$: This model integrates image and tabular data and combines the predictive power of M1 with the interpretability of M2. We use a (non-Bayesian) CNN that determines $\mu(B)$ and BF-VI for the NN without a hidden layer that determines β_1 (see Fig. 6b and c). Both NNs are jointly trained by optimizing the ELBO. The resulting posterior for β_1 differs from the simple logistic regression (see Fig. 7), indicating a diminished effect of age after including the image. Again, e^{β_1} can be interpreted as the factor by which the odds for a lesion to be malignant change when increasing the predictor age by one unit and holding the image constant.

While the main interest of our study is on the posteriors, we also determine the predictive performance on the test set. To quantify and compare the test prediction performances, we look at the achieved log scores (M1: -0.076 , M2: -0.085 , M3: -0.076) and the AUCs with 95% CI (M1: $0.83(0.79, 0.86)$, M2: $0.66(0.61, 0.71)$, M3: $0.82(0.79, 0.85)$). For both measures, higher is better. Interestingly, the image-based models (M1, M3) have higher predictive power than M2, which only uses tabular data. The semi-structured model M3, including tabular and image information, has a similar predictive power compared to M1, which only uses images. The benefit of the semi-structured model here is that it provides interpretable parameters for the tabular data along with uncertainty quantification without losing predictive performance.

4 Summary and outlook

The proposed BF-VI is flexible enough to approximate any posterior in principle without being restricted to variational distributions with known parametric distribution families like Gaussians. In benchmark experiments, BF-VI accurately fits non-trivial posteriors in low-dimensional problems in a BBVI setting. For higher-dimensional models, BF-VI outperforms published results from other NF-VI methods (Dhaka et al. 2021) on the studied benchmark datasets. Still, we observe that the posterior cannot be fitted perfectly in high dimensions by BF-VI, especially since the tails of the approximation are too short. We attribute this limitation to known difficulties in the optimization process and the asymmetry of the KL divergence. These challenges of VI were not in the focus of our study, and we leave it to further research.

To the best of our knowledge, we are the first to demonstrate how BBVI can be used in semi-structured models. We used BF-VI on the public melanoma challenge dataset, integrating image data and tabular data by combining a deep CNN and an interpretable model part. We see a valuable application of BF-VI in models with interpretable parameters, i.e., statistical or semi-structured models where we can model complex posterior distributions of the interpretable parameters. Especially in semi-structured models with deep NN components that cannot be fitted with MCMC, BF-VI allows determining the variational distribution for the interpretable model parts. Moreover, efficient SGD optimizers can be used in BF-VI to fit all model parts jointly. We plan to extend our research on BF-VI for semi-structured models in the future and investigate the quality of the posterior approximations.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10182-024-00497-z>.

Acknowledgements The research of BS was supported by the Novartis Research Foundation (FreeNovation 2019). The research of the OD and DD was partly supported by the Federal Ministry of Education and Research of Germany (BMBF) in the project DeepDoubt (grant no. 01IS19083A). We, further, would like to thank Nadja Klein, Lucas Kook, and Rebekka Axthelm for fruitful discussions.

Funding Open Access funding enabled and organized by Projekt DEAL. This study was funded by Bundesministerium für Bildung, Wissenschaft, Forschung und Technologie (01IS19083A), Novartis Stiftung für Medizinisch-Biologische Forschung.

Declarations

Conflict of interest None.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line

to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Agrawal, A., Sheldon, D.R., Domke, J.: Advances in black-box VI: normalizing flows, importance weighting, and optimization. *Adv. Neural Inf. Process. Syst.* **33**, 17358–17369 (2020)
- Baumann, P.F.M., Hothorn, T., Rügamer, D.: Deep conditional transformation models. In: *Machine Learning and Knowledge Discovery in Databases. Research Track*, pp. 3–18. Springer, Cham (2021)
- Bernštejn, S.: Démonstration du théorème de Weierstrass fondée sur le calcul des probabilités. *Commun. Soc. Math. Kharkov* **13**, 1–2 (1912)
- Bingham, E., Chen, J.P., Jankowiak, M., Obermeyer, F., Pradhan, N., Karaletsos, T., Singh, R., Szerlip, P.A., Horsfall, P., Goodman, N.D.: Pyro: deep universal probabilistic programming. *J. Mach. Learn. Res.* **20**, 28–1286 (2019)
- Blei, D., Ranganath, R., Mohamed, S.: Variational inference: foundations and modern methods. In: *NIPS tutorial* (2016)
- Blei, D.M., Kucukelbir, A., McAuliffe, J.D.: Variational inference: a review for statisticians. *J. Am. Stat. Assoc.* **112**(518), 859–877 (2017)
- Blundell, C., Cornebise, J., Kavukcuoglu, K., Wierstra, D.: Weight uncertainty in neural network. In: *International Conference on Machine Learning*, pp. 1613–1622. PMLR (2015)
- Blundell, C., Cornebise, J., Kavukcuoglu, K., Wierstra, D.: Weight uncertainty in neural network. In: *International Conference on Machine Learning*, pp. 1613–1622. PMLR (2015)
- Bogachev, V.I., Kolesnikov, A.V., Medvedev, K.V.: Triangular transformations of measures. *Sbornik Math.* **196**(3), 309 (2005)
- Buri, M., Curt, A., Steeves, J., Hothorn, T.: Baseline-adjusted proportional odds models for the quantification of treatment effects in trials with ordinal sum score outcomes. *BMC Med. Res. Methodol.* **20**, 1–14 (2020)
- Campanella, G., Kook, L., Häggström, I., Hothorn, T., Fuchs, T.J.: Deep conditional transformation models for survival analysis (2022). arXiv preprint [arXiv:2210.11366](https://arxiv.org/abs/2210.11366)
- Carpenter, B., Gelman, A., Hoffman, M.D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., Riddell, A.: Stan: a probabilistic programming language. *J. Stat. Softw.* (2017). <https://doi.org/10.18637/jss.v076.i01>
- Dhaka, A.K., Catalina, A., Andersen, M.R., Magnusson, M., Huggins, J., Vehtari, A.: Robust, accurate stochastic optimization for variational inference. *Adv. Neural Inf. Process. Syst.* **33**, 10961–10973 (2020)
- Dhaka, A.K., Catalina, A., Welandawe, M., Andersen, M.R., Huggins, J., Vehtari, A.: Challenges and opportunities in high dimensional variational inference. *Adv. Neural Inf. Process. Syst.* **34**, 7787–7798 (2021)
- Dinh, L., Sohl-Dickstein, J., Bengio, S.: Density estimation using real NVP. In: *International Conference on Learning Representations* (2016)
- Durkan, C., Bekasov, A., Murray, I., Papamakarios, G.: Neural spline flows. *Adv. Neural Inf. Process. Syst.* **32**, 7511–7522 (2019)
- Farouki, R.T.: The Bernstein polynomial basis: a centennial retrospective. *Comput. Aided Geom. Des.* **29**(6), 379–419 (2012)
- Hothorn, T., Kneib, T., Bühlmann, P.: Conditional transformation models. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **76**(1), 3–27 (2014)
- Hothorn, T., Moest, L., Buehlmann, P.: Most likely transformations. *Scand. J. Stat.* **45**(1), 110–134 (2018)
- Huggins, J., Kasprzak, M., Campbell, T., Broderick, T.: Validated variational inference via practical posterior error bounds. In: *International Conference on Artificial Intelligence and Statistics*, pp. 1792–1802. PMLR (2020)

- Jaini, P., Selby, K.A., Yu, Y.: Sum-of-squares polynomial flow. In: International Conference on Machine Learning, pp. 3009–3018. PMLR (2019)
- Jordan, M.I., Ghahramani, Z., Jaakkola, T.S., Saul, L.K.: An introduction to variational methods for graphical models. *Mach. Learn.* **37**(2), 183–233 (1999)
- Klein, N., Hothorn, T., Barbanti, L., Kneib, T.: Multivariate conditional transformation models. *Scand. J. Stat.* **49**(1), 116–142 (2019)
- Kook, L., Herzog, L., Hothorn, T., Dürr, O., Sick, B.: Deep and interpretable regression models for ordinal outcomes. *Pattern Recognit.* **122**, 108263 (2022)
- Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. In: Advances in Neural Information Processing Systems, vol. 30 (2017)
- Lohse, T., Rohrmann, S., Faeh, D., Hothorn, T.: Continuous outcome logistic regression for analyzing body mass index distributions. *F1000Research* **6**, 1933 (2017)
- Louizos, C., Welling, M.: Multiplicative normalizing flows for variational Bayesian neural networks. In: International Conference on Machine Learning, pp. 2218–2227. PMLR (2017)
- Magnusson, M., Bürkner, P., Vehtari, A.: PosteriorDB: a set of posteriors for Bayesian inference and probabilistic programming
- Papamakarios, G., Pavlakou, T., Murray, I.: Masked autoregressive flow for density estimation. In: Proceedings of the 31st International Conference on Neural Information Processing Systems, pp. 2335–2344 (2017)
- Ramasinghe, S., Fernando, K., Khan, S., Barnes, N.: Robust normalizing flows using Bernstein-type polynomials (2021). arXiv preprint [arXiv:2102.03509](https://arxiv.org/abs/2102.03509)
- Ranganath, R., Gerrish, S., Blei, D.: Black box variational inference. In: Artificial Intelligence and Statistics, pp. 814–822. PMLR, Cambridge (2014)
- Rezende, D., Mohamed, S.: Variational inference with normalizing flows. In: International Conference on Machine Learning, pp. 1530–1538. PMLR (2015)
- Rügamer, D., Baumann, P.F.M., Kneib, T., Hothorn, T.: Probabilistic time series forecasts with autoregressive transformation models. arXiv preprint (2021)
- Sick, B., Hothorn, T., Dürr, O.: Deep transformation models: tackling complex regression problems with neural network based transformation models. In: 2020 25th International Conference on Pattern Recognition (ICPR), pp. 2476–2481. IEEE (2021)
- Siegfried, S., Hothorn, T.: Count transformation models. *Methods Ecol. Evol.* **11**(7), 818–827 (2020). <https://doi.org/10.1111/2041-210X.13383>
- Van Den Berg, R., Hasenclever, L., Tomczak, J.M., Welling, M.: Sylvester normalizing flows for variational inference. In: 34th Conference on Uncertainty in Artificial Intelligence 2018, UAI 2018, pp. 393–402. Association For Uncertainty in Artificial Intelligence (AUAI) (2018)
- Welandawe, M., Andersen, M.R., Vehtari, A., Huggins, J.H.: Robust, automated, and accurate black-box variational inference (2022). arXiv preprint [arXiv:2203.15945](https://arxiv.org/abs/2203.15945)
- Yao, Y., Vehtari, A., Simpson, D., Gelman, A.: Yes, but did it work?: Evaluating variational inference. In: International Conference on Machine Learning, pp. 5581–5590. PMLR (2018)
- Yao, Y., Vehtari, A., Gelman, A.: Stacking for non-mixing Bayesian computations: the curse and blessing of multimodal posteriors. *J. Mach. Learn. Res.* **23**, 1–45 (2022)

Authors and Affiliations

Oliver Dürr¹  · Stefan Hörtling¹ · Danil Dold¹ · Ivonne Kovylov¹ · Beate Sick^{2,3}

✉ Oliver Dürr
oliver.duerr@htwg-konstanz.de

✉ Beate Sick
beate.sick@uzh.ch; sick@zhaw.ch

Stefan Hörtling
Stefanhoertling@web.de

Danil Dold
danil.dold@htwg-konstanz.de

Ivonne Kovylov
ivonne97ko@gmail.com

¹ IOS, Konstanz University of Applied Sciences, Alfred Wachtel Straße 8, 78462 Konstanz, Germany

² IDP, Zurich University of Applied Sciences, Technikumstrasse 81, 8401 Winterthur, Switzerland

³ EBPI, University of Zurich, Hirschengraben 84, 8001 Zurich, Switzerland