University of Wollongong

Research Online

Centre for Statistical & Survey Methodology
Working Paper Series

Faculty of Engineering and Information
Sciences

2010

Variational Bayesian Inference for Parametric and Nonparametric Regression with Missing Data

C. Faes

University of Wollongong

J. T. Ormerod

University of Wollongong

M. Wand

University of Wollongong, mwand@uow.edu.au

Follow this and additional works at: <https://ro.uow.edu.au/cssmwp>

Recommended Citation

Faes, C.; Ormerod, J. T.; and Wand, M., Variational Bayesian Inference for Parametric and Nonparametric Regression with Missing Data, Centre for Statistical and Survey Methodology, University of Wollongong, Working Paper 07-10, 2010, 30p.
<https://ro.uow.edu.au/cssmwp/58>

Research Online is the open access institutional repository for the University of Wollongong. For further information contact the UOW Library: research-pubs@uow.edu.au

Centre for Statistical and Survey Methodology

The University of Wollongong

Working Paper

07-10

**Variational Bayesian Inference for Parametric and
Nonparametric Regression with Missing Data**

BY C. FAES¹, J.T. ORMEROD² & M.P. WAND³

¹ Interuniversity Institute for Biostatistics and Statistical Bioinformatics, Hasselt University, BE3590 Diepenbeek, BELGIUM

² School of Mathematics and Statistics, University of Sydney, Sydney 2006, AUSTRALIA

³ Centre for Statistical and Survey Methodology, School of Mathematics and Applied Statistics, University of Wollongong, Wollongong 2522, AUSTRALIA

Copyright © 2008 by the Centre for Statistical & Survey Methodology, UOW. Work in progress, no part of this paper may be reproduced without permission from the Centre.

Centre for Statistical & Survey Methodology, University of Wollongong, Wollongong NSW 2522. Phone +61 2 4221 5435, Fax +61 2 4221 4845. Email: anica@uow.edu.au

Variational Bayesian Inference for Parametric and Nonparametric Regression with Missing Data

BY C. FAES¹, J.T. ORMEROD² & M.P. WAND³

¹ *Interuniversity Institute for Biostatistics and Statistical Bioinformatics, Hasselt University, BE3590 Diepenbeek, BELGIUM*

² *School of Mathematics and Statistics, University of Sydney, Sydney 2006, AUSTRALIA*

³ *Centre for Statistical and Survey Methodology, School of Mathematics and Applied Statistics, University of Wollongong, Wollongong 2522, AUSTRALIA*

12th May, 2010

SUMMARY

Bayesian hierarchical models are attractive structures for conducting regression analyses when the data are subject to missingness. However, the requisite probability calculus is challenging and Monte Carlo methods typically are employed. We develop an alternative approach based on deterministic variational Bayes approximations. Both parametric and nonparametric regression are treated. We demonstrate that variational Bayes can achieve good accuracy, but with considerably less computational overhead. The main ramification is fast approximate Bayesian inference in parametric and nonparametric regression models with missing data.

Keywords: Bayesian inference; Directed acyclic graphs; Incomplete data; Mean field approximation; Penalized splines; Variational approximation.

1 Introduction

Bayesian inference for parametric regression has a long history (e.g. Box & Tiao, 1973; Gelman, Carlin, Stern & Rubin, 2004). Mixed model representations of smoothing splines and penalized splines afford Bayesian inference for nonparametric regression (e.g. Ruppert, Wand & Carroll, 2003). Whilst this notion goes back at least to Wahba (1978), recent developments in Bayesian inference methodology, especially Markov Chain Monte Carlo (MCMC) algorithms and software, has led to Bayesian approaches to nonparametric regression becoming routine. See, for example, Crainiceanu, Ruppert & Wand (2005) and Gurrin, Scurrah & Hazelton (2005). There is also a large literature on Bayesian nonparametric regression using regression splines with a variable selection approach (e.g. Denison, Holmes, Mallick & Smith, 2002). The present article deals only with penalized spline nonparametric regression, where hierarchical Bayesian models for nonparametric regression are relatively simple.

When the data are susceptible to missingness a Bayesian approach allows relatively straightforward incorporation of standard missing data models (e.g. Little & Rubin, 2004; Daniels & Hogan, 2008), resulting in a larger hierarchical Bayesian model. Inference via MCMC is simple in principle, but can be costly in processing time. For example, on the third author's laptop computer (Mac OS X; 2.33 GHz processor, 3 GBytes of RAM), obtaining 10000 MCMC samples for a 25-knot penalized spline model, and sample size of 500, takes about 2.6 minutes via the R language (R Development Core Team, 2010) package *BRugs* (Ligges *et al.* 2010). If 30% of the predictor data are reset to be missing completely at random and the appropriate missing data adjustment is made to the model then

10000 MCMC draws takes about 7.3 minutes; representing an approximate three-fold increase. The situation worsens for more complicated nonparametric and semiparametric regression models. MCMC-based inference, via **BRugs**, for the missing data/bivariate smoothing example in Section 7 of Wand (2009) requires about a week on the aforementioned laptop.

This paper is concerned with fast Bayesian parametric and nonparametric regression analysis in situations where some of the data are missing. Speed is achieved by using variational approximate Bayesian inference, often shortened to *variational Bayes*. This is a deterministic approach that yields approximate inference, rather than ‘exact’ inference produced by an MCMC approach. However, as we shall see, the approximations can be very good. An accuracy assessment, described in Section 3.4, showed that variational Bayes achieves better than 80% accuracy for the main model parameters.

Variational Bayes is now part of mainstream Computer Science methodology (e.g. Bishop, 2006) and are used in problems such as speech recognition, document retrieval (e.g. Jordan, 2004) and functional magnetic resonance imaging (e.g. Flandin & Penny, 2007). Recently, they have seen use in statistical problems such as cluster analysis for gene-expression data (Teschendorff *et al.*, 2005) and finite mixture models (McGrory & Titterton, 2007). Ormerod & Wand (2010) contains an exposition on variational Bayes from a statistical perspective. A pertinent feature is their heavy algebraic nature. Even relatively simple models require significant notation and algebra for description of variational Bayes.

To the best of our knowledge, the present article is the first to develop and investigate variational Bayes for regression analysis with missing data. In principle, variational Bayes methods can be used in essentially all missing data regression contexts: e.g. generalized linear models, mixed models, generalized additive models, geostatistical models and their various combinations. It is prudent, however, to start with simpler regression models where the core tenets can be elucidated without excessive notation and algebra. Hence, the present paper treats the simplest parametric and nonparametric regression models: single predictor with homoscedastic Gaussian errors. The full array of missing data scenarios: missing completely at random (MCAR), missing at random (MAR) and missing not at random (MNAR) are treated. For parametric regression with probit missing data mechanisms, we show that variational Bayes is purely algebraic, without the need for quadrature or Monte Carlo-based approximate integration. The nonparametric regression extension enjoys many of the assets of parametric regression, but requires some univariate quadrature. Comparisons with MCMC show quite good accuracy, but with computation in the order of seconds rather than minutes. The upshot is fast approximate Bayesian inference in parametric and nonparametric regression models with missing data.

Section 2 summarises the variational Bayes approach. Inference in the simple linear regression model with missing data is the focus of Section 3. In Section 4 we describe extension to nonparametric regression. Some closing discussion is given in Section 5.

1.1 Notation

If $\mathcal{P}$ is a logical condition then $I(\mathcal{P}) = 1$ if $\mathcal{P}$ is true and $I(\mathcal{P}) = 0$ if $\mathcal{P}$ is false. We use Φ to denote the standard normal distribution function.

Column vectors with entries consisting of sub-scripted variables are denoted by a bold-faced version of the letter for that variable. Round brackets will be used to denote the entries of column vectors. For example $\mathbf{x} = (x_1, \dots, x_n)$ denotes a $n \times 1$ vector with entries $x_1, \dots, x_n$. The element-wise product of two matrices $\mathbf{A}$ and $\mathbf{B}$ is denoted by $\mathbf{A} \odot \mathbf{B}$. We use $\mathbf{1}_d$ to denote the $d \times 1$ column vector with all entries equal to 1. The norm of a column vector $\mathbf{v}$, defined to be $\sqrt{\mathbf{v}^T \mathbf{v}}$, is denoted by $\|\mathbf{v}\|$. Scalar functions applied to

vectors are evaluated element-wise. For example,

$$\Phi(a_1, a_2, a_3) \equiv (\Phi(a_1), \Phi(a_2), \Phi(a_3)).$$

The density function of a random vector $\mathbf{u}$ is denoted by $p(\mathbf{u})$. The conditional density of $\mathbf{u}$ given $\mathbf{v}$ is denoted by $p(\mathbf{u}|\mathbf{v})$. The covariance matrix of $\mathbf{u}$ is denoted by $\text{Cov}(\mathbf{u})$. A $d \times 1$ random vector $\mathbf{x}$ has a Multivariate Normal distribution with parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, denoted by $\mathbf{x} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, if its density function is

$$p(\mathbf{x}) = (2\pi)^{-d/2} |\boldsymbol{\Sigma}|^{-1/2} \exp\{-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\}.$$

A random variable x has an Inverse Gamma distribution with parameters $A, B > 0$, denoted by $x \sim \text{IG}(A, B)$ if its density function is $p(x) = B^A \Gamma(A)^{-1} x^{-A-1} e^{-B/x}$, $x > 0$. If y_i has distribution D_i for each $1 \leq i \leq n$, and the y_i are independent, then we write $y_i \stackrel{\text{ind.}}{\sim} D_i$.

2 Elements of Variational Bayes

Variational Bayes methods are a family of approximate inference techniques based on the notions of minimum Kullback-Leibler divergence and product assumptions on the posterior densities of the model parameters. They are known as *mean field* approximations in the statistical physics literature (e.g. Parisi, 1988). Detailed expositions on variational Bayes may be found in Bishop (2006, Sections 10.1–10.4) and Ormerod & Wand (2010). In this section we describe the elements of variational Bayes.

Consider a generic Bayesian model with parameter vector $\boldsymbol{\theta} \in \Theta$ and observed data vector $\mathbf{y}$. Bayesian inference is based on the posterior density function

$$p(\boldsymbol{\theta}|\mathbf{y}) = \frac{p(\mathbf{y}, \boldsymbol{\theta})}{p(\mathbf{y})}.$$

We will suppose that $\mathbf{y}$ and $\boldsymbol{\theta}$ are continuous random vectors, which conforms with the models in Sections 3 and 4. Let q be an arbitrary density function over Θ . Then the marginal likelihood $p(\mathbf{y})$ satisfies $p(\mathbf{y}) \geq \underline{p}(\mathbf{y}; q)$ where

$$\underline{p}(\mathbf{y}; q) \equiv \exp \int_{\Theta} q(\boldsymbol{\theta}) \log \left\{ \frac{p(\mathbf{y}, \boldsymbol{\theta})}{q(\boldsymbol{\theta})} \right\} d\boldsymbol{\theta}.$$

The gap between $\log\{p(\mathbf{y})\}$ and $\log\{\underline{p}(\mathbf{y}; q)\}$ is known as the *Kullback-Leibler divergence* and is minimized by

$$q_{\text{exact}}(\boldsymbol{\theta}) = p(\boldsymbol{\theta}|\mathbf{y}),$$

the exact posterior density function. However, for most models of practical interest, $q_{\text{exact}}(\boldsymbol{\theta})$ is intractable and restrictions need to be placed on q to achieve tractability. Variational Bayes relies on product density restrictions:

$$q(\boldsymbol{\theta}) = \prod_{i=1}^M q_i(\boldsymbol{\theta}_i) \text{ for some partition } \{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M\} \text{ of } \boldsymbol{\theta}.$$

Under this restriction, the optimal densities (with respect to minimum Kullback-Leibler divergence) can be shown to satisfy

$$q_i^*(\boldsymbol{\theta}_i) \propto \exp\{E_{-\boldsymbol{\theta}_i} \log p(\mathbf{y}, \boldsymbol{\theta})\}, \quad 1 \leq i \leq M, \quad (1)$$

where $E_{-\boldsymbol{\theta}_i}$ denotes expectation with respect to the density $\prod_{j \neq i} q_j(\boldsymbol{\theta}_j)$. Conditions (1) give rise to iterative schemes for obtaining the simultaneous solutions over each member of the partition.

An alternative form for the solutions is

$$q_i^*(\theta_i) \propto \exp\{E_{-\theta_i} \log p(\theta_i|\text{rest})\}, \quad 1 \leq i \leq M, \quad (2)$$

where

$$\text{rest} \equiv \{\mathbf{y}, \theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_M\}$$

is the set containing the rest of the random vectors in the model, apart from θ_i . The distributions $\theta_i|\text{rest}$, $1 \leq i \leq M$, are known as the *full conditionals* in the MCMC literature. Gibbs sampling (e.g. Robert & Casella, 2004) involves successive draws from these full conditionals. We prefer (2) to (1), since it lends itself to considerable simplification because of graph theoretic results which we describe next.

2.1 Directed Acyclic Graphs and Markov Blanket Theory

The missing data regression models of Sections 3 and 4 are hierarchical Bayesian models, and hence can be represented as probabilistic directed acyclic graphs (DAGs). DAGs provide a useful ‘road map’ of the models structure, and aid the algebra required for variational Bayes. Random variables or vectors correspond to *nodes* while *directed edges* (i.e. *arrows*) convey conditional dependence. The observed data components of the DAG are sometimes called *evidence* nodes, whilst the model parameters correspond to *hidden* nodes. Bishop (2006, Chapter 8) and Wasserman (2004, Chapter 17) provide very good summaries of DAGs and their probabilistic properties. Figures 1 and 6 contain DAGs for models considered in the present paper.

The formulation of variational Bayes algorithms greatly benefit from a DAG-related known concept known as *Markov blanket* theory. First we define the Markov blanket of a node on a DAG:

Definition. *The Markov blanket of a node on a DAG is the set of children, parents and co-parents of that node. Two nodes are co-parents if they have at least one child node in common.*

Markov blankets are important in the formulation of variational Bayes algorithms because of:

Theorem (Pearl, 1988). *For each node on a probabilistic DAG, the conditional distribution of the node given the rest of the nodes is the same as the conditional distribution of the node given its Markov blanket.*

For our generic Bayesian example, this means that

$$p(\theta_i|\text{rest}) = p(\theta_i|\text{Markov blanket of } \theta_i).$$

It immediately follows that

$$q_i^*(\theta_i) \propto \exp\{E_{-\theta_i} \log p(\theta_i|\text{Markov blanket of } \theta_i)\}, \quad 1 \leq i \leq M. \quad (3)$$

For large DAGs, such as those in Figure 6, (3) yields considerable algebraic economy. In particular, it shows that the $q_i^*(\theta_i)$ require only local calculations on the model’s DAG.

3 Simple Linear Regression with Missing Predictor Data

In this section we confine attention to the simple linear regression model with homoscedastic Gaussian errors. For complete data on the predictor/response pairs (x_i, y_i) , $1 \leq i \leq n$, this model is

$$y_i = \beta_0 + \beta_1 x_i + \varepsilon_i, \quad \varepsilon_i \stackrel{\text{ind.}}{\sim} N(0, \sigma_\varepsilon^2).$$

We couch this in a Bayesian framework by taking

$$\beta_0, \beta_1 \stackrel{\text{ind.}}{\sim} N(0, \sigma_\beta^2) \quad \text{and} \quad \sigma_\varepsilon^2 \sim \text{IG}(A_\varepsilon, B_\varepsilon)$$

for hyperparameters $\sigma_\beta^2, A_\varepsilon, B_\varepsilon > 0$. Use of these conjugate priors simplifies the variational Bayes algebra. Other priors, such as those described by Gelman (2006), may be used. However, they result in more complicated variational Bayes algorithms.

Now suppose that the predictors are susceptible to missingness. Bayesian inference then requires a probabilistic model for the x_i s. We will suppose that

$$x_i \sim N(\mu_x, \sigma_x^2) \tag{4}$$

and take $\mu_x \sim N(0, \sigma_{\mu_x}^2)$ and $\sigma_x^2 \sim \text{IG}(A_x, B_x)$ for hyperparameters $\sigma_{\mu_x}^2, A_x, B_x > 0$. If normality of the x_i s cannot be reasonably assumed then (4) should be replaced by an appropriate parametric model. The variational Bayes algorithm will need to be changed accordingly. For concreteness and simplicity we will assume that (4) is reasonable for the remainder of the article.

For $1 \leq i \leq n$ let R_i be a binary random variable such that

$$R_i = \begin{cases} 1, & \text{if } x_i \text{ is observed,} \\ 0, & \text{if } x_i \text{ is missing.} \end{cases}$$

Bayesian inference for the regression model parameters differs according to the dependence of the distribution of R_i on the observed data (e.g. Gelman *et al.* 2004, Section 17.2). We will consider three missingness mechanisms:

1. $P(R_i = 1) = p$ for some constant $0 < p < 1$. In this case the missing-data mechanism is independent of the data, and the x_i s are said to be missing completely at random (MCAR). Under MCAR, the observed data are a simple random sample of the complete data.
2. $P(R_i = 1 | \phi_0, \phi_1, y_i) = \Phi(\phi_0 + \phi_1 y_i)$ for parameters $\phi_0, \phi_1 \stackrel{\text{ind.}}{\sim} N(0, \sigma_\phi^2)$ and hyperparameter $\sigma_\phi^2 > 0$. In this case, the missing-data mechanism depends on the observed y_i s but not on the missing x_i s. Inference for the regression parameters β_0, β_1 and σ_ε^2 is unaffected by the ϕ_0 and ϕ_1 or the conditional distribution $R_i | \phi_0, \phi_1, y_i$. The x_i s are said to be missing at random (MAR). In addition, the independence of the priors for (ϕ_0, ϕ_1) from those of the regression parameters means that the missingness is *ignorable* (Little & Rubin, 2004).
3. $P(R_i = 1 | \phi_0, \phi_1) = \Phi(\phi_0 + \phi_1 x_i)$ for parameters $\phi_0, \phi_1 \stackrel{\text{ind.}}{\sim} N(0, \sigma_\phi^2)$ and hyperparameter $\sigma_\phi^2 > 0$. In this case, the missing-data mechanism depends on the unobserved x_i s and inference for the regression parameters β_0, β_1 and σ_ε^2 depends on the ϕ_0 and ϕ_1 and $R_i | \phi_0, \phi_1, y_i$. The x_i s are said to be missing not at random (MNAR).

Define the matrices

$$\mathbf{X} = \begin{bmatrix} 1 & x_1 \\ \vdots & \vdots \\ 1 & x_n \end{bmatrix}, \quad \mathbf{Y} = \begin{bmatrix} 1 & y_1 \\ \vdots & \vdots \\ 1 & y_n \end{bmatrix}, \quad \boldsymbol{\beta} = \begin{bmatrix} \beta_0 \\ \beta_1 \end{bmatrix} \quad \text{and} \quad \boldsymbol{\phi} = \begin{bmatrix} \phi_0 \\ \phi_1 \end{bmatrix}.$$

Then the three missing data models can be summarized as follows:

$$\begin{aligned}
y_i|x_i, \beta, \sigma_\varepsilon^2 &\stackrel{\text{ind.}}{\sim} N(\mathbf{X}\beta, \sigma_\varepsilon^2), & x_i|\mu_x, \sigma_x^2 &\stackrel{\text{ind.}}{\sim} N(\mu_x, \sigma_x^2), \\
\beta &\sim N(\mathbf{0}, \sigma_\beta^2 \mathbf{I}), & \mu_x &\sim N(0, \sigma_{\mu_x}^2), & \sigma_\varepsilon^2 &\sim \text{IG}(A_\varepsilon, B_\varepsilon), & \sigma_x^2 &\sim \text{IG}(A_x, B_x). \\
R_i|\phi, x_i, y_i &\stackrel{\text{ind.}}{\sim} \begin{cases} \text{Bernoulli}(p), & \text{model with } x_i \text{ MCAR,} \\ \text{Bernoulli}[\Phi\{(\mathbf{Y}\phi)_i\}], & \text{model with } x_i \text{ MAR,} \\ \text{Bernoulli}[\Phi\{(\mathbf{X}\phi)_i\}], & \text{model with } x_i \text{ MNAR,} \end{cases} \\
\phi &\sim N(\mathbf{0}, \sigma_\phi^2 \mathbf{I}).
\end{aligned} \tag{5}$$

Of course, for the model with x_i MCAR, the assumption $R_i|\phi \stackrel{\text{ind.}}{\sim} \text{Bernoulli}(p)$ simplifies to $R_i \stackrel{\text{ind.}}{\sim} \text{Bernoulli}(p)$ and ϕ is superfluous.

The following additional notation is useful in the upcoming sections. Let n_{obs} denote the number of observed x_i s and n_{mis} be the number of missing x_i s. Let $\mathbf{x}_{\text{obs}}$ be the $n_{\text{obs}} \times 1$ vector containing the observed x_i s and $\mathbf{x}_{\text{mis}}$ be $n_{\text{mis}} \times 1$ vector containing the missing x_i s. We re-order the data so that the observed data is first. Hence, the full vector of predictors is

$$\mathbf{x} \equiv \begin{bmatrix} \mathbf{x}_{\text{obs}} \\ \mathbf{x}_{\text{mis}} \end{bmatrix}.$$

Finally, let $y_{x_{\text{mis},i}}$ be the value of the response variable corresponding to $x_{\text{mis},i}$.

3.1 Incorporation of Auxiliary Variables

It is now well-established that Bayesian models with probit regression components benefit from the introduction of auxiliary variables. This was demonstrated by Albert & Chib (1993) for inference via Gibbs sampling and by Girolami & Rogers (2006) for variational Bayes inference. Appropriate auxiliary variables are:

$$\begin{aligned}
a_i|\phi &\sim N((\mathbf{Y}\phi)_i, 1) \text{ for the model with } x_i \text{ MAR,} \\
\text{and } a_i|\phi &\sim N((\mathbf{X}\phi)_i, 1) \text{ for the model with } x_i \text{ MNAR.}
\end{aligned} \tag{6}$$

A consequence of (6) is

$$P(R_i = r|a_i) = I(a_i \geq 0)^r I(a_i < 0)^{1-r}, \quad r = 0, 1.$$

As will become clear in Section 3.3, variational Bayes becomes completely algebraic (i.e. without the need for numerical integration or Monte Carlo methods) if auxiliary variables are incorporated into the model.

3.2 Directed Acyclic Graphs Representations

Figure 1 provides DAG summaries of the three missing data models, after the incorporation of the auxiliary variables $\mathbf{a} = (a_1, \dots, a_n)$ given by (6). To enhance clarity, the hyperparameters are suppressed in the DAGs.

The DAGs in Figure 1 show the interplay between the regression parameters and missing data mechanism parameters. For the MCAR model the observed data indicator vector $\mathbf{R} = (R_1, \dots, R_n)$ is completely separate from the rest of the DAG. Delineation between the MAR and MNAR is more subtle, but can be gleaned from the directed edges in the respective DAGs and graph theoretical results. The Markov blanket theorem of Section 2.1 provides one way to distinguish MAR from MNAR. Table 1 lists the Markov

Figure 1: DAGs for the three missing data models for simple linear regression, given by (5). Shaded nodes correspond to the observed data.

blankets for each hidden node under the two missing-data models. Under MAR, there is a separation between the two hidden node sets

$$\{\beta, \sigma_\varepsilon^2, \mathbf{x}_{\text{mis}}, \mu_x, \sigma_x^2\} \quad \text{and} \quad \{a, \phi\}$$

in that their Markov blankets have no overlap. It follows immediately that Bayesian inference for the regression parameters based on Gibbs sampling or variational Bayes is not impacted by the missing-data mechanism. In the MNAR case, this separation does not occur since, for example, the Markov blanket of $\mathbf{x}_{\text{mis}}$ includes $\{a, \phi\}$.

node	Markov blanket under MAR	Markov blanket under MNAR
β	$\{y, \sigma_\varepsilon^2, \mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}}\}$	$\{y, \sigma_\varepsilon^2, \mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}}\}$
σ_ε^2	$\{y, \beta, \mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}}\}$	$\{y, \beta, \mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}}\}$
$\mathbf{x}_{\text{mis}}$	$\{y, \beta, \sigma_\varepsilon^2, \mathbf{x}_{\text{obs}}, \mu_x, \sigma_x^2\}$	$\{y, \beta, \sigma_\varepsilon^2, \mathbf{x}_{\text{obs}}, \mu_x, \sigma_x^2, a, \phi\}$
μ_x	$\{\mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}}, \sigma_x^2\}$	$\{\mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}}, \sigma_x^2\}$
σ_x^2	$\{\mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}}, \mu_x\}$	$\{\mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}}, \mu_x\}$
a	$\{y, R, \phi\}$	$\{\mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}}, R, \phi\}$
ϕ	$\{a, y\}$	$\{\mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}}, a\}$

Table 1: The Markov blankets for each node in the DAGs of Figure 1.

One can also use *d-separation theory* (Pearl, 1988; see also Section 8.2 of Bishop, 2006) to establish that, under MAR,

$$\{\beta, \sigma_\varepsilon^2, \mathbf{x}_{\text{mis}}, \mu_x, \sigma_x^2\} \perp\!\!\!\perp \{a, \phi\} \mid \{y, \mathbf{x}_{\text{obs}}, R\}$$

where $u \perp\!\!\!\perp v \mid w$ denotes conditional independence of u and v given w . The key to this result is the fact that all paths from the nodes in $\{\beta, \sigma_\varepsilon^2, \mathbf{x}_{\text{mis}}, \mu_x, \sigma_x^2\}$ to those in $\{a, \phi\}$ must pass through the y node. In Figure 1 we see that the y node has ‘head-to-tail’ pairs of edges that block the path between a and the regression parameters.

3.3 Approximate Inference via Variational Bayes

We will now provide details on approximate inference for each of the simple linear regression missing data models. As we shall see, variational Bayes boils down to iterative schemes for the parameters of the optimal q densities. The current subsection does little more than listing algorithms for variational Bayes inference. Section 3.4 addresses accuracy of these algorithms.

For a generic random variable v and density function $q(v)$ let

$$\mu_{q(v)} \equiv E_q(v), \quad \text{and} \quad \sigma_{q(v)}^2 \equiv \text{Var}_q(v).$$

Also, in the special case that $q(v)$ is an Inverse Gamma density function we let

$$(A_{q(v)}, B_{q(v)}) \equiv \text{shape and rate parameters of } q(v)$$

In other words, $v \sim \text{IG}(A_{q(v)}, B_{q(v)})$. Note the relationship $\mu_{q(1/v)} = A_{q(v)}/B_{q(v)}$. For a generic random vector $\mathbf{v}$ and density function $q(\mathbf{v})$ let $\boldsymbol{\mu}_{q(\mathbf{v})} \equiv E_q(\mathbf{v})$ and

$$\boldsymbol{\Sigma}_{q(\mathbf{v})} \equiv \text{Cov}_q(\mathbf{v}) = \text{covariance matrix of } \mathbf{v} \text{ under density } q(\mathbf{v}).$$

To avoid notational clutter we will omit the asterisk when applying these definitions to the optimal q^* densities.

3.3.1 MCAR model

For the MCAR model we impose the product density restriction

$$q(\boldsymbol{\beta}, \sigma_\varepsilon^2, \mathbf{x}_{\text{mis}}, \mu_x, \sigma_x^2) = q(\boldsymbol{\beta}, \mu_x) q(\sigma_\varepsilon^2, \sigma_x^2) q(\mathbf{x}_{\text{mis}}).$$

However, d -separation theory (e.g. Section 10.2.5 of Bishop, 2006) can be used to show that *induced* products $q(\boldsymbol{\beta}, \mu_x) = q(\boldsymbol{\beta})q(\mu_x)$ and $q(\sigma_\varepsilon^2, \sigma_x^2) = q(\sigma_\varepsilon^2)q(\sigma_x^2)$ arise in the solutions. Application of (2) for each component of the induced product leads to the optimal densities taking the form

$$\begin{aligned} q^*(\boldsymbol{\beta}) &= \text{Bivariate Normal density,} \\ q^*(\sigma_\varepsilon^2) &= \text{Inverse Gamma density,} \\ q^*(\mathbf{x}_{\text{mis}}) &= \text{product of } n_{\text{mis}} \text{ univariate Normal densities,} \\ q^*(\mu_x) &= \text{univariate normal density} \\ \text{and } q^*(\sigma_x^2) &= \text{Inverse Gamma density.} \end{aligned} \tag{7}$$

The optimal parameters may be obtained iteratively, using Algorithm 1. The updates for $E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X})$ and $E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X})$ are the same for all algorithms in Section 3.3, so we list them here:

$$\begin{aligned} E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) &\leftarrow \begin{bmatrix} \mathbf{1} & \mathbf{x}_{\text{obs}} \\ \mathbf{1} & \boldsymbol{\mu}_{q(\mathbf{x}_{\text{mis}})} \end{bmatrix}, \\ E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X}) &\leftarrow \begin{bmatrix} n & \mathbf{1}^T \mathbf{x}_{\text{obs}} + \mathbf{1}^T \boldsymbol{\mu}_{q(\mathbf{x}_{\text{mis}})} \\ \mathbf{1}^T \mathbf{x}_{\text{obs}} + \mathbf{1}^T \boldsymbol{\mu}_{q(\mathbf{x}_{\text{mis}})} & \|\mathbf{x}_{\text{obs}}\|^2 + \|\boldsymbol{\mu}_{q(\mathbf{x}_{\text{mis}})}\|^2 + n_{\text{mis}} \sigma_{q(\mathbf{x}_{\text{mis}})}^2 \end{bmatrix}. \end{aligned} \tag{8}$$

For the MCAR model, $\log\{\underline{p}(\mathbf{y}; q)\}$ takes the form

$$\begin{aligned} \log \underline{p}(\mathbf{y}; q) &= \frac{1}{2}(n_{\text{mis}} + 3) - (n - \frac{1}{2}n_{\text{mis}}) \log(2\pi) + \frac{n_{\text{mis}}}{2} \log(\sigma_{q(\mathbf{x}_{\text{mis}})}^2) + \frac{1}{2} \log \left| \frac{1}{\sigma_\beta^2} \boldsymbol{\Sigma}_{q(\boldsymbol{\beta})} \right| \\ &\quad - \frac{1}{2\sigma_\beta^2} \{ \|\boldsymbol{\mu}_{q(\boldsymbol{\beta})}\|^2 + \text{tr}(\boldsymbol{\Sigma}_{q(\boldsymbol{\beta})}) \} + \frac{1}{2} \log(\sigma_{q(\mu_x)}^2 / \sigma_{\mu_x}^2) - \frac{1}{2} (\mu_{q(\mu_x)}^2 + \sigma_{q(\mu_x)}^2) / \sigma_{\mu_x}^2 \\ &\quad + A_\varepsilon \log(B_\varepsilon) - A_{q(\sigma_\varepsilon^2)} \log(B_{q(\sigma_\varepsilon^2)}) + \log \Gamma(A_{q(\sigma_\varepsilon^2)}) - \log \Gamma(A_\varepsilon) \\ &\quad + A_x \log(B_x) - A_{q(\sigma_x^2)} \log(B_{q(\sigma_x^2)}) + \log \Gamma(A_{q(\sigma_x^2)}) - \log \Gamma(A_x). \end{aligned}$$

Note that, within each iteration of Algorithm 1, this expression applies only after each of the parameter updates have been made.

Derivation of (7) and Algorithm 1 is discussed in Appendix A.

Initialize: $\mu_{q(1/\sigma_\varepsilon^2)}, \mu_{q(1/\sigma_x^2)} > 0$, $\mu_{q(\beta)}(2 \times 1)$ and $\Sigma_{q(\beta)}(2 \times 2)$.

Cycle:

$$\sigma_{q(x_{\text{mis}})}^2 \leftarrow 1 / \left[\mu_{q(1/\sigma_\varepsilon^2)} + \mu_{q(1/\sigma_x^2)} \left\{ \mu_{q(\beta_1)}^2 + (\Sigma_{q(\beta)})_{22} \right\} \right]$$

for $i = 1, \dots, n_{\text{mis}}$:

$$\mu_{q(x_{\text{mis}}, i)} \leftarrow \sigma_{q(x_{\text{mis}})}^2 \left[\mu_{q(1/\sigma_\varepsilon^2)} \mu_{q(\mu_x)} + \mu_{q(1/\sigma_\varepsilon^2)} \{ y_{x_{\text{mis}}, i} \mu_{q(\beta_1)} - (\Sigma_{q(\beta)})_{12} - \mu_{q(\beta_0)} \mu_{q(\beta_1)} \} \right]$$

update $E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X})$ and $E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X})$ using (8)

$$\Sigma_{q(\beta)} \leftarrow \left\{ \mu_{q(1/\sigma_\varepsilon^2)} E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X}) + \frac{1}{\sigma_\beta^2} \mathbf{I} \right\}^{-1} \quad ; \quad \mu_{q(\beta)} \leftarrow \Sigma_{q(\beta)} \mu_{q(1/\sigma_\varepsilon^2)} E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X})^T \mathbf{y}$$

$$\sigma_{q(\mu_x)}^2 \leftarrow 1 / (n \mu_{q(1/\sigma_\varepsilon^2)} + 1/\sigma_{\mu_x}^2) \quad ; \quad \mu_{q(\mu_x)} \leftarrow \sigma_{q(\mu_x)}^2 \mu_{q(1/\sigma_\varepsilon^2)} (\mathbf{1}^T \mathbf{x}_{\text{obs}} + \mathbf{1}^T \mu_{q(\mathbf{x}_{\text{mis}})})$$

$$B_{q(\sigma_\varepsilon^2)} \leftarrow B_\varepsilon + \frac{1}{2} \|\mathbf{y}\|^2 - \mathbf{y}^T E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \mu_{q(\beta)} + \frac{1}{2} \text{tr} \{ E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X}) (\Sigma_{q(\beta)} + \mu_{q(\beta)} \mu_{q(\beta)}^T) \}$$

$$B_{q(\sigma_x^2)} \leftarrow B_x + \frac{1}{2} (\|\mathbf{x}_{\text{obs}} - \mu_{q(\mu_x)} \mathbf{1}\|^2 + \|\mu_{q(\mathbf{x}_{\text{mis}})} - \mu_{q(\mu_x)} \mathbf{1}\|^2 + n \sigma_{q(\mu_x)}^2 + n_{\text{mis}} \sigma_{q(x_{\text{mis}})}^2)$$

$$\mu_{q(1/\sigma_\varepsilon^2)} \leftarrow (A_x + \frac{1}{2}n) / B_{q(\sigma_\varepsilon^2)} \quad ; \quad \mu_{q(1/\sigma_x^2)} \leftarrow (A_\varepsilon + \frac{1}{2}n) / B_{q(\sigma_x^2)}$$

until the increase in $p(\mathbf{y}; q)$ is negligible.

Algorithm 1: Iterative scheme for obtaining the parameters in the optimal densities $q^*(\beta)$, $q^*(\sigma_\varepsilon^2)$, $q^*(\mu_x)$, $q^*(\sigma_x^2)$ and $q^*(x_{\text{mis}, i})$ for the MCAR simple linear regression model.

3.3.2 MAR model

For the MAR model we impose the product density restriction

$$q(\beta, \sigma_\varepsilon^2, \mathbf{x}_{\text{mis}}, \mu_x, \sigma_x^2, \phi, \mathbf{a}) = q(\beta, \mu_x, \phi) q(\sigma_\varepsilon^2, \sigma_x^2) q(\mathbf{x}_{\text{mis}}) q(\mathbf{a}) \quad (9)$$

As with the MCAR model, we have the induced products $q(\beta, \mu_x, \phi) = q(\beta) q(\mu_x) q(\phi)$ and $q(\sigma_\varepsilon^2, \sigma_x^2) = q(\sigma_\varepsilon^2) q(\sigma_x^2)$. Application of (2) leads to the optimal densities for $\beta, \sigma_\varepsilon^2, \mathbf{x}_{\text{mis}}, \mu_x$ and σ_x^2 have the same form as the MCAR model, given by (7).

The missing-data mechanism parameters have optimal densities

$$\begin{aligned} q^*(\phi) &= N((\mathbf{Y}^T \mathbf{Y} + \frac{1}{\sigma_\phi^2} \mathbf{I})^{-1}, (\mathbf{Y}^T \mathbf{Y} + \frac{1}{\sigma_\phi^2} \mathbf{I})^{-1} \mathbf{Y}^T \mu_{q(\mathbf{a})}) \\ q^*(\mathbf{a}) &= \left[\prod_{i=1}^n \left\{ \frac{I(a_i \geq 0)}{\Phi((\mathbf{Y} \mu_{q(\phi)})_i)} \right\}^{y_i} \left\{ \frac{I(a_i < 0)}{1 - \Phi((\mathbf{Y} \mu_{q(\phi)})_i)} \right\}^{1-y_i} \right] \\ &\quad \times (2\pi)^{-n/2} \exp\{-\frac{1}{2} \|\mathbf{a} - \mathbf{Y} \mu_{q(\phi)}\|^2\}. \end{aligned} \quad (10)$$

The optimal parameters can be found iteratively via Algorithm 2. Appendix A discusses its derivation.

Initialize: $\mu_{q(1/\sigma_\varepsilon^2)}, \mu_{q(1/\sigma_x^2)} > 0$, $\boldsymbol{\mu}_{q(\beta)}(2 \times 1)$ and $\boldsymbol{\Sigma}_{q(\beta)}(2 \times 2)$.
 Cycle:

$$\begin{aligned}
 \sigma_{q(x_{\text{mis}})}^2 &\leftarrow 1 / \left[\mu_{q(1/\sigma_x^2)} + \mu_{q(1/\sigma_\varepsilon^2)} \left\{ \mu_{q(\beta_1)}^2 + (\boldsymbol{\Sigma}_{q(\beta)})_{22} \right\} \right] \\
 \text{for } i = 1, \dots, n_{\text{mis}}: \\
 \mu_{q(x_{\text{mis}}, i)} &\leftarrow \sigma_{q(x_{\text{mis}})}^2 \left[\mu_{q(1/\sigma_x^2)} \mu_{q(\mu_x)} + \mu_{q(1/\sigma_\varepsilon^2)} \left\{ y_{x_{\text{mis}}, i} \mu_{q(\beta_1)} - (\boldsymbol{\Sigma}_{q(\beta)})_{12} - \mu_{q(\beta_0)} \mu_{q(\beta_1)} \right\} \right] \\
 \text{update } E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \text{ and } E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X}) \text{ using (8)} \\
 \boldsymbol{\Sigma}_{q(\beta)} &\leftarrow \left\{ \mu_{q(1/\sigma_\varepsilon^2)} E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X}) + \frac{1}{\sigma_\beta^2} \mathbf{I} \right\}^{-1} \quad ; \quad \boldsymbol{\mu}_{q(\beta)} \leftarrow \boldsymbol{\Sigma}_{q(\beta)} \mu_{q(1/\sigma_\varepsilon^2)} E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X})^T \mathbf{y} \\
 \sigma_{q(\mu_x)}^2 &\leftarrow 1 / (n \mu_{q(1/\sigma_x^2)} + 1/\sigma_{\mu_x}^2) \quad ; \quad \mu_{q(\mu_x)} \leftarrow \sigma_{q(\mu_x)}^2 \mu_{q(1/\sigma_x^2)} (\mathbf{1}^T \mathbf{x}_{\text{obs}} + \mathbf{1}^T \boldsymbol{\mu}_{q(\mathbf{x}_{\text{mis}})}) \\
 B_{q(\sigma_\varepsilon^2)} &\leftarrow B_\varepsilon + \frac{1}{2} \|\mathbf{y}\|^2 - \mathbf{y}^T E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\beta)} + \frac{1}{2} \text{tr} \{ E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X}) (\boldsymbol{\Sigma}_{q(\beta)} + \boldsymbol{\mu}_{q(\beta)} \boldsymbol{\mu}_{q(\beta)}^T) \} \\
 B_{q(\sigma_x^2)} &\leftarrow B_x + \frac{1}{2} (\|\mathbf{x}_{\text{obs}} - \mu_{q(\mu_x)} \mathbf{1}\|^2 + \|\boldsymbol{\mu}_{q(\mathbf{x}_{\text{mis}})} - \mu_{q(\mu_x)} \mathbf{1}\|^2 + n \sigma_{q(\mu_x)}^2 + n_{\text{mis}} \sigma_{q(x_{\text{mis}})}^2) \\
 \mu_{q(1/\sigma_\varepsilon^2)} &\leftarrow (A_\varepsilon + \frac{1}{2}n) / B_{q(\sigma_\varepsilon^2)} \quad ; \quad \mu_{q(1/\sigma_x^2)} \leftarrow (A_x + \frac{1}{2}n) / B_{q(\sigma_x^2)} \\
 \boldsymbol{\Sigma}_{q(\phi)} &\leftarrow (\mathbf{Y}^T \mathbf{Y} + \frac{1}{\sigma_\phi^2} \mathbf{I})^{-1}; \quad \boldsymbol{\mu}_{q(\phi)} \leftarrow \boldsymbol{\Sigma}_{q(\phi)} \mathbf{Y}^T \boldsymbol{\mu}_{q(\mathbf{a})} \\
 \boldsymbol{\mu}_{q(\mathbf{a})} &\leftarrow \mathbf{Y} \boldsymbol{\mu}_{q(\phi)} + (2\mathbf{R} - \mathbf{1}) \odot \frac{(2\pi)^{-1/2} \exp\{-\frac{1}{2}(\mathbf{Y} \boldsymbol{\mu}_{q(\phi)})^2\}}{\Phi((2\mathbf{R} - \mathbf{1}) \odot (\mathbf{Y} \boldsymbol{\mu}_{q(\phi)}))}
 \end{aligned}$$

until the increase in $\underline{p}(\mathbf{y}; q)$ is negligible.

Algorithm 2: Iterative scheme for obtaining the parameters in the optimal densities $q^*(\beta)$, $q^*(\sigma_\varepsilon^2)$, $q^*(\mu_x)$, $q^*(\sigma_x^2)$, $q^*(x_{\text{mis}, i})$ and $q^*(\phi)$ for the MAR simple linear regression model.

For the MAR model, $\log\{\underline{p}(\mathbf{y}; q)\}$ takes the form

$$\begin{aligned}
 \log \underline{p}(\mathbf{y}; q) &= \frac{1}{2}(n_{\text{mis}} + 5) - (n - \frac{1}{2}n_{\text{mis}}) \log(2\pi) + \frac{n_{\text{mis}}}{2} \log(\sigma_{q(x_{\text{mis}})}^2) + \frac{1}{2} \log |\frac{1}{\sigma_\beta^2} \boldsymbol{\Sigma}_{q(\beta)}| \\
 &\quad - \frac{1}{2\sigma_\beta^2} \{\|\boldsymbol{\mu}_{q(\beta)}\|^2 + \text{tr}(\boldsymbol{\Sigma}_{q(\beta)})\} + \frac{1}{2} \log(\sigma_{q(\mu_x)}^2 / \sigma_{\mu_x}^2) - \frac{1}{2} (\mu_{q(\mu_x)}^2 + \sigma_{q(\mu_x)}^2) / \sigma_{\mu_x}^2 \\
 &\quad + A_\varepsilon \log(B_\varepsilon) - A_{q(\sigma_\varepsilon^2)} \log(B_{q(\sigma_\varepsilon^2)}) + \log \Gamma(A_{q(\sigma_\varepsilon^2)}) - \log \Gamma(A_\varepsilon) \\
 &\quad + A_x \log(B_x) - A_{q(\sigma_x^2)} \log(B_{q(\sigma_x^2)}) + \log \Gamma(A_{q(\sigma_x^2)}) - \log \Gamma(A_x) \\
 &\quad + \frac{1}{2} \|\mathbf{Y} \boldsymbol{\mu}_{q(\phi)}\|^2 - \frac{1}{2} \text{tr} \{ \mathbf{Y}^T \mathbf{Y} (\boldsymbol{\mu}_{q(\phi)} \boldsymbol{\mu}_{q(\phi)}^T + \boldsymbol{\Sigma}_{q(\phi)}) \} \\
 &\quad + \mathbf{R}^T \log \{ \Phi(\mathbf{Y} \boldsymbol{\mu}_{q(\phi)}) + (\mathbf{1} - \mathbf{R})^T \log \{ \mathbf{1} - \Phi(\mathbf{Y} \boldsymbol{\mu}_{q(\phi)}) \} \} \\
 &\quad + \frac{1}{2} \log |\frac{1}{\sigma_\phi^2} \boldsymbol{\Sigma}_{q(\phi)}| - \frac{1}{2\sigma_\phi^2} \{\|\boldsymbol{\mu}_{q(\phi)}\|^2 + \text{tr}(\boldsymbol{\Sigma}_{q(\phi)})\}.
 \end{aligned}$$

Note that, within each iteration of Algorithm 2, this expression applies only after each of the parameter updates have been made.

3.3.3 MNAR model

For the simple linear regression model with predictors MNAR model we again impose the product density restriction (9). The regression model parameters have the same distributional forms as the MCAR and MAR cases (given by (7)). However, the parameters of $q^*(x_{\text{mis}})$ depend on the missing-data mechanism parameters.

The missing-data mechanism parameters have the same form as (10), but with $\mathbf{Y}$ replaced by $E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X})$. Also, for $1 \leq i \leq n_{\text{mis}}$, $a_{x_{\text{mis},i}}$ denotes the entry of $\mathbf{a}$ corresponding to $x_{\text{mis},i}$.

Initialize: $\mu_{q(1/\sigma_\varepsilon^2)}, \mu_{q(1/\sigma_x^2)} > 0$, $\boldsymbol{\mu}_{q(\beta)}(2 \times 1)$ and $\boldsymbol{\Sigma}_{q(\beta)}(2 \times 2)$.

Cycle:

$$\sigma_{q(x_{\text{mis}})}^2 \leftarrow 1 / \left[\mu_{q(1/\sigma_x^2)} + \mu_{q(1/\sigma_\varepsilon^2)} \left\{ \mu_{q(\beta_1)}^2 + (\boldsymbol{\Sigma}_{q(\beta)})_{22} \right\} + \mu_{q(\phi_1)}^2 + (\boldsymbol{\Sigma}_{q(\phi)})_{22} \right]$$

for $i = 1, \dots, n_{\text{mis}}$:

$$\begin{aligned} \mu_{q(x_{\text{mis},i})} \leftarrow \sigma_{q(x_{\text{mis}})}^2 & \left[\mu_{q(1/\sigma_x^2)} \mu_{q(\mu_x)} + \mu_{q(1/\sigma_\varepsilon^2)} \left\{ y_{x_{\text{mis},i}} \mu_{q(\beta_1)} - (\boldsymbol{\Sigma}_{q(\beta)})_{12} - \mu_{q(\beta_0)} \mu_{q(\beta_1)} \right\} \right. \\ & \left. + \mu_{q(a_{x_{\text{mis},i}})} \mu_{q(\phi_1)} - (\boldsymbol{\Sigma}_{q(\phi)})_{12} - \mu_{q(\phi_0)} \mu_{q(\phi_1)} \right]. \end{aligned}$$

update $E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X})$ and $E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X})$ using (8)

$$\boldsymbol{\Sigma}_{q(\beta)} \leftarrow \left\{ \mu_{q(1/\sigma_\varepsilon^2)} E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X}) + \frac{1}{\sigma_\beta^2} \mathbf{I} \right\}^{-1} \quad ; \quad \boldsymbol{\mu}_{q(\beta)} \leftarrow \boldsymbol{\Sigma}_{q(\beta)} \mu_{q(1/\sigma_\varepsilon^2)} E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X})^T \mathbf{y}$$

$$\sigma_{q(\mu_x)}^2 \leftarrow 1 / (n \mu_{q(1/\sigma_x^2)} + 1/\sigma_{\mu_x}^2) \quad ; \quad \mu_{q(\mu_x)} \leftarrow \sigma_{q(\mu_x)}^2 \mu_{q(1/\sigma_x^2)} (\mathbf{1}^T \mathbf{x}_{\text{obs}} + \mathbf{1}^T \boldsymbol{\mu}_{q(\mathbf{x}_{\text{mis}})})$$

$$B_{q(\sigma_\varepsilon^2)} \leftarrow B_\varepsilon + \frac{1}{2} \|\mathbf{y}\|^2 - \mathbf{y}^T E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\beta)} + \frac{1}{2} \text{tr} \{ E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X}) (\boldsymbol{\Sigma}_{q(\beta)} + \boldsymbol{\mu}_{q(\beta)} \boldsymbol{\mu}_{q(\beta)}^T) \}$$

$$B_{q(\sigma_x^2)} \leftarrow B_x + \frac{1}{2} (\|\mathbf{x}_{\text{obs}} - \mu_{q(\mu_x)} \mathbf{1}\|^2 + \|\boldsymbol{\mu}_{q(\mathbf{x}_{\text{mis}})} - \mu_{q(\mu_x)} \mathbf{1}\|^2 + n \sigma_{q(\mu_x)}^2 + n_{\text{mis}} \sigma_{q(x_{\text{mis}})}^2)$$

$$\mu_{q(1/\sigma_\varepsilon^2)} \leftarrow (A_\varepsilon + \frac{1}{2}n) / B_{q(\sigma_\varepsilon^2)} \quad ; \quad \mu_{q(1/\sigma_x^2)} \leftarrow (A_x + \frac{1}{2}n) / B_{q(\sigma_x^2)}$$

$$\boldsymbol{\Sigma}_{q(\phi)} \leftarrow \left\{ E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X}) + \frac{1}{\sigma_\phi^2} \mathbf{I} \right\}^{-1} \quad ; \quad \boldsymbol{\mu}_{q(\phi)} \leftarrow \boldsymbol{\Sigma}_{q(\phi)} E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X})^T \boldsymbol{\mu}_{q(\mathbf{a})}$$

$$\boldsymbol{\mu}_{q(\mathbf{a})} \leftarrow E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)} + (2\mathbf{R} - \mathbf{1}) \odot \frac{(2\pi)^{-1/2} \exp\{-\frac{1}{2}(E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)})^2\}}{\Phi((2\mathbf{R} - \mathbf{1}) \odot (E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)})}$$

until the increase in $p(\mathbf{y}; q)$ is negligible.

Algorithm 3: Iterative scheme for obtaining the parameters in the optimal densities $q^*(\beta)$, $q^*(\sigma_\varepsilon^2)$, $q^*(\mu_x)$, $q^*(\sigma_x^2)$, $q^*(x_{\text{mis},i})$ and $q^*(\phi)$ for the MNAR simple linear regression model.

For the MNAR model $\log\{p(\mathbf{y}; q)\}$ is given by

$$\begin{aligned} \log p(\mathbf{y}; q) = & \frac{1}{2}(n_{\text{mis}} + 5) - (n - \frac{1}{2}n_{\text{mis}}) \log(2\pi) + \frac{n_{\text{mis}}}{2} \log(\sigma_{q(x_{\text{mis}})}^2) + \frac{1}{2} \log \left| \frac{1}{\sigma_\beta^2} \boldsymbol{\Sigma}_{q(\beta)} \right| \\ & - \frac{1}{2\sigma_\beta^2} \{ \|\boldsymbol{\mu}_{q(\beta)}\|^2 + \text{tr}(\boldsymbol{\Sigma}_{q(\beta)}) \} + \frac{1}{2} \log(\sigma_{q(\mu_x)}^2 / \sigma_{\mu_x}^2) - \frac{1}{2} (\mu_{q(\mu_x)}^2 + \sigma_{q(\mu_x)}^2) / \sigma_{\mu_x}^2 \\ & + A_\varepsilon \log(B_\varepsilon) - A_{q(\sigma_\varepsilon^2)} \log(B_{q(\sigma_\varepsilon^2)}) + \log \Gamma(A_{q(\sigma_\varepsilon^2)}) - \log \Gamma(A_\varepsilon) \\ & + A_x \log(B_x) - A_{q(\sigma_x^2)} \log(B_{q(\sigma_x^2)}) + \log \Gamma(A_{q(\sigma_x^2)}) - \log \Gamma(A_x) \\ & + \frac{1}{2} \|E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)}\|^2 - \frac{1}{2} \text{tr} \{ E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X}) (\boldsymbol{\mu}_{q(\phi)} \boldsymbol{\mu}_{q(\phi)}^T + \boldsymbol{\Sigma}_{q(\phi)}) \} \\ & + \mathbf{R}^T \log \{ \Phi(E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)}) + (\mathbf{1} - \mathbf{R})^T \log \{ \mathbf{1} - \Phi(E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)}) \} \\ & + \frac{1}{2} \log \left| \frac{1}{\sigma_\phi^2} \boldsymbol{\Sigma}_{q(\phi)} \right| - \frac{1}{2\sigma_\phi^2} \{ \|\boldsymbol{\mu}_{q(\phi)}\|^2 + \text{tr}(\boldsymbol{\Sigma}_{q(\phi)}) \}. \end{aligned}$$

Note that, within each iteration of Algorithm 3, this expression applies only after each of the parameter updates have been made.

3.4 Assessment of Accuracy

We now turn attention to the issue of accuracy of variational Bayes inference for models (5). Algorithms 1–3 provide speedy approximate inference for the model parameters, but come with no guarantees of achieving an acceptable level of accuracy.

Let θ denote a generic univariate model parameter. There are numerous means by which the accuracy of a variational Bayes approximate density $q^*(\theta)$ can be measured with respect to the exact posterior density $p(\theta|\mathbf{y})$. Kullback-Leibler distance is an obvious choice but can be dominated by the tail-behaviour of the densities involved (e.g. Hall, 1987). We recommend working with the L_1 loss, or *integrated absolute error* (IAE) of q^* , given by

$$\text{IAE}(q^*) = \int_{-\infty}^{\infty} |q^*(\theta) - p(\theta|\mathbf{y})| d\theta.$$

This error measure has the attractions of being (a) invariant to monotone transformations on the parameter θ and (b) a scale-independent number between 0 and 2 (e.g. Devroye & Györfi, 1985). The second of these motivates the accuracy measure

$$\text{accuracy}(q^*) = 1 - \{\text{IAE}(q^*) / \sup_{q \text{ a density}} \text{IAE}(q)\} = 1 - \frac{1}{2}\text{IAE}(q^*). \quad (11)$$

Note that $0 \leq \text{accuracy}(q^*) \leq 1$ and will be expressed as a percentage in the examples to follow.

Computation of $\text{accuracy}(q^*)$ is a little challenging, since it depends on the posterior $p(\theta|\mathbf{y})$ that we are trying to avoid by using approximate inference methods. However, MCMC with sufficiently large samples can be used to approximate $p(\theta|\mathbf{y})$ arbitrarily well. The accuracy assessments that we present in this section are based on MCMC samples obtained using **BRugs** (Ligges *et al.* 2010) with a burnin of size 10000. A thinning factor of 5 was applied to post-burnin samples of size 50000. This resulted in MCMC samples of size 10000 for density estimation. Density estimates were obtained using the binned kernel density estimate `bkde()` function in the R package `KernSmooth` (Wand & Ripley, 2009). The bandwidth was chosen using a direct plug-in rule, corresponding to the default version of `dpik()`. These density estimates act as a proxy for the exact posterior densities. For sample sizes as large as 10000 and well-behaved posteriors the quality of these proxies should be quite good. Nevertheless, it must be noted that they are subject to errors inherent in density estimation and bandwidth selection.

Due to space considerations, our accuracy assessments are limited to the MCAR and MNAR models.

3.4.1 MCAR model

We ran a simulation study to assess the accuracy of Algorithm 1. Throughout the simulation study we fixed

$$n = 500, \beta_0 = \beta_1 = 1, \mu_x = \frac{1}{2} \quad \text{and} \quad \sigma_x^2 = \frac{1}{36}. \quad (12)$$

The noise and missing-data levels were varied according to

$$\sigma_\varepsilon \in \{0.05, 0.2, 0.8\} \quad \text{and} \quad p \in \{0.6, 0.8\}. \quad (13)$$

The values of σ_ε correspond to low, medium and high noise scenarios. Having $p = 0.8$ corresponds to moderate (20%) missingness while $p = 0.6$ corresponds to more severe (40%) missingness. The hyperparameters were set at

$$\sigma_\beta^2 = \sigma_{\mu_x}^2 = 10^{-8} \quad \text{and} \quad A_\varepsilon = B_\varepsilon = A_x = B_x = \frac{1}{100}. \quad (14)$$

Algorithm 1 was terminated when the relative increase in $\log\{p(\mathbf{y}|q)\}$ was less than 10^{-10} .

Figure 2 summarises the accuracy results based on 100 data sets simulated according to (12) and (13). For moderate missingness ($p = 0.8$) we see that variational Bayes is very accurate, with all accuracy values above 90% regardless of the noise level. For severe missingness ($p = 0.6$) there is some degradation in the accuracy of variational Bayes, but it stays about 83% for all replications.

Figure 2: Summary of simulation for simple linear regression with predictor MCAR. For each setting, the accuracy values are summarized as a boxplot.

Figure 3 allows appreciation for the accuracy of variational Bayes in this context. It corresponds to a typical realization from the simulation study with $\sigma_\epsilon = 0.2$ and $p = 0.8$.

3.4.2 MNAR model

In our simulation for MNAR model, the missingness was controlled by the two pair of probit coefficients:

$$(\phi_0, \phi_1) = (2.95, -2.95) \quad \text{and} \quad (\phi_0, \phi_1) = (0.85, -1.05).$$

In each case, the probability of missingness increases as a function of the covariate. For the first pair the missingness probability ranges from 0.0 to 0.5 with an average of 0.25. For the second pair the range is 0.2 to 0.58 with an average of 0.39, representing more severe missingness. The hyperparameter for (ϕ_0, ϕ_1) was set at $\sigma_\phi^2 = 10^{-8}$.

Figure 4 summarizes the accuracy results based on 100 simulated data sets while Figure 5 plots the variational Bayes and MCMC approximate posteriors for a typical realization from the simulation study with $\sigma_\epsilon = 0.2$ and $(\phi_0, \phi_1) = (2.95, -2.95)$.

The parameters corresponding to the regression part of the model $(\beta_0, \beta_1, \sigma_\epsilon^2)$ show high accuracy, with almost all accuracy levels above 80%. The accuracy drops considerably when the amount of missing data is large or when the data are noisy. This might be expected since there is a decrease in the amount of information about the parameters. The accuracy of the missing covariates is high in all situations, even when the missing data percentage is very large.

Figure 3: Variational Bayes approximate posteriors for the regression model parameters and three missing x_i s for simple linear regression with predictors MCAR. The regression parameters are $(\beta_0, \beta_1, \sigma_\epsilon) = (1, 1, 0.2)$ and the probability of the x_i s being observed is $p = 0.8$. The vertical lines correspond to the true values of the parameters from which the data were simulated (described in the text). The MCMC posteriors are based on samples of size 10000 and kernel density estimation. The accuracy values correspond to the definition given at (11).

The variational Bayes approximations generally are poor for the missingness mechanism parameters ϕ_0 and ϕ_1 . This is due to strong posterior correlation between ϕ and α in probit auxiliary variable models, as is reported in Section 2.1 of Holmes & Held (2006), for example. This deficiency of variational Bayes is isolated to the lower nodes of the right-most DAG in Figure 1 and can only be remedied through a more elaborate variational approximation – for example, one that allows posterior dependence between ϕ and α . Such elaboration will bring computational costs, which need to be traded off against the importance of making inference about the MNAR parameters. In many applied contexts, these parameters are not of primary interest.

3.5 Speed Comparisons

While running the simulation studies described in Section 3.4 we kept track of the time taken for each model to be fitted. The results are summarized in Table 3.5. The computer involved used the Mac OS X operating system with a 2.33 GHz processor and 3 GBytes of random access memory.

As with most speed comparisons, some caveats need to be taken into account. Firstly, the MCMC and variational Bayes answers were computed using different programming languages. The MCMC model fits were obtained using the BUGS inference engine (Lunn *et al.* 2000) with interfacing via the package BRugs (Ligges, *et al.* 2010) in the R computing environment (R Development Core Team, 2010). The variational Bayes model fits

Figure 4: Summary of simulation for simple linear regression with predictor MNAR. For each setting, the accuracy values are summarized as a boxplot.

	MAR models	MNAR models
MCMC	(5.89,5.84)	(33.8,33.9)
var. Bayes	(0.0849,0.0850)	(0.705,0.790)
ratio	(76.6,78.7)	(59.5,67.8)

Table 2: 99% Wilcoxon confidence intervals based on computation times, in seconds, from the simulation study described in Section 3.4.

were implemented using R. Secondly, no effort was made to tailor MCMC scheme to the models at hand. Thirdly, as detailed in Section 3.4, both methods had arbitrarily chosen stopping criteria. Despite these caveats, Table 3.5 gives an impression of the relative computing times involved if an ‘off-the-shelf’ MCMC implementation is used.

Caveats aside, the results indicate that variational Bayes is at least 60 times faster than MCMC across all models. Hence, a model that takes minutes to run in MCMC takes only seconds with variational Bayes approximation.

4 Nonparametric Regression with Missing Predictor Data

We now describe extension to nonparametric regression with missing predictor data. The essence of this extension is replacement of the linear mean function

$$\beta_0 + \beta_1 x \quad \text{by} \quad f(x)$$

where f is a smooth flexible function. There are numerous approaches to modelling and estimating f . The one which is most conducive to inference via variational Bayes is penalized splines with mixed model representation. This involves the model

$$f(x) = \beta_0 + \beta_1 x + \sum_{k=1}^K u_k z_k(x), \quad u_k \stackrel{\text{ind.}}{\sim} N(0, \sigma_u^2) \quad (15)$$

Figure 5: Variational Bayes approximate posteriors for the regression model parameters and three missing x_i s for simple linear regression with predictors MNAR. The regression parameters are $(\beta_0, \beta_1, \sigma_\epsilon) = (1, 1, 0.2)$ and the probability of x_i being observed is $\Phi(\phi_0 + \phi_1 x_i)$ where $(\phi_0, \phi_1) = (2.95, -2.95)$. The vertical lines correspond to the true values of the parameters from which the data were simulated (described in the text). The MCMC posteriors are based on samples of size 10000 and kernel density estimation. The accuracy values correspond to the definition given at (11).

where the $\{z_k(\cdot) : 1 \leq k \leq K\}$ are an appropriate set of spline basis functions. Several options exist for the z_k . Our preference is suitable transformed O’Sullivan penalized splines (Wand & Ormerod, 2008) since this leads to approximate smoothing splines; which have good boundary and extrapolation properties.

From the graphical model standpoint, moving from parametric regression to non-parametric regression using mixed model-based penalized splines simply involves enlarging the DAGs from parametric regression. Figure 6 shows the nonparametric regression DAGs for the three missing data mechanisms treated in Section 3. Comparison with Figure 1 shows the only difference is the addition of the σ_u^2 node, and replacement of β by (β, u) . Note that (β, u) could be broken up into separate nodes, but the update expressions are simpler if these two random vectors are kept together.

The variational Bayes algorithms for the DAGs in 6 simply involve modification of the Algorithms 1–3 to accommodate the additional nodes and edges. However, the spline basis functions give rise to non-standard forms and numerical integration is required. We will give a detailed account of this extension in the MNAR case only. The MCAR and MAR cases require similar arguments, but are simpler.

Figure 6: DAGs for the three missing data models for nonparametric regression with mixed model-based penalized spline modelling of the regression function; given by (15). Shaded nodes correspond to the observed data.

Define the $1 \times (K + 2)$ vector

$$\mathbf{C}_x \equiv (1, x, z_1(x), \dots, z_K(x))$$

corresponding to evaluation of penalized spline basis functions at an arbitrary location $x \in \mathbb{R}$. Then the optimal densities for the $x_{\text{mis},i}$, $1 \leq i \leq n_{\text{mis}}$, take the form

$$q^*(x_{\text{mis},i}) \propto \exp\left(-\frac{1}{2} \mathbf{C}_{x_{\text{mis},i}} \boldsymbol{\Lambda}_{\text{mis},i} \mathbf{C}_{x_{\text{mis},i}}^T\right) \quad (16)$$

where the $(K + 2) \times (K + 2)$ matrices $\boldsymbol{\Lambda}_{\text{mis},i}$, $1 \leq i \leq n_{\text{mis}}$ correspond to each entry of $\mathbf{x}_{\text{mis}} = (x_{\text{mis},1}, \dots, x_{\text{mis},n_{\text{mis}}})$ but does not depend on $x_{\text{mis},i}$. A derivation of (16) and expressions for the $\boldsymbol{\Lambda}_{\text{mis},i}$ are given in Appendix B.

The right-hand side of (16) does not have a closed-form integral, so numerical integration is required to obtain the normalizing factors and required moments. We will take a basic quadrature approach. In the interests of computational efficiency, we use the same quadrature grid over all $1 \leq i \leq n_{\text{mis}}$. Let

$$\mathbf{g} = (g_1, \dots, g_M)$$

be an equally-spaced grid of size M in $\mathbb{R}$. An example of numerical integration via quadrature is

$$\int_{-\infty}^{\infty} z_1(x) dx \approx \sum_{j=1}^M w_j z_1(g_j) = \mathbf{w}^T \mathbf{z}_1(\mathbf{g})$$

where $\mathbf{w} = (w_1, \dots, w_M)$ is vector of quadrature weights. Examples of $\mathbf{w}$ for common quadrature schemes are

$$\mathbf{w} = \begin{cases} \frac{1}{2} \delta \times (1, 2, 2, \dots, 2, 2, 1), & \text{for the trapezoidal rule,} \\ \frac{1}{3} \delta \times (1, 4, 2, 4, 2, \dots, 2, 4, 2, 4, 1), & \text{for Simpson's rule,} \end{cases}$$

where $\delta = (g_M - g_1)/(M - 1)$ is the distance between successive grid-points. Next, define the $M \times (K + 2)$ matrix:

$$\mathbf{C}_g \equiv \begin{bmatrix} 1 & g_1 & z_1(g_1) & \cdots & z_K(g_1) \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & g_M & z_1(g_M) & \cdots & z_K(g_M) \end{bmatrix} = \begin{bmatrix} \mathbf{C}_{g_1} \\ \vdots \\ \mathbf{C}_{g_M} \end{bmatrix}. \quad (17)$$

For a given quadrature grid $\mathbf{g}$, $\mathbf{C}_g$ contains the totality of basis function evaluations required for variational Bayes updates.

For succinct statement of quadrature approximations to $E_{q(x_{\text{mis}})}(C)$ and $E_{q(x_{\text{mis}})}(C^T C)$ the following additional matrix notation is useful:

$$\mathbf{Q}_g \equiv [\exp(-\frac{1}{2} \mathbf{C}_{g_j} \mathbf{\Lambda}_{\text{mis},i} \mathbf{C}_{g_j}^T)]_{1 \leq i \leq n_{\text{mis}}, 1 \leq j \leq M} \quad \text{and} \quad \mathbf{C} \equiv \begin{bmatrix} \mathbf{C}_{\text{obs}} \\ \mathbf{C}_{\text{mis}} \end{bmatrix}$$

where $\mathbf{C}_{\text{obs}}$ corresponds to the $\mathbf{x}_{\text{obs}}$ component of $\mathbf{C}$ and $\mathbf{C}_{\text{mis}}$ corresponds to the $\mathbf{x}_{\text{mis}}$ component of $\mathbf{C}$. Clearly

$$E_{q(x_{\text{mis}})}(\mathbf{C}) \equiv \begin{bmatrix} \mathbf{C}_{\text{obs}} \\ E_{q(x_{\text{mis}})}(\mathbf{C}_{\text{mis}}) \end{bmatrix} \quad \text{and} \quad E_{q(x_{\text{mis}})}(\mathbf{C}^T \mathbf{C}) = \mathbf{C}_{\text{obs}}^T \mathbf{C}_{\text{obs}} + E_{q(x_{\text{mis}})}(\mathbf{C}_{\text{mis}}^T \mathbf{C}_{\text{mis}}).$$

Then we have the following efficient quadrature approximations:

$$E_{q(x_{\text{mis}})}(\mathbf{C}_{\text{mis}}) \approx \frac{\mathbf{Q}_g \text{diag}(\mathbf{w}) \mathbf{C}_g}{\mathbf{1}^T \otimes (\mathbf{Q}_g \mathbf{w})} \quad \text{and} \quad E_{q(x_{\text{mis}})}(\mathbf{C}_{\text{mis}}^T \mathbf{C}_{\text{mis}}) \approx \mathbf{C}_g^T \text{diag} \left(\sum_{i=1}^{n_{\text{mis}}} \frac{(\mathbf{e}_i^T \mathbf{Q}_g) \odot \mathbf{w}}{\mathbf{e}_i^T \mathbf{Q}_g \mathbf{w}} \right) \mathbf{C}_g$$

with $\mathbf{e}_i$ denoting the $n_{\text{mis}} \times 1$ vector with 1 in the i th position and zeroes elsewhere. Since there are exponentials in entries of $\mathbf{Q}_g$ some care needs to be taken to avoid overflow and underflow. Working with logarithms is recommended.

Algorithm 4 chronicles the iterative scheme for nonparametric regression with predictors MNAR. The lower bound on the marginal log-likelihood is

$$\begin{aligned} p(\mathbf{y}; q) = & \left(\frac{1}{2} n_{\text{mis}} - n \right) \log(2\pi) + \frac{1}{2} (K + 5 + n_{\text{mis}}) - \log(\sigma_\beta^2) - \frac{1}{2\sigma_\beta^2} [\|\mu_{q(\beta)}\|^2 + \text{tr}(\Sigma_{q(\beta)})] \\ & + \frac{1}{2} \log |\Sigma_{q(\beta, \mathbf{u})}| + \frac{1}{2} \log \{ \sigma_{q(\mu_x)}^2 / \sigma_{\mu_x}^2 \} - \frac{1}{2\sigma_{\mu_x}^2} \{ \mu_{q(\mu_x)}^2 + \sigma_{q(\mu_x)}^2 \} - \frac{\mathbf{Q}_g \text{diag}(\mathbf{w}) \log(\mathbf{Q}_g)}{\mathbf{1}^T \otimes (\mathbf{Q}_g \mathbf{w})} \\ & + A_\varepsilon \log(B_\varepsilon) - \log \Gamma(A_\varepsilon) - A_{q(\sigma_\varepsilon^2)} \log(B_{q(\sigma_\varepsilon^2)}) + \log(A_{q(\sigma_\varepsilon^2)}) \\ & + A_u \log(B_u) - \log \Gamma(A_u) - A_{q(\sigma_u^2)} \log(B_{q(\sigma_u^2)}) + \log(A_{q(\sigma_u^2)}) \\ & + A_x \log(B_x) - \log \Gamma(A_x) - A_{q(\sigma_x^2)} \log(B_{q(\sigma_x^2)}) + \log(A_{q(\sigma_x^2)}) \\ & + \frac{1}{2} \|E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)}\|^2 - \frac{1}{2} \text{tr} \left[E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X}) (\boldsymbol{\mu}_{q(\phi)} \boldsymbol{\mu}_{q(\phi)}^T + \Sigma_{q(\phi)}) \right] \\ & + \mathbf{R}^T \log \Phi(E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)}) + (\mathbf{1} - \mathbf{R})^T \log \{ \mathbf{1} - \Phi(E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)}) \} \\ & + \frac{1}{2} \log \left| \frac{1}{\sigma_\phi^2} \Sigma_{q(\phi)} \right| - \frac{1}{2\sigma_\phi^2} \left\{ \|\boldsymbol{\mu}_{q(\phi)}\|^2 + \text{tr}(\Sigma_{q(\phi)}) \right\} \end{aligned}$$

4.1 Illustration

Our first illustration involves data simulated according to

$$y_i \sim N(f(x_i), \sigma_\varepsilon^2), \quad f(x) = \sin(4\pi x), \quad x_i \sim N\left(\frac{1}{2}, \frac{1}{36}\right) \quad \text{and} \quad \sigma_\varepsilon^2 = 0.35, \quad 1 \leq i \leq 300$$

and with 20% of the x_i s removed completely at random. This simulation setting, with identical parameters, was also used in Wand (2009).

We applied the MCAR analogue of Algorithm 4 and compared the results with MCMC fitting via BRugs. The penalized splines used the truncated linear spline basis with 30 knots: $z_k(x) = (x - \kappa_k)_+$, $1 \leq k \leq 30$, with the knots equally-spaced over the range of the observed x_i s. Truncated linear splines were used to allow straightforward coding in BUGS. If a comparison with MCMC is not being done then O'Sullivan splines are recommended for variational Bayesian inference in this context. The hyperparameters were set at the values given in (14). The MCMC sampling involved a burnin of size 20000, and a thinning factor of 20 applied to post-burnin samples of size 200000 resulting in samples of size 10000 being retained for inference. In addition, we used the over-relaxed

Set M , the size of the quadrature grid, and g_1 and g_M , the quadrature grid limits. The interval (g_1, g_M) should contain each of the observed x_i s. Obtain $\mathbf{g} = (g_1, \dots, g_M)$ where $g_j = g_1 + (j-1)\delta$, $1 \leq j \leq M$, and $\delta = (g_M - g_1)/(M-1)$. Obtain the quadrature weights $\mathbf{w} = (w_1, \dots, w_M)$ and set $\mathbf{C}_g$ using (17). Initialize: $\mu_{q(1/\sigma_\varepsilon^2)}, \mu_{q(1/\sigma_x^2)} > 0$, $\mu_{q(\mu_x)}$, $\mu_{q(\beta, \mathbf{u})}((K+2) \times 1)$, $\Sigma_{q(\beta, \mathbf{u})}((K+2) \times (K+2))$, $\mu_{q(\phi)}(2 \times 1)$, $\Sigma_{q(\phi)}(2 \times 1)$ and $\mu_{q(a)}(n \times 1)$. Cycle:

update $\Lambda_{\text{mis}, i}$, $1 \leq i \leq n_{\text{mis}}$, using (18) and (19)

$$\mathbf{Q}_g \leftarrow [\exp(-\frac{1}{2} \mathbf{C}_{g_j} \Lambda_{\text{mis}, i} \mathbf{C}_{g_j}^T)]_{1 \leq i \leq n_{\text{mis}} \atop 1 \leq j \leq M} \quad ; \quad E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{C}) \leftarrow \begin{bmatrix} \mathbf{C}_{\text{obs}} \\ \frac{\mathbf{Q}_g \text{diag}(\mathbf{w}) \mathbf{C}_g}{\mathbf{1}^T \otimes (\mathbf{Q}_g \mathbf{w})} \end{bmatrix}$$

for $i = 1, \dots, n_{\text{mis}}$:

$$\mu_{q(x_{\text{mis}, i})} \leftarrow \{E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{C}_{\text{mis}})\}_{i2} \quad ; \quad \sigma_{q(x_{\text{mis}, i})}^2 \leftarrow \frac{\mathbf{Q}_g \text{diag}(\mathbf{w}) (\mathbf{g} - \mu_{q(x_{\text{mis}, i})} \mathbf{1})^2}{\mathbf{1}^T \otimes (\mathbf{Q}_g \mathbf{w})}$$

$$E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{C}^T \mathbf{C}) \leftarrow \mathbf{C}_{\text{obs}}^T \mathbf{C}_{\text{obs}} + \mathbf{C}_g^T \text{diag} \left(\sum_{i=1}^{n_{\text{mis}}} \frac{(\mathbf{e}_i^T \mathbf{Q}_g) \odot \mathbf{w}}{\mathbf{e}_i^T \mathbf{Q}_g \mathbf{w}} \right) \mathbf{C}_g$$

$$\Sigma_{q(\beta, \mathbf{u})} \leftarrow \left\{ \mu_{q(1/\sigma_\varepsilon^2)} E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{C}^T \mathbf{C}) + \frac{1}{\sigma_\beta^2} \mathbf{I} \right\}^{-1} \quad ; \quad \mu_{q(\beta, \mathbf{u})} \leftarrow \Sigma_{q(\beta, \mathbf{u})} \mu_{q(1/\sigma_\varepsilon^2)} E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{C})^T \mathbf{y}$$

$$\sigma_{q(\mu_x)}^2 \leftarrow 1 / (n \mu_{q(1/\sigma_x^2)} + 1/\sigma_x^2) \quad ; \quad \mu_{q(\mu_x)} \leftarrow \sigma_{q(\mu_x)}^2 \mu_{q(1/\sigma_x^2)} (\mathbf{1}^T \mathbf{x}_{\text{obs}} + \mathbf{1}^T \mu_{q(\mathbf{x}_{\text{mis}})})$$

$$B_{q(\sigma_\varepsilon^2)} \leftarrow B_\varepsilon + \frac{1}{2} \|\mathbf{y}\|^2 - \mathbf{y}^T E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{C}) \mu_{q(\beta, \mathbf{u})} + \frac{1}{2} \text{tr} \{ E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{C}^T \mathbf{C}) (\Sigma_{q(\beta, \mathbf{u})} + \mu_{q(\beta, \mathbf{u})} \mu_{q(\beta, \mathbf{u})}^T) \}$$

$$B_{q(\sigma_u^2)} \leftarrow B_u + \frac{1}{2} \{ \|\mu_{q(\mathbf{u})}\|^2 + \text{tr}(\Sigma_{q(\mathbf{u})}) \}$$

$$B_{q(\sigma_x^2)} \leftarrow B_x + \frac{1}{2} \left(\|\mathbf{x}_{\text{obs}} - \mu_{q(\mu_x)} \mathbf{1}\|^2 + \|\mu_{q(\mathbf{x}_{\text{mis}})} - \mu_{q(\mu_x)} \mathbf{1}\|^2 + n \sigma_{q(\mu_x)}^2 + \sum_{i=1}^{n_{\text{mis}}} \sigma_{q(x_{\text{mis}, i})}^2 \right)$$

$$\mu_{q(1/\sigma_\varepsilon^2)} \leftarrow (A_\varepsilon + \frac{1}{2}n)/B_{q(\sigma_\varepsilon^2)} \quad ; \quad \mu_{q(1/\sigma_x^2)} \leftarrow (A_x + \frac{1}{2}n)/B_{q(\sigma_x^2)} \quad ; \quad \mu_{q(1/\sigma_u^2)} \leftarrow (A_u + \frac{1}{2}K)/B_{q(\sigma_u^2)}$$

$$\Sigma_{q(\phi)} \leftarrow \{ E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X}) + \frac{1}{\sigma_\phi^2} \mathbf{I} \}^{-1} \quad ; \quad \mu_{q(\phi)} \leftarrow \Sigma_{q(\phi)} E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X})^T \mu_{q(a)}$$

$$\mu_{q(a)} \leftarrow E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \mu_{q(\phi)} + (2\mathbf{R} - \mathbf{1}) \odot \frac{(2\pi)^{-1/2} \exp\{-\frac{1}{2} (E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \mu_{q(\phi)})^2\}}{\Phi((2\mathbf{R} - \mathbf{1}) \odot (E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \mu_{q(\phi)}))}$$

until the increase in $p(\mathbf{y}; q)$ is negligible.

Algorithm 4: Iterative scheme for obtaining the parameters in the optimal densities $q^*(\beta, \mathbf{u})$, $q^*(\sigma_\varepsilon^2)$, $q^*(\sigma_u^2)$, $q^*(\mu_x)$, $q^*(\sigma_x^2)$, $q^*(x_{\text{mis}, i})$ and $q^*(\phi)$ for the MNAR nonparametric regression model.

form of MCMC (Neal, 1998). In BRugs this involves setting `overRelax=TRUE` in the `modelUpdate()` function. Using these settings, all chains appeared to behave reasonably well.

The resulting posterior densities for the model parameters and four randomly chosen missing x_i values are shown in Figure 7. The vertical lines correspond to the true values,

except σ_u^2 where ‘truth’ is not readily defined. Good to excellent accuracy of variational Bayes is apparent for all posterior densities. There is some noticeable discordance in the case of σ_u^2 . This is perhaps due to some lack of identifiability for this parameter.

Figure 7: Variational Bayes approximate posteriors for the regression model parameters and four missing x_i s for nonparametric regression with predictors MCAR. The vertical lines correspond to the true values of the parameters from which the data were simulated (described in the text). The MCMC posteriors are based on samples of size 10000 and kernel density estimation. The accuracy values correspond to the definition given at (11).

A novel aspect of this example is the multimodality of the posteriors for the $x_{\text{mis},i}$. This arises from the periodic nature of f , since more than one x conforms with a particular y . It is noteworthy that the variational Bayes approximations are able to handle this multimodality quite well

We then applied Algorithm 4 to data simulated according to

$$y_i \sim N(f(x_i), \sigma_\epsilon^2), \quad f(x) = \sin(4\pi x^2), \quad x_i \sim N\left(\frac{1}{2}, \frac{1}{36}\right) \quad \text{and} \quad \sigma_\epsilon^2 = 0.35, \quad 1 \leq i \leq 500$$

and the observed predictor indicators generated according to

$$R_i \sim \text{Bernoulli}(\Phi(\phi_0 + \phi_1 x_i)) \quad \text{with } \phi_0 = 3 \text{ and } \phi_1 = -3.$$

The hyperparameters were as in (14) and $\sigma_\phi^2 = 10^{-8}$. We also ran an MCMC analysis using **BRugs**. The spline basis functions and MCMC sample sizes were the same as those used in the MCAR example. Figure 8 shows resulting posterior density functions. As with the parametric regression examples, variational Bayes is seen to have good to excellent performance for all parameters except ϕ_0 and ϕ_1 .

Our last example involves two variables from **Ozone** data-frame (source: Breiman & Friedman, 1985) in the R package **mlbench** (Leisch & Dimitriadou, 2009). The response variable is daily maximum one-hour-average ozone level and the predictor variable is daily temperature (degrees Fahrenheit) at El Monte, California, USA. The **Ozone** data-frame is such that 5 of the response values are missing and 137 of the predictor values are missing. So that we could apply the methodology of the current section, directly we omitted the 5 records for which the response was missing. This resulted in a sample size of $n = 361$ with $n_{\text{mis}} = 137$ missing predictor values.

Preliminary checks shown the normality assumption for the predictors and errors, along with homoscedasticity, to be quite reasonable. We then assumed MNAR nonparametric regression model and fed the stardized data into to Algorithm 4. MCMC fitting of the same model via **BRugs** was also done for comparison. The results were then transformed to the original scale. Figure 9 shows resulting posterior density functions approximations.

Figure 8: Variational Bayes approximate posteriors for the regression model parameters and five missing x_i s for nonparametric regression with predictors MNAR. The vertical lines correspond to the true values of the parameters from which the data were simulated (described in the text). The MCMC posteriors are based on samples of size 10000 and kernel density estimation. The accuracy values correspond to the definition given at (11).

In Figure 10 the fitted function estimates for all three examples are shown. Good agreement is seen between variational Bayes and MCMC.

Finally, it is worth noting that these three penalized spline examples had much bigger speed increases for variational Bayes compared with MCMC in BUGS. The total elapsed time for the variational Bayes analysis was 75 seconds. For BRugs, with the MCMC sample sizes described above, the three examples required 15.5 hours to run. This corresponds to a speed-up in the order of several hundreds.

5 Discussion

We have derived variational Bayes algorithms for fast approximate inference in parametric and nonparametric regression with missing predictor data. The central finding of this paper is that, for using regression models with missing predictor data, variational Bayes inference achieves good to excellent accuracy for the main parameters of interest. Poor accuracy is realized for the missing data mechanism parameters. As we note at the end of Section 3.4.2, better accuracy for these auxiliary parameters maybe achievable with a more elaborate variational scheme – in situations where they are of interest. The nonparametric regression examples illustrate that variational Bayes approximates multi-modal posterior densities with a high degree of accuracy.

The article has been confined to single predictor models so that the main ideas could be maximally elucidated. Numerous extensions could be made relatively straightforwardly, based on the methodology developed here. Examples include: missing response data, multiple regression and additive models and additive mixed models.

Figure 9: Variational Bayes approximate posteriors for the regression model parameters and four missing x_i s for nonparametric regression applied to the ozone data with predictors MNAR. The MCMC posteriors are based on samples of size 10000 and kernel density estimation. The accuracy values correspond to the definition given at (11). Summary of nonparametric regression for ozone data with with predictor MNAR.

Figure 10: Posterior mean functions and corresponding pointwise 95% credible sets for all three nonparametric regression examples. The grey curves correspond to MCMC-based inference, whilst the black curves correspond to variational Bayesian inference.

Appendix A: Derivation of Algorithms 1, 2 and 3

The derivations for Algorithms 1, 2 and 3 are similar to each other. Hence, we only give the full derivation for the most complicated one – Algorithm 3. Adjustments for Algorithms 1 and 2 are relatively simple.

The full conditionals

As a first step, we determine the full conditionals of each node in the right-most DAG of Figure 1. These involve manipulations that have become standard in the Gibbs sampling

literature (e.g. Robert & Casella, 2004). Note that the calculations can be made simpler by appealing to the Markov blanket theorem and the content of Table 1. For example,

$$\begin{aligned} p(\beta|\text{rest}) &= p(\beta|\text{Markov blanket of } \beta) = p(\beta|\mathbf{y}, \sigma_\varepsilon^2, \mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}}) \\ &\propto p(\mathbf{y}|\beta, \sigma_\varepsilon^2, \mathbf{x}_{\text{mis}}, \mathbf{x}_{\text{obs}})p(\beta) \propto \exp\left(-\frac{1}{2\sigma_\varepsilon^2}\|\mathbf{y} - \mathbf{X}\beta\|^2 - \frac{1}{2\sigma_\beta^2}\|\beta\|^2\right). \end{aligned}$$

Algebraic manipulation of the exponent leads to

$$p(\beta|\text{rest}) \sim N\left(\left(\frac{1}{\sigma_\varepsilon^2}\mathbf{X}^T\mathbf{X} + \frac{1}{\sigma_\beta^2}\mathbf{I}\right)^{-1}\frac{1}{\sigma_\varepsilon^2}\mathbf{X}^T\mathbf{y}, \left(\frac{1}{\sigma_\varepsilon^2}\mathbf{X}^T\mathbf{X} + \frac{1}{\sigma_\beta^2}\mathbf{I}\right)^{-1}\right).$$

Continuing in this fashion we obtain:

$$\begin{aligned} p(\sigma_\varepsilon^2|\text{rest}) &\sim \text{IG}(A_\varepsilon + \tfrac{1}{2}n, B_\varepsilon + \tfrac{1}{2}\|\mathbf{y} - \mathbf{X}\beta\|^2), \\ p(\mu_x|\text{rest}) &\sim N\left(\frac{(\mathbf{1}^T\mathbf{x})/\sigma_x^2}{n/\sigma_x^2 + 1/\sigma_{\mu_x}^2}, \frac{1}{n/\sigma_x^2 + 1/\sigma_{\mu_x}^2}\right), \\ p(\sigma_x^2|\text{rest}) &\sim \text{IG}(A_x + \tfrac{1}{2}n, B_x + \tfrac{1}{2}\|\mathbf{x} - \mu_x\mathbf{1}\|^2), \\ p(x_{\text{mis}}|\text{rest}) &\sim \text{product of } N\left(\frac{\mu_x/\sigma_x^2 + (\beta_1/\sigma_\varepsilon^2)(y_{x_{\text{mis},i}} - \beta_0) + \phi_1(a_i - \phi_0)}{1/\sigma_x^2 + \beta_1^2/\sigma_\varepsilon^2 + \phi_1^2}, \frac{1}{1/\sigma_x^2 + \beta_1^2/\sigma_\varepsilon^2 + \phi_1^2}\right) \\ &\quad \text{densities over } 1 \leq i \leq n_{\text{mis}}, \\ p(\phi|\text{rest}) &\sim N\left((\mathbf{X}^T\mathbf{X} + \frac{1}{\sigma_\phi^2}\mathbf{I})^{-1}\mathbf{X}^T\mathbf{y}, (\mathbf{X}^T\mathbf{X} + \frac{1}{\sigma_\phi^2}\mathbf{I})^{-1}\right) \end{aligned}$$

and
$$p(\mathbf{a}|\text{rest}) = (2\pi)^{-n/2} \prod_{i=1}^n \left[\left\{ \frac{I(a_i \geq 0)}{\Phi((\mathbf{X}\phi)_i)} \right\}^{R_i} \left\{ \frac{I(a_i < 0)}{1 - \Phi((\mathbf{X}\phi)_i)} \right\}^{1-R_i} \exp\left\{-\frac{1}{2}(a_i - \phi_0 - \phi_1 x_i)^2\right\} \right].$$

Expressions for $E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X})$ and $E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T\mathbf{X})$

The correspond to the updates given at (8).

Derivation:

The results for $E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X})$ and the (1,1), (1,2) and (2,1) entries of $E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T\mathbf{X})$ follow from

$$E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{x}) = \begin{bmatrix} \mathbf{x}_{\text{obs}} \\ E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{x}_{\text{mis}}) \end{bmatrix} = \begin{bmatrix} \mathbf{x}_{\text{obs}} \\ \boldsymbol{\mu}_{q(\mathbf{x}_{\text{mis}})} \end{bmatrix}.$$

The (2,2) entry of $E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T\mathbf{X})$ is

$$E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{x}^T\mathbf{x}) = \{E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{x})\}^T E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{x}) + \text{tr}\{\text{Cov}_q(\mathbf{x})\} = \|\mathbf{x}_{\text{obs}}\|^2 + \|\boldsymbol{\mu}_{q(\mathbf{x}_{\text{mis}})}\|^2 + \sum_{i=1}^{n_{\text{mis}}} \text{Var}_q(x_{\text{mis},i}).$$

The stated result then follows from the fact that $\text{Var}_q(x_{\text{mis},i}) = \sigma_{q(\mathbf{x}_{\text{mis}})}^2$ for each $1 \leq i \leq n_{\text{mis}}$.

Expression for $q^*(\beta)$

$$q^*(\beta) \sim N(\boldsymbol{\mu}_{q(\beta)}, \boldsymbol{\Sigma}_{q(\beta)})$$

where

$$\boldsymbol{\Sigma}_{q(\beta)} = \left\{ \mu_{q(1/\sigma_\varepsilon^2)} E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T\mathbf{X}) + \frac{1}{\sigma_\beta^2}\mathbf{I} \right\}^{-1} \quad \text{and} \quad \boldsymbol{\mu}_{q(\beta)} = \boldsymbol{\Sigma}_{q(\beta)} \mu_{q(1/\sigma_\varepsilon^2)} E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X})^T \mathbf{y}.$$

Derivation:

With $\mathbf{V} \equiv (\frac{1}{\sigma_\varepsilon^2} \mathbf{X}^T \mathbf{X} + \frac{1}{\sigma_\beta^2} \mathbf{I})^{-1}$ we have

$$\begin{aligned} q^*(\boldsymbol{\beta}) &\propto \exp\{E_{-\boldsymbol{\beta}} \log p(\boldsymbol{\beta}|\text{rest})\} \\ &\propto \exp\left(E_{q(\sigma_\varepsilon^2)q(\mathbf{x}_{\text{mis}})}\left\{-\frac{1}{2}[\{\boldsymbol{\beta} - \mathbf{V}(1/\sigma_\varepsilon^2)\mathbf{X}^T \mathbf{y}\}^T \mathbf{V}^{-1}\{\boldsymbol{\beta} - \mathbf{V}(1/\sigma_\varepsilon^2)\mathbf{X}^T \mathbf{y}\}]\right\}\right) \\ &\propto \exp\left(E_{q(\sigma_\varepsilon^2)q(\mathbf{x}_{\text{mis}})}\left\{-\frac{1}{2}[\boldsymbol{\beta}^T\{(1/\sigma_\varepsilon^2)\mathbf{X}^T \mathbf{X} + (1/\sigma_\beta^2)\}\boldsymbol{\beta} - 2\boldsymbol{\beta}^T(1/\sigma_\varepsilon^2)\mathbf{X}^T \mathbf{y}]\right\}\right) \\ &\propto \exp\left(-\frac{1}{2}[\boldsymbol{\beta}^T\{\mu_{q(1/\sigma_\varepsilon^2)}E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X}) + (1/\sigma_\beta^2)\}\boldsymbol{\beta} - 2\boldsymbol{\beta}^T\mu_{q(1/\sigma_\varepsilon^2)}E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X})^T \mathbf{y}]\right). \end{aligned}$$

The result then follows after completion of the square in the exponent.

Expression for $q^*(\sigma_\varepsilon^2)$

$$q^*(\sigma_\varepsilon^2) \sim \text{IG}(A_\varepsilon + \frac{1}{2}n, B_{q(\sigma_\varepsilon^2)})$$

where

$$B_{q(\sigma_\varepsilon^2)} = B_\varepsilon + \frac{1}{2}\|\mathbf{y}\|^2 - \mathbf{y}^T E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X})\mu_{q(\boldsymbol{\beta})} + \frac{1}{2}\text{tr}\{E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X})(\boldsymbol{\Sigma}_{q(\boldsymbol{\beta})} + \boldsymbol{\mu}_{q(\boldsymbol{\beta})}\boldsymbol{\mu}_{q(\boldsymbol{\beta})}^T)\}.$$

Note that $\mu_{q(1/\sigma_\varepsilon^2)} = (A_x + \frac{1}{2}n)/B_{q(\sigma_\varepsilon^2)}$.

Derivation:

From the expression for $p(\sigma_\varepsilon^2|\text{rest})$, application of (2) leads to

$$q^*(\sigma_\varepsilon^2) \sim \text{IG}(A_\varepsilon + \frac{1}{2}n, B_\varepsilon + \frac{1}{2}E_{q(\mathbf{x}_{\text{mis}})q(\boldsymbol{\beta})}\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2).$$

Noting that

$$\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2 = \|\mathbf{y}\|^2 - 2\mathbf{y}^T \mathbf{X}\boldsymbol{\beta} + \text{tr}(\mathbf{X}^T \mathbf{X}\boldsymbol{\beta}\boldsymbol{\beta}^T).$$

and taking expectations with respect to $q(\mathbf{x}_{\text{mis}})q(\boldsymbol{\beta})$ we obtain the stated result.

Expression for $q^*(\mu_x)$

$$q^*(\mu_x) \sim N(\mu_{q(\mu_x)}, \sigma_{q(\mu_x)}^2)$$

where

$$\sigma_{q(\mu_x)}^2 = 1/(n\mu_{q(1/\sigma_x^2)} + 1/\sigma_{\mu_x}^2) \quad \text{and} \quad \mu_{q(\mu_x)} = \sigma_{q(\mu_x)}^2 \mu_{q(1/\sigma_x^2)}(\mathbf{1}^T \mathbf{x}_{\text{obs}} + \mathbf{1}^T \boldsymbol{\mu}_{q(\mathbf{x}_{\text{mis}})}).$$

Derivation:

This is similar to the derivation of $q^*(\boldsymbol{\beta})$, but is simpler.

Expression for $q^*(\sigma_x^2)$

$$q^*(\sigma_x^2) \sim \text{IG}(A_x + \frac{1}{2}n, B_{q(\sigma_x^2)})$$

where

$$B_{q(\sigma_x^2)} = B_x + \frac{1}{2}(\|\mathbf{x}_{\text{obs}} - \mu_{q(\mu_x)}\mathbf{1}\|^2 + \|\boldsymbol{\mu}_{q(\mathbf{x}_{\text{mis}})} - \mu_{q(\mu_x)}\mathbf{1}\|^2 + n\sigma_{q(\mu_x)}^2 + n_{\text{mis}}\sigma_{q(\mathbf{x}_{\text{mis}})}^2).$$

Note that $\mu_{q(1/\sigma_x^2)} = (A_x + \frac{1}{2}n)/B_{q(\sigma_x^2)}$.

Derivation:

From the result for the full conditional of σ_x^2 :

$$q^*(\sigma_x^2) \sim \text{IG}(A_x + \frac{1}{2}n, B_x + \frac{1}{2}E_q\|\mathbf{x} - \mu_x\mathbf{1}\|^2).$$

To obtain an explicit expression for $E_q\|\mathbf{x} - \mu_x\mathbf{1}\|^2$ first note that

$$\|\mathbf{x} - \mu_x\mathbf{1}\|^2 = \|\mathbf{x}_{\text{obs}} - \mu_x\mathbf{1}\|^2 + \|\mathbf{x}_{\text{mis}} - \mu_x\mathbf{1}\|^2.$$

Then

$$E_q \|\mathbf{x}_{\text{obs}} - \mu_x \mathbf{1}\|^2 = \|\mathbf{x}_{\text{obs}} - \mu_{q(\mu_x)} \mathbf{1}\|^2 + n_{\text{obs}} \sigma_{q(\mu_x)}^2.$$

Also, by independence of $\mathbf{x}_{\text{mis}}$ and μ_x (under the q densities)

$$E_q \|\mathbf{x}_{\text{mis}} - \mu_x \mathbf{1}\|^2 = \|\mu_{q(\mathbf{x}_{\text{mis}})} - \mu_{q(\mu_x)} \mathbf{1}\|^2 + n_{\text{mis}} (\sigma_{q(\mathbf{x}_{\text{mis}})}^2 + \sigma_{q(\mu_x)}^2).$$

Combining we get

$$E_q \|\mathbf{x} - \mu_x \mathbf{1}\|^2 = \|\mathbf{x}_{\text{obs}} - \mu_{q(\mu_x)} \mathbf{1}\|^2 + \|\mu_{q(\mathbf{x}_{\text{mis}})} - \mu_{q(\mu_x)} \mathbf{1}\|^2 + n \sigma_{q(\mu_x)}^2 + n_{\text{mis}} \sigma_{q(\mathbf{x}_{\text{mis}})}^2.$$

The stated result then follows immediately.

Expression for $q^*(x_{\text{mis},i})$

For $1 \leq i \leq n_{\text{mis}}$:

$$q^*(x_{\text{mis},i}) \sim N(\mu_{q(x_{\text{mis},i})}, \sigma_{q(x_{\text{mis}})}^2)$$

where

$$\sigma_{q(x_{\text{mis}})}^2 = 1 / \left[\mu_{q(1/\sigma_x^2)} + \mu_{q(1/\sigma_\varepsilon^2)} \left\{ \mu_{q(\beta_1)}^2 + (\Sigma_{q(\beta)})_{22} \right\} + \mu_{q(\phi_1)}^2 + (\Sigma_{q(\phi)})_{22} \right]$$

and

$$\begin{aligned} \mu_{q(x_{\text{mis},i})} = \sigma_{q(x_{\text{mis}})}^2 & \left[\mu_{q(1/\sigma_x^2)} \mu_{q(\mu_x)} + \mu_{q(1/\sigma_\varepsilon^2)} \left\{ y_{x_{\text{mis},i}} \mu_{q(\beta_1)} - (\Sigma_{q(\beta)})_{12} - \mu_{q(\beta_0)} \mu_{q(\beta_1)} \right\} \right. \\ & \left. + \mu_{q(a_{x_{\text{mis},i}})} \mu_{q(\phi_1)} - (\Sigma_{q(\phi)})_{12} - \mu_{q(\phi_0)} \mu_{q(\phi_1)} \right]. \end{aligned}$$

Derivation:

First note that, since

$$q^*(\mathbf{x}_{\text{mis}}) \propto \exp\{E_{-\mathbf{x}_{\text{mis}}} \log p(\mathbf{x}_{\text{mis}} | \text{rest})\},$$

$q^*(\mathbf{x}_{\text{mis}})$ inherits the product structure of $p(\mathbf{x}_{\text{mis}} | \text{rest})$. Therefore, it suffices to work with $p(x_{\text{mis},i} | \text{rest})$ and $q^*(x_{\text{mis},i})$. Let

$$v \equiv 1 / (1/\sigma_x^2 + \beta_1^2/\sigma_\varepsilon^2 + \phi_1^2).$$

Then

$$\begin{aligned} q^*(x_{\text{mis},i}) & \propto \exp\{E_{-x_{\text{mis},i}} \log p(x_{\text{mis},i} | \text{rest})\} \\ & \propto \exp \left\{ E_{-x_{\text{mis},i}} \left(-\frac{1}{2v} [x_{\text{mis},i} - v\{\mu_x/\sigma_x^2 + (\beta_1/\sigma_\varepsilon^2)(y_{x_{\text{mis},i}} - \beta_0) + \phi_1(a_{x_{\text{mis},i}} - \phi_0)\}]^2 \right) \right\} \\ & \propto \exp \left(-\frac{1}{2} [x_{\text{mis},i}^2 E_q(1/\sigma_x^2 + \beta_1^2/\sigma_\varepsilon^2 + \phi_1^2) \right. \\ & \quad \left. - 2x_{\text{mis},i} E_q\{\mu_x/\sigma_x^2 + (\beta_1/\sigma_\varepsilon^2)(y_{x_{\text{mis},i}} - \beta_0) + \phi_1(a_{x_{\text{mis},i}} - \phi_0)\}] \right) \end{aligned}$$

where the E_q s in the previous expression are assumed to be over the joint q -densities of the respective random variables. The first of these expectations is

$$E_q(1/\sigma_x^2 + \beta_1^2/\sigma_\varepsilon^2 + \phi_1^2) = \mu_{q(1/\sigma_x^2)} + \mu_{q(1/\sigma_\varepsilon^2)} \left\{ \mu_{q(\beta_1)}^2 + (\Sigma_{q(\beta)})_{22} \right\} + \mu_{q(\phi_1)}^2 + (\Sigma_{q(\phi)})_{22}$$

whilst the second one is

$$\begin{aligned} E_q\{\mu_x/\sigma_x^2 + (\beta_1/\sigma_\varepsilon^2)(y_{x_{\text{mis},i}} - \beta_0) + \phi_1(a_{x_{\text{mis},i}} - \phi_0)\} = \\ \mu_{q(1/\sigma_x^2)} \mu_{q(\mu_x)} + \mu_{q(1/\sigma_\varepsilon^2)} \left\{ y_{x_{\text{mis},i}} \mu_{q(\beta_1)} - (\Sigma_{q(\beta)})_{12} - \mu_{q(\beta_0)} \mu_{q(\beta_1)} \right\} \\ + \mu_{q(a_{x_{\text{mis},i}})} \mu_{q(\phi_1)} - (\Sigma_{q(\phi)})_{12} - \mu_{q(\phi_0)} \mu_{q(\phi_1)}. \end{aligned}$$

The result then follows after completion of the square in the exponent.

Expression for $q^*(\phi)$

$$q^*(\phi) \sim N(\boldsymbol{\mu}_{q(\phi)}, \boldsymbol{\Sigma}_{q(\phi)})$$

where

$$\boldsymbol{\Sigma}_{q(\phi)} = \{E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}^T \mathbf{X}) + (1/\sigma_\phi^2) \mathbf{I}\}^{-1} \quad \text{and} \quad \boldsymbol{\mu}_{q(\phi)} = \boldsymbol{\Sigma}_{q(\phi)} E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X})^T \boldsymbol{\mu}_{q(\mathbf{a})}.$$

Derivation:

This is similar to the derivation of $q^*(\beta)$, but is simpler.

Expression for $q^*(\mathbf{a})$ and its mean

$$q^*(\mathbf{a}) = \left[\prod_{i=1}^n \left\{ \frac{I(a_i \geq 0)}{\Phi((\mathbf{X} \boldsymbol{\mu}_{q(\phi)})_i)} \right\}^{R_i} \left\{ \frac{I(a_i < 0)}{1 - \Phi((\mathbf{X} \boldsymbol{\mu}_{q(\phi)})_i)} \right\}^{1-R_i} \right] \times (2\pi)^{-n/2} \exp\left\{-\frac{1}{2} \|\mathbf{a} - \mathbf{X} \boldsymbol{\mu}_{q(\phi)}\|^2\right\}.$$

From this we get

$$\boldsymbol{\mu}_{q(\mathbf{a})} = E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)} + \frac{(2\pi)^{-1/2} (2\mathbf{R} - \mathbf{1}) \odot \exp\left\{-\frac{1}{2} (E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)})^2\right\}}{\Phi((2\mathbf{R} - \mathbf{1}) \odot (E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)}))}.$$

Derivation:

If $R_i = 1$ then

$$\begin{aligned} q^*(a_i) &\propto \exp E_{q(\mathbf{x}_{\text{mis}}, \phi)}\left[-\frac{1}{2} \{a_i - (\mathbf{X} \phi)_i\}^2\right], \quad a_i \geq 0 \\ &\propto \exp\left(-\frac{1}{2} \left[a_i - \{E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)}\}_i\right]^2\right), \quad a_i \geq 0. \end{aligned}$$

Hence, if $R_i = 1$, then

$$q^*(a_i) = \frac{(2\pi)^{-1/2} \exp\left(-\frac{1}{2} [a_i - \{E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)}\}_i]^2\right)}{\Phi(\{E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)}\}_i)}, \quad a_i \geq 0.$$

Similarly, if $R_i = 0$, then

$$q^*(a_i) = \frac{(2\pi)^{-1/2} \exp\left(-\frac{1}{2} [a_i - \{E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)}\}_i]^2\right)}{1 - \Phi(\{E_{q(\mathbf{x}_{\text{mis}})}(\mathbf{X}) \boldsymbol{\mu}_{q(\phi)}\}_i)}, \quad a_i < 0.$$

The expression for $q^*(\mathbf{a})$ is obtained after applying these results to each factor in $p(\mathbf{a}|\text{rest})$.

The $\boldsymbol{\mu}_{q(\mathbf{a})}$ expression follows from the following results concerning means of truncated normal density functions:

$$\int_0^\infty x \frac{(2\pi)^{-1/2} \exp\{-\frac{1}{2}(x - \mu)^2\}}{\Phi(\mu)} dx = \mu + \frac{(2\pi)^{-1/2} e^{-\mu^2/2}}{\Phi(\mu)}$$

and

$$\int_{-\infty}^0 x \frac{(2\pi)^{-1/2} \exp\{-\frac{1}{2}(x - \mu)^2\}}{1 - \Phi(\mu)} dx = \mu - \frac{(2\pi)^{-1/2} e^{-\mu^2/2}}{\Phi(-\mu)}.$$

Appendix B: Derivation of (16) and Expression for $\Lambda_{\text{mis},i}$

First note that $p(x_{\text{mis},i}|\text{rest})$ satisfies

$$p(x_{\text{mis},i}|\text{rest}) \propto \exp \left\{ -\frac{1}{2}\sigma_\varepsilon^{-2}\mathbf{C}_{x_{\text{mis},i}} \begin{bmatrix} \beta \\ \mathbf{u} \end{bmatrix} \begin{bmatrix} \beta \\ \mathbf{u} \end{bmatrix}^T \mathbf{C}_{x_{\text{mis},i}} + \sigma_\varepsilon^{-2}y_{x_{\text{mis},i}} \mathbf{C}_{x_{\text{mis},i}} \begin{bmatrix} \beta \\ \mathbf{u} \end{bmatrix} \right. \\ \left. -\frac{1}{2}\sigma_x^{-2}x_{\text{mis},i}^2 + \sigma_x^{-2}\mu_x x_{\text{mis},i} - \frac{1}{2}\phi_1^2 x_{\text{mis},i}^2 + (a_{x_{\text{mis},i}}\phi_1 - \phi_0\phi_1) x_{\text{mis},i} \right\}$$

where, for $1 \leq i \leq n_{\text{mis}}$, $a_{x_{\text{mis},i}}$ is the entry of $\mathbf{a}$ corresponding to $x_{\text{mis},i}$. The expression for $q(x_{\text{mis},i})$ is then given by

$$q(x_{\text{mis},i}) \propto \exp \left\{ -\frac{1}{2}\mu_{q(1/\sigma_\varepsilon^2)} \mathbf{C}_{x_{\text{mis},i}} (\boldsymbol{\mu}_{q(\beta,\mathbf{u})} \boldsymbol{\mu}_{q(\beta,\mathbf{u})}^T + \boldsymbol{\Sigma}_{q(\beta,\mathbf{u})}) \mathbf{C}_{x_{\text{mis},i}}^T + \mu_{q(1/\sigma_\varepsilon^2)} y_{x_{\text{mis},i}} \mathbf{C}_{x_{\text{mis},i}} \boldsymbol{\mu}_{q(\beta,\mathbf{u})} \right. \\ \left. -\frac{1}{2}\mu_{q(1/\sigma_x^2)} x_{\text{mis},i}^2 + \mu_{q(1/\sigma_x^2)} \mu_q(\mu_x) x_{\text{mis},i} - \frac{1}{2}(\mu_{q(\phi_1)}^2 + \{\boldsymbol{\Sigma}_{q(\phi)}\}_{22}) x_{\text{mis},i}^2 \right. \\ \left. + \left(\mu_{q(a_{x_{\text{mis},i}})} \mu_{q(\phi_1)} - \mu_{q(\phi_0)} \mu_{q(\phi_1)} - \{\boldsymbol{\Sigma}_{q(\phi)}\}_{12} \right) x_{\text{mis},i} \right\}.$$

Expression (16) then follows, using the $\mathbf{C}_x$ notation and the definitions

$$\Lambda_{\text{mis},i} \equiv \Lambda_1 + \Lambda_{2,\text{mis},i} + \Lambda_{3,\text{mis},i} \quad (18)$$

with

$$\Lambda_1 \equiv \mu_{q(1/\sigma_\varepsilon^2)} (\boldsymbol{\mu}_{q(\beta,\mathbf{u})} \boldsymbol{\mu}_{q(\beta,\mathbf{u})}^T + \boldsymbol{\Sigma}_{q(\beta,\mathbf{u})}),$$

$$\Lambda_{2,\text{mis},i} \equiv \begin{bmatrix} 0 & -\mu_{q(1/\sigma_x^2)} \mu_q(\mu_x) & -\mu_{q(1/\sigma_x^2)} y_{x_{\text{mis},i}} \boldsymbol{\mu}_{q(\mathbf{u})}^T \\ -\mu_{q(1/\sigma_\varepsilon^2)} y_{x_{\text{mis},i}} \boldsymbol{\mu}_{q(\beta_1)} & \mu_{q(1/\sigma_x^2)} & 0 \\ -\mu_{q(1/\sigma_\varepsilon^2)} y_{x_{\text{mis},i}} \boldsymbol{\mu}_{q(\mathbf{u})} & 0 & 0 \end{bmatrix}$$

$$\text{and } \Lambda_{3,\text{mis},i} \equiv \begin{bmatrix} 0 & \mu_{q(\phi_0)} \mu_{q(\phi_1)} + \{\boldsymbol{\Sigma}_{q(\phi)}\}_{12} & 0 \\ \mu_{q(\phi_0)} \mu_{q(\phi_1)} + \{\boldsymbol{\Sigma}_{q(\phi)}\}_{12} & \mu_{q(\phi_1)}^2 + \{\boldsymbol{\Sigma}_{q(\phi)}\}_{22} & 0 \\ -\mu_{q(a_{x_{\text{mis},i}})} \mu_{q(\phi_1)} & 0 & 0 \end{bmatrix}. \quad (19)$$

References

- Albert, J.H. & Chib, S. (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, **88**, 669–679.
- Bishop, C.M. (2006). *Pattern Recognition and Machine Learning*. New York: Springer.
- Box, G.P. & Tiao, G.C. (1973). *Bayesian Inference in Statistical Analysis*. Reading, Massachusetts: Addison Wesley.
- Breiman, L. & Friedman, J.H. (1985). Estimating optimal transformations for multiple regression and correlation. *Journal of the American Statistical Association*, **80**, 580–598.

- Crainiceanu, C., Ruppert, D. & Wand, M.P. (2005). Bayesian analysis for penalized spline regression using WinBUGS. *Journal of Statistical Software*, Volume 14, Article 14, 1–24.
- Daniels, M.J. & Hogan, J.W. (2008). *Missing Data in Longitudinal Studies: Strategies for Bayesian Modeling and Sensitivity Analysis*. Boca Raton, Florida: Chapman and Hall (CRC Press).
- Denison, D., Holmes, C., Mallick, B. & Smith, A. (2002). *Bayesian methods for nonlinear classification and regression*. Chichester, UK: Wiley.
- Devroye, L. & Györfi, L. (1985). *Density Estimation: The L_1 View*. New York: Wiley.
- Flandin, G. & Penny, W.D. (2007). Bayesian fMRI data analysis with sparse spatial basis function priors. *NeuroImage*, **34**, 1108–1125.
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis*, **1**, 515–533.
- Gelman, A., Carlin, J.B., Stern, H.S. & Rubin, D.B. (2004). *Bayesian Data Analysis*. Boca Raton, Florida: Chapman and Hall.
- Girolami, M. & Rogers, S. (2006). Variational Bayesian multinomial probit regression. *Neural Computation*, **18**, 1790–1817.
- Gurrin, L.C., Scurrah, K.J. & Hazelton, M.L. (2005). Tutorial in biostatistics: spline smoothing with linear mixed models. *Statistics in Medicine*, **24**, 3361–3381.
- Hall, P. (1987). On Kullback-Leibler loss and density estimation. **15**, 1491–1519.
- Holmes, C.C. & Held, L. (2006). Bayesian auxiliary variable models for binary and multinomial regression. *Bayesian Analysis*, **1**, 145–168.
- Jordan, M.I. (2004). Graphical models. *Statistical Science*, **19**, 140–155.
- Leisch, F. & Dimitriadou, E. (2009). mlbench 1.1–6. Machine Learning Benchmark Problems. R package. <http://cran.r-project.org>
- Ligges, U., Thomas, A., Spiegelhalter, D., Best, N., Lunn, D., Rice, K. & Sturtz, S. (2010). BRugs 0.5 OpenBUGS and its R/S-PLUS interface BRugs. R package. <http://www.stats.ox.ac.uk/pub/RWin/bin/windows/contrib/2.10>
- Little, R.J. & Rubin, D.B. (2004). *Statistical Analysis with Missing Data, Second Edition*. New York: John Wiley & Sons.
- Lunn, D.J., Thomas, A., Best, N. & Spiegelhalter, D. (2000). WinBUGS – a Bayesian modelling framework: concepts, structure, and extensibility. *Statistics and Computing*, **10**, 325–337.
- McGrory, C.A. & Titterton, D.M. (2007). Variational approximations in Bayesian model selection for finite mixture distributions. *Computational Statistics and Data Analysis*, **51**, 5352–5367.
- Neal, R. (1998). Suppressing random walks in Markov chain Monte Carlo using ordered over-relaxation. *Learning in Graphical Models*, Jordan, M.I. Ed. Dordrecht: Kluwer

Academic Publishers, pp. 205–230.

Ormerod, J.T. & Wand, M.P. (2010). Explaining variational approximations. *The American Statistician*, in press.

Parisi, G. (1988). *Statistical Field Theory*. Redwood City, California: Addison-Wesley.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. San Mateo, California: Morgan Kaufmann.

R Development Core Team (2010). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, <http://www.R-project.org>

Robert, C.P. & Casella, G. (2004). *Monte Carlo Statistical Methods, Second Edition*. New York: Springer-Verlag.

Ruppert, D., Wand, M.P. & Carroll, R.J. (2003). *Semiparametric Regression*. New York: Cambridge University Press.

Teschendorff, A.E., Wang, Y., Barbosa-Morais, N.L., Brenton, J.D. & Caldas C. (2005). A variational Bayesian mixture modelling framework for cluster analysis of gene-expression data. *Bioinformatics*, **21**, 3025–3033.

Wahba, G. (1978). Improper priors, spline smoothing and the problem of guarding against model errors in regression. *Journal of the Royal Statistical Society, Series B*, **40**, 364–372.

Wand, M.P. (2009). Semiparametric regression and graphical models. *Australian and New Zealand Journal of Statistics*, **51**, 9–41.

Wand, M.P. & Ormerod, J.T. (2008). On O’Sullivan penalised splines and semiparametric regression. *Australian and New Zealand Journal of Statistics*, **50**, 179–198.

Wand, M.P. & Ripley, B.D. (2009). KernSmooth 2.23. Functions for kernel smoothing corresponding to the book: Wand, M.P. & Jones, M.C. (1995) “Kernel Smoothing”. R package. <http://cran.r-project.org>

Wasserman, L. (2004). *All of Statistics*. New York: Springer.