University of Wollongong

Research Online

Centre for Statistical & Survey Methodology
Working Paper Series

Faculty of Engineering and Information
Sciences

2008

Semiparametric regression and graphical models

M. P. Wand

University of Wollongong, mwand@uow.edu.au

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

Recommended Citation

Wand, M. P., Semiparametric regression and graphical models, Centre for Statistical and Survey Methodology, University of Wollongong, Working Paper 18-08, 2008, 32p.
<https://ro.uow.edu.au/cssmwp/15>

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

18-08

Semiparametric regression and graphical models

Matt P. Wand

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

Semiparametric regression and graphical models

BY M. P. WAND¹

University of Wollongong

1st August, 2008

SUMMARY

Semiparametric regression models that use spline basis functions with penalisation have graphical models representations. This link is more powerful than previously established mixed model representations of semiparametric regression since a bigger class of models can be accommodated. Complications such as missingness and measurement error are more naturally handled within the graphical models architecture. Directed acyclic graphs, also known as Bayesian networks, play a prominent role. Graphical models-based Bayesian “inference engines”, such as BUGS and VIBES, facilitate fitting and inference. Underlying these are Markov chain Monte Carlo schemes and recent developments in variational approximation theory and methodology.

Keywords: Additive models; Bayesian networks; BUGS; Directed acyclic graphs; Markov chain Monte Carlo; Measurement error models; Missing data; Mixed models; Penalised splines; Variational approximation; Variational inference; VIBES.

1 Introduction

The main thrust of two of my publications from five years ago, Ruppert, Wand & Carroll (2003) and Wand (2003), is that *mixed models* are a very useful framework for carrying out *semiparametric regression* analyses. The thrust of this paper is that the more general *graphical models* framework is also very useful for semiparametric regression, especially when the problem is *non-standard*.

Semiparametric regression is an embellishment of parametric regression that uses *penalised spline* basis functions to achieve greater flexibility. Ruppert *et al.* (2003) survey the field up to about 2002. The mixed model aspects of semiparametric regression, and antecedents such as *smoothing splines*, have been known for some time (e.g. Wahba, 1978). However, the advent of formal mixed model software in the 1990s led to a surge in research on mixed model approaches to semiparametric regression, mainly in the last decade. A recent survey of semiparametric regression for the period 2003–2007, to be published as Ruppert, Wand & Carroll (2008), revealed more than 150 research articles making use of the mixed model-based semiparametric regression. Sophisticated semiparametric regression analyses are now being routinely carried out with the ‘work’ being done by established mixed model software such as `lme()` and `PROC MIXED`, or with BUGS if a Bayesian approach is adopted. Recent examples are Harezlak, Ryan, Giedd & Lange (2005) and Crainiceanu, Diggle & Rowlingson (2008).

In many applied situations, however, complications such as non-Gaussian response and missingness prevent the use of standard mixed models methodology and software. I will argue that *graphical models* are a better vehicle for fitting and inference in this case. Software for graphical models is less mature but, in 2008, some reasonable options exist. In this article I make use of BUGS (Lunn *et al.* 2000) and VIBES (Bishop, Spiegelhalter & Winn, 2003); each of which are *Bayesian inference engines* built upon graphical models architecture. Methodology and software of this type is an ongoing active area of research.

¹School of Mathematics and Applied Statistics, University of Wollongong, Wollongong 2522, Australia.
e-mail: mwand@uow.edu.au

Graphical models is currently a very vibrant area of research, although much of it is taking place in Computer Science rather than Statistics. The recent book *Pattern Recognition and Machine Learning* by C.M. Bishop (2006) states that “graphical models have emerged as a general framework for describing and applying probabilistic models” in the areas of Machine Learning and Pattern Recognition.

The central thesis of the present article is that (non-standard) semiparametric regression can be embedded in graphical models architecture and benefit from ongoing graphical models research. There is also the potential for new applications for semiparametric regression in areas of research that are intrinsically graphical models-based. Some examples are causal inference, social networks and phylogenetic trees.

The proposed marriage of semiparametric regression and graphical models is in keeping with a current general trend, that is seeing ideas being exchanged between Statistics and Computer Science much more freely than in earlier days of each discipline. The foreword of a recent special issue of *Statistical Science* on Bayesian statistics describes “the dissolving of the frontier between Statistics and Computer Science” (Casella & Robert, 2004). The special issue contains two review articles, Jordan (2004) and Titterton (2004), of recent Computer Science literature involving Bayesian statistics. Each of these have been strong influences on the present article. In 2006 *Statistica Sinica* had a special issue titled *Challenges in Statistical Machine Learning*.

Semiparametric regression has already benefited from other areas of Computer Science. One spectacular example is boosting (Schapire, 1990; Freund; 1995; Freund & Schapire; 1996). Tutz & Binder (2006) describe the evolution from boosting as a means to improve classification procedures to a powerful tool for semiparametric regression analysis and provide relevant references. Also see Buehlmann & Hothorn (2007) and accompanying discussion. Kernel machine research (e.g. Schölkopf & Smola, 2002) is another area for which there is a great deal of common ground; see the recent Statistics articles by Pearce & Wand (2006), Wahba (2006) and Hastie & Zhu (2006). To date, there seems to have been very little interplay between graphical models and semiparametric regression. A rare example of such interplay is Liang, Truong & Wong (2001) who use graphical models in their Bayesian nonparametric regression procedure.

Section 2 summarises semiparametric regression, focussing on mixed model and hierarchical Bayesian representations. In Section 3 a brief summary of graphical models is provided. A graphical models viewpoint of semiparametric regression is put forward in Section 4. Section 5 then pays special attention to non-standard variants of semiparametric regression. Both Sections 4 and 5 work with Bayesian inference engines based on Markov chain Monte Carlo and BUGS software. Section 6 describes an alternative type of inference engine, based on *variational approximation*. A case study involving relative cancer mapping, when some auxiliary data are missing, is described in Section 7. In Section 8 I add some brief discussion on what might be potentially new areas of application for semiparametric regression, in light of this article’s central thesis. Concluding remarks are given in Section 9.

1.1 Notation and Conventions

Column vectors with entries consisting of subscripted variables are denoted by a bold-faced version of the letter for that variable. For example, the vector containing $x_1, \dots, x_n$ is denoted by $\mathbf{x}$. Scalar functions applied to vectors are evaluated element-wise. For example, $\tanh(a_1, a_2, a_3) = (\tanh(a_1), \tanh(a_2), \tanh(a_3))$.

The density function of a random vector $\mathbf{x}$ is denoted by $[\mathbf{x}]$. The conditional density of $\mathbf{y}$ given $\mathbf{x}$ is denoted by $[\mathbf{y}|\mathbf{x}]$. 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 $[x] \sim B^A \Gamma(A)^{-1} x^{-A-1} e^{-B/x}$, $x > 0$. For a general random vector $\mathbf{v}$, $\mathbf{v} \sim (\boldsymbol{\mu}, \boldsymbol{\Sigma})$ is shorthand for $E(\mathbf{v}) = \boldsymbol{\mu}$ and $\text{Cov}(\mathbf{v}) = \boldsymbol{\Sigma}$, the covariance matrix of $\mathbf{v}$. If, for $1 \leq i \leq n$, y_i has

distribution D_i and the y_i are independent then I write $y_i \stackrel{\text{ind.}}{\sim} D_i$.

There are several directed graphs in this article. I use the same conventions as Bishop (2006). Random nodes are denoted by open circles. Non-random nodes are shown as small solid circles. Observed (“evidence”) nodes are distinguished from parameter (“hidden”) nodes using shading.

All Bayesian models are fitted using standardised versions of continuous variables. Unless otherwise stated, Markov chain Monte Carlo examples use a burnin period of 5000 iterations and then retain 5000 iterations. They are then thinned by a factor of 5, resulting in samples of size 1000 being retained for inference.

2 Semiparametric Regression

Three examples of semiparametric regression models are

$$[y_i] \stackrel{\text{ind.}}{\sim} \text{Bernoulli}[\text{logit}^{-1}\{\beta_1 x_{1i} + f_2(x_{2i}) + f_{34}(x_{3i}, x_{4i})\}], \quad 1 \leq i \leq n, \quad (1)$$

$$[y_i] \stackrel{\text{ind.}}{\sim} \text{Poisson}[\exp\{\beta_0(x_{1i}) + \beta_1(x_{1i})x_{2i}\}], \quad 1 \leq i \leq n, \quad (2)$$

$$[y_{ij}|u_{i,\text{sbj}}] \stackrel{\text{ind.}}{\sim} N(u_{i,\text{sbj}} + f_1(x_{1i}) + \beta_2^T \mathbf{x}_{2i}, \sigma_\varepsilon^2), \quad u_{i,\text{sbj}} \stackrel{\text{ind.}}{\sim} N(0, \sigma_{\text{sbj}}^2), \quad 1 \leq j \leq n_i, \quad 1 \leq i \leq m. \quad (3)$$

Model (1) is extension of the generalised additive model paradigm that allows nonparametric bivariate components. In the case where (x_{3i}, x_{4i}) correspond to geographic position then (1) is sometimes called a *geoadditive model* (e.g. Kammann & Wand, 2003). In Model (2), β_0 and β_1 are smooth functions of the x_1 variable. This model is known as a Poisson *varying coefficient model*. Model (3) is usually called an *additive mixed model* since it represents the fusion of an additive model and a linear mixed model.

An example data set that might benefit from (3) is shown in Figure 1. The source and description of the data is given in Section 4. A question of interest is how the response variable, spinal bone mineral density, differs among the 4 ethnicity groups. However, the analysis is complicated by (a) the non-linear effect of age, and (b) correlation arising from repeated measurements on the same subject.

In the mixed model approach to semiparametric regression nonparametric functional relationships are handled through modelling mechanisms such as:

$$f(x) = \beta_0 + \beta_1 x + \sum_{k=1}^K u_{k,\text{spl}} z_k(x), \quad u_{k,\text{spl}} \text{ i.i.d } N(0, \sigma_{\text{spl}}^2). \quad (4)$$

Here $z_1, \dots, z_K$ are a set of spline basis functions. The simplest example is $z_k(x) = (x - \kappa_k)_+$ for some knot sequence $\kappa_1, \dots, \kappa_K$. However, more sophisticated options now exist: see, e.g. Wood (2003), Welham *et al.* (2006) and Wand & Ormerod (2008). Most of the spline bases described in these three references are in accordance with the classical nonparametric regression method known as *smoothing splines* (e.g. Wahba, 1990; Eubank, 1999). This approach is extendable to multivariate functions using either radial basis functions (e.g. Wood, 2003; Ruppert *et al.*, 2003) or tensor products (e.g. Wood, 2006).

The upshot of (4) is that most semiparametric regression models are expressible as

$$E(\mathbf{y}|\mathbf{u}) = g(\mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u}), \quad \mathbf{u} \sim (\mathbf{0}, \mathbf{G}). \quad (5)$$

Here g is a scalar ‘link’ function. The fixed effects term, $\mathbf{X}\boldsymbol{\beta}$, handles covariates that enter the model linearly while the random effects component $\mathbf{Z}\mathbf{u}$, and corresponding covariance matrix $\mathbf{G}$, handles non-linear effects, random subject effects and other

Figure 1: Data on spinal bone mineral density versus age, data broken down according to ethnicity of the subjects. Points for the same subject are connected by lines.

spatial correlation structure. There will often be other parameters arising, for example, in the variance structure (e.g. $\mathbf{R} = \text{Cov}(\mathbf{y}|\mathbf{u})$) but I will ignore this in the current discussion.

The hierarchical Bayesian version of (5) takes the form

$$\begin{aligned} [\mathbf{y}|\boldsymbol{\beta}, \mathbf{u}] &= f_1(\mathbf{y}; \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u}); & [\mathbf{u}|\mathbf{G}] &= f_2(\mathbf{u}; \mathbf{G}) \\ [\boldsymbol{\beta}] &= f_3(\boldsymbol{\beta}; \mathbf{A}_{\boldsymbol{\beta}}); & [\mathbf{G}] &= f_4(\mathbf{G}; \mathbf{A}_{\mathbf{G}}) \end{aligned} \quad (6)$$

where $\mathbf{A}_{\boldsymbol{\beta}}$ and $\mathbf{A}_{\mathbf{G}}$ are hyper-parameters, $f_1, \dots, f_4$ are fixed conditional density functions and $[v|w]$ denotes the conditional density of v given w . Inference is based on posteriors for parameters of interest; in particular

$$[\boldsymbol{\beta}|\mathbf{y}], \quad [\mathbf{u}|\mathbf{y}] \quad \text{and} \quad [\mathbf{G}|\mathbf{y}].$$

In semiparametric regression it is very rare that analytical solutions for these posteriors exist and approximation methods need to be employed. MCMC approximation via the BUGS software (e.g. Lunn *et al.* 2000) often provides satisfactory solutions.

Semiparametric regression, especially via mixed models and hierarchical Bayesian approaches, is now a major branch of Statistics. In Ruppert, Wand & Carroll (2008) we reviewed literature on the topic for the period 2003-2007 and found about 300 papers with connections to semiparametric regression. About 100 of these were in non-Statistics journals. Applications include quantitative trait prediction (Gianola, Fernando & Stella 2006), modelling of on-line auctions (Jank & Shmueli, 2007) and disease mapping (Crainiceanu, Diggle & Rowlingson, 2008).

Graphical models, described in the next section, can be used for both frequentist and Bayesian statistical models. For the remainder of this article I will restrict attention to

Bayesian semiparametric regression. This is in keeping with the graphical model software used in the examples.

3 Graphical Models

Graphical models is a relatively young branch of mathematics that combines ideas from graph theory and probability. Sometimes known as *probabilistic graphical models*, they facilitate visualisation of probability models. Graph-theoretic results have been established for determining conditional independence relationships, and devising efficient algorithms for inference. The fundamental components of a graph are *nodes* and *edges* which link pairs of nodes. Directed graphs add an arrow-head to each link, conveying a *parent-child* relationship. Several examples of graphs are given later in this section.

Even though there are instances of graphical models in probabilistic and statistical contexts going back several decades (e.g. Wright, 1934; Besag, 1974; Geman & Geman, 1984) the 1980s saw the emergence of substantive theoretical advances and its use in applications. Much of this occurred outside of mainstream Statistics, and was mainly driven by applications in Machine Learning and Pattern Recognition. Pearl (1988) is a watershed book on modern graphical models. It was soon followed by several others: Whittaker (1990), Jensen (1996), Lauritzen (1996), Castillo, Gutiérrez & Hadi (1997), Jordan (1999) and Cowell, Dawid, Lauritzen & Spiegelhalter (1999). The preparation of this paper has been aided by summaries of the field contained in Jordan (2004), Wasserman (2004, Chapter 17) and Bishop (2006, Chapter 8) and each are highly recommended as background reading for the current article. I also adopt the conventions of Bishop (2006) for displaying graphs.

There are two main types of graphical models: *directed acyclic graphs* (DAGs), also known as *Bayesian networks*, and *undirected graphs*, also known as *Markov random fields*. Of these, DAGs are more immediately relevant to semiparametric regression and attention will be restricted to this sub-class of graphical models.

Figure 2: A *directed acyclic graph* (DAG) involving five random variables: $x_1, \dots, x_5$.

An elementary example DAG is given in Figure 2. The $x_1, \dots, x_5$ are random variables

corresponding to each of the nodes. The joint distribution of the x_i s defined by this graph takes the form

$$[x_1, x_2, x_3, x_4, x_5] = \prod_{i=1}^5 [x_i | \text{parents of } x_i] = [x_1][x_2][x_3|x_1, x_2][x_4|x_2][x_5|x_3, x_4].$$

More generally, a DAG with N nodes corresponding to the random vectors $\mathbf{x}_1, \dots, \mathbf{x}_N$ has its joint distribution given by

$$[\mathbf{x}_1, \dots, \mathbf{x}_N] = \prod_{k=1}^N [\mathbf{x}_k | \text{parents of } \mathbf{x}_k].$$

Of particular relevance to this article is the DAG representation of hierarchical Bayesian models. Consider a Bayesian version of simple linear regression

$$[y_i | \beta_0, \beta_1, \sigma^2] \stackrel{\text{ind.}}{\sim} N(\beta_0 + \beta_1 x_i, \sigma^2), \quad 1 \leq i \leq n, \quad (7)$$

$$[\beta_0] \sim N(\mu_{\beta_0}, \sigma_{\beta_0}^2), \quad [\beta_1] \sim N(\mu_{\beta_1}, \sigma_{\beta_1}^2), \quad [\sigma^2] \sim \text{IG}(A, B).$$

Then Figure 3 shows (7) represented as a DAG. Constant nodes, corresponding to the hyperparameters and x_i s, are shown as small solid circles. The shading of the y_i nodes indicates the replacement by observed values. Bayesian inference involves conditioning on these nodes to obtain posterior densities. For example, the posterior density of the slope parameter β_1 is $[\beta_1 | y_1, \dots, y_n]$.

Figure 3: DAG representation of a hierarchical Bayesian simple linear regression model. Shaded nodes correspond to observed data.

More compact versions of Figure 3 are shown in Figure 4. In Figure 4 (a) we introduce a *plate*, shown here as a rectangle, for the x_i and y_i nodes. The plate convention is that all

subscripted nodes inside the plate represent several nodes corresponding to the subscript ranging from 1 to the number in the bottom right-hand corner of the plate. Panel (b) of Figure 4 suppresses the constant nodes, and conveys only the essential probabilistic structure of the model. In Figure 4 (c) we replace the y_i nodes by a single node for the random vector $\mathbf{y} = (y_1, \dots, y_n)$ and the β_0 and β_1 nodes by one for $\boldsymbol{\beta} = (\beta_0, \beta_1)$. This graph hides the fact that the y_i are conditionally independent given the parameters, but provides a particularly succinct summary of (7).

Figure 4: Compact graphical representations of a hierarchical Bayesian logistic regression model. Panel (a) uses the plate convention for the subscripted nodes. In panel (b) the constant nodes are suppressed. Panel (c) treats the vectors $\mathbf{y} = (y_1, \dots, y_n)$ and $\boldsymbol{\beta} = (\beta_0, \beta_1)$ as entities.

The nodes corresponding the parameters of a hierarchical Bayesian model are often referred to as *hidden* in graphical models literature. The observed data correspond to *evidence* nodes. In Figure 4 (c) the evidence node is

$$\mathcal{E} = \{\mathbf{y}\}$$

and the set of hidden nodes is

$$\mathcal{H} = \{\boldsymbol{\beta}, \sigma^2\}.$$

Bayesian inference relies upon

$$[\mathcal{H}|\mathcal{E}] = \frac{[\mathcal{H}, \mathcal{E}]}{[\mathcal{E}]} = \frac{[\mathbf{y}, \boldsymbol{\beta}, \sigma^2]}{[\mathbf{y}]} = \frac{[\mathbf{y}|\boldsymbol{\beta}, \sigma^2][\boldsymbol{\beta}][\sigma^2]}{\int_0^\infty \int_{\mathbb{R}^2} [\mathbf{y}|\boldsymbol{\beta}, \sigma^2][\boldsymbol{\beta}][\sigma^2] d\boldsymbol{\beta} d\sigma^2},$$

the posterior density of the parameters given the data.

For general hierarchical Bayesian models the probability calculus required to make inference about parameters of inference can be aided by DAG representation and graph theoretic results. An example of such a result is that, conditional on its parents, each node is independent of the rest of the graph except for its descendants.

Another example is *Markov blanket* theory (Pearl, 1988). The Markov blanket of a node is defined to be the set of its parents, co-parents and children. The Markov blanket of a node separates it from the rest of the graph in that conditioning on it renders that node independent from the rest of the graph. An illustration of the use of Markov blankets is provided in Section 4. There is also the theory of *d-separation* (Geiger, Verma & Pearl, 1990), which provides necessary and sufficient conditions for two sets of nodes in a DAG to be independent after conditioning on a third set of nodes.

DAG representation of hierarchical Bayesian models has had profound influence on Bayesian inference since the early 1990s. As pointed out in Jordan (2004), systematic application of graph-theoretic algorithms to Bayesian inference problems has led to so-called Bayesian “inference engines” (Cowell, *et al.*, 1999) and is exploited by the popular BUGS software.

Numerous packages in the R language are concerned with graphical models and are summarised on the web-site *CRAN Task View: gRaphical Models in R*. As I write, this web-site has address `cran.r-project.org/web/views/gR.html`.

4 Graphical Models Viewpoint of Semiparametric Regression

As discussed in Section 2, many semiparametric regression analyses may be couched in the framework of hierarchical Bayesian models. These, in turn, have natural representations as DAGs. As an illustration, consider the Bayesian additive mixed model for the data shown in Figure 1. It consists of longitudinal measurements on the spinal bone mineral density (SBMD) of a cohort of young female subjects (source: Bachrach *et al.* 1999). The number of subjects is $m = 230$. Let n_i , $1 \leq i \leq m$, denote the number of measurements for the i th subject. One question of interest concerns differences in mean SBMD among the 4 ethnic groups, Asian, Black, Hispanic and White, after accounting for age. An appropriate model is the *Bayesian additive mixed model*

$$\begin{aligned} [y_{ij} | \boldsymbol{\beta}, \mathbf{u}_{\text{sbj}}, \mathbf{u}_{\text{spl}}, \sigma_{\text{sbj}}^2, \sigma_{\text{spl}}^2, \sigma_{\varepsilon}^2] &\stackrel{\text{ind.}}{\sim} N(\boldsymbol{\beta}_{\mathbf{x}}^T \mathbf{x}_i + u_{i,\text{sbj}} + f(\text{age}_{ij}; \beta_0, \beta_1, \sigma_{\text{spl}}^2), \sigma_{\varepsilon}^2), \\ [\mathbf{u}_{\text{sbj}} | \sigma_{\text{sbj}}^2] &\sim N(\mathbf{0}, \sigma_{\text{sbj}}^2 \mathbf{I}), \quad [\mathbf{u}_{\text{spl}} | \sigma_{\text{spl}}^2] \sim N(\mathbf{0}, \sigma_{\text{spl}}^2 \mathbf{I}), \quad [\beta_0, \beta_1] \sim N(\mathbf{0}, \sigma_{\beta}^2 \mathbf{I}) \end{aligned} \quad (8)$$

$$[\boldsymbol{\beta}_{\mathbf{x}}] \sim N(\mathbf{0}, \sigma_{\beta}^2 \mathbf{I}), \quad [\sigma_{\text{sbj}}^2] \sim \text{IG}(A_{\text{sbj}}, B_{\text{sbj}}), \quad [\sigma_{\text{spl}}^2] \sim \text{IG}(A_{\text{spl}}, B_{\text{spl}}), \quad [\sigma_{\varepsilon}^2] \sim \text{IG}(A_{\varepsilon}, B_{\varepsilon}).$$

Here y_{ij} denotes the j th ($1 \leq j \leq n_i$) SBMD measurement on subject i ($1 \leq i \leq m$), age_{ij} is the age in years at which y_{ij} was recorded, $\mathbf{x}_i = (1, \text{black}_i, \text{hispanic}_i, \text{white}_i)$ where black_i , hispanic_i and white_i are indicator variables for ethnicity (Asian ethnicity is taken as the baseline). In addition, $u_{i,\text{sbj}}$ i.i.d. $N(0, \sigma_{\text{sbj}}^2)$ are random subject intercepts, and the ε_{ij} i.i.d. $N(0, \sigma_{\varepsilon}^2)$, independent of the $u_{i,\text{sbj}}$'s, account for within-subject variability. We will model the smooth function for the age effect using penalised splines:

$$f(\text{age}; \beta_0, \beta_1, \sigma_{\text{spl}}^2) = \beta_0 + \beta_1 \text{age} + \sum_{k=1}^K u_k z_k(\text{age}), \quad u_{k,\text{spl}} \text{ i.i.d. } N(0, \sigma_{\text{spl}}^2)$$

where the z_k are the spline basis functions described in Wand & Ormerod (2008). Graphical representation of (8) is displayed in Figure 5. In panel (a) the predictors and hyper-parameters are included as constant nodes. These are suppressed in Figure 5 (b).

In graphical models phraseology $\mathcal{E} = \{\mathbf{y}\}$ is the evidence node and

$$\mathcal{H} = \{\boldsymbol{\beta}, \mathbf{u}_{\text{sbj}}, \mathbf{u}_{\text{spl}}, \sigma_{\text{sbj}}^2, \sigma_{\text{spl}}^2, \sigma_{\varepsilon}^2\} \quad (9)$$

is the set of the hidden nodes. We wish to learn

$$[\mathcal{H} | \mathcal{E}] = [\boldsymbol{\beta}, \mathbf{u}_{\text{sbj}}, \mathbf{u}_{\text{spl}}, \sigma_{\text{sbj}}^2, \sigma_{\text{spl}}^2, \sigma_{\varepsilon}^2 | \mathbf{y}]. \quad (10)$$

The probability calculus required to obtain (10) is somewhat difficult since one gets stuck with intractable integrals that arise from integrating out the variance components. Useful inferential statements, such as credible intervals for the components of $\boldsymbol{\beta}$, are therefore burdensome via direct calculation. *Gibbs sampling* (e.g. Robert & Casella, 2004), a special case of MCMC, circumvents this problem by providing samples of arbitrary size from (10). There are several Gibbs sampling options due to the various ways in which $\mathcal{H}$ can be partitioned. For the partition corresponding to the nodes of Figure 5, and ordering as in (9), the Gibbs sampling strategy is:

Initialise: $\boldsymbol{\beta}^{[0]}, \mathbf{u}_{\text{sbj}}^{[0]}, \mathbf{u}_{\text{spl}}^{[0]}, (\sigma_{\text{sbj}}^2)^{[0]}, (\sigma_{\text{spl}}^2)^{[0]}, (\sigma_{\varepsilon}^2)^{[0]}$.

Cycle: $g = 1, \dots, B + G$:

Figure 5: DAG representation of the hierarchical Bayesian model for the spinal bone mineral density data (source: Bachrach et al., 1999). In panel (a) hyperparameters and the vector of age values (taken to be deterministic) are shown as small solid circles. In panel (b) these non-random nodes are suppressed.

$$\begin{aligned}
\beta^{[g]} &\sim [\beta^{[g-1]} | \mathbf{u}_{\text{sbj}}^{[g-1]}, \mathbf{u}_{\text{spl}}^{[g-1]}, (\sigma_{\text{sbj}}^2)^{[g-1]}, (\sigma_{\text{spl}}^2)^{[g-1]}, (\sigma_\epsilon^2)^{[g-1]}, \mathbf{y}], \\
\mathbf{u}_{\text{sbj}}^{[g]} &\sim [\mathbf{u}_{\text{sbj}}^{[g-1]} | \beta^{[g]}, \mathbf{u}_{\text{spl}}^{[g-1]}, (\sigma_{\text{sbj}}^2)^{[g-1]}, (\sigma_{\text{spl}}^2)^{[g-1]}, (\sigma_\epsilon^2)^{[g-1]}, \mathbf{y}], \\
\mathbf{u}_{\text{spl}}^{[g]} &\sim [\mathbf{u}_{\text{spl}}^{[g-1]} | \beta^{[g]}, \mathbf{u}_{\text{sbj}}^{[g]}, (\sigma_{\text{sbj}}^2)^{[g-1]}, (\sigma_{\text{spl}}^2)^{[g-1]}, (\sigma_\epsilon^2)^{[g-1]}, \mathbf{y}], \\
(\sigma_{\text{sbj}}^2)^{[g]} &\sim [(\sigma_{\text{sbj}}^2)^{[g-1]} | \beta^{[g]}, \mathbf{u}_{\text{sbj}}^{[g]}, \mathbf{u}_{\text{spl}}^{[g]}, (\sigma_{\text{spl}}^2)^{[g-1]}, (\sigma_\epsilon^2)^{[g-1]}, \mathbf{y}], \\
(\sigma_{\text{spl}}^2)^{[g]} &\sim [(\sigma_{\text{spl}}^2)^{[g-1]} | \beta^{[g]}, \mathbf{u}_{\text{sbj}}^{[g]}, \mathbf{u}_{\text{spl}}^{[g]}, (\sigma_{\text{sbj}}^2)^{[g]}, (\sigma_\epsilon^2)^{[g-1]}, \mathbf{y}], \\
(\sigma_\epsilon^2)^{[g]} &\sim [(\sigma_\epsilon^2)^{[g-1]} | \beta^{[g]}, \mathbf{u}_{\text{sbj}}^{[g]}, \mathbf{u}_{\text{spl}}^{[g]}, (\sigma_{\text{sbj}}^2)^{[g]}, (\sigma_{\text{spl}}^2)^{[g]}, \mathbf{y}].
\end{aligned} \tag{11}$$

For sufficiently high value of the burnin sample size B , the draws

$$\boldsymbol{\beta}^{[g]}, \mathbf{u}_{\text{sbj}}^{[g]}, \mathbf{u}_{\text{spl}}^{[g]}, (\sigma_{\text{sbj}}^2)^{[g]}, (\sigma_{\text{spl}}^2)^{[g]}, (\sigma_{\varepsilon}^2)^{[g]}, \quad B+1 \leq g \leq B+G,$$

are a sample of size G from (10) and can be used for inference. Implementation of this Gibbs sampling scheme requires the *full conditionals* sampling distributions

$$[\boldsymbol{\beta}|\text{rest}], [\mathbf{u}_{\text{sbj}}|\text{rest}], [\mathbf{u}_{\text{spl}}|\text{rest}], [\sigma_{\text{sbj}}^2|\text{rest}], [\sigma_{\text{spl}}^2|\text{rest}] \text{ and } [\sigma_{\varepsilon}^2|\text{rest}] \quad (12)$$

where “rest” denotes the nodes in the graph apart from the node appearing before the vertical bar. Determination of (12) benefits from the Markov blanket result stated in Section 3:

$$[\text{node}|\text{rest}] = [\text{node}|\text{Markov blanket of node}].$$

The Markov blanket of σ_{sbj}^2 is shown in Figure 6.

Figure 6: Markov blanket of the node σ_{sbj}^2 .

For this node we then have

$$\begin{aligned} [\sigma_{\text{sbj}}^2|\text{rest}] &= [\sigma_{\text{sbj}}^2|\text{Markov blanket of } \sigma_{\text{sbj}}^2] = [\sigma_{\text{sbj}}^2|\mathbf{u}_{\text{sbj}}] \propto [\mathbf{u}_{\text{sbj}}|\sigma_{\text{sbj}}^2][\sigma_{\text{sbj}}^2] \\ &\propto (\sigma_{\text{sbj}}^2)^{-m/2} \exp\{-\|\mathbf{u}_{\text{sbj}}\|^2/(2\sigma_{\text{spl}}^2)\} (\sigma_{\text{sbj}}^2)^{-A_{\text{sbj}}-1} \exp(-B_{\text{sbj}}/\sigma_{\text{spl}}^2) \\ &= (\sigma_{\text{sbj}}^2)^{-(A_{\text{sbj}}+m/2)-1} \exp\{-(B_{\text{sbj}} + \frac{1}{2}\|\mathbf{u}_{\text{sbj}}\|^2)/\sigma_{\text{spl}}^2\} \\ &\sim \text{IG}(A_{\text{sbj}} + \frac{1}{2}m, B_{\text{sbj}} + \frac{1}{2}\|\mathbf{u}_{\text{sbj}}\|^2). \end{aligned}$$

Continuing in this fashion we obtain the set of full conditionals as:

$$\begin{aligned} [\boldsymbol{\beta}|\text{rest}] &\sim N((\mathbf{X}^T\mathbf{X} + \frac{\sigma_{\varepsilon}^2}{\sigma_{\beta}^2}\mathbf{I})^{-1}\mathbf{X}^T(\mathbf{y} - \mathbf{Z}_{\text{sbj}}\mathbf{u}_{\text{sbj}} - \mathbf{Z}_{\text{spl}}\mathbf{u}_{\text{spl}}), \sigma_{\varepsilon}^2(\mathbf{X}^T\mathbf{X} + \frac{\sigma_{\varepsilon}^2}{\sigma_{\beta}^2}\mathbf{I})^{-1}), \\ [\mathbf{u}_{\text{sbj}}|\text{rest}] &\sim N((\mathbf{Z}_{\text{sbj}}^T\mathbf{Z}_{\text{sbj}} + \frac{\sigma_{\varepsilon}^2}{\sigma_{\text{sbj}}^2}\mathbf{I})^{-1}\mathbf{Z}_{\text{sbj}}^T(\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}_{\text{spl}}\mathbf{u}_{\text{spl}}), \sigma_{\varepsilon}^2(\mathbf{Z}_{\text{sbj}}^T\mathbf{Z}_{\text{sbj}} + \frac{\sigma_{\varepsilon}^2}{\sigma_{\text{sbj}}^2}\mathbf{I})^{-1}), \\ [\mathbf{u}_{\text{spl}}|\text{rest}] &\sim N((\mathbf{Z}_{\text{spl}}^T\mathbf{Z}_{\text{spl}} + \frac{\sigma_{\varepsilon}^2}{\sigma_{\text{spl}}^2}\mathbf{I})^{-1}\mathbf{Z}_{\text{spl}}^T(\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}_{\text{sbj}}\mathbf{u}_{\text{sbj}}), \sigma_{\varepsilon}^2(\mathbf{Z}_{\text{spl}}^T\mathbf{Z}_{\text{spl}} + \frac{\sigma_{\varepsilon}^2}{\sigma_{\text{spl}}^2}\mathbf{I})^{-1}), \\ [\sigma_{\text{sbj}}^2|\text{rest}] &\sim \text{IG}(A_{\text{sbj}} + \frac{1}{2}m, B_{\text{sbj}} + \frac{1}{2}\|\mathbf{u}_{\text{sbj}}\|^2), \\ [\sigma_{\text{spl}}^2|\text{rest}] &\sim \text{IG}(A_{\text{spl}} + \frac{1}{2}K, B_{\text{spl}} + \frac{1}{2}\|\mathbf{u}_{\text{spl}}\|^2) \text{ and} \\ [\sigma_{\varepsilon}^2|\text{rest}] &\sim \text{IG}(A_{\varepsilon} + \frac{1}{2}\sum_{i=1}^m n_i, B_{\varepsilon} + \frac{1}{2}\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta} - \mathbf{Z}_{\text{sbj}}\mathbf{u}_{\text{sbj}} - \mathbf{Z}_{\text{spl}}\mathbf{u}_{\text{spl}}\|^2). \end{aligned} \quad (13)$$

Approximate Bayesian inference can then proceed via implementation of (11) and (13). Implementation in the R language (R Core Development Team, 2008) is very straightforward. However, BUGS offers even more immediate results. Note that in the Windows version of BUGS, known as WinBUGS, there is the option to specify the model using a graphical model drawing facility. Figure 7 is a screen-shot of the graph used for fitting (8) in BUGS.

Figure 7: Screen-shot of DAG drawn in WinBUGS for specifying the semiparametric regression model applied to the spinal bone mineral density data.

Figure 8 summarises the MCMC output and subsequent Bayesian inference for the parameters in (8). In keeping with previously published analyses, a statistically significant difference is found between Black and Asian females in terms of mean spinal bone mineral density. Note that, in Figure 8, σ_{sbj} is replaced by the (*effective*) degrees of freedom for the non-linear age effect (e.g. Buja, Hastie & Tibshirani, 1989).

The fitted curves $\hat{f}(\text{age})$, together with 95% pointwise credible sets, is shown in Figure 8.

5 Non-standard Semiparametric Regression

The biggest gains from a graphical models viewpoint of semiparametric regression are realised when the setting is a non-standard one. In ‘standard’ semiparametric regression the response variable is approximately Gaussian and all data are cleanly observed. However, in many applications, the data do not conform with this bucolic state of affairs, and the analyst has to deal with the likes of categorical response variables, outliers, missingness and measurement error. In this section I focus on three aspects of non-standard semiparametric regression: non-Gaussian response, predictors subject to missingness and predictors subject to measurement error; and view them through the graphical models prism.

Figure 8: Summary of MCMC-based inference for parameters in the fitted model for the spinal bone mineral density data. The columns are: parameter, trace plot of MCMC sample, plot of sample against 1-lagged sample, sample autocorrelation function, kernel estimates posterior density and basic numerical summaries.

5.1 Non-Gaussian response

As is well-known, generalisation of regression models to non-Gaussian response variables has its challenges. In semiparametric regression it essentially entails the extension from linear mixed models to generalised linear mixed models. If taking a frequentist likelihood-based approach then an immediate consequence is intractable integrals. If a Bayesian/MCMC approach is taken then the full conditionals are no longer standard distributions like those appearing in (13). More elaborate MCMC schemes, such as Metropolis-Hastings and adaptive rejection sampling (e.g. Robert & Casella, 2004) are required. Metropolis-Hastings algorithms can also benefit from the graphical models viewpoint and Markov blanket theory. As pointed out by Jordan (2004): “factors that do not appear in the Markov blanket of a set of variables being considered in a proposed update can be neglected.”

Zhao, Coull, Staudenmayer & Wand (2006) recently studied semiparametric regression for the case where the response variable distribution is in the one-parameter exponential family. Implementation in BUGS was described, which then means that such semiparametric regression models have a DAG representation. For example, the logistic additive mixed model fitted to data on respiratory infection in Indonesian children in Section 4.1 of Zhao *et al.* (2006) has the graphical representation shown in Figure 10. This DAG is very similar to the one in Figure 5 for the spinal bone mineral density example. The main difference is that y is now a binary rather than Gaussian node.

Figure 9: MCMC-based estimate of the non-linear age effect in the spinal bone mineral density example. The dashed lines correspond to pointwise 95% credible sets.

Figure 10: DAG representation of the Bayesian logistic additive mixed model applied to data on respiratory infection in Indonesian children in Zhao et al. (2006).

However, many other non-Gaussian response models of interest fall outside the one-parameter exponential family structure. Examples from the semiparametric regression

literature include Negative Binomial (e.g. Thurston, Wand & Weincke, 2000), Efron's Double Exponential Family (e.g. Nott, 2006), Beta (e.g. Branscum, Johnson & Thurmond, 2007), Student's t (e.g. Staudenmayer, Lake & Wand, 2008), and Generalised Extreme Value (e.g. Yee & Stephenson; Padoan & Wand, 2008) distributions. All can be embedded within a graphical models framework. Current joint research with Jennifer K. Marley is investigating BUGS fitting of non-Gaussian response semiparametric regression models such as these.

As the response becomes less standard then the suitability of established mixed model software is less likely and more general software packages such as BUGS are about the only current option. Essentially, this means that ordinary mixed model architectures are inadequate and that more general graphical model architectures are required. The next two subsections provide even more potent illustrations of this state of affairs.

5.2 Predictors subject to missingness

Consider the simple nonparametric regression setting

$$y_i = f(x_i) + \varepsilon_i, \quad \varepsilon_i \text{ i.i.d. } N(0, \sigma_\varepsilon^2), \quad 1 \leq i \leq n, \quad (14)$$

for a smooth function f , and assume that the x_i s can be modelled as coming from a normal distribution with mean μ_x and variance σ_x^2 . Suppose, however, that some of the x_i s are missing. An appropriate hierarchical Bayesian model for this situation is

$$\begin{aligned} [y_i | x_i, \boldsymbol{\beta}, \mathbf{u}, \sigma_\varepsilon^2] &\stackrel{\text{ind.}}{\sim} N\left(\beta_0 + \beta_1 x_i + \sum_{k=1}^K u_k z_k(x_i), \sigma_\varepsilon^2\right), \\ [\mathbf{u} | \sigma_u^2] &\sim N(\mathbf{0}, \sigma_u^2 \mathbf{I}), \quad [x_i | \mu_x, \sigma_x^2] \text{ i.i.d. } N(\mu_x, \sigma_x^2), \\ [\boldsymbol{\beta}] &\sim N(\mathbf{0}, \sigma_\beta^2 \mathbf{I}), \quad [\mu_x] \sim N(0, \sigma_{\mu_x}^2), \\ [\sigma_u^2] &\sim \text{IG}(A_u, B_u), \quad [\sigma_\varepsilon^2] \sim \text{IG}(A_\varepsilon, B_\varepsilon), \quad [\sigma_x^2] \sim \text{IG}(A_x, B_x). \end{aligned} \quad (15)$$

Let $\mathbf{x}^{\text{obs}}$ be the vector of observed x_i s and $\mathbf{x}^{\text{mis}}$ be the missing values. Then the observed data, or evidence nodes, are

$$\mathcal{E} = \{\mathbf{y}, \mathbf{x}^{\text{obs}}\}$$

while the parameters, or hidden nodes, are

$$\mathcal{H} = \{\boldsymbol{\beta}, \mathbf{u}, \mathbf{x}^{\text{mis}}, \sigma_u^2, \sigma_\varepsilon^2, \mu_x, \sigma_x^2\}.$$

The DAG for (15) is given in Figure 11.

Bayesian inference requires

$$[\mathcal{H} | \mathcal{E}] = [\boldsymbol{\beta}, \mathbf{u}, \mathbf{x}^{\text{mis}}, \sigma_u^2, \sigma_\varepsilon^2, \mu_x, \sigma_x^2 | \mathbf{y}, \mathbf{x}^{\text{obs}}].$$

Note the extra layer of complexity imposed by missingness since

$$[\boldsymbol{\beta}, \mathbf{u}, \sigma_u^2, \sigma_\varepsilon^2, \mu_x, \sigma_x^2 | \mathbf{y}, \mathbf{x}^{\text{obs}}] = \frac{\int [\mathbf{y} | \mathbf{x}, \boldsymbol{\beta}, \sigma_\varepsilon^2] [\mathbf{x} | \mu_x, \sigma_x^2] [\boldsymbol{\beta}] [\mathbf{u} | \sigma_u^2] [\sigma_u^2] [\sigma_\varepsilon^2] [\mu_x] [\sigma_x^2] d\mathbf{x}^{\text{mis}}}{\int [\mathbf{y} | \mathbf{x}, \boldsymbol{\beta}, \sigma_\varepsilon^2] [\mathbf{x} | \mu_x, \sigma_x^2] [\boldsymbol{\beta}] [\mathbf{u} | \sigma_u^2] [\sigma_u^2] [\sigma_\varepsilon^2] [\mu_x] [\sigma_x^2] d\mathcal{H}}$$

now involves integration over the missing data vector $\mathbf{x}^{\text{mis}}$.

I tested BUGS fitting of (15) to simulated data with

$$n = 300, \quad f(x) = \sin(4\pi x), \quad \mu_x = \frac{1}{2}, \quad \sigma_x^2 = \frac{1}{36}, \quad \sigma_\varepsilon^2 = 0.35 \quad (16)$$

and 20% of the x_i s missing completely at random. The hyperparameters were set to be

$$\sigma_\beta^2 = \sigma_{\mu_x}^2 = 10^8, \quad A_u = B_u = A_\varepsilon = B_\varepsilon = A_x = B_x = \frac{1}{100}. \quad (17)$$

Figure 11: Graphical representation of the penalised spline nonparametric regression model with the predictor subject to missingness. Shading corresponds to the observed, or evidence, nodes.

Since the spline basis functions for the missing x_i s have to be computed inside BUGS I used truncated lines:

$$z_k(x) = (x - \kappa_k)_+ \quad \text{with} \quad \kappa_k = \{(K+1-k) \min(x_i^{\text{obs}}) + k \max(x_i^{\text{obs}})\} / (K+1), \quad 1 \leq k \leq K$$

and $K = 25$. The relevant BUGS code is

```
model
{
  for(i in 1:nObs)
  {
    muObs[i] <- beta0 + betal*xObs[i] + inprod(u[],ZxObs[i,])
    yxObs[i] ~ dnorm(muObs[i],tauEps)
    xObs[i] ~ dnorm(muX,tauX)
  }
  for(i in 1:nMis)
  {
    muMis[i] <- beta0 + betal*xMis[i] + inprod(u[],ZxMis[i,])
    yxMis[i] ~ dnorm(muMis[i],tauEps)
    xMis[i] ~ dnorm(muX,tauX)
  }
  for (k in 1:K)
  {
    for (i in 1:nMis)
    {
      ZxMis[i,k] <- (xMis[i]-knots[k])*step(xMis[i]-knots[k])
    }
    u[k] ~ dnorm(0,tauU)
  }
  beta0 ~ dnorm(0,1.0E-8) ; betal ~ dnorm(0,1.0E-8)
  muX ~ dnorm(0,1.0E-8) ; tauX ~ dgamma(0.01,0.01)
  tauU ~ dgamma(0.01,0.01) ; tauEps ~ dgamma(0.01,0.01)
}
```

where, for example, $y_{xObs} []$ is the vector of y_i values that have an observed x_i partner.

The upper panels of Figure 12 summarises the MCMC output produced by BUGS for the nodes μ_x , σ_x and σ_ε . The true values (16) from which the data were simulated are shown as vertical dashed lines in the posterior density plots. The lowest panel monitors the effective degrees of freedom for estimation of f . The chains are seen to be reasonably well-behaved.

parameter	trace	lag 1	acf	density	summary
μ_x					posterior mean: 0.479 95% credible interval: (0.458, 0.499)
σ_x					posterior mean: 0.16 95% credible interval: (0.146, 0.175)
σ_ε					posterior mean: 0.329 95% credible interval: (0.302, 0.357)
degrees of freedom for f					posterior mean: 14.1 95% credible interval: (12, 16.6)

Figure 12: Summary of MCMC-based inference for parameters in the missing predictor nonparametric regression model. The columns are: parameter, trace plot of MCMC sample, plot of sample against 1-lagged sample, sample autocorrelation function, kernel estimates posterior density and basic numerical summaries. The vertical dashed lines in the density plots correspond to the true values of the parameters according to the simulation set-up.

Figure 13 shows the estimate of f as well as pointwise 95% credible intervals. The missing data, known from simulation but hidden from the methodology, are shown as grey circles.

Lastly, we study some of the output for the hidden node $\mathbf{x}^{\text{mis}}$. Five components were chosen at random and the MCMC summaries are shown in Figure 14. Interestingly, the posterior densities of some of the $\mathbf{x}^{\text{mis}}$ components are multimodal. This arises from the periodic nature of the underlying signal. Knowledge about the ordinate manifests in the posterior of x_i^{mis} as two or three clumps of probability mass corresponding, roughly, to vertical slicing of f at that ordinate.

We close this subsection by noting that a moderate amount of research on missingness for mixed model-based semiparametric now exists. References include French & Wand (2004), Chen & Ibrahim (2006), Geraci & Bottai (2006) and Yuan & Little (2007).

5.3 Predictors subject to measurement error

In the previous subsection the x_i s in (15) were subject to missingness. Now suppose instead that they are subject to measurement error. In this case, rather than observing x_i we observe

$$w_i = x_i + z_i, \quad 1 \leq i \leq n, \quad (18)$$

where the z_i i.i.d. $N(0, \sigma_z^2)$ and independent of the x_i s. The contamination variance, σ_z^2 , is assumed to be known.

This is an instance of nonparametric regression with measurement error. Carroll *et al.* (2006) is a recent survey of this and related topics. A hierarchical Bayesian model for (14)

Figure 13: MCMC-based estimate of f in the missing predictor nonparametric regression model. The dashed lines correspond to pointwise 95% credible sets. The grey points are those for which the x values were missing and not used by the fitting procedure.

and (18) is

$$\begin{aligned}
 [y_i | x_i, \boldsymbol{\beta}, \mathbf{u}, \sigma_\varepsilon^2] &\stackrel{\text{ind.}}{\sim} N\left(\beta_0 + \beta_1 x_i + \sum_{k=1}^K u_k z_k(x_i), \sigma_\varepsilon^2\right), \\
 [\mathbf{u} | \sigma_u^2] &\sim N(\mathbf{0}, \sigma_u^2 \mathbf{I}), \quad [x_i | \mu_x, \sigma_x^2] \stackrel{\text{ind.}}{\sim} N(\mu_x, \sigma_x^2), \quad [w_i | x_i] \stackrel{\text{ind.}}{\sim} N(x_i, \sigma_z^2), \\
 [\boldsymbol{\beta}] &\sim N(\mathbf{0}, \sigma_\beta^2 \mathbf{I}), \quad [\mu_x] \sim N(0, \sigma_{\mu_x}^2), \\
 [\sigma_u^2] &\sim \text{IG}(A_u, B_u), \quad [\sigma_\varepsilon^2] \sim \text{IG}(A_\varepsilon, B_\varepsilon), \quad [\sigma_x^2] \sim \text{IG}(A_x, B_x).
 \end{aligned} \tag{19}$$

The observed data, or set of evidence nodes, is

$$\mathcal{E} = \{\mathbf{y}, \mathbf{w}\}$$

where $\mathbf{w}$ is the vector of w_i s. The set of parameters, or hidden nodes, is

$$\mathcal{H} = \{\boldsymbol{\beta}, \mathbf{u}, \mathbf{x}, \sigma_u^2, \sigma_\varepsilon^2, \mu_x, \sigma_x^2\}.$$

Graphical representation of (19) is shown in Figure 15.

BUGS fitting of (19) was tested using the parameter settings given by (16) and (17) and with σ_z set to be 0.1. As for the missing data example spline basis functions have to be computed inside BUGS, so I used truncated line basis functions with knots

$$\kappa_k = \{(K + 1 - k) \min(x_i) + k \max(x_i)\} / (K + 1), \quad 1 \leq k \leq K,$$

(which depend on the hidden x node) and $K = 20$. The BUGS code is:

parameter	trace	lag 1	acf	density	summary
x_{10}^{mis}					posterior mean: 0.535 95% credible interval: (0.123, 0.712)
x_{18}^{mis}					posterior mean: 0.424 95% credible interval: (0.287, 0.856)
x_{27}^{mis}					posterior mean: 0.523 95% credible interval: (0.163, 0.733)
x_{44}^{mis}					posterior mean: 0.493 95% credible interval: (0.202, 0.767)
x_{59}^{mis}					posterior mean: 0.471 95% credible interval: (0.234, 0.783)

Figure 14: Summary of MCMC-based inference for five randomly chosen missing predictors in the missing predictor nonparametric regression model. The columns are: missing predictor, trace plot of MCMC sample, plot of sample against 1-lagged sample, sample autocorrelation function, kernel estimates posterior density and basic numerical summaries. The vertical dashed lines in the density plots correspond to the true values of the predictors for the simulation.

Figure 15: Graphical representation of the penalised spline nonparametric regression model with the predictor subject to measurement error. Shading corresponds to the observed, or evidence, nodes.

```

model
{
  for(i in 1:n)
  {
    x[i] ~ dnorm(muX,tauX)
    w[i] ~ dnorm(x[i],tauZ)
    mu[i] <- beta0 + beta1*x[i] + inprod(u[],Z[i,])
    y[i] ~ dnorm(mu[i],tauEps)
  }
  for (k in 1:K)
  {
    knots[k] <- ((K+1-k)*ranked(x[],1)+k*ranked(x[],n))/(K+1)
    for (i in 1:n)
    {
      Z[i,k] <- (x[i]-knots[k])*step(x[i]-knots[k])
    }
    u[k] ~ dnorm(0,tauU)
  }
  beta0 ~ dnorm(0,1.0E-8) ; beta1 ~ dnorm(0,1.0E-8)
  muX ~ dnorm(0,1.0E-8) ; tauX ~ dgamma(0.01,0.01)
  tauU ~ dgamma(0.01,0.01) ; tauEps ~ dgamma(0.01,0.01)
}

```

The upper panels of Figure 16 is the analogue of Figure 12 for current measurement error example. Once again, the chains are seen to be reasonably well-behaved and true parameters are inside the 95% credible sets.

Figure 16: Summary of MCMC-based inference for parameters in the nonparametric regression measurement error model. The columns are: parameter, trace plot of MCMC sample, plot of sample against 1-lagged sample, sample autocorrelation function, kernel estimates posterior density and basic numerical summaries. The vertical dashed lines in the density plots correspond to the true values of the parameters according to the simulation set-up.

Figure 17 shows the estimate of f as well as pointwise 95% credible intervals. The grey circles are the unobserved (x_i, y_i) pairs which, because this is a simulation study, are known. The curve estimate is seen to be quite reasonable despite having to adjust for contamination of the x_i s.

Models of type (19) were first formulated by Berry, Ruppert & Carroll (2002).

Figure 17: MCMC-based estimate of f in the nonparametric regression with measurement error model. The shaded region corresponds to pointwise 95% credible sets. The black points are the observed (w, y) pairs (contaminated data). The grey points are the unobserved (x, y) pairs (clean data).

6 Variational Inference Engines

Each of the examples in the previous two sections were performed using an MCMC-based Bayesian inference engine, namely BUGS. However, MCMC is computationally intensive and inference can be quite slow. The example in Section 5.3 involving measurement error took about a day to run on my laptop computer. An alternative to MCMC, which offers the possibility of much faster approximate inference, is *variational approximation*. So-called *variational inference machines* have emerged in recent years for conducting inference in DAG models. The most prominent is VIBES (Variational Inference for BayESian networks); authored by Bishop, Spiegelhalter & Winn (2003). Several others are described in Murphy (2007), including an in-progress successor to VIBES named Infer.NET (Minka, Winn, Guiver & Kannan, 2008) (current web-site: research.microsoft.com/mlp/ml/infer/infer.htm). An illustration of VIBES is given later in this section. Before that I will provide a brief description of variational approximation.

Variational approximation is an alternative to MCMC that is gathering steam as a means of making inference in complex models when the latter becomes untenable. Most contemporary literature on variational approximation for graphical models is in Computer Science rather than Statistics. Review articles that summarise contemporary variational inference are Jordan, Ghahramani, Jaakkola & Saul (1999), Jordan (2004), Titterton (2004) and Bishop (2006).

The essence of variational approximation is the use of variational forms for non-linear functions. An example is

$$\log(x) = \min_{\xi > 0} \{ \xi x - \log(\xi) - 1 \}, \quad \text{for all } x > 0.$$

The fact that $\xi x - \log(\xi) - 1$ is linear in x for every value of the *variational parameter* $\xi > 0$ allows for simplifications of expressions involving the logarithmic function. The value of ξ can then be chosen to make the approximation as accurate as possible.

An instructive example of variational inference is simple Bayesian logistic regression:

$$[y_i | \beta_0, \beta_1] \stackrel{\text{ind.}}{\sim} \frac{e^{y_i(\beta_0 + \beta_1 x_i)}}{1 + e^{\beta_0 + \beta_1 x_i}}, \quad [\beta_0, \beta_1] \sim N(\mathbf{0}, \sigma_\beta^2 \mathbf{I}), \quad 1 \leq i \leq n \quad (20)$$

where $y_i \in \{0, 1\}$ is the i th realisation of a binary response variables and x_i is the corresponding predictor. This is a special case of an example given in Jaakkola & Jordan (2000). Inference about the slope parameter requires

$$[\beta_1 | \mathbf{y}] \propto e^{-\beta_1^2 / (2\sigma_\beta^2)} \int_{-\infty}^{\infty} \exp \left(\sum_{i=1}^n [\beta_0 y_i - \log\{1 + \exp(\beta_0 + \beta_1 x_i)\}] - \beta_0^2 / (2\sigma_\beta^2) \right) d\beta_0. \quad (21)$$

The presence of $-\log\{1 + \exp(\beta_0 + \beta_1 x_i)\}$ in the exponent of the integrand makes the integrals irreducible. However, we can make use of the variational form

$$-\log(1 + e^x) = \max_{\xi \in \mathbb{R}} \{A(\xi)x^2 + B(\xi)x + C(\xi)\} \quad \text{for all } x \in \mathbb{R} \quad (22)$$

where

$$A(\xi) = -\tanh(\xi/2)/(4\xi), \quad B(\xi) = -1/2 \text{ and } C(\xi) = \xi/2 - \log(1 + e^\xi) + \xi \tanh(\xi/2)/4$$

(Jaakkola & Jordan, 2000). Figure 18 is a graphical representation of (22); where the function $-\log(1 + e^x)$ is seen to be the maximum of a family of parabolae.

Figure 18: Variational representation of the function $-\log(1 + e^x)$, as the maximum of a family of parabolae.

In (21) one can then replace

$$-\log\{1 + \exp(\beta_0 + \beta_1 x_i)\} \quad \text{by} \quad A(\xi_i)x^2 + B(\xi_i)x + C(\xi_i), \quad 1 \leq i \leq n. \quad (23)$$

This entails the introduction of a vector of n variational parameters $\xi = (\xi_1, \dots, \xi_n)$. For any choice of $\xi \in \mathbb{R}^n$ one can solve the posterior density problem analytically and arrive at the following family of solutions:

$$[\beta_1 | \mathbf{y}; \xi] \sim N(\mu(\xi), \sigma^2(\xi)), \quad \xi \in \mathbb{R}^n,$$

where

$$\mu(\xi) \equiv \frac{(2n\bar{\lambda}(\xi) + \sigma_\beta^{-2})(\mathbf{x}^T \mathbf{y} - \bar{x}/2)}{(2n\bar{\lambda}(\xi) + \sigma_\beta^{-2})\{2(\mathbf{x}^2)^T \lambda(\xi) + \sigma_\beta^{-2} - 4\{\lambda(\xi)^T \mathbf{x}\}\}}$$

$$\text{and } \sigma^2(\xi) \equiv [2(\mathbf{x}^2)^T \lambda(\xi) + \sigma_\beta^{-2} - 4\{\lambda(\xi)^T \mathbf{x}\}^2 / \{2n\bar{\lambda}(\xi) + \sigma_\beta^{-2}\}]^{-1},$$

with $\lambda(\xi) \equiv \tanh(\xi/2)/(4\xi)$ and $\bar{\lambda}(\xi) \equiv \frac{1}{n} \sum_{i=1}^n \lambda(\xi_i)$. The variational parameters ξ should then be chosen to make the approximation (23) as accurate as possible. This involves maximisation of the lower bounds on the left-hand side of

$$A(\xi_i)x^2 + B(\xi_i)x + C(\xi_i) \leq -\log\{1 + \exp(\beta_0 + \beta_1 x_i)\}, \quad 1 \leq i \leq n.$$

An Expectation Maximisation-type iterative scheme can be devised for carrying this out (Jaakkola & Jordan, 2000). Let $[\beta_0, \beta_1 | \mathbf{y}; \xi] \sim N(\boldsymbol{\mu}(\xi), \Sigma(\xi))$ be the variational approximation to $[\beta_0, \beta_1 | \mathbf{y}]$ based on ξ . Then, with $\mathbf{y} = (y_1, \dots, y_n)$ and $\mathbf{X} = [1 \ x_i]_{1 \leq i \leq n}$, the algorithm:

Cycle:

1. $\Sigma(\xi) \leftarrow [\sigma_\beta^{-2} \mathbf{I} + 2\mathbf{X}^T \text{diag}\{\lambda(\xi)\} \mathbf{X}]^{-1}$
2. $\boldsymbol{\mu}(\xi) \leftarrow \Sigma(\xi) \mathbf{X}^T (\mathbf{y} - \frac{1}{2} \mathbf{1})$
3. $\xi \leftarrow \sqrt{\text{diagonal}[\mathbf{X}\{\Sigma(\xi) + \boldsymbol{\mu}(\xi)\boldsymbol{\mu}(\xi)^T\} \mathbf{X}^T]}$

usually leads to rapid convergence to the optimum.

I compared the posterior distribution approximations for β_0 and β_1 obtained via this variational approach to data on 223 birthweight measurements (grammes) and occurrence of *bronchopulmonary dysplasia* (source: Pagano & Gauvreau, 1993). Throughout this example I work with the standardised version of the birthweights rather than the original birthweight values with $\sigma_\beta^2 = 10^8$. Figure 19 shows the variational approximations to $[\beta_0 | \mathbf{y}]$, $[\beta_1 | \mathbf{y}]$ and $[\beta_0, \beta_1 | \mathbf{y}]$. As a benchmark I obtained 1 million realisations from the posteriors using MCMC and BUGS and constructed kernel density estimates using direct plug-in bandwidth selectors (available in the R packages `KernSmooth` and `ks`). We see from Figure 19 that the Jaakkola & Jordan (2000) variational approximations are reasonable, but not extremely accurate.

While this example, based on (22), provides an illustration of variational approximation it should be pointed out that many other methods exist. A common general approach to variational approximation involves the theory of Kullback-Liebler divergence; see, for example, Titterton (2004) and Bishop (2006, Section 10.1). It should also be pointed out that other analytic approximations exist and can be used for approximate inference in DAG models. A particularly simple and popular one is *Laplace approximation* which, for example, is used by Spiegelhalter & Lauritzen (1990) in DAG models and Breslow & Clayton (1993) in generalised linear mixed models.

Current joint research with John T. Ormerod involves a Kullback-Liebler divergence where the lower bound on the likelihood is reduced to the calculation of n univariate integrals, which are calculated numerically using adaptive Gauss-Hermite quadrature. The results for model (20) applied to the bronchopulmonary dysplasia data are shown in Figure 20. The accuracy is seen to be very good in this case, and considerably better than that of Jaakkola & Jordan (2000)

Figure 19: Assessment of the accuracy of the Jaakkola & Jordan (2000) variational approximation method. The dashed curves indicate the variational approximation to $[\beta_0|\mathbf{y}]$, $[\beta_1|\mathbf{y}]$ and $[\beta_0, \beta_1|\mathbf{y}]$ for the Bayesian logistic regression fit to data on birthweights and occurrence of bronchopulmonary dysplasia. The solid curves are approximations to the posteriors using MCMC samples of size one million (obtained using BUGS).

Figure 20: Assessment of the accuracy variational approximation method arising from current joint research with John T. Ormerod . The dashed curves indicate the variational approximation to $[\beta_0|\mathbf{y}]$, $[\beta_1|\mathbf{y}]$ and $[\beta_0, \beta_1|\mathbf{y}]$ for the Bayesian logistic regression fit to data on birthweights and occurrence of bronchopulmonary dysplasia. The solid curves are approximations to the posteriors using MCMC samples of size one million (obtained using BUGS).

I also tested the use of VIBES for Bayesian semiparametric regression by getting it to fit (8) to the spinal bone mineral density data. Figure 21 is a screen-shot of the specified model in VIBES, obtained using its graph drawing interface.

Approximate posterior densities for 4 of the model parameters are shown in Figure 22. The regression coefficients, corresponding to the indicators for Black and Hispanic, have narrower posteriors compared with those obtained using MCMC via BUGS. The posteriors for the standard deviation parameters, σ_{spl} and σ_{ε} , are quite close to the BUGS answers. The fitted age effects, shown in Figure 22 are also similar to those obtained via BUGS. However, the credible interval bars cannot be produced from the VIBES output.

Variational inference machines, such as VIBES, is a young and emerging field. It has the potential to yield satisfactory solutions to complex graphical model inferential problems much more quickly than what is current being achieved via MCMC. There is also the question of statistical properties of variational approximations to quantities such as maximum likelihood estimators and posterior densities. Jordan (2004) states that “variational inference is still in its infancy” and cites Tatikonda & Jordan (2004) for early work on asymptotics for variational approximation. Several other relevant references are listed in Section 3.3 of Jordan (2004). In the Statistics literature, pioneering work

Figure 21: Screen-shot of DAG drawn in VIBES for specifying the semiparametric regression model applied to the spinal bone mineral density data.

Figure 22: Approximate posterior densities for a selection of parameters in the VIBES fit of an additive mixed model to the spinal bone mineral density data. The MCMC-based approximations are shown for comparison.

on variational approximation theory has been undertaken by D.M. Titterington and co-authors. Examples of published work to date include Hall, Humphreys & Titterington (2002) and Wang & Titterington (2004, 2006)

Figure 23: Bayes estimates for the age effect arising from a VIBES fit of an additive mixed model to the spinal bone mineral density data.

7 Example: Relative Cancer Mapping with Missingness

We applied the new Ormerod & Wand variational approximation methodology, mentioned in the previous section, to some real data for which semiparametric regression in the face of missingness is appropriate. The data, corresponding to a female cancer study in Cape Cod, Massachusetts, USA, are described in French & Wand (2004). Of primary interest for these data is *relative cancer mapping* where the geographical variation of a certain cancer type, relative to other cancers, is assessed. In this illustration the cancer type of interest is lung cancer. Define

$$\text{lungCancer}_i = \begin{cases} 1 & \text{female } i \text{ has lung cancer} \\ 0 & \text{female } i \text{ has other type of cancer} \end{cases}$$

for $1 \leq i \leq n$ where $n = 2540$ is the number of females in the study. For each female in the study we observe the longitude (lon_i) and latitude (lat_i) values of her residence, her age in years (age_i) and whether or not she has ever smoked (smoked_i). An exception is the 15.4% of females for which smoking information is missing.

To account for age and smoking, as well as missingness in smoking, we entertained

the hierarchical Bayesian additive models

$$[\text{lungCancer}_i | \text{smoked}_i, \beta_0, \beta_{\text{smk}}, \beta_{\text{age}}, \boldsymbol{\beta}_{\text{geo}}, \mathbf{u}_{\text{age}}, \mathbf{u}_{\text{geo}}, \sigma_{\text{age}}^2, \sigma_{\text{geo}}^2] \\ \stackrel{\text{ind.}}{\sim} \text{Bernoulli}[\text{logit}^{-1}\{\beta_0 + \beta_{\text{smk}} \text{smoked}_i + f(\text{age}_i; \beta_{\text{age}}, \sigma_{\text{age}}^2) + g(\text{lon}_i, \text{lat}_i; \boldsymbol{\beta}_{\text{geo}}, \sigma_{\text{geo}}^2)\}]$$

$$[\text{smoked}_i | \tilde{\beta}_0, \tilde{\beta}_{\text{age}}, \tilde{\boldsymbol{\beta}}_{\text{geo}}, \tilde{\mathbf{u}}_{\text{age}}, \tilde{\mathbf{u}}_{\text{geo}}, \tilde{\sigma}_{\text{age}}^2, \tilde{\sigma}_{\text{geo}}^2] \\ \stackrel{\text{ind.}}{\sim} \text{Bernoulli}[\text{logit}^{-1}\{\tilde{\beta}_0 + \tilde{f}(\text{age}_i; \tilde{\beta}_{\text{age}}, \tilde{\sigma}_{\text{age}}^2) + \tilde{g}(\text{lon}_i, \text{lat}_i; \tilde{\boldsymbol{\beta}}_{\text{geo}}, \tilde{\sigma}_{\text{geo}}^2)\}].$$

The univariate functions f and $\tilde{f}$ are handled analogously to that for age in the spinal bone mineral density example of Section 4. The bivariate functions g and $\tilde{g}$ use penalised thin plate splines as described in Chapter 13 of Ruppert, Wand & Carroll (2003). Figure 24 provides a graphical description of this model. Our variational approximations for the missing smoking data node are similar in nature to those described in Section 5.2 of Jaakkola & Jordan (2000).

Figure 24: Graphical representation of the semiparametric regression missing data model for the relative cancer mapping example.

The functional components of our variational approximation fit to the above model are shown in Figure 25. The upper panels of Figure 25 show the effects of age and geographical location on lung cancer occurrence (relative to other cancer types). The age curve is monotonic as expected. The geographical fit suggests a ‘hot spot’ around -70.4° longitude and 41.65° latitude. The lower panels are the fitted effects of age and geography on smoking status. Some geographical variability in smoking status is apparent.

Figure 25: Ormerod & Wand variational approximation fit to the female lung cancer data. The displayed fit for each predictor corresponds to slices of the fitted model with the other predictors set to their medians. Panel (a): estimate of the effect of age on the relative probability of lung cancer. Panel (b): estimate of the effect of geography on the relative probability of lung cancer. Panel (c): estimate of the effect of age on the probability of ever being a smoker. Panel (d): estimate of the effect of geography on the probability of ever being a smoker.

Also, there is a somewhat surprising decline in the age curve after about 75 years. However, there is also a high degree of variability (not shown) in these function estimates for high ages.

8 Potential for New Semiparametric Regression Applications

A final advantage of the graphical models viewpoint of semiparametric regression is that it brings the latter field closer to other areas of research that rely heavily on graphical model theory and methodology. Examples include social networks (e.g. Wasserman & Faust, 1994), causal inference (e.g. Cox & Wermuth, 2001; van der Laan & Robins, 2003), hidden Markov models (e.g. Cappé, Moulines & Ryden 2005) and phylogenetic trees (e.g. Jordan, 2004).

Synergistic development of this type has been recently witnessed as result of mixed model representations of semiparametric regression. Semiparametric regression methodology is now very much a part of longitudinal data analysis (e.g. Fitzmaurice, Davidian, Verbeke & Molenberghs, 2008), spatial statistics (e.g. Hennerfeind, Brezger & Fahrmeir, 2006; Crainiceanu, Diggle & Rowlingson, 2008) and analysis of complex sample surveys (e.g. Briedt & Opsomer, 2008). There is great potential for similar outcomes in the graphical models realm.

9 Concluding Remarks

I have explained why I believe graphical models to be a useful structure for semiparametric regression analysis. Particular attention has been paid to non-standard situations where there is more to gain from the graphical models viewpoint. As theory, methodology and software for graphical models continue to be developed I envisage sophisticated semiparametric regression analyses becoming more routine and streamlined by taking advantage of graphical models representations.

Acknowledgements

I would like to thank the editors for inviting me to write this paper. I am grateful to John Ormerod for his assistance with the variational approximation examples to he and Raymond Carroll for their valuable feedback. This research was partially supported by Australian Research Council Discovery Project DP0877055.

References

- Bachrach, L.K., Hastie, T., Wang, M.-C., Narasimhan, B. & Marcus, R. (1999). Bone mineral acquisition in healthy Asian, Hispanic, Black and Caucasian youth. A longitudinal study. *Journal of Clinical Endocrinology and Metabolism*, **84**, 4702–12.
- Berry, S.A., Carroll, R.J. & Ruppert, D. (2002). Bayesian smoothing and regression splines for measurement error problems. *Journal of the American Statistical Association*, **97**, 160–169.
- Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society, Series B*, **36**, 192–236.
- Bishop, C.M. *Pattern Recognition and Machine Learning*. New York: Springer.
- Bishop, C.M., Spiegelhalter, D.J. & Winn, J. (2003). VIBES: A variational inference engine for Bayesian networks. In *Advances in Neural Information Processing Systems* (S. Becker, S. Thrun and K. Obermayer, eds.) **15**, 793–800, Cambridge, Massachusetts: MIT Press.
- Branscum, A. J., Johnson, W. O. & Thurmond, M.C. (2007). Bayesian beta regression: Applications to household expenditure data and genetic distance between foot-and-mouth disease viruses. *Australian and New Zealand Journal of Statistics*, **49**, 287–301.
- Breidt, F.J. & Opsomer, J.D. (2008). Nonparametric and semiparametric estimation in complex surveys. In: *Sample Surveys: Theory, Methods and Inference, Handbook of Statistics*, Vol. 29, C.R. Rao and D. Pfeffermann (Editors), Amsterdam: North Holland.
- Breslow, N.E. and Clayton, D.G. (1993). Approximate inference in generalized linear mixed models. *Journal of the American Statistical Association*, **88**, 9–25.
- Bühlmann, P. & Hothorn, T. (2007). Boosting algorithms: regularization, prediction and model fitting. *Statistical Science*, **22**, 477–505.

- Buja, A. and Hastie, T. and Tibshirani, R. (1989). Linear smoothers and additive models. *The Annals of Statistics*, **17**, 453–510.
- Cappé, O., Moulines, E. & Ryden, T. (2005). *Inference in Hidden Markov Models*. New York: Springer.
- Carroll, R.J., Ruppert, D., Stefanski, L.A. & Crainiceanu, C.M. (2006). *Measurement Error in Nonlinear Models (Second Edition)*. Boca Raton, Florida: Chapman & Hall/CRC.
- Casella, G. & Robert, C.P. (2004). Introduction to the special issue: Bayes then and now. *Statistical Science*, **19**, 1–2.
- Castillo, E., Gutiérrez, J.M. & Hadi, A.S. (1997). *Expert Systems and Probabilistic Network Models*. New York: Springer.
- Chen, Q.X. & Ibrahim, J.G. (2006). Semiparametric models for missing covariate and response data in regression models. *Biometrics*, **62**, 177–184.
- Cowell, R.G., Dawid, A.P., Lauritzen, S.L. and Spiegelhalter, D.J. (1999). *Probabilistic Networks and Expert Systems*. New York: Springer.
- Cox, D.R. & Wermuth, N. (2001). Causal inference and statistical fallacies. *International Encyclopedia of the Social and Behavioral Sciences*. P.B. Bates and N.J. Smelser (eds). Amsterdam: Elsevier, **3**, 1554–1661.
- Crainiceanu, C.M., Diggle, P.J. & Rowlingson, B. (2008). Bivariate binomial spatial modelling Loa loa prevalence in tropical Africa (with discussion). *Journal of the American Statistical Association*, **103**, 21–43.
- Eubank, R.L. (1999). *Nonparametric Regression and Spline Smoothing (Second Edition)*. New York: Marcel Dekker.
- Fitzmaurice, G., Davidian, M., Verbeke, G. & Molenberghs, G. (Eds.) (2008). *Longitudinal Data Analysis: A Handbook of Modern Statistical Methods*. Boca Raton, Florida: Chapman & Hall/CRC.
- French, J.L. & Wand, M.P. (2004). Generalized additive models for cancer mapping with incomplete covariates. *Biostatistics*, **5**, 177–191.
- Freund, Y. (1995). Boosting a weak learning algorithm by majority. *Information and Computation*, **121**, 256–285.
- Freund, Y. & Schapire, R.E. (1996). Experiments with a new boosting algorithm. In *Machine Learning: Proceedings of the Thirteenth International Conference*, San Francisco: Morgan Kaufman, pp. 148–156.
- Geiger, D., Verma, T.S. and Pearl, J. (1990). Identifying independence in Bayesian networks. *Networks*, **20**, 507–534.
- Geman, S. & Geman, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Pattern Analysis and Machine Intelligence*, **6**, 721–741.
- Geraci, M. & Bottai, M. (2006). Use of auxiliary data in semi-parametric spatial regression with nonignorable missing responses. *Statistical Modelling*, **6**, 321–336.

- Gianola, D., Fernando, R.L. & Stella, A. (2006). Genomic-assisted prediction of genetic value with semiparametric procedures. *Genetics*, **173**, 1761–1776.
- Hall, P., Humphreys, K. & Titterton, D.M. (2002). On the adequacy of variational lower bound functions for likelihood-based inference in Markovian models with missing values. *Journal of the Royal Statistical Society, Series B*, **64**, 549–564.
- Harezlak, J., Ryan, L.M., Giedd, J.N. & Lange, N. (2005). Individual and population penalized regression splines for accelerated longitudinal designs. *Biometrics*, **61**, 1037–1048.
- Hastie, T. & Zhu, J. (2006). Comment on paper by Moguerza and Muñoz. *Statistical Science*, **21**, 352–357.
- Hennerfeind, A., Brezger, A. & Fahrmeir, L. (2006). Geoadditive survival models. *Journal of the American Statistical Association*, **101**, 1065–1075.
- Jaakkola, T.S. & Jordan, M.I. (2000). Bayesian parameter estimation via variational methods. *Statistics and Computing*, **10**, 25–37.
- Jank, W. & Shmueli, G. (2007). Modelling concurrency of events in on-line auctions via spatiotemporal semiparametric models. *Applied Statistics*, **56**, 1–27.
- Jensen, F.V. (1996). *An Introduction to Bayesian Networks*. London: UCL Press.
- Jordan, M.I. (1999). *Learning in Graphical Models*. Cambridge, Massachusetts: MIT Press.
- Jordan, M.I. (2004). Graphical models. *Statistical Science*, **19**, 140–155.
- Jordan, M.I., Ghahramani, Z., Jaakkola, T.S. & Saul, L.K. (1999). An introduction to variational methods for graphical models. *Machine Learning*, **37**, 183–233.
- Kamman, E.E. & Wand, M.P. (2003). Geoadditive models. *Applied Statistics*, **52**, 1–18.
- Lauritzen, S.L. (1996). *Graphical Models*. Oxford, UK: Clarendon Press.
- Liang, F.M., Truong, Y.K. and Wong, W.H. (2001). Automatic Bayesian model averaging for linear regression and applications in Bayesian curve fitting. *Statistica Sinica*, **4**, 1005–1029.
- 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.
- Minka, T., Winn, J., Guiver, G. & Kannan, A. (2008). *Infer.Net*, Microsoft Research Cambridge, UK.
- Murphy, K. (2007). Software for graphical models: a review. *International Society for Bayesian Analysis Bulletin*, **14**, 13–15.
- Nott, D. (2006). Semiparametric estimation of mean and variance functions for non-Gaussian data. *Computational Statistics*, **21**, 603–620.
- Padoan, S.A. and Wand, M.P. (2008). Mixed-model based additive models for sample extremes. *Statistics and Probability Letters*, doi:10.1016/j.spl.2008.04.009

- Pagano, M. & Gauvreau, K. (1993). *Principles of Biostatistics*. Florence, Kentucky: Duxbury.
- Pearce, N.D. and Wand, M.P. (2006). Penalized splines and reproducing kernel methods. *The American Statistician*, **60**, 233–240.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. San Mateo, California: Morgan Kaufmann.
- R Development Core Team (2008). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0. 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.
- Ruppert, D., Wand, M.P. & Carroll, R.J. (2008). Semiparametric regression during 2003–2007. *Journal of the American Statistical Association*, to appear.
- Schapire, R.E. (1990). The strength of weak learnability. *Machine Learning*, **5**, 197–227.
- Schölkopf, B. & Smola, A.J. (2002). *Learning with Kernels*. Cambridge USA: MIT Press.
- Spiegelhalter, D.J. & Lauritzen, S.L. (1990). Sequential updating of conditional probabilities on directed graphical structures. *Networks*, **20**, 579–605.
- Staudenmayer, J., Lake, E.E. & Wand, M.P. (2008). Robustness for general design mixed models using the t -distribution. *Statistical Modelling*, to appear.
- Tatikonda, S. & Jordan, M.I. (2002). Loopy belief propagation and Gibbs measures. In *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence*, 493–500, San Mateo, California: Morgan Kaufmann.
- Thurston, S., Wand, M.P. & Weincke, J. (2000). Negative binomial additive models, *Biometrics*, **56**, 139–144.
- Titterton, D.M. (2004). Bayesian methods for neural networks and related models. *Statistical Science*, **19**, 128–139.
- Tutz, G. & Binder, H. (2006). Generalized additive modelling with implicit variable selection by likelihood based boosting. *Biometrics*, **62**, 961–971.
- van der Laan, M.J. & Robins, J.M. (2003). *Unified Methods for Censored Longitudinal Data and Causality*. New York: Springer.
- 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.
- Wahba, G. (1990). *Spline Models for Observational Data*. Philadelphia: SIAM.
- Wand, M.P. (2003). Smoothing and mixed models. *Computational Statistics*, **18**, 223–249.
- 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.

Wang, B. & Titterton, D.M. (2004). Convergence properties of a general algorithm for calculating variational Bayesian estimates for a normal mixture model. *Bayesian Analysis*, **1**, 625–650.

Wang, B. & Titterton, D.M. (2006). Lack of consistency of mean field and variational Bayes approximations for state space models. *Neural Processing Letters*, **20**, 151–170.

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

Wasserman, S. & Faust, K. (1994). *Social Networks Analysis: Methods and Applications*. Cambridge UK: Cambridge University Press.

Welham, S.J., Cullis, B.R., Kenward, M.G. & Thompson, R. (2007). A comparison of mixed model splines for curve fitting. *Australian and New Zealand Journal of Statistics*, **49**, 1–23.

Whittaker, J. (1990). *Graphical Models in Applied Multivariate Statistics*. New York: Wiley.

Wright, S. (1934). The method of path coefficients. *Annals of Mathematical Statistics*, **5**, 161–215.

Wood, S.N. (2003). Thin plate regression splines. *Journal of the Royal Statistical Society, Series B*, **65**, 95–114.

Yee, T.W. & Stephenson, A.G. (2007). Vector generalized linear and additive extreme value models. *Extremes*, **10**, 1–19.

Yuan, Y. & Little, R.J.A. (2007). Parametric and semiparametric model-based estimates of the finite population mean for two-stage cluster samples with item nonresponse. *Biometrics*, **63**, 1172–1180.

Zhao, Y., Staudenmayer, J., Coull, B.A. and Wand, M.P. (2006). General design Bayesian generalized linear mixed models. *Statistical Science*, **21**, 35–51.