



Density estimation via Bayesian inference engines

M. P. Wand¹ · J. C. F. Yu¹

Received: 11 March 2021 / Accepted: 24 September 2021 / Published online: 1 November 2021
© Springer-Verlag GmbH Germany, part of Springer Nature 2021

Abstract

We explain how effective automatic probability density function estimates can be constructed using contemporary Bayesian inference engines such as those based on no-U-turn sampling and expectation propagation. Extensive simulation studies demonstrate that the proposed density estimates have excellent comparative performance and scale well to very large sample sizes due to a binning strategy. Moreover, the approach is fully Bayesian and all estimates are accompanied by point-wise credible intervals. An accompanying package in the R language facilitates easy use of the new density estimates.

Keywords Expectation propagation · Mixed model-based penalized splines · No-U-turn sampler · Semiparametric mean field variational Bayes · Slice sampling

1 Introduction

Bayesian inference engines have become established as an important paradigm for inference in arbitrarily large and complex graphical models. Software platforms such as Infer.NET (Minka 2018) and Stan (Carpenter et al. 2017) are instances of such Bayesian inference engines. They deliver approximate Bayesian inference, with varying degrees of inferential accuracy, by calling upon contemporary approaches such as expectation propagation, Hamiltonian Monte Carlo and variational approximation. The purpose of this short article is to show that effective and scalable probability density function estimation, or density estimation for short, can be achieved using Bayesian inference engines. We provide easy access for users of the R statistical computing environment (R Core Team 2018) via a package named `densEst-Bayes` (Wand 2021).

Even though density estimators such as the histogram have had a presence in statistics and data analysis for most of its history, automatic density estimation started as a major area of research in the early 1980s when computing power

✉ J. C. F. Yu
jamescfyu@gmail.com

¹ University of Technology Sydney, Sydney, Australia

aided its feasibility. Practical methodology, usually involving kernel density estimation with a data-driven bandwidth choice, such as Rudemo (1982), Bowman (1984) and Sheather and Jones (1991) was accompanied by deep theoretical analysis such as Hall and Marron (1987). Several other proposals ensued, many of which are summarized in Chapter 3 of Wand and Jones (1995). A more recent proposal of this general type is due to Botev et al. (2010), in which kernel density estimation is combined with diffusion theory to yield an advanced plug-in type bandwidth selector. A simulation study given there demonstrates superior practical performance compared with earlier proposals.

In a separate literature, starting mainly in the early 1990s, practical methodology for inference in Bayesian graphical models emerged as a major area of activity. The most prominent approach is Markov chain Monte Carlo which aims to produce samples from the posterior density functions of hidden nodes (parameters and latent variables) in a graphical model. By the mid-1990s the BUGS Bayesian inference engine (e.g., Lunn et al. 2009) had emerged and, for the first time, data analysts could perform approximate Bayesian inference for an arbitrarily complicated Bayesian graphical model by doing little more than specifying the model and inputting the data. The last 25 years has seen various refinements of this paradigm. An interesting review of the state-of-affairs in the mid-2000s is provided by Murphy (2007). Since that time two new major Bayesian inference engines have emerged: *Infer.NET* (Minka 2018) and *Stan* (Carpenter et al. 2017). The former of these is distinguished by the fact that its main approaches to approximate Bayesian inference are deterministic, rather than based on Monte Carlo sampling, with expectation propagation (e.g., Minka 2001) and mean field variational Bayes (e.g., Wainwright and Jordan 2008) being the underlying principles called upon. The *Stan* Bayesian inference engine uses Hamiltonian Monte Carlo and a variant known as the no-U-turn sampler (Hoffman and Gelman 2014) to obtain samples from the posterior density functions of hidden nodes. In an area with close ties to density estimation: nonparametric regression and various extensions, Luts et al. (2018) and Harezlak et al. (2018) provide several illustrations of approximate Bayesian inference via *Infer.NET* and *Stan* respectively.

The more fundamental problem of automatic probability density function estimation via Bayesian inference engines is the focus here. The crux of our approach is to express the density estimation problem as a Poisson nonparametric regression problem. This involves replacement of the original data by bin counts on a fine equally spaced grid (Eilers and Marx 1996) as detailed in Sect. 2.1. Bayesian Poisson nonparametric regression using mixed model representations of low-rank smoothing splines (e.g., Ruppert et al. 2009) can be expressed as a Bayesian graphical model and is easy to feed into a Bayesian inference engine. The conversion of the density estimation problem to that of fitting a Poisson nonparametric regression model also has the advantage of scaling well to massive sample sizes since the only cost for large sample sizes is the binning step. Once the input data have been binned the remaining operations are unaffected by sample size. In Sect. 3, we report the results of a simulation study that demonstrates Bayesian inference engine density estimation to be very accurate in comparison with existing methods.

Various Bayesian approaches to density estimation have been proposed over the past few decades and articles on the topic number in the dozens. Broad themes include use of continuous-time stochastic processes (e.g., Leonard 1978; Lenk 1988) and nonparametric Bayes discrete-time stochastic process structures (e.g., Escobar and West 1995; Petrone 1999). Unlike the Poisson nonparametric regression/low-rank smoothing spline model considered here, see (1) in Sect. 2, the majority of these approaches are not amenable to immediate implementation in a Bayesian inference engine.

A Bayesian inference engine estimator at any particular abscissa has a corresponding variability measure – nominally in the form of a 95% credible interval. This entails the option of adding a variability band around the plotted density estimate which has the advantage of providing a visualization of the sample variability. Interpretation of variability bands requires caution since they are based on *point-wise* credible intervals. The problem of obtaining *simultaneous* credible interval bands, in the spirit of Sun and Loader (1994), is not explored here. In Sect. 3.3, we demonstrate that the empirical coverages of the credible intervals produced by Bayesian inference engine density estimation is somewhat conservative but usually meets advertized coverage levels. The density estimation literature contains numerous proposals for the construction of confidence intervals (e.g., Hall and Titterton 1988; Chen 1996; Giné and Nickl 2010) for frequentist inference concerning density function values. The methodology is generally of a high technical level, with delicate asymptotic arguments and practical implementation hindered by *ad hoc* smoothing parameter choice. In contrast, Bayesian inference engine density estimation provides credible intervals in a simple and natural way.

Full details of our approach and some examples are presented in Sect. 2. In Sect. 3, we report the results of simulation studies concerned with evaluation of our proposal. We describe the `densEstBayes` R package in Sect. 4. Conclusions are drawn in Sect. 5.

2 Approach

We start by describing our Bayesian inference engine approach to density estimation in generic form. There are various choices to be made such as the actual Bayesian inference engine to use and auxiliary parameters such as the number of spline basis functions and Bayesian model hyperparameters. These choices are discussed in Sects. 2.2 and 2.3. A key feature of the approach is conversion of the density estimation problem to a Poisson nonparametric regression problem. Its justification is given in Sect. 2.1.

The following notation is needed to describe the approach. A random variable v has an Inverse Gamma distribution with shape parameter κ and scale parameter λ if and only if its density function is

$$p(v) = 1/\{\lambda^\kappa \Gamma(\kappa)\} v^{-\kappa-1} \exp(-v/\lambda), \quad v > 0,$$

and $p(v) = 0$ for $v \leq 0$. The statement $v_i \stackrel{\text{ind.}}{\sim} \mathcal{D}_i, 1 \leq i \leq n$, means each of the random variables v_i are independent with distribution \mathcal{D}_i .

Given a univariate random sample x_1, \dots, x_n , the generic approach to obtaining an estimate of the sample's probability density function is:

1. Linearly transform the $x_i, 1 \leq i \leq n$, to the unit interval.
2. Replace the $x_i, 1 \leq i \leq n$, by bin counts on a fine equally-spaced grid of size M over the unit interval. Let $(g_\ell, c_\ell), 1 \leq \ell \leq M$, denote the grid point/grid count pairs. The choice of M is discussed in Sect. 2.3.
3. Fit the Bayesian mixed model-based penalized spline model:

$$\begin{aligned}
 c_\ell | \beta_0, \beta_1, u_1, \dots, u_K &\stackrel{\text{ind.}}{\sim} \text{Poisson} \left\{ \exp \left(\beta_0 + \beta_1 g_\ell + \sum_{k=1}^K u_k z_k(g_\ell) \right) \right\}, \\
 \beta_0, \beta_1 &\stackrel{\text{ind.}}{\sim} N(0, \sigma_\beta^2), \quad u_1, \dots, u_K | \sigma^2 \stackrel{\text{ind.}}{\sim} N(0, \sigma^2), \\
 \sigma^2 | a &\sim \text{Inverse-Gamma} \left(\frac{1}{2}, 1/a \right), \quad a \sim \text{Inverse-Gamma} \left(\frac{1}{2}, 1/s_\sigma^2 \right)
 \end{aligned} \tag{1}$$

via some Bayesian inference engine. Its choice is discussed in Sect. 2.2. Choice of the spline basis $\{z_k(\cdot) : 1 \leq k \leq K\}$ and hyperparameters $\sigma_\beta, \sigma_\sigma > 0$ is covered in Sect. 2.3.

4. For any $x \in [0, 1]$, the density estimate of the transformed data is

$$\hat{p}(x) = C^{-1} \left\{ \text{posterior mean of } \exp \left(\beta_0 + \beta_1 x + \sum_{k=1}^K u_k z_k(x) \right) \right\} \tag{2}$$

and C is chosen to ensure that $\int_0^1 \hat{p}(x') dx' = 1$. Point-wise credible intervals to accompany the $\hat{p}(x), 0 \leq x \leq 1$, are readily available from Step 3. Details are given in Sect. 2.4.

5. Linearly transform the density estimate and corresponding credible intervals to the original data units.

Note that

$$\sigma^2 | a \sim \text{Inverse-Gamma} \left(\frac{1}{2}, 1/a \right), \quad a \sim \text{Inverse-Gamma} \left(\frac{1}{2}, 1/s_\sigma^2 \right)$$

is equivalent to the standard deviation parameter σ having a Half Cauchy prior density function with scale parameter s_σ :

$$p(\sigma) = 2/[\pi s_\sigma \{1 + (\sigma/s_\sigma)^2\}], \quad \sigma > 0.$$

The use of the auxiliary variable a in (1) aids the construction of approximate Bayesian inference schemes such as those discussed in Sects. 2.2.1 and 2.2.3.

The density estimate produced by steps 1.-5. takes the form of an exponentiated cubic spline, where the coefficients are subject to a roughness penalty. The essence of this general approach goes back, at least, to Boneva et al. (1971). Several articles,

such as Wahba (1975) and Good and Gaskins (1980), have built on this general paradigm. The class of density estimates presented in this section is of the same ilk, but uses low-rank smoothing splines and takes advantage of the Bayesian inference engine revolution.

2.1 Justification for use of Poisson nonparametric regression

Conversion of the density estimation problem to a Poisson nonparametric regression via binning over a fine grid is a relatively old trick, and is explained and used in Section 8 of Eilers and Marx (1996) for a version of penalized spline-based nonparametric regression. The justification hinges upon an equivalence between Poisson and Multinomial maximum likelihood estimators as explained in Section 13.4.4 of Bishop et al. (2007).

2.2 Choice of Bayesian inference engine

Potentially, the only difficult step in Bayesian Poisson nonparametric density estimation is fitting the Bayesian model (1). Established Bayesian inference engines such BUGS, Infer.NET, JAGS (Plummer 2003) and Stan essentially remove this difficulty. At the time of this writing the main costs are computing time and the occasional need for chain diagnostics. Refinements of these packages and improved future Bayesian inference engines will continue to make the fitting of (1) faster and more routine. Another option for fitting (1) is self-implementation of one of the very many approximate Bayesian inference schemes in the literature. For example, the “stepping out” slice sampling strategy of Neal (2003) has a particularly simple implementation and, if programmed in a low-level language, can be reasonably fast compared with the general purpose Markov chain Monte Carlo schemes used by established Bayesian inference engines.

In our exploration and demonstration of the efficacy of density estimation via Bayesian inference engines, we settled on four approaches, and these feature in the numerical evaluations given in Sect. 3. We now provide some details on each of these four approaches.

2.2.1 Expectation propagation

Expectation propagation is a class of deterministic approximations of the joint posterior density function of the parameters in a graphical model, based on notions such as Bethe free energy and expectation constraints. A theoretical framework for expectation propagation is provided by Heskes et al. (2005), where the problem is expressed in Lagrangian form. However, the underlying optimization problem is challenging due to its non-convex saddle point nature. The most common strategy for obtaining practical solutions is to use iterative *message passing* on an appropriate *factor graph* with so-called *damping* adjustments. Expectation propagation message passing updates correspond to Kullback–Leibler projections onto particular exponential family density functions

and, intuitively, are iterative moment-matching operations. Kim and Wand (2018) provide the algorithmic details of expectation propagation for various generalized, linear and mixed models.

Model (1) is a special case of the models treated in Kim and Wand (2018). Let $\beta \equiv (\beta_0, \beta_1)$ and $\mathbf{u} \equiv (u_1, \dots, u_K)$ be the coefficient vectors and $\mathbf{c} \equiv (c_1, \dots, c_M)$ be the vector of bin counts. The factor graph for expectation propagation approximation of the posterior density functions

$$p(\beta, \mathbf{u} | \mathbf{c}) \quad \text{and} \quad p(\sigma^2 | \mathbf{c}),$$

is formed by noting the following algebraic truism:

$$p(\mathbf{c}, \beta, \mathbf{u}, \sigma^2, a) = p(a) p(\sigma^2 | a) \left\{ \prod_{k=1}^K \int_{-\infty}^{\infty} p(\tilde{u}_k | \sigma^2) \delta\left(\tilde{u}_k - \mathbf{e}_{k+2}^T \begin{bmatrix} \beta \\ \mathbf{u} \end{bmatrix}\right) d\tilde{u}_k \right\} \\ \times \left\{ \int_{\mathbb{R}^2} p(\tilde{\beta}) \delta\left(\tilde{\beta} - \mathbf{E}_2^T \begin{bmatrix} \beta \\ \mathbf{u} \end{bmatrix}\right) d\tilde{\beta} \right\} \left\{ \prod_{\ell=1}^M \int_{-\infty}^{\infty} p(c_\ell | \alpha_\ell) \delta\left(\alpha_\ell - \mathbf{h}_\ell^T \begin{bmatrix} \beta \\ \mathbf{u} \end{bmatrix}\right) d\alpha_\ell \right\}. \tag{3}$$

Here, δ denotes the univariate Dirac delta function, δ denotes the bivariate Dirac delta function, \mathbf{e}_j is the $(2 + K) \times 1$ vector with 1 in the j th entry and all other entries equal to zero, \mathbf{E}_2 is the $(2 + K) \times 2$ matrix with \mathbf{I}_2 in first two rows and all other entries equal to zero and $\mathbf{h}_\ell \equiv (1, g_\ell, z_1(g_\ell), \dots, z_K(g_\ell))$. Figure 1 is a factor graph representation of (3). The factors are shown as solid rectangles and stochastic variables as circles, with edges joining stochastic variables to the factors that include them. The integral signs in (3) are ignored and Kim and Wand (2018) use the phrase *derived variable* factor graph to make this distinction from regular factor graphs, with the $\alpha_\ell = \mathbf{h}_\ell^T [\beta^T \ \mathbf{u}^T]^T$ being examples of derived variables.

Bayesian density estimation via expectation propagation proceeds by updating messages passed between each of the neighboring nodes on Fig. 1 and iteration until convergence. Full details are in Kim and Wand (2018). The message updates required evaluation of versions of the following non-analytic integral functions:

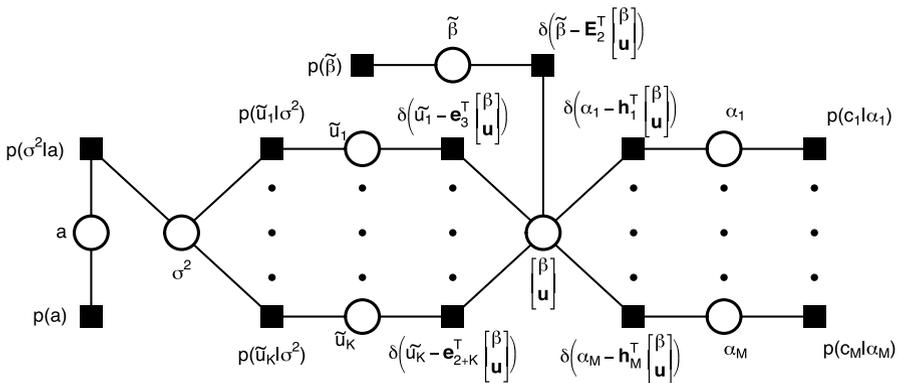


Fig. 1 Derived variable factor graph corresponding to the representation of $p(\mathbf{c}, \beta, \mathbf{u}, \sigma^2, a)$ given by (3)

$$\begin{aligned}
 \mathcal{A}(p, q, r, s, t, u) &= \int_{-\infty}^{\infty} \frac{x^p \exp(qx - rx^2) dx}{(x^2 + sx + t)^u}, \\
 \mathcal{B}(p, q, r, s, t, u) &= \int_{-\infty}^{\infty} \frac{x^p \exp\{qx - re^x - se^x/(t + e^x)\} dx}{(t + e^x)^u} \\
 \text{and } \mathcal{C}(p, q, r) &= \int_{-\infty}^{\infty} x^p \exp(qx - rx^2 - e^x) dx
 \end{aligned} \tag{4}$$

with various restrictions on the parameters as detailed in Section 2.1 of Kim and Wand (2018). Inversion of the function $\log - \text{digamma}$, where $\text{digamma}(x) \equiv \frac{d}{dx} \log\{\Gamma(x)\}$, is also required. All other calculations are algebraic.

2.2.2 No U-turn sampling

No U-turn sampling, due to Hoffman and Gelman (2014), is a Markov chain Monte Carlo scheme that fine tunes Hamiltonian Monte Carlo sampling. Hamiltonian Monte Carlo, also known as hybrid Monte Carlo, dates back to the mid-1980s statistical physics literature. Its naming, and introduction to the mainstream statistics literature, is due to Neal (2011). As explained there, Hamiltonian Monte Carlo is a version of Metropolis–Hastings-based Markov chain Monte Carlo. The traditional approach to producing Metropolis–Hastings proposal distributions involves random walks. As explained in Neal (2011), the posterior distribution space can be explored more efficiently when Metropolis–Hastings proposal distributions are instead produced using the principles of Hamiltonian dynamics. Full details are provided by this landmark article.

The essence of no U-turn sampling is adaptive choice of parameters that are inherent to Hamiltonian Monte Carlo, such as step size and number of steps, via the introduction of slice variables. No U-turn sampling is the default and preferred algorithm in the Stan Bayesian inference engine for obtaining samples from the posterior distributions of hidden nodes in a graphical model. In recent years no U-turn sampling has established itself as a durable and high-quality Markov chain Monte Carlo scheme. Almost, all of the Bayesian semiparametric regression examples in Harezlak et al. (2018) use no U-turn sampling. Its availability within the R package `rstan` means that (1) can be embedded within the R computing environment via just a few lines of code.

2.2.3 Semiparametric mean field variational Bayes

Mean field variational Bayes aims to achieve approximate Bayesian inference for (1) via a product density restriction approximations such as

$$p(\boldsymbol{\beta}, \mathbf{u}, \sigma^2, a | \mathbf{c}) \approx q(\boldsymbol{\beta}, \mathbf{u})q(\sigma^2)q(a). \tag{5}$$

and choosing the q -density functions to minimize the Kullback–Leibler divergence of the right-hand side of (5) from the left-hand side. However, the form of the optimal q -density of the coefficients does not admit a closed form. A practical remedy is the pre-specification

$q(\boldsymbol{\beta}, \mathbf{u})$ is a $N(\boldsymbol{\mu}_{q(\boldsymbol{\beta}, \mathbf{u})}, \boldsymbol{\Sigma}_{q(\boldsymbol{\beta}, \mathbf{u})})$ density function

for some mean vector $\boldsymbol{\mu}_{q(\boldsymbol{\beta}, \mathbf{u})}$ and covariance matrix $\boldsymbol{\Sigma}_{q(\boldsymbol{\beta}, \mathbf{u})}$, followed by Kullback–Leibler minimization subject to this restriction. This augmentation of mean field variational Bayes has various names such as *fixed-form variational Bayes* and *non-conjugate variational message passing*. Rohde and Wand (2016) make a case for the term *semiparametric mean field variational Bayes*, and we use that label here. Model (1) is a special case of the Poisson additive mixed models treated in Section 3.1 of Luts and Wand (2015). The Poisson response and $r = 1$ special case of Algorithm 1 in Luts and Wand (2015) leads to fast approximate Bayesian density estimation.

2.2.4 Slice sampling

For random variables $s_1 \in \mathbb{R}$, $s_2 > 0$, as well as random vectors \mathbf{s}_3 and \mathbf{s}_4 , of the same dimension, let

$$x|s_1, s_2, \mathbf{s}_3, \mathbf{s}_4 \sim \mathcal{H}(s_1, s_2, \mathbf{s}_3, \mathbf{s}_4)$$

denote that the random variable x , conditional on $(s_1, s_2, \mathbf{s}_3, \mathbf{s}_4)$ has density function

$$p(x|s_1, s_2, \mathbf{s}_3, \mathbf{s}_4) \propto \exp \left\{ s_1 x - x^2 / (2s_2) - \mathbf{1}^T \exp(x\mathbf{s}_3 + \mathbf{s}_4) \right\}, \quad -\infty < x < \infty, \tag{6}$$

where $\mathbf{1}$ denotes a vector of ones having the same number of rows as \mathbf{s}_3 and \mathbf{s}_4 and $\exp(x\mathbf{s}_3 + \mathbf{s}_4)$ is evaluated element-wise. Then scalar Gibbs sampling for (1) is such that draws are required from either density functions of the form (6) or Inverse Gamma density functions. The latter is trivial and the former is relatively easy if one uses the “stepping out” slice sampling approach of Neal (2003). Let $\mathbf{C} \equiv [\mathbf{X} \ \mathbf{Z}]$ and define \mathbf{C}_j to be the j th column of \mathbf{C} and \mathbf{C}_{-j} to be the matrix \mathbf{C} with its j th column removed. Similarly, define

$$\begin{bmatrix} \boldsymbol{\beta} \\ \mathbf{u} \end{bmatrix}_{-j} \quad \text{to be} \quad \begin{bmatrix} \boldsymbol{\beta} \\ \mathbf{u} \end{bmatrix} \quad \text{with its } j\text{th entry removed.}$$

If G denotes the total number of samples, including the warm-up, then a suitable slice sampling within Gibbs sampling scheme is (after e.g., setting an initial value for $(\sigma^2)^{[0]}$):

For $g = 1, \dots, G$:

$$\mathbf{v} \leftarrow [\sigma_{\boldsymbol{\beta}}^2 \mathbf{1}_2^T, (\sigma^2)^{[g-1]} \mathbf{1}_K^T]^T$$

For $j = 1, \dots, 2 + K$:

$$\begin{bmatrix} \boldsymbol{\beta} \\ \mathbf{u} \end{bmatrix}_j^{[g]} \sim \mathcal{H} \left((\mathbf{C}^T \mathbf{c})_j, \mathbf{v}_j, \mathbf{C}_j, \mathbf{C}_{-j} \begin{bmatrix} \boldsymbol{\beta} \\ \mathbf{u} \end{bmatrix}_{-j}^{[g-1]} \right)$$

$$a^{[g]} \sim \text{Inverse-Gamma} \left(1, 1/(\sigma^2)^{[g-1]} + 1/s_{\sigma}^2 \right),$$

$$(\sigma^2)^{[g]} \sim \text{Inverse-Gamma} \left(\frac{1}{2}(K + 1), \frac{1}{2} \|\mathbf{u}^{[g]}\|^2 + 1/a^{[g]} \right).$$

where $\mathbf{c} = (c_1, \dots, c_M)$ is the vector of bin counts.

After omission of the warm-up samples, the R retained samples

$$\begin{bmatrix} \beta \\ \mathbf{u} \end{bmatrix}^{[g]}, \quad (\sigma^2)^{[g]}, \quad 1 \leq g \leq R,$$

can be used for construction of a Bayesian density estimate and point-wise credible intervals. In the Sect. 2.6 examples and Sect. 3 simulation studies we used a warm-up of size 100 and $R = 1000$ retained samples.

2.3 Choice of auxiliary parameters

Full specification of the Bayesian inference engine-based density estimator requires choice of the basis functions, hyperparameters and binning grid size. For most density functions that arise in applications the choices of these auxiliary parameters have very little effect on the estimate. We provide good default settings here. If the density function has intricate features and the sample size is very large to the extent that these features can be estimated reasonably then some adjustment to these defaults may be required.

For the spline basis functions $\{z_k(\cdot) : 1 \leq k \leq K\}$, we use cubic canonical O'Sullivan splines as described in Section 4 of Wand and Ormerod (2008). The default number of basis functions is $K = 50$. The number of grid points used for binning is defaulted to $M = 401$ which is in keeping with Table 1 of Hall and Wand (1996). Linear binning (e.g., Hall and Wand 1996) is used, followed by rounding to the nearest integer, to get the bin counts c_ℓ for use in the Poisson nonparametric regression model. The default hyperparameter values are $\sigma_\beta = s_\sigma = 1000$, assuming the transformation of the input data to the unit interval has taken place, corresponding to approximate noninformativity.

For the Monte Carlo-based approaches we ran a pilot study to test the effect of the warm-up and retained sample sizes on density estimation accuracy. For the no-U-turn sampler, we found that a warm-up of length 1000 with 1000 retained samples was adequate without significant degradation of accuracy. For the slice sampling a warm-up of length 100, followed by 1000 retained samples, was found to be adequate. For expectation propagation and semiparametric mean field variational Bayes, the default stopping criterion is the relative change in $E_q(1/\sigma^2)$ falling below 10^{-5} .

2.4 Pointwise credible interval construction

Point-wise credible intervals are a simple by-product of the Bayesian inference engine output. Suppose that $x_0 \in [0, 1]$ is a typical abscissae of interest. In the case of no-U-turn sampling the samples from the posterior density functions of the β_j and u_k can be used to form a sample corresponding to the $\hat{p}(x_0)$ according to the form

given by (2). For $0 < \alpha < 1$, an approximate $100(1 - \alpha)\%$ credible interval for $p(x_0)$ has upper and lower limits corresponding to the $\alpha/2$ and $(2 - \alpha)/2$ sample quantiles of this sample.

For expectation propagation and semiparametric mean field variational Bayes, we instead have $\mu_{q(\beta,u)}$ and $\Sigma_{q(\beta,u)}$ as Bayesian inference engine outputs. If we let

$$\mathcal{L}(x_0) \equiv [1 \ x_0 \ z_1(x_0) \ \dots \ z_K(x_0)]^T$$

be the vector of basis function evaluations at x_0 then, with Φ denoting the $N(0, 1)$ cumulative distribution function,

$$\mathcal{L}(x_0)^T \mu_{q(\beta,u)} \pm \Phi^{-1}((2 - \alpha)/2) \sqrt{\mathcal{L}(x_0)^T \Sigma_{q(\beta,u)} \mathcal{L}(x_0)}$$

is an approximate $100(1 - \alpha)\%$ credible interval for the linear form

$$\beta_0 x_0 + \beta_1 x_0 + \sum_{k=1}^K u_k z_k(x_0).$$

Simple manipulations then lead to an approximate $100(1 - \alpha)\%$ credible interval for the $p(x_0)$.

2.5 Pre-processing options

If the input data are strongly skewed or contain gross outliers then some pre-processing may be worthwhile. The fourth example of the upcoming Fig. 2 applies a logarithmic transform to the input data. Bayesian inference engine density estimation is applied to these transformed data. The estimate is back-transformed for graphical display.

2.6 Examples

Figure 2 provides four examples of Bayesian inference density estimation for the following univariate data sets:

- ages in years at first inauguration of the 29 presidents of the United States of America who have held office during 1900–2021;
- maximum daily temperature in degrees Fahrenheit in Melbourne, Australia, for the 101 days that followed a very hot day, defined to be 95 degrees Fahrenheit or higher, during 1981–1990;
- time intervals in minutes between all 3,507 adjacent pairs of eruptions of the Old Faithful Geyser in Yellowstone National Park, U.S.A, during 2011, obtained from The Geyser Observation and Study Association web-site (www.geyserstudy.org);
- incomes of 7,201 United Kingdom citizens for the year 1975, divided by average income. The source of these data is the Economic and Social Research Council Data Archive at the University of Essex, United Kingdom.

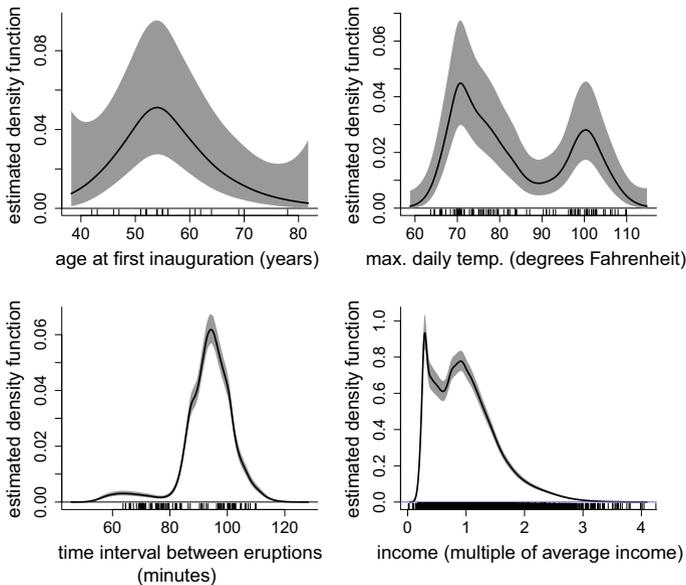


Fig. 2 Examples of Bayesian inference engine density estimation, with semiparametric mean field variational Bayes used in each case. The tick marks at the base of each plot show the data and shaded regions correspond to point-wise 95% credible intervals. Top left: the data are the ages (years) at first inauguration of the 28 U.S. presidents whom have held office during 1900–2021. Top right: the data are maximum daily temperature (degrees Fahrenheit) in Melbourne, Australia, for the 101 days that followed a very hot day, defined to be 95 degrees Fahrenheit or higher, during 1981–1990. Bottom left: the data are time intervals (minutes) between all 3,507 adjacent pairs of eruptions of the Old Faithful Geyser during the 2011. Bottom right: the data are incomes of 7,201 United Kingdom citizens for the year 1975. The data have been divided by average income

Since the data sets increase in size from the top left panel to the bottom right panel the 95% credible intervals become narrower. The last three density estimates having interesting bimodal structure. For the maximum daily temperatures in Melbourne the bimodality is explained by the southerly buster phenomenon, which often produces a dramatic temperature drop after a very hot day in southern Australia.

3 Evaluation

The new density estimation strategies described in Sect. 2 add to a large field of existing automatic density estimators. We now investigate how they compare in terms of accuracy and computing time.

3.1 Density estimation accuracy

We ran a large simulation study involving 3 sample sizes, 10 true density functions and 6 automatic density function estimators. The sample sizes are

$n \in \{100, 1000, 10000\}$, the true density functions are density numbers 1–10 in Table 1 of Marron and Wand (1992). The density estimation methods are (in order of development):

kernel density estimation with bandwidth chosen according to least squares cross-validation (Rudemo 1982; Bowman 1984),
 kernel density estimation with bandwidth chosen according a direct plug-in strategy as described in Section 3.6.1 of Wand and Jones (1995),
 the diffusion kernel density estimator of Botev et al. (2010),
 and the four types of Bayesian inference engine density estimators described in Sect. 2—involving each of expectation propagation, no-U-turn sampling, semiparametric mean field variational Bayes and slice sampling.

Estimation accuracy of a generic density estimate \hat{p} was measured using the accuracy score

$$\text{accuracy}(\hat{p}) = 100 \left(1 - \frac{1}{2} \int_{-\infty}^{\infty} |\hat{p}(x) - p_{\text{true}}(x)| dx \right)$$

which uses the fact that the L_1 error $\int_{-\infty}^{\infty} |\hat{p}(x) - p_{\text{true}}(x)| dx$ is a scale-free number between 0 and 2 and linearly transforms this error measure to an accuracy percentage. The simulation study was run over 1000 replications. For each pair of methods, the accuracy paired difference samples were analyzed using visual inspections of side-by-side box plots and Wilcoxon confidence intervals. The main findings were as follows:

- The four Bayesian inference engine approaches were such that there was very little practical differences between them in terms of accuracy. Further investigations have revealed that the choice of Bayesian inference engine has a negligible effect on the density estimate in terms of visual appearance.
- The diffusion kernel density estimator usually dominated the other kernel density estimates in terms of accuracy. For some settings such as Marron–Wand density number 5 there were pronounced practical improvements of the diffusion kernel density estimator compared with ordinary kernel density estimation with least squares cross-validation bandwidth choice.
- In 29 out of the 30 settings both the no-U-turn-based and slice sampling-based density estimator exhibited a statistically significant better accuracy than the diffusion kernel density estimator. Although for most of these settings the practical improvement was negligible, there were some practical improvements in a few settings. Figure 3 and its accompanying discussion describes these improvements.

Figure 3 shows some of the practical advantages of Bayesian inference engine density estimation over a state-of-the art approach. The upper panels are for sample sizes of $n = 100$ and estimation of the third Marron–Wand density function, which is strongly skewed. The scatterplot in the upper left panel of Fig. 3 shows that,

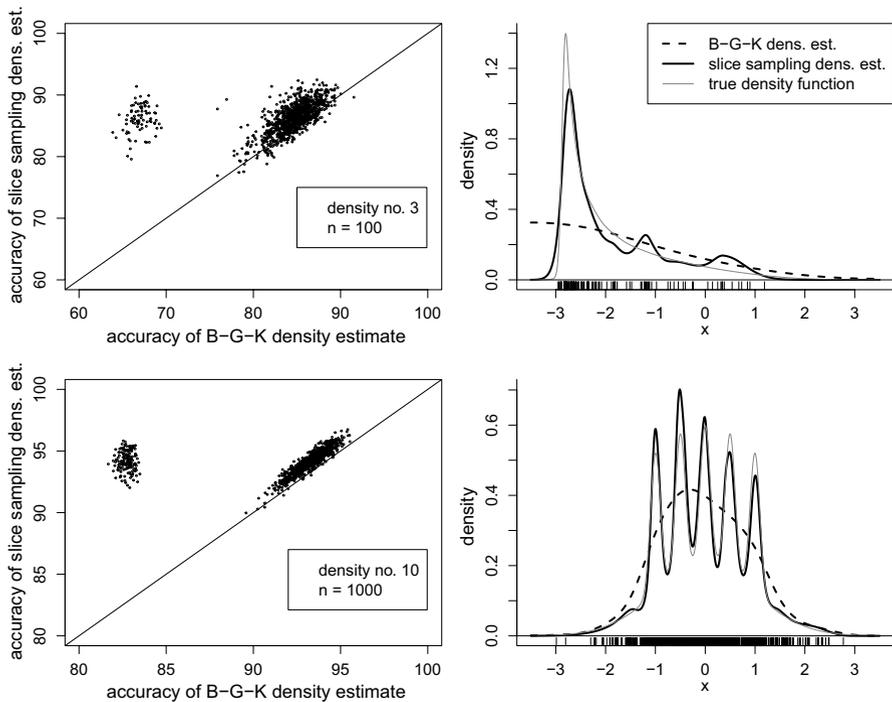


Fig. 3 Simulation study results that demonstrate improvements of Bayesian inference engine density estimation over an existing approach. Top row: Estimation of Marron–Wand density number 3 with a sample size of $n = 100$. The top left panel scatterplot shows accuracy values of estimates based on slice sampling versus those based on the Botev–Grotowski–Kroese approach. The 1:1 line is also plotted. The upper left cluster represents a significant practical improvement of the Bayesian inference engine approach and contains 8.1% of the replications. The top right shows example density estimates corresponding to this cluster. Bottom row: Analogous to the bottom row but for Marron–Wand density number 10 with $n = 1000$. The upper left cluster contains 14.9% of the replications

whilst most of the time the two approaches have very similar accuracies, about 8% of the replication are such that the diffusion kernel density estimate suffers from a pronounced drop in accuracy — corresponding to the cluster above the 1:1 line. The top right panel shows typical estimates from this cluster, with the diffusion kernel density estimate over-smoothing. The lower panels of Fig. 3 tell a similar story for samples of size $n = 1000$ and estimation of the tenth Marron–Wand density function, which is claw-shaped.

3.2 Computing time

We kept track of the computing times in the simulation study described in the previous subsection. The default auxiliary parameter values described in Sect. 2.3 were used. The simulations were run on a MacBook Air laptop computer with a 2.2 gigahertz processor and 8 gigabytes of random access memory. Of course, the computing

Table 1 Percentiles for the computing times in seconds for each of the Bayesian inference engine approaches used in the simulation study described in Sect. 3.1

	10th percentile	50th percentile	90th percentile
Expect. propagation	19	31	78
No-U-turn sampling	6.9	12	23
Semipar. MFVB	0.076	0.21	1.2
Slice sampling	13	13	15

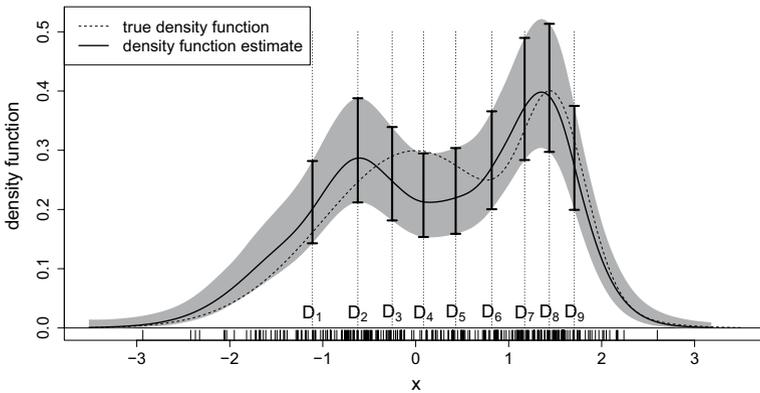


Fig. 4 Density function given at (7) and the semiparametric mean field variational Bayes estimate based on a random sample of size $n = 250$. The vertical line segments and notches indicate 95% credible intervals for $p_{\text{true}}(D_j)$ at each of the deciles $D_j, 1 \leq j \leq 9$

times are impacted by choices such as warm-up length and hardware specifications. Nonetheless, the results presented here give an idea of computing time using typical early 2020s personal hardware, as well as comparative performance.

Table 1 provides the 10th, 50th and 90th quantiles of the computing times in seconds for each approach. Semiparametric mean field variational Bayes is the fastest by far and usually returns an estimate in less than half a second. However, as explained in Sect. 3.4, this speed has to be counterbalanced against occasional convergence failure problems. Despite our implementation in a low-level language, expectation propagation can be quite slow due to the large number of numerical integrations that it requires. The no-U-turn and slice sampling approaches typically take about 10 to 15 s. The computing time for slice sampling is less variable, with a 90th percentile of 15 s compared with 24 s for no-U-turn sampling.

3.3 Bayesian inferential accuracy

In a second simulation study we investigated the degree to which the point-wise credible sets produced by the proposed Bayesian density estimates meet their advertized coverage levels. Figure 4 shows the semiparametric mean field

variational Bayes-based density estimate from on a sample size of $n = 1000$ generated from the eighth density function in Table 1 of Marron and Wand (1992):

$$p_{\text{true}}(x) = 3 \left[\exp(-x^2/2) + \exp\left\{-9\left(x - \frac{3}{2}\right)^2/2\right\} \right] / (4\sqrt{2\pi}). \tag{7}$$

We focused on inference for each of $p_{\text{true}}(D_1), \dots, p_{\text{true}}(D_9)$, where D_1, \dots, D_9 are the population deciles of p_{true} . Figure 4 shows the locations of the D_j along with 95% credible intervals for each $p_{\text{true}}(D_j)$ using an $n = 250$ Bayesian density estimate via the semiparametric mean field variational Bayes approach. In this example, eight of the nine credible intervals cover the true density function value. The exception is $p_{\text{true}}(D_4)$, which is not quite covered by its 95% credible interval.

To assess coverage accuracy, for sample sizes $n = 100$, $n = 1000$ and $n = 10,000$, we generated 1000 random samples and obtained density estimates using each of the three Bayesian inference engines described in Sect. 2.2. Table 2 shows the empirical coverage percentages.

We see from Table 2 that, in almost every case, the empirical coverage level exceeds the 95% advertized coverage level. This indicates that inference based on the proposed Bayesian inference engine density estimators is honest in that it delivers on what it promises. However, with an average empirical coverage of 97.6% it is apparent that the estimators over-deliver compared with their 95% advertizement. The conclusion from this limited study concerning Bayesian accuracy of the proposed Bayesian inference engine density estimators is that they are honest although they err on the side of conservatism.

Table 2 Empirical coverage percentages for the Bayesian inference problem conveyed by Fig. 4 involving 95% credible intervals at each of the deciles of the density function given at (7)

		D_1	D_2	D_3	D_4	D_5	D_6	D_7	D_8	D_9
$n = 100$	EP	98.4	98.7	98.8	98.8	98.0	93.2	98.3	95.8	92.6
	NUTS	97.7	98.4	99.0	99.2	97.5	93.9	98.8	92.7	95.5
	SMFVB	97.5	98.0	99.0	99.2	97.6	90.6	98.6	90.2	95.5
	slice	97.6	98.3	99.0	99.5	97.5	93.9	98.7	92.2	95.1
$n = 1000$	EP	98.3	98.8	98.9	98.9	98.7	96.1	98.9	98.4	96.9
	NUTS	98.4	98.3	98.7	98.8	98.9	95.7	98.7	96.5	97.3
	SMFVB	98.5	98.3	98.5	98.9	98.9	94.7	99.0	96.7	97.7
	slice	98.5	98.6	98.8	98.6	98.9	95.3	98.8	97.1	97.8
$n = 10,000$	EP	98.3	98.6	98.4	98.8	98.2	97.5	98.3	98.1	98.4
	NUTS	97.1	98.4	98.0	98.5	97.7	97.9	98.5	98.8	97.9
	SMFVB	97.3	98.3	97.9	98.9	97.7	97.7	98.6	98.9	98.2
	slice	97.2	98.4	97.7	98.6	98.0	97.7	98.5	98.6	98.1

The methods are no U-turn sampling (NUTS), expectation propagation (EP), semiparametric mean field variational Bayes (SMFVB) and slice sampling (slice). The number of replications was 1000

3.4 Numerical issues

The semiparametric mean field variational Bayes approach involves fixed point iteration to find the minimum Kullback–Leibler divergence $\mu_{q(\beta,u)}$ and $\Sigma_{q(\beta,u)}$ parameters. For the default of 50 spline basis functions there are 1430 free parameters in this search. Often convergence is successful and rapid. However, despite efforts to obtain good starting values, the semiparametric mean field variational Bayes approach failed to converge for 13.6% of the 30,000 data sets in the Sect. 3.1 simulation study. Therefore, further numerical analytic research is required to make semiparametric mean field variational Bayes more practical. Expectation propagation converged in almost all 30,000 data sets but sometimes could be quite slow taking as long as several minutes on a 2020s laptop. Despite being a deterministic alternative to Monte Carlo sampling, the integrals that arise when expectation propagation is applied to model (1) require quadrature and are also quite numerous. This leads to expectation propagation often being considerably slower than the Monte Carlo-based approaches. Further research is warranted for speeding up expectation propagation to acceptable levels.

4 Accompanying R package

An R package that accompanies this article is available on the Comprehensive R Archive Network (<https://www.R-project.org>) and named `densEstBayes` (Wand 2021). Once installed, the following few commands illustrate its default use:

```
library(densEstBayes) ; x <- rnorm(1000)
dest <- densEstBayes(x) ; plot(dest) ; rug(x)
```

The other arguments of the `densEstBayes()` function, named `method` and `control`, respectively, allow for different Bayesian inference engines to be specified and auxiliary parameters to be controlled. In the version of `densEstBayes` that is available at the time of this writing the Bayesian inference engines are Hamiltonian Monte Carlo, no-U-turn sampling, semiparametric mean field variational Bayes and slice sampling. The last of these is the default method due to it achieving a good balance in terms of accuracy performance, numerical reliability and speed.

The `densEstBayes` package is accompanied by a vignette which provides fuller details on its use. The vignette PDF file is opened via the command `densEstBayesVignette()`.

5 Conclusions

We have demonstrated that Bayesian inference engines based on the Poisson non-parametric regression and mixed model-based penalized splines offer a competitive class of density estimators, and sometimes lead to noticeable improvements in accuracy compared with state-of-the-art approaches. Moreover, Bayesian inference

engines are accompanied by principled variability bands which enhance graphical display. Our recommended default Bayesian inference engine takes about 10–15 s to compute on early 2020s laptop computers, which is a reasonable price to pay given its attractive attributes. Future Bayesian inference engines and hardware enhancements offer the prospect of further improvement.

Acknowledgements We are grateful for assistance from Eman Alfaifi. This research was supported by Australian Research Council Discovery Project DP140100441.

References

- Bishop, Y.M.M., Fienberg, S.E., Holland, P.W.: *Discrete Multivariate Analysis: Theory and Practice*. Springer, New York (2007)
- Boneva, L.I., Kendall, D., Stefanov, I.: Spline transformations: three new diagnostic aids for the statistical data-analyst (with discussion). *J. Roy. Stat. Soc. B* **33**, 1–71 (1971)
- Botev, Z.I., Grotowski, J.F., Kroese, D.P.: Kernel density estimation via diffusion. *Ann. Stat.* **38**, 2916–2957 (2010)
- Bowman, A.W.: An alternative method of cross-validation for the smoothing of density estimates. *Biometrika* **71**, 353–360 (1984)
- 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.* **76**(1), 1–32 (2017)
- Chen, S.X.: Empirical likelihood confidence intervals for nonparametric density estimation. *Biometrika* **83**, 329–341 (1996)
- Escobar, M.D., West, M.: Bayesian density estimation and inference using mixtures. *J. Am. Stat. Assoc.* **90**, 577–588 (1995)
- Eilers, P.H.C., Marx, B.D.: Flexible smoothing with B-splines and penalties (with discussion). *Stat. Sci.* **11**, 89–121 (1996)
- Giné, E., Nickl, R.: Confidence bands in density estimation. *Ann. Stat.* **38**, 1122–1170 (2010)
- Good, I.J., Gaskins, R.A.: Density estimation and bump-hunting by the penalized likelihood method exemplified by scattering and meteorite data (with discussion). *J. Am. Stat. Assoc.* **75**, 42–73 (1980)
- Hall, P., Marron, J.S.: Extent to which least-squares cross-validation minimised integrated square error in nonparametric density estimation. *Probab. Theory Relat. Fields* **74**, 567–581 (1987)
- Hall, P., Titterton, D.M.: On confidence bands in nonparametric density estimation and regression. *J. Multivar. Anal.* **27**, 228–254 (1988)
- Hall, P., Wand, M.P.: On the accuracy of binned kernel density estimators. *J. Multivar. Anal.* **56**, 165–184 (1996)
- Harezlak, J., Ruppert, D., Wand, M.P.: *Semiparametric Regression with R*. Springer, New York (2018)
- Heskes, T., Opper, M., Wiegerinck, W., Winther, O., Zoeter, O.: Approximate inference techniques with expectation constraints. *J. Stat. Mech: Theory Exp.* **11**, 1–24 (2005)
- Hoffman, M.D., Gelman, A.: The no-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* **15**, 1593–1623 (2014)
- Kim, A.S.I., Wand, M.P.: On expectation propagation for generalised, linear and mixed models. *Aust. N. Z. J. Stat.* **60**, 75–102 (2018)
- Lenk, P.J.: The logistic normal distribution for Bayesian, nonparametric, predictive densities. *J. Am. Stat. Assoc.* **83**, 509–516 (1988)
- Leonard, T.: Density estimation, stochastic processes, and prior information. *J. R. Stat. Soc. B* **40**, 113–146 (1978)
- Lunn, D., Spiegelhalter, D., Thomas, A., Best, N.: The BUGS project: evolution, critique and future directions. *Stat. Med.* **28**, 3049–3067 (2009)
- Luts, J., Wand, M.P.: Variational inference for count response semiparametric regression. *Bayesian Anal.* **10**, 991–1023 (2015)
- Luts, J., Wang, S.S.J., Ormerod, J.T., Wand, M.P.: Semiparametric regression analysis via Infer.NET. *J. Stat. Softw.* **87**(2), 1–37 (2018)
- Marron, J.S., Wand, M.P.: Exact mean integrated squared error. *Ann. Stat.* **20**, 712–736 (1992)

- Minka, T.P.: Expectation propagation for approximate Bayesian inference. In: Breese, J.S., Koller, D. (eds) Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence, pp. 362–369. Burlington, Massachusetts: Morgan Kaufmann (2001)
- Minka, T., Winn, J., Guiver, J., Zayov, Y., Fabian, D., Bronskill: Infer.NET 0.3, Microsoft Research Cambridge. <http://dotnet.github.io/infer> (2018)
- Murphy, K.: Software for graphical models: a review. *Int. Soc. Bayesian Anal. Bull.* **14**, 13–15 (2007)
- Neal, R.: Slice sampling (with discussion). *Ann. Stat.* **31**, 705–767 (2003)
- Neal, R.: MCMC using Hamiltonian dynamics. In: Brooks, S., Gelman, A., Jones, G.L., Meng, X.-L. (eds.) *Handbook of Markov Chain Monte Carlo*, pp. 113–162. CRC Press, Boca Raton, Florida (2011)
- Petrone, S.: Bayesian density estimation using Bernstein polynomials. *Can. J. Stat.* **27**, 105–120 (1999)
- Plummer, M.: JAGS: a program for analysis of Bayesian graphical models using Gibbs sampling. In: Hornik, K., Leisch, F., Zeileis, A. (eds) *Proceedings of the 3rd International Workshop on Distributed Statistical Computing* (2003)
- R Core Team: R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org> (2018)
- Rohde, D., Wand, M.P.: Semiparametric mean field variational Bayes: general principles and numerical issues. *J. Mach. Learn. Res.* **17**(172), 1–47 (2016)
- Rudemo, M.: Empirical choice of histograms and kernel density estimators. *Scand. J. Stat.* **9**, 65–78 (1982)
- Ruppert, D., Wand, M.P., Carroll, R.J.: Semiparametric regression during 2003–2007. *Electr. J. Stat.* **3**, 1193–1256 (2009)
- Sheather, S.J., Jones, M.C.: A reliable data-based bandwidth selection method for kernel density estimation. *J. Roy. Stat. Soc. B* **53**, 683–690 (1991)
- Sun, J., Loader, C.R.: Simultaneous confidence bands for linear regression and smoothing. *Ann. Stat.* **22**, 1328–1345 (1994)
- Wainwright, M.J., Jordan, M.I.: Graphical models, exponential families, and variational inference. *Found. Trends Mach. Learn.* **1**, 1–305 (2008)
- Wahba, G.: Interpolating spline methods for density estimation I. Equi-spaced knots. *Ann. Stat.* **3**, 30–48 (1975)
- Wand, M.P., Jones, M.C.: *Kernel Smoothing*. Chapman and Hall, London (1995)
- Wand, M.P., Ormerod, J.T.: On semiparametric regression with O’Sullivan penalized splines. *Aust. N. Z. J. Stat.* **50**, 179–198 (2008)
- Wand, M.P.: **densEstBayes**: Density estimation via Bayesian inference algorithms. R package version 1.0. <http://cran.r-project.org> (2021)

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.