

Rao-Blackwellization for Bayesian Variable Selection and Model Averaging in Linear and Binary Regression: A Novel Data Augmentation Approach

Joyee Ghosh* and Merlise A. Clyde†

Abstract

Choosing the subset of covariates to use in regression or generalized linear models is a ubiquitous problem. The Bayesian paradigm addresses the problem of model uncertainty by considering models corresponding to all possible subsets of the covariates, where the posterior distribution over models is used to select models or combine them via Bayesian model averaging (BMA). Although conceptually straightforward, BMA is often difficult to implement in practice, since either the number of covariates is too large for enumeration of all subsets, calculations cannot be done analytically, or both. For orthogonal designs with the appropriate choice of prior, the posterior probability of any model can be calculated without having to enumerate the entire model space and scales linearly with the number of predictors, p . In this article we extend this idea to a much broader class of non-orthogonal design matrices. We propose a novel method which augments the observed non-orthogonal design by at most p new rows to obtain a design matrix with orthogonal columns and generate the “missing” response variables in a data augmentation algorithm. We show that our data augmentation approach keeps the original posterior distribution of interest unaltered, and develop methods to construct Rao-Blackwellized estimates of several quantities of interest, including posterior model probabilities of any model, which may not be available from an ordinary Gibbs sampler. Our method can be used for BMA in linear regression and binary regression with non-orthogonal design matrices in conjunction with independent “spike and slab” priors with a continuous prior component that is a Cauchy or other heavy tailed distribution that may be represented as a scale mixture of normals. We provide simulated and real examples to illustrate the methodology. Supplemental materials for the manuscript are available online.

Key words: Gibbs sampling; MCMC; Missing data; Model probability; Model uncertainty; Orthogonal design; Posterior probability; Rao-Blackwell, Variable selection

*Joyee Ghosh is Assistant Professor, Department of Statistics and Actuarial Science, The University of Iowa, Iowa City, IA 52242-1409. Email joyee-ghosh@uiowa.edu

†Merlise A. Clyde is Associate Professor, Department of Statistical Science, Duke University, Durham, NC 27705. Email clyde@stat.duke.edu

1 Introduction

Linear models remain among the most popular methods for modeling relationships between a Gaussian response variable and a set of explanatory variables and provide a foundation for many nonparametric regression methods with the appropriate selection of basis vectors, such as splines, wavelets and kernel regression (Clyde and George 2004). The generalization of the Gaussian linear model to other exponential families leads to the familiar generalized linear models, e.g. probit or logistic regression for binary outcomes. For many applications, it is routine to collect data on many variables of interest and entertain a wide variety of possible models. Computational advances have facilitated the Bayesian treatment of such problems, where the Bayesian approach to model uncertainty proceeds by treating the model as an unknown parameter, and assigns prior probabilities to all possible models under consideration. Given the observed data, these probabilities are updated via Bayes theorem to obtain posterior probabilities of models, which may be used in the selection of a model via decision theoretic approaches taking into consideration other costs (Fouskakis et al. 2009) or in Bayesian model averaging (BMA) using the full joint posterior distribution (see Draper (1995); Hoeting et al. (1999) or Clyde and George (2004) for reviews and additional references).

While it is straightforward to formulate the model uncertainty problem under the Bayesian paradigm, its implementation often becomes non-trivial in large problems, particularly with highly correlated predictors. Heaton and Scott (2010) compare several recent stochastic search algorithms and note that traditional Markov chain Monte Carlo (MCMC) methods for sampling from the posterior distribution of models often fail to reach regions of the model space with high posterior probability. For identifying high probability models, they argue in favor of using alternative stochastic search algorithms (Berger and Molina 2005; Scott and Carvalho 2008; Clyde et al. 2010) that rely on having analytic expressions for marginal likeli-

hoods, such as with Zellner’s g -prior (Zellner 1986) or mixtures of g -priors (Zellner and Siow 1980; Liang et al. 2008); the renormalized likelihoods are used to estimate posterior model probabilities and marginal inclusion probabilities in lieu of (noisy) Monte Carlo frequencies. While these methods are better at identifying high probability models, both Heaton and Scott (2010) and Clyde et al. (2010) note substantial disagreement in estimates such as posterior inclusion probabilities between MCMC methods and these alternative stochastic search methods in high dimensional problems. Clyde and Ghosh (2010) prove that estimators based on renormalized likelihoods lead to biased estimates of inclusion probabilities and other quantities under BMA, which in turn may result in higher mean squared errors. Alternative choices are independent proper “spike and slab” variable selection priors (Ishwaran and Rao 2005) which lead to generalized ridge regression estimates. Despite progress over the last decade, there is clearly still a need for improved algorithms and estimators in the variable selection/model averaging problem.

In the case of linear regression with designs matrices with orthogonal columns and independent normal priors on regression coefficients, one can obtain closed form expressions for posterior probabilities of models and marginal posterior inclusion probabilities (conditional on σ^2), without explicit enumeration of the entire model space of 2^p models. Such “orthogonal designs” arise naturally in the context of designed experiments or wavelet regression, where $p = n$. The computational advantages from orthogonality have been exploited by Chipman et al. (1997); Clyde et al. (1998); Clyde and George (1999, 2000); Johnstone and Silverman (2005), among others, to provide computationally efficient estimators in non-enumerable model spaces that provide adaptive shrinkage with outstanding frequentist properties. For general design matrices, Clyde et al. (1996); Clyde and Parmigiani (1996) transform the non-orthogonal design to an orthogonal design and developed efficient algorithms for Bayesian model averaging (BMA) based on importance sampling. Their approach led to better performance in terms of prediction in many cases, but does not readily lend itself for variable

selection as models are defined in terms of combinations of the original predictor variables. In this paper, we propose a novel method that augments a non-orthogonal design with additional rows, such that the columns of the resulting “complete” design matrix are mutually orthogonal. Unlike Clyde et al. (1996) the Orthogonal Data Augmentation algorithm presented in Section 2 permits both model averaging and model selection in terms of the original variables. The response variables corresponding to the newly introduced rows are treated as missing data and are sampled using MCMC algorithms. Exploiting properties of orthogonal designs, in Section 2.2 we construct “Rao-Blackwellized” (RB) estimates of posterior model probabilities by marginalizing over the missing responses and prove that the RB estimates have smaller variances than the ergodic averages. A key feature of the ODA formulation is the use of Rao-Blackwellization to provide a simple method for estimating the mass of unsampled models. In Section 3 we discuss how to select the augmented design and its properties. We compare the ODA algorithm and RB estimator to other methods in Section 4 using the simulation design of Nott and Kohn (2005), for which $p = 15$, so that enumeration is feasible and estimates can be compared to the truth. We show that our estimates of model probabilities can outperform Monte Carlo estimates and provide accurate estimates of the unsampled mass. The ODA formulation allows extensions to heavier tailed prior distributions on regression coefficients constructed as scale mixtures of normals. In Section 5, we show how to implement ODA with independent Cauchy priors and compare the results to the multivariate Zellner-Siow Cauchy prior (Zellner and Siow 1980). In Section 6 we apply our methods to the ozone data analyzed by Friedman and Silverman (1989) and more recently by Liang et al. (2008), with a model space of dimension 2^{44} that prohibits enumeration. Empirically, we show that ODA does as well or better than BMA with the Zellner g prior, Zellner-Siow Cauchy prior or shrinkage methods using the lasso (Tibshirani 1996) or horseshoe prior (Carvalho et al. 2010). Finally, using the latent variable formulations, we may extend the class of models to binary regression. We illustrate the ODA method for

probit regression in Section 7 using the well known Pima Indian diabetes dataset. In Section 8 we conclude with possible extensions of our method.

2 Orthogonal Data Augmentation

The basis of the popular Expectation-Maximization (EM) algorithm of Dempster et al. (1977) or the Data Augmentation (DA) algorithm of Tanner and Wong (1987) (and their many extensions) is to augment the observed data with missing data, so that the resulting complete data model is much simpler. Letting $\mathbf{Y}_o \equiv (Y_1, Y_2, \dots, Y_{n_o})^T$ denote the observed response vector of length n_o , we may augment it with a vector of length n_a of unknown missing response variables \mathbf{Y}_a , so that the complete response data is denoted as $\mathbf{Y}_c^T = (\mathbf{Y}_o^T, \mathbf{Y}_a^T)$, with length $n_c = n_o + n_a$. In the regression context, we must also specify a design matrix for the missing \mathbf{Y}_a . For fractional factorial experiments, there is a natural choice for this design matrix based on the remaining fraction from the full factorial experiment. We propose a method that can be used to construct a design matrix for the more general observational setting.

To begin, let $\mathbf{X}_o = [\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_p]$ denote the $n_o \times (p+1)$ observed design where \mathbf{x}_0 is an $n \times 1$ vector of ones representing an intercept. Without loss of generality we assume that the remaining p predictor variables \mathbf{x}_j for $j > 0$ have been centered so that they are orthogonal to the intercept ($\mathbf{x}_j^T \mathbf{x}_0 = 0$) and have been rescaled by dividing by their standard deviation (using n_o , rather than $n_o - 1$ in the denominator). This standardization leads to the diagonal elements of $\mathbf{X}_o^T \mathbf{X}_o$ being the sample size n_o and the off-diagonal elements of $\mathbf{X}_o^T \mathbf{X}_o$ being n_o times the correlation of the columns of \mathbf{X}_o . Such standardizations are commonly used in ridge regression or other shrinkage methods such as the lasso so that coefficients may be interpreted on the same scale. The normal linear model using the full design matrix may be

written as

$$\mathbf{Y}_o = \mathbf{X}_o \boldsymbol{\beta} + \boldsymbol{\epsilon}, \quad (1)$$

where $\boldsymbol{\beta} = (\beta_0, \beta_1, \dots, \beta_p)^T$ is a vector of regression coefficients, $\boldsymbol{\epsilon} \sim \mathbf{N}(\mathbf{0}, \sigma^2 \mathbf{I}_{n_o})$, \mathbf{I}_{n_o} is the identity matrix of dimension n_o and σ^2 is the variance.

For the variable selection problem, models may be represented by a binary vector $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_p)^T \in \{0, 1\}^p \equiv \boldsymbol{\Gamma}$, where $\gamma_j = 0$ implies that variable \mathbf{x}_j has been removed from the model or equivalently that $\gamma_j = 0 \Leftrightarrow \beta_j = 0$, so that under model $\boldsymbol{\gamma}$

$$\mathbf{Y}_o \mid \boldsymbol{\beta}_\boldsymbol{\gamma}, \sigma^2, \boldsymbol{\gamma} \sim \mathbf{N}(\mathbf{X}_{o\boldsymbol{\gamma}} \boldsymbol{\beta}_\boldsymbol{\gamma}, \sigma^2 \mathbf{I}_{n_o}) \quad (2)$$

where $\mathbf{X}_{o\boldsymbol{\gamma}}$ is $n \times (p_\boldsymbol{\gamma} + 1)$ design matrix and $\boldsymbol{\beta}_\boldsymbol{\gamma}$ is the $p_\boldsymbol{\gamma} + 1$ dimensional vector of non-zero model specific regression coefficients; by default, the intercept is always assumed to be included in $\mathbf{X}_{o\boldsymbol{\gamma}}$. To transform the problem of model search with a general design matrix to one with an orthogonal design, we augment the rows of the observed design matrix \mathbf{X}_o with a $n_a \times (p + 1)$ design matrix \mathbf{X}_a , such that the resulting “complete” design matrix

$$\mathbf{X}_c = \begin{bmatrix} \mathbf{X}_o \\ \mathbf{X}_a \end{bmatrix} \quad (3)$$

has orthogonal columns, $\mathbf{X}_c^T \mathbf{X}_c = \mathbf{X}_o^T \mathbf{X}_o + \mathbf{X}_a^T \mathbf{X}_a = \mathbf{D}$, where $\mathbf{D} = \text{diag}(\delta_0, \delta_1, \dots, \delta_p)$ is a diagonal matrix with $\delta_j > 0$. For now, we assume that such a matrix will exist and defer discussion of how to select \mathbf{X}_a until Section 3. Under model $\boldsymbol{\gamma}$, given model specific parameters $(\boldsymbol{\beta}_\boldsymbol{\gamma}, \sigma^2)$ and the design \mathbf{X}_a , \mathbf{Y}_a has the same distribution as the observed data \mathbf{Y}_o and is independent of \mathbf{Y}_o , leading to the complete data model

$$\mathbf{Y}_c \mid \boldsymbol{\beta}_\boldsymbol{\gamma}, \sigma^2, \boldsymbol{\gamma} \sim \mathbf{N}(\mathbf{X}_{c\boldsymbol{\gamma}} \boldsymbol{\beta}_\boldsymbol{\gamma}, \sigma^2 \mathbf{I}_{n_c}). \quad (4)$$

We begin with limiting conjugate prior distributions for β_0 and σ^2 and conjugate prior distributions for β_γ

$$p(\beta_0) \propto 1 \quad (5a)$$

$$\beta_j \mid \sigma^2, \lambda_j, \gamma_j \stackrel{ind}{\sim} \mathbf{N} \left(0, \sigma^2 \frac{\gamma_j}{\lambda_j} \right) \text{ for } j = 1, \dots, p \quad (5b)$$

$$p(\sigma^2) \propto 1/\sigma^2 \quad (5c)$$

where β_j is degenerate at 0 if $\gamma_j = 0$. As both β_0 and σ^2 are assumed to be included in all models, this and arguments based on orthogonal parametrizations and invariance to scale and location transformation have been used to justify this objective choice in many variable selection applications, see Liang et al. (2008); Berger et al. (1998) for more details.

Using independent Bernoulli prior distributions on the inclusion indicators γ_j

$$p(\boldsymbol{\gamma} \mid \boldsymbol{\pi}) = \prod_{j=1}^p \pi_j^{\gamma_j} (1 - \pi_j)^{1-\gamma_j} \quad (6)$$

the posterior distribution of $\boldsymbol{\gamma}$ given σ^2 and the complete data \mathbf{Y}_c also has an independent product Bernoulli distribution

$$p(\boldsymbol{\gamma} \mid \mathbf{Y}_c, \sigma^2) = \prod_{j=1}^p \left\{ p(\gamma_j = 1 \mid \mathbf{Y}_c, \sigma^2)^{\gamma_j} (1 - p(\gamma_j = 1 \mid \mathbf{Y}_c, \sigma^2))^{1-\gamma_j} \right\} \quad (7a)$$

with posterior inclusion probability $\rho_j(\mathbf{Y}_c, \sigma^2, \lambda_j) \equiv p(\gamma_j = 1 \mid \mathbf{Y}_c, \sigma^2)$ expressed via posterior odds as

$$\begin{aligned} \rho_j(\mathbf{Y}_c, \sigma^2, \lambda_j) &\equiv O_j(\mathbf{Y}_c, \sigma^2, \lambda_j) / (1 + O_j(\mathbf{Y}_c, \sigma^2, \lambda_j)) \quad (7b) \\ O_j(\mathbf{Y}_c, \sigma^2, \lambda_j) &\equiv \frac{\pi_j}{1 - \pi_j} \left(\frac{\lambda_j}{\delta_j + \lambda_j} \right)^{1/2} \exp \left\{ \frac{1}{2} \frac{\delta_j}{\delta_j + \lambda_j} \frac{\widehat{\beta}_j^2}{\sigma^2 \delta_j} \right\} \text{ for } j = 1, \dots, p. \end{aligned}$$

The estimate $\hat{\beta}_j$ is the j^{th} element of $\hat{\boldsymbol{\beta}} = (\mathbf{X}_c^T \mathbf{X}_c)^{-1} \mathbf{X}_c^T \mathbf{Y}_c$, the vector of OLS regression coefficients of $\boldsymbol{\beta}$ using the complete data $\mathbf{Y}_c, \mathbf{X}_c$. Because $\mathbf{X}_c^T \mathbf{X}_c$ is diagonal, the inversion is trivial. The quantity, $\hat{\beta}_j^2 \delta_j / \sigma^2$ in the exponential term of the posterior odds is the square of the Z-score for testing the hypothesis $\beta_j = 0$, thus the posterior probabilities provide a way of converting frequentist test statistics into the posterior probabilities of hypotheses via Bayes Theorem.

Predictions under orthogonal designs with BMA are greatly simplified. To predict a future \mathbf{Y}^* observed at \mathbf{X}^* , the mean of the posterior predictive distribution

$$E(\mathbf{Y}^* | \mathbf{Y}_c, \sigma^2) = \mathbf{x}_0^* \hat{\beta}_0 + \sum_{j=1}^p \mathbf{x}_j^* p(\gamma_j = 1 | \mathbf{Y}_c, \sigma^2) \frac{\delta_j}{\delta_j + \lambda_j} \hat{\beta}_j \quad (8)$$

takes the form of a multiple shrinkage estimator with linear shrinkage of the OLS estimate induced from the normal prior distribution given that β_j is non-zero, and nonlinear data-dependent shrinkage from the posterior inclusion probability $p(\gamma_j = 1 | \mathbf{Y}_c, \sigma^2)$ that arises from the uncertainty of whether β_j is non-zero. Note that in the orthogonal case, the summation is over p terms as opposed to 2^p terms in BMA with non-orthogonal designs. Of course, while \mathbf{Y}_a is not necessary under orthogonal designs, σ^2 is generally unknown. In the context of wavelets with no missing data ($\mathbf{Y}_c = \mathbf{Y}_o$), Clyde et al. (1998) use robust estimates of σ^2 from the data. As using a “plug-in” estimate for σ^2 may underestimate uncertainty, they suggest using a Gibbs sampler and calculate RB estimates of inclusion probabilities and fitted values. In the next section, we introduce a two-component Gibbs sampler based on the idea of orthogonal data augmentation that allows us to incorporate uncertainty in \mathbf{Y}_a and σ^2 .

2.1 ODA Algorithm

Given the model for the complete data (4) and the prior specification in (5a-5c) and (6), we may construct a two block Gibbs sampler with $[\sigma^2, \mathbf{Y}_a]$ in one block and the vector $[\boldsymbol{\gamma}]$ in the other using the following sequence of distributions:

$$1/\sigma^2 \mid \boldsymbol{\gamma}, \mathbf{Y}_o \sim \text{G} \left(\frac{n_o - 1}{2}, \frac{(\mathbf{Y}_o^T \mathbf{Y}_o - \tilde{\boldsymbol{\beta}}_{o\boldsymbol{\gamma}}^T (\mathbf{X}_{o\boldsymbol{\gamma}}^T \mathbf{X}_{o\boldsymbol{\gamma}} + \Lambda_\boldsymbol{\gamma}) \tilde{\boldsymbol{\beta}}_{o\boldsymbol{\gamma}})}{2} \right) \quad (9a)$$

$$\mathbf{Y}_a \mid \sigma^2, \boldsymbol{\gamma}, \mathbf{Y}_o \sim \text{N} \left(\mathbf{X}_{a\boldsymbol{\gamma}} \tilde{\boldsymbol{\beta}}_{o\boldsymbol{\gamma}}, \sigma^2 (\mathbf{I}_{n_a} + \mathbf{X}_{a\boldsymbol{\gamma}} (\mathbf{X}_{o\boldsymbol{\gamma}}^T \mathbf{X}_{o\boldsymbol{\gamma}} + \Lambda_\boldsymbol{\gamma})^{-1} \mathbf{X}_{a\boldsymbol{\gamma}}^T) \right) \quad (9b)$$

$$\boldsymbol{\gamma}_j \mid \mathbf{Y}_a, \sigma^2, \mathbf{Y}_o \stackrel{ind}{\sim} \text{Ber}(\rho_j(\mathbf{Y}_c, \sigma^2, \lambda_j)) \text{ for } j = 1, 2, \dots, p, \text{ where} \quad (9c)$$

$$\tilde{\boldsymbol{\beta}}_{o\boldsymbol{\gamma}} = (\mathbf{X}_{o\boldsymbol{\gamma}}^T \mathbf{X}_{o\boldsymbol{\gamma}} + \Lambda_\boldsymbol{\gamma})^{-1} \mathbf{X}_{o\boldsymbol{\gamma}}^T \mathbf{Y}_o \quad (10)$$

is the posterior mean for $\boldsymbol{\beta}_\boldsymbol{\gamma}$ under model $\boldsymbol{\gamma}$ and observed data \mathbf{Y}_o , $\rho_j(\mathbf{Y}_c, \sigma^2, \lambda_j)$ is given by equation (7b), and $\Lambda_\boldsymbol{\gamma}$ is a $(p_\boldsymbol{\gamma} + 1) \times (p_\boldsymbol{\gamma} + 1)$ diagonal matrix with first entry $\lambda_0 = 0$ and the remaining $p_\boldsymbol{\gamma}$ entries are given by the subset of λ_j where $\gamma_j = 1$. The block $[\mathbf{Y}_a, \sigma^2]$ is sampled by drawing σ^2 from $p(\sigma^2 \mid \boldsymbol{\gamma}, \mathbf{Y}_o)$, and then drawing \mathbf{Y}_a from $p(\mathbf{Y}_a \mid \boldsymbol{\gamma}, \sigma^2, \mathbf{Y}_o)$. Alternatively the joint distribution could be decomposed by drawing first $\mathbf{Y}_a \mid \boldsymbol{\gamma}, \mathbf{Y}_o$ and then drawing $\sigma^2 \mid \mathbf{Y}_c, \boldsymbol{\gamma}$. In practice we noted no difference, but prefer the former order as the draws for σ^2 are independent of \mathbf{Y}_a given $\boldsymbol{\gamma}$ and are from the full conditional in the collapsed sampler where \mathbf{Y}_a has been integrated out. Because of the independence of $\boldsymbol{\gamma}$ given $[\mathbf{Y}_c, \sigma^2]$, the entire block may be drawn in parallel.

The Markov chain may be used to “marginalize” over the distribution of \mathbf{Y}_a , so that we account for uncertainty in \mathbf{Y}_a appropriately, however, \mathbf{X}_a has been held fixed throughout.

This in fact does not alter the posterior distribution $p(\boldsymbol{\beta}_\gamma, \sigma^2, \gamma \mid \mathbf{X}_o, \mathbf{Y}_o)$ as

$$\begin{aligned}
& p(\boldsymbol{\beta}_\gamma, \sigma^2, \gamma \mid \mathbf{X}_o, \mathbf{X}_a, \mathbf{Y}_o) \\
&= \int p(\mathbf{Y}_a, \boldsymbol{\beta}_\gamma, \sigma^2, \gamma \mid \mathbf{X}_o, \mathbf{X}_a, \mathbf{Y}_o) d\mathbf{Y}_a \\
&= \frac{\int p(\mathbf{Y}_o \mid \boldsymbol{\beta}_\gamma, \sigma^2, \gamma, \mathbf{X}_o) p(\mathbf{Y}_a \mid \boldsymbol{\beta}_\gamma, \sigma^2, \gamma, \mathbf{X}_a) p(\boldsymbol{\beta}_\gamma, \sigma^2 \mid \gamma) p(\gamma) d\mathbf{Y}_a}{\sum_{\gamma \in \Gamma} \iiint p(\mathbf{Y}_o \mid \boldsymbol{\beta}_\gamma, \sigma^2, \gamma, \mathbf{X}_o) p(\mathbf{Y}_a \mid \boldsymbol{\beta}_\gamma, \sigma^2, \gamma, \mathbf{X}_a) p(\boldsymbol{\beta}_\gamma, \sigma^2 \mid \gamma) p(\gamma) d\mathbf{Y}_a d\boldsymbol{\beta}_\gamma d\sigma^2} \\
&= \frac{p(\mathbf{Y}_o \mid \boldsymbol{\beta}_\gamma, \sigma^2, \gamma, \mathbf{X}_o) p(\boldsymbol{\beta}_\gamma, \sigma^2 \mid \gamma) p(\gamma)}{\sum_{\gamma \in \Gamma} \iint p(\mathbf{Y}_o \mid \boldsymbol{\beta}_\gamma, \sigma^2, \gamma, \mathbf{X}_o) p(\boldsymbol{\beta}_\gamma, \sigma^2 \mid \gamma) p(\gamma) d\boldsymbol{\beta}_\gamma d\sigma^2} \\
&= p(\boldsymbol{\beta}_\gamma, \sigma^2, \gamma \mid \mathbf{X}_o, \mathbf{Y}_o) \tag{11}
\end{aligned}$$

is the distribution of interest. We now show how Rao-Blackwellization may be used with the ODA sampler to reduce Monte Carlo variation in estimates.

2.2 Rao-Blackwellized Estimates of Quantities of Interest

Ergodic averages given the output of MCMC may be used to estimate functions of interest, such as model probabilities, inclusion probabilities and predictions. Let $\mathbf{1}(\gamma = \gamma^*)$ denote the indicator function for the event $\gamma = \gamma^*$, then the ergodic average or Monte Carlo (MC) frequency estimator of the probability of model γ^* is $\widehat{p}^{\text{MC}}(\gamma^* \mid \mathbf{Y}_o) = \frac{1}{K} \sum_{k=1}^K \mathbf{1}(\gamma^{(k)} = \gamma^*)$, where K is the number of draws from the Markov chain. Similarly posterior inclusion probabilities may be estimated by taking an ergodic average of $\mathbf{1}(\gamma_j^{(k)} = 1)$ and in general for a function $t(\gamma)$,

$$\widehat{t}^{\text{MC}}(\gamma) = \frac{1}{K} \sum_{k=1}^K t(\gamma^{(k)}). \tag{12}$$

“Rao-Blackwellization” of estimates (Gelfand and Smith 1990) has been suggested as a way to construct improved estimates by replacing $t(\gamma^{(k)})$ with its expectation given other components in the sampler, and then applying the ergodic average to marginalize over the

components. Applying this to the ODA sampler, leads to RB estimates of $t(\boldsymbol{\gamma})$

$$\widehat{t}^{\text{RB}}(\boldsymbol{\gamma}) = \frac{1}{K} \sum_{k=1}^K \mathbb{E}[t(\boldsymbol{\gamma}^{(k)}) \mid \mathbf{Y}_o, \mathbf{Y}_a^{(k)}, \sigma^{2(k)}] = \frac{1}{K} \sum_{k=1}^K \mathbb{E}[t(\boldsymbol{\gamma}^{(k)}) \mid \mathbf{Y}_c^{(k)}, \sigma^{2(k)}]. \quad (13)$$

Because of the product Bernoulli posterior for $\boldsymbol{\gamma}$ conditional on \mathbf{Y}_a, σ^2 , analytic expressions for the conditional expectation of the indicators of models, individual inclusion indicators or predictions is trivial. For ODA, the RB estimators of model probabilities and inclusion probabilities take the form

$$\widehat{p}^{\text{RB}}(\boldsymbol{\gamma} \mid \mathbf{Y}_o) = \frac{1}{K} \sum_{k=1}^K \prod_{j=1}^p \rho_j(\mathbf{Y}_c^{(k)}, \sigma^{2(k)}, \lambda_j)^{\gamma_j} (1 - \rho_j(\mathbf{Y}_c^{(k)}, \sigma^{2(k)}, \lambda_j))^{1-\gamma_j} \quad (14a)$$

$$\widehat{p}^{\text{RB}}(\gamma_j = 1 \mid \mathbf{Y}_o) = \frac{1}{K} \sum_{k=1}^K \rho_j(\mathbf{Y}_c^{(k)}, \sigma^{2(k)}, \lambda_j) \quad (14b)$$

where $\rho_j(\mathbf{Y}_c, \sigma^2, \lambda_j)$ is the marginal inclusion probability given $\mathbf{Y}_c, \sigma^2, \lambda_j$ (7b). For a matrix \mathbf{X} (at observed or new data points), RB estimates for model averaging of $\mathbf{X}\boldsymbol{\beta}$ are given by

$$\widehat{\mathbf{E}}^{\text{RB}}(\beta_j \mid \mathbf{Y}_o) = \frac{1}{K} \sum_{k=1}^K \rho_j(\mathbf{Y}_c^{(k)}, \sigma^{2(k)}, \lambda_j) \frac{\delta_j}{\delta_j + \lambda_j} \widehat{\beta}_j^{(k)} \quad (14c)$$

$$\widehat{\mathbf{E}}^{\text{RB}}(\mathbf{X}\boldsymbol{\beta} \mid \mathbf{Y}_o) = \mathbf{x}_0 \frac{1}{K} \sum_{k=1}^K \widehat{\beta}_0^{(k)} + \sum_{j=1}^p \mathbf{x}_j \widehat{\mathbf{E}}^{\text{RB}}(\beta_j \mid \mathbf{Y}_o), \quad (14d)$$

where $\widehat{\beta}_j^{(k)}$ is the least squares estimate of β_j from the complete data $\mathbf{Y}_c^{(k)}, \mathbf{X}_c$ for the k^{th} iteration.

The RB estimate in (14a) may be used to estimate the posterior model probability for any model, including those that are not sampled in the MCMC (whose ergodic average is 0). This gives us a simple way of estimating the remaining unsampled posterior mass. Let Γ_K denote the set of unique models in the K draws from the ODA sampler. Because the estimates of model probabilities given by equation (14a) sum to 1 for the entire model space

Γ , we can estimate the total posterior probability of unsampled models by

$$\widehat{p}^{RB}(\Gamma - \Gamma_K | \mathbf{Y}_o) = \sum_{\gamma \in \Gamma - \Gamma_K} \widehat{p}^{RB}(\gamma | \mathbf{Y}_o) = 1 - \sum_{\gamma \in \Gamma_K} \widehat{p}^{RB}(\gamma | \mathbf{Y}_o). \quad (15)$$

RB estimates are commonly believed to reduce the variance over ergodic averages, however, Liu et al. (1994) and Geyer (1995) show that this is not always true when the dependence in the Markov chain is taken into account. For the two block ODA algorithm, however, Rao-Blackwellization is guaranteed to result in variance reduction:

Theorem 1. *Under the assumption that the Markov chain induced by the ODA sampler has reached its stationary distribution*

$$\begin{aligned} \text{var}[\widehat{p}^{RB}(\gamma | \mathbf{Y}_o)] &\leq \text{var}[\widehat{p}^{MC}(\gamma | \mathbf{Y}_o)] \\ \text{var}[\widehat{p}^{RB}(\gamma_j = 1 | \mathbf{Y}_o)] &\leq \text{var}[\widehat{p}^{MC}(\gamma_j = 1 | \mathbf{Y}_o)] \end{aligned}$$

Proof. The ODA algorithm is a 2-component Gibbs sampler with components $[\boldsymbol{\gamma}]$ and $[\mathbf{Y}_a, \sigma^2]$, which is also referred to as “data augmentation”. The proof follows immediately from Theorem 4.1 of Liu et al. (1994), who show that for such a Gibbs sampler, the RB estimator leads to a reduction in variance compared to the Monte Carlo estimate in estimating $\mathbb{E}[t(\boldsymbol{\gamma}) | \mathbf{Y}_o]$ for any scalar function $t(\boldsymbol{\gamma})$. \square

Liu et al. (1994) and Liu (1994) also contrast collapsing and grouping as ways of reducing Monte Carlo variation and recommend collapsing in general. A collapsed Gibbs sampler results from integrating out one or more components of a Gibbs sampler analytically. Grouping or blocking refers to drawing two or more components together from their joint conditional distribution, rather than drawing them individually from their respective conditional distributions. Because of the conjugate priors for $\boldsymbol{\beta}_\gamma, \sigma^2$, the marginal posterior distribution for

γ is available as

$$p(\gamma | \mathbf{Y}_o) = \frac{p(\mathbf{Y}_o | \gamma)p(\gamma | \pi)}{\sum_{\gamma \in \Gamma} p(\mathbf{Y}_o | \gamma)p(\gamma | \pi)}, \text{ where} \quad (16)$$

$$p(\mathbf{Y}_o | \gamma) \propto |\Lambda_\gamma|^{+1/2} |\mathbf{X}_{\sigma\gamma}^T \mathbf{X}_{\sigma\gamma} + \Lambda_\gamma|^{-1/2} \left(\|\mathbf{Y}_o - \mathbf{X}_{\sigma\gamma} \tilde{\boldsymbol{\beta}}_{\sigma\gamma}\|^2 + \tilde{\boldsymbol{\beta}}_{\sigma\gamma}^T \Lambda_\gamma \tilde{\boldsymbol{\beta}}_{\sigma\gamma} \right)^{-\frac{n-1}{2}}, \quad (17)$$

$|\Lambda_\gamma|^+$ is the determinant of the lower positive definite $p_\gamma \times p_\gamma$ block of Λ_γ obtained by excluding the first row and column, and $\tilde{\boldsymbol{\beta}}_{\sigma\gamma}$ is the posterior mean of $\boldsymbol{\beta}_\gamma$ under the observed data only given by equation (10). As it is not possible to sum over the models in Γ for large problems, a collapsed Gibbs sampler may be used to draw the components of γ from the Bernoulli full conditional distributions $p(\gamma_j | \boldsymbol{\gamma}_{(j)}, \mathbf{Y}_o)$, where $\boldsymbol{\gamma}_{(j)}$ is $\boldsymbol{\gamma}$ with the j^{th} component removed. This Gibbs sampler is equivalent to the popular integrated Stochastic Search Variable Selection (SSVS) algorithm introduced by George and McCulloch (1997) and may also be viewed as a collapsed Gibbs sampler induced by integrating out \mathbf{Y}_a and σ^2 in the ODA algorithm. Liu (1994) shows that collapsing reduces the norm of the forward operator of the Gibbs sampler over that of a grouped Gibbs sampler, although this does not imply faster convergence or better mixing for non-reversible chains (Liu 1994; Liu et al. 1994). In general, one must balance computational ease, speed and mixing. By introducing the missing \mathbf{Y}_a, σ^2 , we are able to generate $\boldsymbol{\gamma}$ in one block, while in the collapsed case, we must propose each γ_j conditional on the remaining components. One iteration of the collapsed Gibbs sampler requires p marginal likelihood evaluations and the solution of equation (10) to obtain a new model; thus standard implementations ¹ of SSVS are computationally more expensive than a single iteration of ODA, even with the generation of the latent \mathbf{Y}_a and σ^2 .

¹George and McCulloch (1997) discuss efficient updating under g -priors

3 Existence and Choice of Augmented Designs

Our goal is to find a matrix \mathbf{X}_a given the observed matrix \mathbf{X}_o such that $\mathbf{X}_c^T \mathbf{X}_c = \mathbf{X}_o^T \mathbf{X}_o + \mathbf{X}_a^T \mathbf{X}_a = \mathbf{D}$ is diagonal. Furthermore, \mathbf{X}_a must be real and $\mathbf{X}_a^T \mathbf{X}_a$ should be positive semi-definite (psd), which implies that it has real, non-negative eigenvalues. There is a long rich history describing conditions on the possible eigenvalues of the sums of two symmetric (Hermitian) matrices (see Fulton 2000, for a summary and key results). The earliest significant results are due to Weyl (1912) and may be restated for our purposes in the following lemma:

Lemma 1. *Let \mathbf{A} , \mathbf{O} and \mathbf{D} be three real symmetric matrices of dimension d with ordered eigenvalues $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_d$, $o_1 \geq o_2 \dots \geq o_d$ and $\delta_1 \geq \delta_2 \geq \dots \delta_d$, respectively. Then for $\mathbf{D} = \mathbf{A} + \mathbf{O}$, $\delta_{i+j-1} \leq \alpha_i + o_j$ whenever $i + j - 1 \leq d$.*

Given \mathbf{D} and $\mathbf{X}_o^T \mathbf{X}_o$, the theorem below suggests how we may now construct \mathbf{X}_a .

Theorem 2. *Let $\mathbf{O} \equiv \mathbf{X}_o^T \mathbf{X}_o$ denote a psd symmetric matrix with ordered eigenvalues $o_1 \geq o_2 \geq \dots \geq o_{p+1} \geq 0$ and \mathbf{D} be a diagonal matrix with elements $\delta_1 \geq \delta_2 \geq \dots \geq \delta_{p+1}$. If $\delta_i \geq o_1$ for $1 \leq i \leq p+1$, then a real augmented design matrix \mathbf{X}_a of dimension $(p+1) \times (p+1)$ always exists.*

Proof. As both \mathbf{O} and \mathbf{D} are symmetric with real eigenvalues, $\mathbf{A} \equiv \mathbf{D} - \mathbf{X}_o^T \mathbf{X}_o$ is a real symmetric matrix with spectral decomposition $\mathbf{U} \boldsymbol{\alpha} \mathbf{U}^T$ for \mathbf{U} a $(p+1) \times (p+1)$ orthogonal matrix and diagonal matrix $\boldsymbol{\alpha}$ of ordered eigenvalues $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_{p+1}$. From Lemma 1, $\delta_{i+j-1} - o_j \leq \alpha_i$ whenever $i + j - 1 \leq p + 1$, thus in order for \mathbf{A} to be positive semi-definite, $\delta_{i+j-1} \geq o_j$ for $i + j - 1 \leq p + 1$, and in particular for $j = 1$, $\delta_i \geq o_1$ for $1 \leq i \leq p + 1$. Setting $\mathbf{X}_a = \mathbf{U} \boldsymbol{\alpha}^{1/2} \mathbf{U}^T$, the symmetric matrix square root of \mathbf{A} , ensures that $\mathbf{X}_a^T \mathbf{X}_a \equiv \mathbf{A}$ is psd and that \mathbf{X}_a is real. \square

3.1 Choice of D

If $\delta_i = o_1$ for $i = 1, \dots, p+1$, then the eigenvalues of $\mathbf{X}_a^T \mathbf{X}_a$ will be $o_1 - o_i$ and there will be $r = p+1 - m$ non-zero eigenvalues, where m is the number of eigenvalues of \mathbf{O} that are equal to o_1 . The construction of $\mathbf{X}_a^T \mathbf{X}_a$ above reduces to a $(p+1) \times (p+1)$ rank r matrix, $\mathbf{X}_a^T \mathbf{X}_a = \sum_{i=m+1}^{p+1} (o_1 - o_i) \mathbf{u}_i \mathbf{u}_i^T$ where \mathbf{u}_i are the eigenvectors of \mathbf{U} . If all the eigenvalues of $\mathbf{X}_o^T \mathbf{X}_o$ are distinct then this implies $\mathbf{X}_a^T \mathbf{X}_a$ is of rank p .

Since the columns of \mathbf{X}_o have been standardized, the choice of equal diagonal elements of D , $\delta = \delta_1 = \dots = \delta_{p+1}$ is not unreasonable as the columns of \mathbf{X}_o are all on the same scale. Furthermore, this choice leads to the same sample precision for all β_γ under the complete data. In the special case that the original matrix has orthogonal columns, the eigenvalues $o_i = n_o$, thus the augmented design under the choice $\delta = o_1$ leads to a rank 0 matrix, hence augmentation is unnecessary.

Large δ increases the column sum of squares for \mathbf{X}_c and hence will lead to less Monte Carlo variation for β_γ under the augmented data posterior. However, this also inflates the column sum of squares of \mathbf{X}_a as $\mathbf{X}_o^T \mathbf{X}_o$ is fixed. As a consequence, the augmented design points may be large in magnitude in contrast to the observed points, which increases the leverage of the augmented points.

Theorem 3. *Let $\mathbf{X}_a = \mathbf{U} \boldsymbol{\alpha}^{1/2} \mathbf{U}^T$ where $\mathbf{U} \boldsymbol{\alpha} \mathbf{U}^T = \delta \mathbf{I} - \mathbf{X}_o^T \mathbf{X}_o$ for $\delta = o_1 + \epsilon$ and o_1 the maximum eigenvalue of $\mathbf{X}_o^T \mathbf{X}_o$. Then $\epsilon = 0$ minimizes the leverage of the design points in \mathbf{X}_a .*

Proof. The leverages under the complete design are the diagonal elements of

$$\mathbf{X}_c (\mathbf{X}_c^T \mathbf{X}_c)^{-1} \mathbf{X}_c^T = \delta^{-1} \begin{bmatrix} \mathbf{X}_o \mathbf{X}_o^T & \mathbf{X}_o \mathbf{X}_a^T \\ \mathbf{X}_a \mathbf{X}_o^T & \mathbf{X}_a \mathbf{X}_a^T \end{bmatrix}$$

with lower block, $\delta^{-1} \mathbf{X}_a \mathbf{X}_a^T = \delta^{-1} \mathbf{U} \boldsymbol{\alpha}^{1/2} \mathbf{U}^T \mathbf{U} \boldsymbol{\alpha}^{1/2} \mathbf{U}^T = \delta^{-1} \mathbf{U} \boldsymbol{\alpha} \mathbf{U}^T$. The leverage for row i

simplifies to $\sum_j (1 - o_j/\delta) u_{ij}^2$. Since $\delta \geq o_1$, the leverages are minimized for $\delta = o_1$. \square

While ideally we set $\delta = o_1$, in practice this choice sometimes leads to computationally unstable solutions with a small negative value for the smallest eigenvalue of $\mathbf{X}_\alpha^T \mathbf{X}_\alpha$. To ensure that $\mathbf{X}_\alpha^T \mathbf{X}_\alpha$ is positive definite we use $\epsilon = 0.001$. This definition of \mathbf{D} will be assumed throughout the rest of the paper.

3.2 Invariance

In Theorem 2, we took \mathbf{X}_α to be the symmetric square root of $\mathbf{D} - \mathbf{X}_o^T \mathbf{X}_o$, however, any matrix square root of $\mathbf{D} - \mathbf{X}_o^T \mathbf{X}_o$ may be used to create an augmented design. We characterize the possible solutions and show that the distribution of the augmented \mathbf{Y}_α is invariant under the choice of square root.

Lemma 2. *Let \mathcal{S} denote the set of $(p+1) \times (p+1)$ matrix square roots of $\mathbf{A} = \mathbf{D} - \mathbf{X}_o^T \mathbf{X}_o$ for $\mathbf{A} > 0$ such that for $\mathbf{S} \in \mathcal{S}$, $\mathbf{S}^T \mathbf{S} = \mathbf{A}$. Let \mathcal{O}_{p+1} denote the group of $(p+1) \times (p+1)$ orthogonal matrices with group action (\cdot) matrix multiplication. Then any $\mathbf{S}^* \in \mathcal{S}$ may be written as $\mathbf{S}^* = \mathbf{O}\mathbf{S}$ for some $\mathbf{O} \in \mathcal{O}_{p+1}$ and \mathbf{S} in \mathcal{S} .*

Proof. The orthogonal group “acts” on the left of \mathcal{S} (see Eaton 1983, page 186 definition 6.1), that is for all $\mathbf{S} \in \mathcal{S}$ and $\mathbf{O}_1, \mathbf{O}_2 \in \mathcal{O}_{p+1}$, $(\mathbf{O}_1 \mathbf{O}_2) \cdot \mathbf{S} = \mathbf{O}_1 \cdot (\mathbf{O}_2 \cdot \mathbf{S})$ with identity group element \mathbf{I}_{p+1} , $\mathbf{I}_{p+1} \cdot \mathbf{S} = \mathbf{S}$. Then proposition 6.1 of Eaton (1983) may be used to show that $\mathbf{O}\mathbf{S}$ is one-to-one and onto from \mathcal{S} to \mathcal{S} . \square

Theorem 4. *Let $\mathbf{Y}_\alpha \sim N(\mathbf{S}\boldsymbol{\beta}, \sigma^2 \mathbf{I})$ where $\mathbf{S} \in \mathcal{S}$, the set of matrix square roots of $\delta \mathbf{I} - \mathbf{X}_o^T \mathbf{X}_o$. Then the distribution of \mathbf{Y}_α is invariant to the choice of \mathbf{S} .*

Proof. For $\mathbf{O} \in \mathcal{O}_{p+1}$, $\mathbf{Y}_\alpha^* \equiv \mathbf{O}\mathbf{Y}_\alpha \stackrel{d}{=} \mathbf{O}\mathbf{S}\boldsymbol{\beta} + \mathbf{O}\boldsymbol{\epsilon}$ where $\boldsymbol{\epsilon} \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$. Since $\mathbf{S}^* = \mathbf{O}\mathbf{S} \in \mathcal{S}$, $\mathbf{Y}_\alpha^* \sim N(\mathbf{S}^* \boldsymbol{\beta}, \sigma^2 \mathbf{I})$. \square

The above result says that if \mathbf{Y}_a is an augmentation such that $\mathbf{Y}_a \sim \mathbf{N}(\mathbf{S}\boldsymbol{\beta}, \sigma^2\mathbf{I})$, then for any other matrix square root of $\delta\mathbf{I} - \mathbf{X}_o^T\mathbf{X}_o$, say \mathbf{S}^* , (where $\mathbf{S}^* = \mathbf{O}\mathbf{S}$ for some $\mathbf{O} \in \mathcal{O}_{p+1}$), the corresponding augmentation \mathbf{Y}_a^* (obtained under left-multiplication by the orthogonal matrix \mathbf{O}) is also normally distributed as $\mathbf{N}(\mathbf{S}^*\boldsymbol{\beta}, \sigma^2\mathbf{I})$. Because it is fast to compute, we take \mathbf{X}_a to be the Cholesky decomposition of $\mathbf{D} - \mathbf{X}_o^T\mathbf{X}_o$ for the remainder of the paper. Results for sensitivity analysis using two other choices of \mathbf{X}_a based on singular value decompositions are reported in the Supplemental Materials. These results suggest that the ODA algorithm is not sensitive to the choice of \mathbf{X}_a .

4 Simulation Study

We compare ODA and other MCMC samplers for data generated using the simulation design in Nott and Kohn (2005) (similar to Raftery et al. (1997); Fernández et al. (2001)) with a sample size $n = 50$ and 15 predictors. The first column of the design matrix \mathbf{X} is a column of ones and the next 10 columns, $\mathbf{x}_1, \dots, \mathbf{x}_{10}$, are generated as independent $\mathbf{N}(0, 1)$ variables. The last 5 columns, $\mathbf{x}_{11}, \dots, \mathbf{x}_{15}$ depend on the first five predictors $[\mathbf{x}_{11}, \dots, \mathbf{x}_{15}] = [\mathbf{x}_1, \dots, \mathbf{x}_5] (0.3, 0.5, 0.7, 0.9, 1.1)^T (1, 1, 1, 1, 1) + \mathbf{E}$ where \mathbf{E} is a 50×5 matrix of independent $\mathbf{N}(0, 1)$ random variables. This induces strong correlations among the last five variables and moderate correlations between them and the first five predictors. Given \mathbf{X} , the response variable is generated as $\mathbf{Y}_o = \mathbf{X}\boldsymbol{\beta}^* + \boldsymbol{\epsilon}$, where $\boldsymbol{\beta}^* = (4, 2, 0, 0, 0, -1, 0, 1.5, 0, 0, 0, 1, 0, 0.5, 0, 0)^T$ and $\boldsymbol{\epsilon} \sim \mathbf{N}(0, 2.5^2\mathbf{I})$.

With 15 predictors and conjugate prior distributions, the model space of $2^{15} = 32,768$ possible models may be enumerated easily to provide the true posterior distribution under model averaging. For all methods we use a design matrix \mathbf{X}_o obtained by standardizing \mathbf{X} as described in Section 2 and set $\lambda_j = 1$ in the Gaussian prior distribution for the coefficients in (5b). If the predictors were uncorrelated, this choice would be equivalent to a unit-

information prior (Kass and Raftery 1995). To complete the prior specification, we take $\pi_j = 1/2$ corresponding to an uniform prior distribution on the model space.

We compare ODA with two alternative MCMC algorithms. The first is SSVS, which is a collapsed version of ODA, described previously. Because the SSVS algorithm and the related MC³ algorithm of Raftery et al. (1997) utilize one component at a time updates, high correlation among predictors may make it difficult to escape local modes, leading to poor mixing in practice. To improve mixing, we use a Metropolis-Hastings algorithm based on a mixture kernel that randomly selects a γ_j and switches γ_j to $1 - \gamma_j$ thus adding or deleting a variable as in the MC³ proposal (Raftery et al. 1997), combined with a random swap proposal that randomly exchanges a predictor in the current model with one that is not. This simple random swap (RS) algorithm greatly aids in escaping local modes (Denison et al. 1998; Nott and Green 2004; Clyde et al. 2010) and often performs as well as the more complicated Swendsen-Wang method of Nott and Green (2004). All MCMC methods utilize the same likelihood for the observed data and prior distributions given above for ODA, so that all methods should provide samples from the same target distribution.

For SSVS and RS we compute estimates of posterior model probabilities and inclusion probabilities based on i) Monte Carlo frequencies of visits to models (MC) and ii) marginal likelihoods of the unique models in the sample renormalized (RM) over the set of unique sampled models (Γ_K),

$$\hat{p}^{\text{RM}}(\gamma | \mathbf{Y}_o) = \frac{p(\mathbf{Y}_o | \gamma)p(\gamma)}{\sum_{\gamma \in \Gamma_K} p(\mathbf{Y}_o | \gamma)p(\gamma)} I(\gamma \in \Gamma_K), \quad \hat{\mathbf{p}}^{\text{RM}}(\gamma_j = 1) = \sum_{\gamma \in \Gamma_K} \gamma_j \hat{p}^{\text{RM}}(\gamma | \mathbf{Y}_o). \quad (18)$$

4.1 Comparisons of Algorithms and Estimators

We compared the three algorithms based on a fixed running time as a single “iteration” of each of the samplers is not equivalent: ODA updates γ as one block with a much simpler full conditional, but requires the generation of the $p + 1$ dimensional vector \mathbf{Y}_a and the

scalar σ^2 ; one iteration of SSVS cycles through the p components of γ with p marginal likelihood evaluations for the full conditionals; RS changes one or two components of γ and requires a single marginal likelihood evaluation in the Metropolis-Hastings ratio. Running each algorithm for one hour² results in 450,000, 46,667, and 700,000 iterations for ODA, SSVS, and RS, respectively; all simulations were done in R on a Dell workstation with 8 3.2 GHz Intel Xeon CPUs. Ideally one should run MCMC algorithms long enough to ensure convergence of model probabilities and marginal inclusion probabilities, however, this would require the number of MCMC iterations to be much more than the size of the model space. As this is computationally prohibitive for large model spaces, the typical number of draws is only a small fraction of the dimension of the model space. To reproduce such a scenario but where we can still evaluate the exact posterior distribution, we also run the algorithms for one minute, resulting in 6,400, 667, and 10,000 draws (19.5%, 2%, 30.5% of the model space) for ODA, SSVS and RS, respectively. For comparison, enumeration of the model space took 2.7 minutes. To compare how well the different estimators and algorithms estimate model probabilities and inclusion probabilities, we use $\text{MSE} = \|\hat{\mathbf{p}} - \mathbf{p}\|^2/d$ where $\hat{\mathbf{p}}$ is an estimate of the vector of probabilities, \mathbf{p} is the value obtained under enumeration of the model space, and d is the length of the vector, $d = 2^p$ for the model probabilities and $d = p$ for inclusion probabilities. For each method, we ran the MCMC algorithm 100 times with random starting values, and computed the average MSE, taking the mean over the 100 runs (Table 1).

The results in Table 1 indicate that the Monte Carlo estimates from ODA are better than MC estimates from the other algorithms for the same running time. Likewise using the renormalized marginal likelihoods (RM) is best under the ODA chain. The root MSE results show that even within the short run we may estimate inclusion probabilities within plus or minus 0.02 under ODA RB. The Rao-Blackwellization leads to a reduction in MSE over the MC estimates for both short and long runs, as the theory suggests, however, the renormalized

²All algorithms were programmed in native R

Square Root Average MSE	Short Runs	Long Runs	Short Runs	Long Runs
	Incl. Prob. ($\times 100$)	Incl. Prob. ($\times 100$)	Model Prob. ($\times 10^4$)	Model Prob. ($\times 10^4$)
ODA MC	1.28	0.15	0.83	0.09
ODA RB	1.11	0.13	0.39	0.05
ODA RM	2.43	0.12	0.66	0.02
SSVS MC	1.98	0.23	2.32	0.27
SSVS RM	6.48	0.73	3.22	0.14
RS MC	2.03	0.24	1.42	0.17
RS RM	3.02	0.14	0.92	0.02

Table 1: Nott-Kohn simulation example: Comparison of square root of the average mean squared error for estimation of the 15 posterior inclusion probabilities (Incl. Prob.) and 2^{15} posterior model probabilities (Model Prob.) using Monte Carlo frequencies (MC), Rao-Blackwellized (RB) and renormalized marginal likelihoods (RM) under short runs (1 minute) and long runs (1 hour) of each of the ODA, SSVS, and RS algorithms. The calculations are based on 100 replicates of each algorithm.

estimates are better only in the long runs. As the RM estimates are Fisher consistent (they equal the true value when the population is enumerated), the MSE goes to zero as the number of unique models converges to $|\Gamma|$. The long runs of ODA, SSVS, and RS sample 8484, 4018 and 8274 unique models (on average), corresponding to 99.5, 96.27, and 99.4% of the posterior mass (respectively). However, for the short runs ODA, SSVS and RS sample 84.9, 52.1, and 80.2% of the mass, respectively, and the superior performance of the RM estimates of inclusion probabilities is diminished over the MC estimates for all algorithms. These results confirm those in Clyde and Ghosh (2010), who show that renormalization may result in larger MSE over MC estimates when the number of iterations is small relative to the model space, because of the inherent bias in RM estimators. As ODA RB has better performance compared to MC and RM in the short runs, this should provide improved estimates for non-enumerable model spaces where a smaller fraction of the posterior mass is sampled. We now turn to estimation of the unsampled mass, which may provide some guidance when ODA RB is preferable to ODA RM and whether one should continue running

the MCMC.

4.2 Estimation of Unsampled Posterior Mass

We can use the RB estimate in (15) to estimate the remaining posterior mass for the Nott-Kohn simulations and compare them to the exact value under enumeration. Using a short preliminary sample from the MCMC, George and McCulloch (1997) constructed an estimate of the normalizing constant C , for the posterior distribution of γ given in (16), $p(\gamma | \mathbf{Y}_o) = Cp(\mathbf{Y}_o | \gamma)p(\gamma | \boldsymbol{\pi})$, which in turn leads to an estimate of the unsampled mass. We found that the average bias was negligible for both estimators, although the George and McCulloch estimator exhibits much greater run-to-run variability (See Figure 1 in Supplemental Materials).

This ODA estimate of unsampled mass exhibits a slight tendency towards underestimation; a possible explanation is included in the Supplemental Materials. Although the bias is negligible, we provide two alternatives to correct it. One option is to run an independent Markov chain, and calculate the estimate in equation (15), based on (\mathbf{Y}_a, σ^2) generated from this new chain. We refer to this unbiased estimate as the Independent Rao-Blackwellized (IRB) estimate. Another approach that does not require additional simulation is to split randomly the MCMC samples from the original chain into two halves. The collection of (\mathbf{Y}_a, σ^2) from the first half can be used to estimate the probabilities of models from the second half and vice-versa; adding the two estimates provides an estimate of the sampled mass from the entire chain, which is then subtracted from one to estimate the remaining mass. We call this the Rao-Blackwell Split (RB-Split) estimate. The Supplemental Materials contain an expanded version of this section with a detailed comparison of all estimators. Among the methods, IRB is the best, but at the expense of running a second independent chain. As a compromise that reduces bias, but does not increase computational complexity, we recommend the RB Split method.

A key feature of the ODA method for estimating probabilities is that it does not require the marginal likelihoods to be available in closed form unlike the method of George and McCulloch (1997) to provide estimates of the unsampled mass. For enumerable problems where marginal likelihoods are not available in closed form, ODA clearly has an advantage in estimating not only the unsampled probability but also model probabilities for *all* models, as we illustrate in Section 7. We first show how to extend the method to allow alternatives to Gaussian prior distributions, another important example where marginal likelihoods are not available.

5 ODA with Scale Mixtures of Normals

Normal prior distributions on the non-zero coefficients as in (5b) have been widely used for Bayesian variable selection, as the conjugate prior permits tractable marginal likelihoods used by collapsed MCMC algorithms that sample only from the model space. While normal priors are appealing from a computational perspective, Jeffreys (1961) rejected the normal prior for hypothesis testing of a mean from a single normal population and recommended Cauchy priors instead. There are two potential problems with Gaussian prior distributions. First, posterior probabilities are sensitive to the prior variance, where large variances may lead to false rejection of the hypothesis that the coefficient is non-zero, the well known “Lindley’s” or “Bartlett’s” paradox. Second, the normal prior does not have bounded influence, thus even as the t statistic increases to infinity, the prior mean of zero exerts influence on the marginal likelihood so that Bayes factors in favor of the alternative do not go to infinity but are instead bounded. The lack of bounded influence also affects the posterior mean, as the linear shrinkage of the sample mean is independent of the data. Liang et al. (2008) show that these problems arise in the regression context using the normal g -prior (Zellner 1986) and prove that scale mixtures of Zellner’s g -prior, obtained by placing a prior distribution

on g resolve these problems. These mixtures of g -priors are multivariate priors with the same prior correlation structure as in the observed data, which may not be desirable in the presence of multicollinearity. In the orthogonal wavelet setting, Clyde and George (2000) use scale mixtures of independent normal distributions for both error distributions as well as wavelet coefficients, which provides both robustness to using a prior mean of zero and robustness to outliers in the data. Johnstone and Silverman (2004, 2005) also demonstrate the superior performance of mixtures of a point mass and heavy-tailed prior distributions compared to mixtures of a point mass and Gaussian prior distributions. Scale mixtures of normal distributions have also seen a resurgence of interest for shrinkage estimators (priors with no explicit probability at zero) (Carvalho et al. 2010; Polson and Scott 2010) where MAP estimates are easy to obtain. These scale mixtures may be easily incorporated in the ODA algorithm by placing a prior distribution on the λ_j in (5b), leading to heavy-tailed prior distributions like the Cauchy, Student- t , double exponential, or horseshoe for the non-zero regression coefficients. We illustrate the algorithm for Cauchy prior distributions below.

5.1 ODA for Cauchy Prior Distributions

We add another layer of hierarchy to the normal prior in (5b):

$$\lambda_j \stackrel{iid}{\sim} \text{G}(\alpha/2, \alpha/2) \quad j = 1, \dots, p, \quad (19)$$

which leads to a marginal distribution for β_j given γ_j and σ^2 that is Student- t with location and scale parameters 0 and σ^2 , and α degrees of freedom. In particular, the choice $\alpha = 1$ corresponds to $\beta_j \mid \gamma_j, \sigma^2 \stackrel{iid}{\sim} \text{C}(0, \sigma^2 \gamma_j)$. Collapsing over γ_j , the prior distribution for β_j is a mixture of a point mass at zero and a Cauchy distribution, a robust version of the “spike-and-slab” prior.

With Cauchy prior distributions, marginal likelihoods may no longer be expressed in

closed form and hence explicit expressions for posterior probabilities even where enumeration is feasible do not exist. We may add a step to sample the λ_j 's, leading to the following ODA-Cauchy algorithm. For the complete data model in (4) and prior distributions in (5a-5c), (6) and (19) generate σ^2 , \mathbf{Y}_a , $\boldsymbol{\gamma}$ and $\boldsymbol{\lambda}$ from the following sequence of distributions:

1. $(\sigma^2, \mathbf{Y}_a) \mid \boldsymbol{\gamma}, \boldsymbol{\lambda}, \mathbf{Y}_o$ from equations (9a - 9b)
2. $\gamma_j \mid \mathbf{Y}_c, \sigma^2, \boldsymbol{\lambda}$ for $j = 1, \dots, p$ from equation (9c)
3. $\beta_j \mid \mathbf{Y}_c, \sigma^2, \boldsymbol{\gamma}, \boldsymbol{\lambda} \stackrel{ind}{\sim} \text{N} \left(\widehat{\beta}_j \gamma_j \frac{\delta_j}{\lambda_j + \delta_j}, \sigma^2 \gamma_j (\lambda_j + \delta_j)^{-1} \right)$ where $\widehat{\beta}_j = \mathbf{x}_{c_j}^T \mathbf{Y}_c / \mathbf{x}_{c_j}^T \mathbf{x}_{c_j}$ for $j = 1, 2, \dots, p$
4. $\lambda_j \mid \mathbf{Y}_c, \sigma^2, \boldsymbol{\gamma}, \beta_j \stackrel{ind}{\sim} \text{G} \left(\frac{\alpha + \gamma_j}{2}, \frac{\alpha + \gamma_j \beta_j^2 / \sigma^2}{2} \right)$ for $j = 1, 2, \dots, p$

Note, the full conditional distribution of the scale parameter vector, $\boldsymbol{\lambda}$, is not available in closed form when $\boldsymbol{\beta}$ is integrated out and hence we draw β_j in step 3 to facilitate the drawing of $\boldsymbol{\lambda}$, but integrate out $\boldsymbol{\beta}$ from steps 1 and 2 in order to exploit the collapsed structure of the original ODA algorithm. We give a proof in the supplemental materials that the stationary distribution for this ODA sampler is the desired posterior distribution of $(\sigma^2, \mathbf{Y}_a, \boldsymbol{\gamma}, \boldsymbol{\lambda})$. The RB estimates for ODA-Cauchy are obtained by averaging over the K draws of \mathbf{Y}_a , σ^2 and $\boldsymbol{\lambda}$.

5.2 Comparison of Different Prior Distributions

In this section we use the Nott-Kohn simulations to compare the performance of model averaging with ODA-Cauchy (ODA with independent Cauchy priors) and ODA-Normal (ODA with normal priors) for parameter estimation under squared error loss. In addition, we compare the independent priors in ODA to BMA using the multivariate g -prior (Zellner 1984) with $g = n_o$ (a unit-information prior), the multivariate Zellner-Siow (Zellner and Siow 1980) Cauchy prior distribution, and BMA using BIC to approximate model probabilities. The

latter three estimates were computed under enumeration of the model space using the R package BAS on CRAN (Clyde 2010).

We generate 100 simulated datasets for the Nott-Kohn example and compare the different prior distributions based on their sum of squared errors (SSE) in estimating β , $SSE = \sum_{j=1}^p (\beta_j - \tilde{\beta}_j)^2$, where β_j is specified in Section 4 and $\tilde{\beta}_j$ is the BMA estimate of β_j . BAS is a sampling without replacement algorithm that is guaranteed to enumerate the model space for number of predictors less than 20-25 (Clyde et al. 2010). As the number of predictors is 15, we ran BAS for 2^{15} iterations for enumeration, providing exact posterior means under BMA for the g -prior. BAS uses a one-dimensional Laplace approximation to calculate the marginal likelihood for Zellner-Siow prior (Liang et al. 2008). ODA-Cauchy and ODA-Normal are run for 2^{15} iterations after discarding a burn-in of 5,000. CPU times for running ODA and BAS are not directly comparable currently as ODA is written currently in interpreted R code, while the sampling algorithm in the BAS package is written in C/FORTRAN and makes extensive use of the BLAS library. We plan to make an R package available and will recode in C/Fortran to speed up the calculations and expect that “an iteration” of ODA to be close to that of the BAS package as both algorithms require calculation of the OLS estimates, and use independent Bernoulli distributions to generate models. Currently for each of the 100 simulated datasets, ODA takes around 5 minutes, whereas the other prior distributions take around 0.2-0.3 seconds using the package BAS.

For each simulated dataset, we identified the algorithm that had the smallest SSE, and computed the relative efficiency of each algorithm to the one with the minimum SSE, $EFF(\text{algorithm}_k) = SSE(\text{algorithm}_k) / \min(SSE)$. Boxplots of the relative efficiencies are shown in Figure 1. ODA-normal and ODA-Cauchy have boxplots that are most concentrated around one, indicating that they are often producing the smallest SSE, and at other times performing almost as well as the prior with the smallest SSE. For this simulation study, there is little difference among Zellner’s g -prior and the Zellner-Siow Cauchy prior. Similarly the




Figure 1: Boxplots of relative efficiency ($SSE(\text{algorithm } k)/\min(SSE)$ for each simulated data set) for the Nott-Kohn simulation study for BIC, g -prior, Zellner-Siow Multivariate Cauchy prior, Independent Cauchy, Independent Normal

results for the independent normal and Cauchy priors based on ODA are very close, thus most of the gain is from the generalized ridge estimates based on the independent prior used in ODA. The mean SSE calculated over 100 replicates for ODA-Normal, ODA-Cauchy, Zellner-Siow, g -prior, and BIC are 2.08, 2.09, 2.34, 2.40, and 2.48 respectively.

6 Example: Ozone Data

We illustrate the performance of ODA using the ozone data (Friedman and Silverman 1989). There are 330 observations and 8 meteorological variables, where the response variable is ground level ozone in Los Angeles. Considering the same second order interactions and square terms, as in Liang et al. (2008), there are 44 predictors in all, leading to a model space of 2^{44} models.

For ODA-Normal we may calculate the exact marginal likelihoods of the sampled models, which provides a diagnostic to check convergence for posterior model probabilities. We run

ODA-Normal for one million iterations and discard the first 200,000 as burn-in. For model spaces as large as this one, most of the MC estimates of model probabilities are around $1/K$, where K is the number of MCMC iterations. Figure 2 highlights this problem and contrasts the true marginal likelihoods with their RB and MC estimates. Ideally the points should be tightly clustered around the 45° line. While RB offers an improvement over the MC estimates, RB also exhibits considerable Monte Carlo variation. This example illustrates the difficulty of estimating model probabilities in higher dimensions, which is still an open problem. In problems like this, we do not advocate using the RB estimates of model probabilities for selecting the highest probability model or a set of top models. The ODA algorithm is still useful for implementing BMA with heavy-tailed prior distributions in large model spaces as this one, and provides reliable estimates of marginal inclusion probabilities for variable selection or predictions under BMA.




Figure 2: Comparison of Rao-Blackwellized (RB) and Monte Carlo (MC) estimates of (renormalized) marginal likelihoods of top 1,000 models for the ozone data for ODA-Normal with 1 million iterations.

We explore the predictive performance of the ODA method with normal and Cauchy prior distributions. Following Liang et al. (2008), we leave out a randomly chosen half of the data as training data and predict the ozone concentration for the remaining half ($n^* = n_o/2$) using BMA estimates for ODA-Normal and ODA-Cauchy. For comparison we also include BMA predictions under the g -prior (with $g = n_o/2$) and the Zellner-Siow prior using the R package BAS, the lasso (Tibshirani 1996; Efron et al. 2004) using the R package lars (Hastie and Efron 2007) and the Bayesian Horseshoe (Carvalho et al. 2010) using the R package monomvn (Gramacy 2010). The penalty in the lasso may be viewed as the log of a double exponential prior, which is a scale mixture of normals, while the horseshoe prior is yet another scale mixture of normals. While both result in nonlinear shrinkage rules, neither include a mixture component that puts positive probability on zero. We compare methods on the basis of the predictive root mean square error (RMSE) defined as $\sqrt{\sum_{i=1}^{n^*} (Y_i - \hat{Y}_i)^2 / (n^*)}$ where Y_1, \dots, Y_{n^*} are the response variables in the randomly selected test sample and \hat{Y}_i is the BMA estimate of Y_i .

We run the ODA algorithms for 2^{15} iterations after discarding a 5,000 burn-in and run BAS for $2^{15} + 5000$ iterations. Empirically, predictions under model averaging appear to converge faster than model probabilities in these large spaces, thus we use a shorter MCMC run. The entire procedure is repeated 10 times to have 10 different sets of training and test samples. The resulting RMSE from different prior distributions (and algorithms) are quite similar, with the between sample variation dominating. As before the best methods are ODA, with root average MSE (RAMSE) of 4.098 for ODA-Cauchy, 4.102 for ODA-Normal, 4.131 for horseshoe, 4.153 for lasso and 4.168 for both the Zellner g -prior and Zellner-Siow priors. Overall, the ODA methods are competitive with some of the best available methods.

7 Probit Regression

We now show how to extend the ODA framework to binary regression models. Letting \mathbf{Z}_o denote the binary response variable of length n_o , a probit regression model may be expressed as

$$P(Z_{oi} = 1 \mid \mathbf{x}_{\gamma i}, \boldsymbol{\beta}_\gamma, \boldsymbol{\gamma}) = \Phi(\mathbf{x}_{\gamma i}^T \boldsymbol{\beta}_\gamma), \quad i = 1, 2, \dots, n_o \quad (20)$$

where Φ is the normal cumulative distribution function (cdf). Using the method of Albert and Chib (1993), one can augment the observed data $\{\mathbf{Z}_o, \mathbf{X}_o\}$ by latent data \mathbf{Y}_o such that:

$$Y_{oi} \mid \boldsymbol{\beta}_\gamma \stackrel{ind}{\sim} \mathbf{N}(\mathbf{x}_{\gamma i}^T \boldsymbol{\beta}_\gamma, 1) \quad (21)$$

where $Z_{oi} = \mathbf{1}_{(Y_{oi} > 0)}$. Using the prior specifications given in (5b) with $\sigma^2 = 1$ and $\lambda_j = 1$ for $j = 1, \dots, p$ and (6), we can construct the ODA-Probit sampler, described in the Supplemental Materials.

To illustrate the application of the ODA-Probit algorithm, we use the Pima Indians Diabetes dataset available in the MASS library in R. The response variable is an indicator of whether or not the women had diabetes according to World Health Organization criteria. There are seven explanatory variables: npreg, glu, bp, skin, bmi, ped and age. We provide a variable key in the Supplemental Materials.

We run ten ODA-Probit chains for 300,000 iterations after a burn-in of 5,000, with random starting values. On average, 39 (± 1) unique models were sampled, with an estimate of the remaining mass of 1.2×10^{-5} (range $2.2 \times 10^{-6} - 1.9 \times 10^{-5}$ over the 10 chains), using the RB-split estimate. Inclusion probabilities were stable over the 10 runs, with estimates of marginal inclusion probabilities $p(\boldsymbol{\gamma}_{\text{npreg}} = 1) = 0.947$, (sd = 0.002), $p(\boldsymbol{\gamma}_{\text{glu}} = 1) = 1.000$ (sd = 1.254×10^{-29}), $p(\boldsymbol{\gamma}_{\text{bp}} = 1) = 0.075$ (sd = 0.0002), $p(\boldsymbol{\gamma}_{\text{skin}} = 1) = 0.099$ (sd = 0.001), $p(\boldsymbol{\gamma}_{\text{bmi}} = 1) = 0.997$ (sd = 0.0004), $p(\boldsymbol{\gamma}_{\text{ped}} = 1) = 0.969$ (sd = 0.001), and $p(\boldsymbol{\gamma}_{\text{age}} = 1) = 0.389$

(sd = 0.003). For generalized linear models, analytic expressions for marginal likelihoods are not available in closed form; thus even for an enumerable model space of dimension 2^7 , there is no way of obtaining closed form expressions for true model probabilities or marginal inclusion probabilities. Here we can enumerate all 2^7 models and estimate the model probabilities using the ODA RB estimates for *any* model. We observe that npreg, glu, bmi and ped are identified as important covariates, which were also flagged as important in a paper by Holmes and Held (2006). Although their prior distribution for β is slightly different from ours, Holmes and Held (2006) reported the inclusion probability for age was 0.129, suggesting doubt about the importance of age, while Ripley (1996) found that the best AIC model included age, dropping only bp and skin. Looking more closely at the distribution of $\rho_{\text{age}}^{\text{RB}}$ from one of the chains, we found that it is highly bimodal, with support near zero and one. Images of the model space show that age is actually included in 2 out of the top 4 models (Figure 3), with the highest probability model under ODA-Probit corresponding to the second best model in terms of AIC and the best AIC model corresponding to the second highest posterior probability model. The Bayes factor for comparing the top two models is 1.11, suggesting that these two models receive similar support from the data. In this case, model averaging may be preferable to model selection. Splitting the data as in Ripley (1996), and refitting the ODA-Probit model with the 200 training observations to predict for the remaining 332, we find that the misclassification rate under ODA-Probit with model averaging is 18.7 %, which is better than the highest probability model under ODA-Probit (19.4%), the best AIC model (19.9%) or the best BIC model (20.2%).

8 Discussion

In the majority of problems outside Gaussian regression with conjugate prior distributions, the marginal likelihood of models and hence model probabilities are not analytically




Figure 3: Top 20 models for the Pima Indians Diabetes data. Each row corresponds to a model, with the highest probability model at the top, and each column corresponds to a variable, with black regions indicating variables that are excluded from the model. The y-axis is the log Bayes factor for comparing each model to the 20th best model out of all (2^7) models. Log Bayes factors for comparing any other two models may be found by subtraction.

tractable, even for relatively small model spaces. The orthogonal data augmentation algorithm provides Rao-Blackwellized estimates of model probabilities, inclusion probabilities and predictions which provide improvements over current algorithms with normal priors, but more importantly allow one to consider more robust prior specifications, such as independent Cauchy priors. For linear regression ODA-Cauchy provides results that are equivalent or better than the “gold-standard” Zellner-Siow Cauchy prior, lasso or horseshoe. The ODA algorithm for probit regression provides Rao-Blackwellized estimates of model probabilities and inclusion probabilities so that one does not have to rely on asymptotic approximations, such as BIC, for model probabilities.

There are several extensions of the current method that are possible. For illustration purposes, we have used a fixed uniform prior on the model space throughout; this assumption may be relaxed by placing a prior distribution, for example a Beta distribution, on π_j and updating π_j from its full conditional. The restriction to independent priors on the β_j may also be relaxed. For normal and mixture of normal prior distributions for β_γ that are obtained from the distribution of coefficients in the full model and conditioning on a subset of β being

equal to zero (conditionally compatible priors such as the g -prior or Zellner-Siow Cauchy prior), we may incorporate the prior precision Φ in the solution of the augmented design $\mathbf{X}_o^T \mathbf{X}_o + \mathbf{X}_a^T \mathbf{X}_a + \Phi = \mathbf{D}$. For mixtures of g -priors, $\Phi = g^{-1} \mathbf{X}_o^T \mathbf{X}_o$, so this would lead to updating $\mathbf{X}_a^T \mathbf{X}_a$ after sampling g . Similarly for the $p > n$ case, the incorporating the prior precision into the solution for \mathbf{X}_a will lead to a full rank solution.

Data-augmentation with parameter expansion has been shown to improve convergence in many cases (Liu and Wu 1999; Meng and van Dyk 1999; Hobert and Marchev 2008). We experimented with several parameter expansion schemes based on scale changes for the latent data (not reported here), but did not see any improvement. Liu and Wu (1999) suggest that if the set of transformation forms a locally compact group and the prior on the expansion parameter corresponds to Haar measure, that there is a well defined parameter expansion algorithm that is optimal in terms of convergence. For ODA a natural choice that preserves the invariance of the augmented data is to consider the group of orthogonal rotations \mathcal{O}_{p+1} as the expansion parameter. Generating random orthogonal matrices, however, will increase computational complexity of the algorithm; standard algorithms for generating orthogonal matrices are of order $O(p^3)$ although Genz (1998) describes methods using butterfly matrices that are of order $O(\log(p)p^2)$. It remains to be seen if improvements in convergence offset computational requirements, however, by allowing the augmented design matrix to be random, we may be able to further reduce the leverage of the augmented cases.

Acknowledgments

This research was supported by National Institutes of Health grants NIH/NIEHS 5T32ES007018, NIH/NIEHS P30 ES10126, 1-RC1HL099863-01, and NIH R01-HL090559-01 and National Science Foundation grants DMS-0406115 and DMS-0422400. The authors thank Steve MacEachern and Scott Schmidler for fruitful discussions.

Supplemental Materials

Algorithm, proof, variable key: Detailed algorithm for ODA-Probit, proof of proposition for ODA-Cauchy and ODA-Probit, and variable key for Pima Indians Diabetes dataset. (supp.pdf)

R Code: R code for the ODA algorithm. (code.zip)

References

- Albert, J. H. and Chib, S. (1993), “Bayesian analysis of binary and polychotomous response data,” *Journal of the American Statistical Association*, 88, 669–679.
- Berger, J. O. and Molina, G. (2005), “Posterior model probabilities via path-based pairwise priors,” *Statistica Neerlandica*, 59, 3–15.
- Berger, J. O., Pericchi, L. R., and Varshavsky, J. A. (1998), “Bayes factors and marginal distributions in invariant situations,” *Sankhya, Ser. A*, 60, 307–321.
- Carvalho, C. M., Polson, N. G., and Scott, J. G. (2010), “The horseshoe estimator for sparse signals,” *Biometrika*, to appear, asq017, doi:10.1093/biomet/asq017.
- Chipman, H. A., Kolaczyk, E. D., and McCulloch, R. E. (1997), “Adaptive Bayesian wavelet shrinkage,” *Journal of the American Statistical Association*, 92, 1413–1421.
- Clyde, M., DeSimone, H., and Parmigiani, G. (1996), “Prediction via orthogonalized model mixing,” *Journal of the American Statistical Association*, 91, 1197–1208.
- Clyde, M. and George, E. I. (1999), “Empirical Bayes estimation in wavelet nonparametric regression,” in *Bayesian Inference in Wavelet Based Models*, Springer-Verlag, pp. 309–322.
- Clyde, M. and George, E. I. (2000), “Flexible empirical Bayes estimation for wavelets,” *Journal of the Royal Statistical Society, Series B*, 62, 681–698.
- Clyde, M. and George, E. I. (2004), “Model uncertainty,” *Statistical Science*, 19, 81–94.
- Clyde, M., Ghosh, J., and Littman, M. (2010), “Bayesian adaptive sampling for variable selection and model averaging,” *Journal of Computational and Graphical Statistics*, to appear.
- Clyde, M. and Parmigiani, G. (1996), “Orthogonalizations and prior distributions for orthogonalized model mixing,” in *Modelling and Prediction Honoring Seymour Geisser*, Springer-Verlag.

- Clyde, M., Parmigiani, G., and Vidakovic, B. (1998), “Multiple shrinkage and subset selection in wavelets,” *Biometrika*, 85, 391–401.
- Clyde, M. A. (2010), *BAS: Bayesian Adaptive Sampling for Bayesian Model Averaging*, r package version 0.90.
- Clyde, M. A. and Ghosh, J. (2010), “A note on the bias in estimating posterior probabilities in variable selection,” Technical Report 2010-11, Duke University.
- Dempster, A., Laird, N. M., and Rubin, D. (1977), “Maximum likelihood from incomplete data via the EM algorithm,” *Journal of the Royal Statistical Society, Series B*, 39, 1–22.
- Denison, D. G. T., Mallick, B. K., and Smith, A. F. M. (1998), “Automatic Bayesian curve fitting,” *Journal of the Royal Statistical Society, Series B*, 60, 333–350.
- Draper, D. (1995), “Assessment and propagation of model uncertainty (with discussion),” *Journal of the Royal Statistical Society, Series B*, 57, 45–70.
- Eaton, M. L. (1983), *Multivariate Statistics: A Vector Space Approach*, John Wiley & Sons.
- Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004), “Least angle regression,” *The Annals of Statistics*, 32, 407–499.
- Fernández, C., Ley, E., and Steel, M. F. (2001), “Benchmark priors for Bayesian model averaging,” *Journal of Econometrics*, 100, 381–427.
- Fouskakis, D., Ntzoufras, I., and Draper, D. (2009), “Bayesian variable selection using cost-adjusted BIC, with application to cost-effective measurement of quality of health care,” *Ann. Appl. Statist.*, 3, 663–690.
- Friedman, J. H. and Silverman, B. W. (1989), “Flexible parsimonious smoothing and additive modeling,” *Technometrics*, 31, 3–39.
- Fulton, W. (2000), “Eigenvalues, invariant forms, highest weights, and Schubert calculus,” *Bulletin of the American Mathematical Society*, 37, 209–249.
- Gelfand, A. E. and Smith, A. F. M. (1990), “Sampling-based approaches to calculating marginal densities,” *Journal of the American Statistical Association*, 85, 398–409.
- Genz, A. (1998), “Methods for Generating Random Orthogonal Matrices,” in *Monte Carlo and Quasi-Monte Carlo Methods*, eds. H. Niederreiter and J. Spanier, Springer-Verlag, pp. 199–213.
- George, E. I. and McCulloch, R. E. (1997), “Approaches for Bayesian variable selection,” *Statistica Sinica*, 7, 339–374.
- Geyer, C. J. (1995), “Conditioning in Markov chain Monte Carlo,” *Journal of Computational and Graphical Statistics*, 4, 148–154.

- Gramacy, R. B. (2010), *monomvn: Estimation for multivariate normal and Student-t data with monotone missingness*, r package version 1.8-3.
- Hastie, T. and Efron, B. (2007), *lars: Least Angle Regression, Lasso and Forward Stagewise*, r package version 0.9-7.
- Heaton, M. and Scott, J. (2010), “Bayesian computation and the linear model,” in *Frontiers of Statistical Decision Making and Bayesian Analysis*, eds. M.-H. Chen, D. K. Dey, P. Mueller, D. Sun, and K. Ye.
- Hobert, J. P. and Marchev, D. (2008), “A theoretical comparison of the data-augmentation, marginal augmentation and PX-DA algorithms,” *Annals of Statistics*, 36, 532–554.
- Hoeting, J. A., Madigan, D., Raftery, A. E., and Volinsky, C. T. (1999), “Bayesian model averaging: a tutorial (with discussion),” *Statistical Science*, 14, 382–401, corrected version at <http://www.stat.washington.edu/www/research/online/hoeting1999.pdf>.
- Holmes, C. C. and Held, L. (2006), “Bayesian auxiliary variable models for binary and multinomial regression,” *Bayesian Analysis*, 1, 145–168.
- Ishwaran, H. and Rao, J. (2005), “Spike and slab variable selection: Frequentist and Bayesian strategies,” *Annals of Statistics*, 33, 730–773.
- Jeffreys, H. (1961), *Theory of Probability*, Oxford Univ. Press.
- Johnstone, I. and Silverman, B. (2004), “Needles and straw in haystacks: Empirical Bayes estimates of possibly sparse sequences,” *Annals of Statistics*, 32, 1594–1649.
- Johnstone, I. and Silverman, B. (2005), “Empirical Bayes selection of wavelet thresholds,” *Annals of Statistics*, 33, 1700–1752.
- Kass, R. E. and Raftery, A. E. (1995), “Bayes factors,” *Journal of the American Statistical Association*, 90, 773–795.
- Liang, F., Paulo, R., Molina, G., Clyde, M., and Berger, J. (2008), “Mixtures of g -priors for Bayesian variable selection,” *Journal of the American Statistical Association*, 103, 410–423.
- Liu, J. S. (1994), “The collapsed Gibbs sampler in Bayesian computations with application to a gene regulation problem,” *Journal of the American Statistical Association*, 89, 958–966.
- Liu, J. S., Wong, W. H., and Kong, A. (1994), “Covariance structure of the Gibbs sampler with applications to the comparisons of estimators and augmentation schemes,” *Biometrika*, 81, 27–40.
- Liu, J. S. and Wu, Y. N. (1999), “Parameter expansion for data augmentation,” *Journal of the American Statistical Association*, 94, 1264–1274.

- Meng, X. L. and van Dyk, D. A. (1999), “Seeking efficient data augmentation schemes via conditional and marginal augmentation,” *Biometrika*, 86, 301–320.
- Nott, D. J. and Green, P. J. (2004), “Bayesian variable selection and the Swendsen-Wang algorithm,” *Journal of Computational and Graphical Statistics*, 13, 141–157.
- Nott, D. J. and Kohn, R. (2005), “Adaptive sampling for Bayesian variable selection,” *Biometrika*, 92, 747–763.
- Polson, N. G. and Scott, J. G. (2010), “Shrink globally, act locally: sparse Bayesian regularization and prediction,” in *Bayesian Statistics 9*, eds. J. M. Bernardo, M. J. Bayarri, J. O. Berger, A. P. Dawid, D. Heckerman, A. F. M. Smith, and M. West, Oxford, UK: Oxford Univ. Press.
- Raftery, A. E., Madigan, D., and Hoeting, J. A. (1997), “Bayesian model averaging for linear regression models,” *Journal of the American Statistical Association*, 92, 179–191.
- Ripley, B. D. (1996), *Pattern Recognition and Neural Networks*, Cambridge University Press.
- Scott, J. G. and Carvalho, C. M. (2008), “Feature-inclusion stochastic search for Gaussian graphical models,” *Journal of Computational and Graphical Statistics*, 17, 790–808.
- Tanner, M. A. and Wong, W. H. (1987), “The calculation of posterior distributions by data augmentation,” *Journal of the American Statistical Association*, 82, 528–540.
- Tibshirani, R. (1996), “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society, Series B*, 58, 267–288.
- Weyl, H. (1912), “Das asymptotische Verteilungsgesetz der Eigenwerte lineare partieller Differentialgleichungen,” *Mathematische Annalen*, pp. 441–479.
- Zellner, A. (1984), “Posterior odds ratios for regression hypotheses: General considerations and some specific results,” in *Basic Issues in Econometrics*, ed. A. Zellner, Univ. of Chicago Press, pp. 275–305.
- Zellner, A. (1986), “On assessing prior distributions and Bayesian regression analysis with g -prior distributions,” in *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti*, North-Holland/Elsevier, pp. 233–243.
- Zellner, A. and Siow, A. (1980), “Posterior odds ratios for selected regression hypotheses,” in *Bayesian Statistics: Proceedings of the First International Meeting held in Valencia (Spain)*, pp. 585–603.