Published in final edited form as:

Stat Probab Lett. 2011 August 1; 81(8): 1056–1062. doi:10.1016/j.spl.2011.02.029.

Sparse Variational Analysis of Linear Mixed Models for Large Data Sets

Artin Armagan and

Department of Statistical Science, Duke University, Durham, NC 27708

David Dunson

Department of Statistical Science, Duke University, Durham, NC 27708

Artin Armagan: artin@stat.duke.edu; David Dunson: dunson@stat.duke.edu

Abstract

It is increasingly common to be faced with longitudinal or multi-level data sets that have large numbers of predictors and/or a large sample size. Current methods of fitting and inference for mixed effects models tend to perform poorly in such settings. When there are many variables, it is appealing to allow uncertainty in subset selection and to obtain a sparse characterization of the data. Bayesian methods are available to address these goals using Markov chain Monte Carlo (MCMC), but MCMC is very computationally expensive and can be infeasible in large p and/or large n problems. As a fast approximate Bayes solution, we recommend a novel approximation to the posterior relying on variational methods. Variational methods are used to approximate the posterior of the parameters in a decomposition of the variance components, with priors chosen to obtain a sparse solution that allows selection of random effects. The method is evaluated through a simulation study, and applied to an epidemiological application.

Keywords

Mixed-effects model; Variational approximations; Shrinkage estimation

1. Introduction

It is often of interest to fit a hierarchical model in settings involving large numbers of predictors (p) and/or large sample size (n). For example, in a large prospective epidemiology study, one may obtain longitudinal data for tens of thousands of subjects, while also collecting ~ 100 predictors. Even in more modest studies, involving thousands of subjects, the number of predictors collected is often large. Unfortunately, current methods for inference in mixed effects models are not designed to accommodate large p and/or large n . This article proposes a method for obtaining sparse approximate Bayes inferences in such problems using variational methods (Jordan et al., 1999; Jaakkola and Jordan, 2000).

For concreteness we focus on the linear mixed effects (LME) model (Laird and Ware, 1982), though the proposed methods can be applied directly in many other hierarchical

© 2011 Elsevier B.V. All rights reserved.

Correspondence to: Artin Armagan, artin@stat.duke.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

models. When considering LMEs in settings involving moderate to large p , it is appealing to consider methods that encourage sparse estimation of the random effects covariance matrix. There are a variety of methods available in the literature, including approaches based on Bayesian methods implemented with MCMC (Chen and Dunson, 2003; Kinney and Dunson, 2007; Frühwirth-Schnatter and Tüchler, 2008) and methods based on fast shrinkage estimation (Foster et al., 2007).

Frequentist procedures encounter convergence problems (Pennell and Dunson, 2007) and MCMC based methods tend to be computationally intensive and not to scale well as p and/or n increases. The methods relying on stochastic search variable selection (SSVS) algorithms (George and McCulloch, 1997) face difficulties when p increases beyond ~ 30 in linear regression applications, with the computational burden substantially greater in hierarchical models involving random effects selection. Approaches have been proposed to make MCMC implementations of hierarchical models feasible in large data sets (Huang and Gelman, 2008; Pennell and Dunson, 2007). However, these approaches do not solve the large p problem or allow sparse estimation or selection of random effects covariances. In addition, the algorithms are still time consuming to implement.

Model selection through shrinkage estimation has gained much popularity since the Lasso of Tibshirani (1996). Similar shrinkage effects were later obtained through hierarchical modeling of the regression coefficients in the Bayesian paradigm. A few examples of these are Tipping (2001); Bishop and Tipping (2000); Figueiredo (2003); Park and Casella (2008). Most approaches have relied on *maximum a posteriori* (MAP) estimation. MAP estimation produces a sparse point estimate with no measure of uncertainty, motivating MCMC and variational methods.

It would be appealing to have a fast approach that could be implemented much more rapidly in cases involving moderate to large data sets and numbers of variables, while producing sparse estimates and allowing approximate Bayesian inferences on predictor effects. In particular, it would be very appealing to have an approximation to the marginal posteriors instead of simply obtaining a point estimate. Basing inferences on point estimates does not account for uncertainty in the estimation process, and hence is not useful in applications, such as epidemiology.

One possibility is to rely on a variational approximation to the posterior distribution (Jordan et al., 1999; Jaakkola and Jordan, 2000; Bishop and Tipping, 2000). Within this framework, we develop a method for sparse covariance estimation relying on a decomposition and the use of heavy-tailed priors in a related manner to Tipping (2001), though they did not consider estimation of covariance matrices.

2. Variational inference

Except in very simple conjugate models, the marginal likelihood of the data is not available analytically. As an alternative to MCMC and Laplace approximations (Tierney and Kadane, 1986), a lower-bound on marginal likelihoods may be obtained via variational methods (Jordan et al., 1999) yielding approximate posterior distributions on the model parameters. Let $\boldsymbol{\theta}$ be the vector of all unobserved quantities in the model and $\mathbf{y}$ be the observed data. Given a density $q(\boldsymbol{\theta})$, the marginal log-likelihood can be decomposed as

$$\log p(\mathbf{y}) = \underbrace{\int q(\boldsymbol{\theta}) \log \frac{p(\mathbf{y}, \boldsymbol{\theta})}{q(\boldsymbol{\theta})} d\boldsymbol{\theta}}_{\text{L}} + \text{KL}(q \| p), \quad (1)$$

where $p(\mathbf{y}, \boldsymbol{\theta})$ is an unnormalized posterior density of $\boldsymbol{\theta}$ and $\text{KL}(q\|p)$ denotes the Kullback-Leibler divergence between the true posterior $p(\boldsymbol{\theta}|\mathbf{y})$ and $q(\boldsymbol{\theta})$. Since this quantity is strictly non-negative and is equal to 0 only when $p(\boldsymbol{\theta}|\mathbf{y}) = q(\boldsymbol{\theta})$, the first term in (1) constitutes a lower-bound on $\log p(\mathbf{y})$. Maximizing the first term in the right hand side of (1) is equivalent to minimizing the second term in the right hand side, suggesting that $q(\boldsymbol{\theta})$ is an approximation to the posterior density $p(\boldsymbol{\theta}|\mathbf{y})$.

Following Bishop and Tipping (2000) we consider a factorized form

$$q(\boldsymbol{\theta}) = \prod_i q_i(\boldsymbol{\theta}_i), \quad (2)$$

where $\boldsymbol{\theta}_i$ is a sub-vector of $\boldsymbol{\theta}$ and there are no restrictions on the form of $q_i(\boldsymbol{\theta}_i)$. Then the lower-bound can be maximized with respect to $q_i(\boldsymbol{\theta}_i)$ yielding an iterative procedure

$$q_i(\boldsymbol{\theta}_i) = \frac{\exp\langle \log p(\mathbf{y}, \boldsymbol{\theta}) \rangle_{\boldsymbol{\theta}_{j \neq i}}}{\int \exp\langle \log p(\mathbf{y}, \boldsymbol{\theta}) \rangle_{\boldsymbol{\theta}_{j \neq i}} d\boldsymbol{\theta}_i}, \quad (3)$$

where $\langle \cdot \rangle_{\boldsymbol{\theta}_{j \neq i}}$ denotes the expectation with respect to the distributions $q_j(\boldsymbol{\theta}_j)$ for $j \neq i$. Due to conjugacy, these expectations will be easily evaluated yielding standard distributions for $q_i(\boldsymbol{\theta}_i)$ with parameters expressed in terms of the moments of $\boldsymbol{\theta}_j$. Thus the procedure will consist of initializing the expectations required and re-iterating through them updating the expectations with respect to the densities provided by (3).

Due to the factorized form in (2), information pertaining to the dependence structure among $\boldsymbol{\theta}_j$ s is lost with these approximations. This comes as a trade-off to the computational advantages. It is this factorized form and the use of conjugate priors that help us identify $q(\boldsymbol{\theta}_i)$ as well-known densities, making the required moments readily available. However, considering that this factorization will be realized in a block form, we will be able to preserve the dependence structure within the blocks, i.e. within each $\boldsymbol{\theta}_j$. Potentially, the approximation may become less accurate with increasing dimension if there are important dependencies among the $\boldsymbol{\theta}_j$ s.

In the following text G, IG, IW, N, U, W will respectively denote gamma, inverse gamma, inverse Wishart, normal, uniform and Wishart distributions.

3. The Model

3.1. The Standard Model

Suppose there are n subjects under study, with n_i observations for the i th subject. For subject i at observation j , let y_{ij} denote the response, let $\mathbf{x}_{ij}$ and $\mathbf{z}_{ij}$ denote $p \times 1$ and $q \times 1$ vectors for predictors. Then, a linear mixed effects model can be written as

$$\mathbf{y}_i = \mathbf{X}_i \boldsymbol{\alpha} + \mathbf{Z}_i \boldsymbol{\beta}_i + \boldsymbol{\varepsilon}_i, \quad (4)$$

where $\mathbf{y}_i = (y_{i1}, \dots, y_{in_i})'$, $\mathbf{X}_i = (\mathbf{x}_{i1}, \dots, \mathbf{x}_{in_i})'$, $\mathbf{Z}_i = (\mathbf{z}_{i1}, \dots, \mathbf{z}_{in_i})'$, $\boldsymbol{\alpha}$ is a $p \times 1$ vector of unknown fixed effects, $\boldsymbol{\beta}_i$ is a $q \times 1$ vector unknown subject-specific random effects with $\boldsymbol{\beta}_i \sim N(\mathbf{0}, \mathbf{D})$, and the elements of the residual vector, $\boldsymbol{\varepsilon}_i$, are $N(\mathbf{0}, \sigma^2 \mathbf{I})$.

Given the formulation in (4), the joint density of the observations given the model parameters can be written as

$$p(\mathbf{y}|\alpha, \beta, \sigma^2) = (2\pi\sigma^2)^{-\sum_{i=1}^n n_i/2} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{i=1}^n \sum_{j=1}^{n_i} (y_{ij} - \mathbf{x}'_{ij}\alpha - \mathbf{z}'_{ij}\beta_i)^2 \right\}. \quad (5)$$

For conditionally-conjugate priors, inference is straightforward both in exact (via MCMC) and approximate (via variational methods) cases. To save space, the variational approximations to the posteriors of the model parameters will be given only for the shrinkage model explained in the following section.

3.2. The Shrinkage Model

Let $\mathbf{D} = \mathbf{A}\mathbf{B}\mathbf{A}$ where $\mathbf{B}$ is a symmetric, positive-definite matrix and $\mathbf{A} = \text{diag}(\lambda_1, \dots, \lambda_q)$ with $\lambda_k \in \mathbb{R}$ with no further restrictions. This decomposition is not unique yet is sufficient to guarantee the positive-semidefiniteness of $\mathbf{D}$.

Let us re-write (4) as

$$\mathbf{y}_i = \mathbf{X}_i\alpha + \mathbf{Z}_i\mathbf{A}\mathbf{b}_i + \epsilon_i, \quad (6)$$

where $\mathbf{b}_i \sim N(\mathbf{0}, \mathbf{B})$ and λ_k acts as a scaling factor on the k th row and column of the random effects covariance $\mathbf{D}$. Although the parameterization is redundant, it has been noted that the redundant parameterizations are often useful for computational reasons and for inducing new classes of priors with appealing properties (Gelman, 2006). The incorporation of λ_k allows greater control on adaptive predictor-dependent shrinkage, with values $\lambda_k \approx 0$ (along with small corresponding diagonals in $\mathbf{B}$) leading to the k th predictor being effectively excluded from the random effects component of the model through setting the values in the k th row and column of $\mathbf{D}$ close to zero. This maintains the positive-semidefinite constraint. One issue with redundant parameterization is the lack of identifiability in a frequentist sense, i.e. it will lead to a likelihood which comprises multiple ridges along possible combinations of $\mathbf{A}$ and $\mathbf{B}$. This does not create insurmountable difficulties for Bayesian procedures as a non-flat prior should take care of this problem. When MCMC is used, the sampling of λ_k and $\mathbf{B}$ would occur along these ridges where the prior assigns a positive density. Multiple modes will exist due to the fact that each λ_k may take either sign.

The variational procedure used will converge to one of multiple exchangeable modes which live on the aforementioned ridges in the posterior. The tracking of the lower-bound plays an important role to stop the iterative procedure. As the lower-bound stops its monotonic increase (according to some preset criterion), we stop the procedure and assume that any further change in λ_k and $\mathbf{B}$ will not change inferences as we are moving along one of these ridges.

The joint density of the observations now can be written as

$$p(\mathbf{y}|\alpha, \mathbf{A}, \mathbf{b}, \sigma^2) = (2\pi\sigma^2)^{-\sum_{i=1}^n \frac{n_i}{2}} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{i=1}^n \sum_{j=1}^{n_i} (y_{ij} - \mathbf{x}'_{ij}\alpha - \mathbf{z}'_{ij}\mathbf{A}\mathbf{b}_i)^2 \right\} \quad (7)$$

where $\mathbf{b} = (\mathbf{b}'_1, \dots, \mathbf{b}'_n)'$.

3.2.1. Priors and Posteriors—After this decomposition the joint (conditional) density of the observations remains almost identical, replacing β by $\mathbf{A}\mathbf{b}_i$ in (7). Notice that λ_k and b_{ik}

are interchangeable which is going to allow us to model λ_k as *redundant* random-effects coefficients.

We will use independent t priors for α_k and λ_k due to its shrinkage reinforcing quality. This will be accomplished through the scale mixtures of normals (West, 1987) due to the conjugacy properties, i.e. $\alpha_k \sim N(0, a_k)$ and $\lambda_k \sim N(0, v_k)$ where $\alpha_k^{-1}, v_k^{-1} \sim G(\eta_0, \zeta_0)$. Under this setup, we would hope that those α_k and λ_k corresponding to insignificant fixed and random effects would shrink towards the neighborhood of zero. This will allow us to obtain a much smaller set of fixed effects for prediction purposes as well as a much more compact covariance structure on the random-effects coefficients. The usefulness of this approach will be especially emphasized in high dimensional problems. We also set $\sigma^2 \sim \text{IG}(c_0, d_0)$ and $\mathbf{B} \sim \text{IW}(n_0, \Psi_0)$.

The approximate marginal posterior distributions of the parameters, using (3), are obtained

$$q(\alpha) \stackrel{d}{=} \mathbf{N}(\widehat{\alpha},$$

$$\widehat{\mathbf{A}}), q(\sigma^2) \stackrel{d}{=} \text{IG}(\widehat{c},$$

$$\widehat{d}), q(\mathbf{b}_i) \stackrel{d}{=} \mathbf{N}(\widehat{\mathbf{b}}_i,$$

$$\widehat{\mathbf{B}}_i), q(\lambda) \stackrel{d}{=} \mathbf{N}(\widehat{\lambda},$$

$$\widehat{\mathbf{V}}), q(\mathbf{B}) \stackrel{d}{=} \text{IW}(\widehat{n},$$

$$\widehat{\Psi}), q(\alpha_k^{-1}) \stackrel{d}{=} G(\widehat{\eta},$$

$$\text{as } \widehat{\zeta}_k \text{ and } q(v_k^{-1}) \stackrel{d}{=} G(\eta^*, \zeta_k^*) \text{ where}$$

$$\widehat{\alpha} = \langle \sigma^{-2} \rangle \widehat{\mathbf{A}} \sum_{i=1}^n \mathbf{X}_i' (\mathbf{y}_i - \mathbf{Z}_i \langle \Lambda \rangle \langle \mathbf{b}_i \rangle),$$

$$\widehat{\mathbf{A}} = \langle \sigma^{-2} \rangle^{-1} \left(\sum_{i=1}^n \mathbf{X}_i' \mathbf{X}_i + \langle \sigma^{-2} \rangle^{-1} \langle \mathbf{A}^{-1} \rangle \right)^{-1},$$

$$\widehat{c} = \sum_{i=1}^n n_i / 2 + c_0,$$

$$\widehat{d} = \sum_{i=1}^n \left(\mathbf{y}_i' \mathbf{y}_i - 2 \langle \alpha \rangle' \mathbf{X}_i' \mathbf{y}_i - 2 \langle \mathbf{b}_i \rangle' \langle \Lambda \rangle \mathbf{Z}_i' \mathbf{y}_i + \sum_{k=1}^p \mathbf{x}_{ik}' \langle \alpha \alpha' \rangle \mathbf{x}_{ik} + \sum_{k=1}^q \mathbf{Z}_{ik}' \langle \lambda \lambda' \rangle \bullet \langle \mathbf{b}_i \mathbf{b}_i' \rangle \mathbf{Z}_{ik} + 2 \langle \alpha \rangle' \mathbf{X}_i' \mathbf{Z}_i \langle \Lambda \rangle \langle \mathbf{b}_i \rangle \right) / 2 + d_0,$$

$$\widehat{\mathbf{b}}_i = \langle \sigma^{-2} \rangle \widehat{\mathbf{B}}_i \langle \Lambda \rangle \mathbf{Z}_i' (\mathbf{y}_i - \mathbf{X}_i \langle \alpha \rangle),$$

$$\widehat{\mathbf{B}}_i = \langle \sigma^{-2} \rangle^{-1} \left(\langle \lambda \lambda' \rangle \bullet \mathbf{Z}_i' \mathbf{Z}_i + \langle \sigma^{-2} \rangle^{-1} \langle \mathbf{B}^{-1} \rangle \right)^{-1},$$

$$\widehat{\lambda}_i = \langle \sigma^{-2} \rangle \widehat{\mathbf{V}} \sum_{i=1}^n \text{diag}(\langle \mathbf{b}_i \rangle) \mathbf{Z}_i' (\mathbf{y}_i - \mathbf{X}_i \langle \alpha \rangle),$$

$$\widehat{\mathbf{V}} = \langle \sigma^{-2} \rangle^{-1} \left(\sum_{i=1}^n \langle \mathbf{b}_i \mathbf{b}_i' \rangle \bullet \mathbf{Z}_i' \mathbf{Z}_i + \langle \sigma^{-2} \rangle^{-1} \langle \mathbf{V}^{-1} \rangle \right)^{-1},$$

$$\widehat{n} = n + n_0, \widehat{\Psi} = \sum_{i=1}^n \langle \mathbf{b}_i \mathbf{b}_i' \rangle + \Psi_0,$$

$$\widehat{\eta} = 1/2 + \eta_0, \widehat{\zeta}_k = \langle \alpha_k^2 \rangle / 2 + \zeta_0,$$

$$\eta^* = 1/2 + \eta_0, \zeta_k^* = \langle \zeta_k^2 \rangle / 2 + \zeta_0,$$

$\lambda = \text{diag}(\Lambda)$, $\mathbf{A} = \text{diag}(a_k : k = 1, \dots, p)$, $\mathbf{V} = \text{diag}(v_k : k = 1, \dots, q)$, $(\bullet)$ denotes the Hadamard product and $\text{diag}(\cdot)$, depending on its argument, either builds a vector from the diagonal elements of a matrix or builds a diagonal matrix using the components of a vector as the diagonal elements of that matrix.

$$\begin{aligned}\langle \alpha \rangle &= \widehat{\alpha}, \langle \alpha \alpha' \rangle = \widehat{\mathbf{A}} + \widehat{\alpha} \widehat{\alpha}', \\ \langle \mathbf{b}_i \rangle &= \widehat{\mathbf{b}}_i, \langle \mathbf{b}_i \mathbf{b}_i' \rangle = \widehat{\mathbf{B}}_i + \widehat{\mathbf{b}}_i \widehat{\mathbf{b}}_i', \\ \langle \sigma^{-2} \rangle &= \widehat{c} / \widehat{d}, \langle \boldsymbol{\Lambda} \rangle = \text{diag}(\widehat{\boldsymbol{\lambda}}), \\ \langle \lambda \lambda' \rangle &= \widehat{\mathbf{V}} + \widehat{\boldsymbol{\lambda}} \widehat{\boldsymbol{\lambda}}',\end{aligned}$$

The required moments are $\langle a_k^{-1} \rangle = \widehat{\eta} / \widehat{\zeta}_k$, $\langle v_k^{-1} \rangle = \eta^* / \zeta_k^*$, $\langle \mathbf{B}^{-1} \rangle = \widehat{n} \widehat{\boldsymbol{\Phi}}^{-1}$. The iterative procedure that (3) implies is then to initiate these moments and cycle through them until convergence is reached.

After simplifications, the expression for the lower-bound, L, is given by

$$\begin{aligned}L = \frac{1}{2} \left\{ - \sum_{i=1}^n n_i \log(2\pi) + q(n+1) + p \log |\mathbb{V}(\alpha)| + \log |\mathbb{V}(\lambda)| + \sum_{i=1}^n \log |\mathbb{V}(\mathbf{b}_i)| + q(\widehat{n} - n_0) \log 2 + \log \frac{|\boldsymbol{\Psi}_0|^{n_0}}{|\widehat{\boldsymbol{\Psi}}|^{\widehat{n}}} \right\} \\ + \log \frac{\Gamma_q(\widehat{n}/2)}{\Gamma_q(n_0/2)} \\ + \log \frac{d_0^{c_0}}{\widehat{d}^{\widehat{c}}} \\ + \log \frac{\Gamma(\widehat{c})}{\Gamma(c_0)} \\ + \sum_{j=1}^p \log \frac{\zeta_0^{\eta_0}}{\zeta_j^{\eta_j}} \\ + \sum_{j=1}^q \log \frac{\zeta_0^{\eta_0}}{\zeta_j^{\eta_j}} \\ + p \log \frac{\Gamma(\widehat{\eta})}{\Gamma(\eta_0)} \\ + q \log \frac{\Gamma(\eta^*)}{\Gamma(\eta_0)}.\end{aligned}\quad (8)$$

4. Simulations

Focusing on the model of Section 3.2, we assess the performance in estimating the fixed effects and random effects covariance. We specify two alternatives for the number of subjects, $n = \{400, 2000\}$, three alternatives for the number of potential covariates, $p = \{4, 20, 60\}$, $q = p$, and three alternatives for the underlying sparsity corresponding respectively to the number of potential covariates, $p' = \{.75p, .50p, .25p\}$, $q' = p'$, where p' and q' denote the number of active covariates in the underlying model. We generate $n_i = 8$, $i = 1, \dots, n$ observations per subject and set $\alpha_{1:p'} = \mathbf{1}$ and $\alpha_{(p'+1):p} = \mathbf{0}$. In addition, we let $x_{ij1} = 1$, $\mathbf{x}_{ij(2:p)} \sim N_{p-1}(\mathbf{0}, \mathbf{C})$ and $\mathbf{C} \sim W(p-1, \mathbf{I}_{p-1})$; $\mathbf{z}_{ij} = \mathbf{x}_{ij}$; $\sigma^2 \sim U(1, 3)$; $\mathbf{D}_{1:q' \times 1:q'} \sim W(q', \mathbf{I}_{q'})$ and the rest of the entries along the dimensions $(q' + 1) : q$ are 0.

We run the variational procedure both for the standard and shrinkage models on 100 independently generated data sets. For the standard model, the priors are specified as $\boldsymbol{\alpha} \sim N(\mathbf{0}, \mathbf{A}_0)$, $\sigma^{-2} \sim G(c_0, d_0)$ and $\mathbf{D} \sim IW(n_0, \boldsymbol{\Psi}_0)$ where $\boldsymbol{\alpha}_0 = \mathbf{0}$, $\mathbf{A}_0 = 1000\mathbf{I}$ (Chen and Dunson, 2003), $c_0 = 0.1$, $d_0 = 0.001$ (Smith and Kohn, 2002), $n_0 = q$ and $\boldsymbol{\Phi}_0 = \mathbf{I}$ to reflect our vague prior information on $\boldsymbol{\alpha}$, σ^{-2} and $\mathbf{D}$. All these priors are proper yet specify very vague information about the parameters relative to the likelihoods that are observed in this simulation. For the shrinkage model, we choose $c_0 = \eta_0 = 0.1$, $d_0 = \zeta_0 = 0.001$, $n_0 = q$ and $\boldsymbol{\Psi}_0 = \mathbf{I}$ to express our vague information on σ^{-2} , a_k^{-1} , v_k^{-1} and $\mathbf{B}$. It is important that we refrain from using improper priors on the higher level parameters, i.e. a_k^{-1} , v_k^{-1} , for

meaningful marginal likelihoods as the limiting cases of these conjugate priors will lead to the impropriety of the posterior distribution and consequently the decomposition in (1) will lose its meaning.

The boxplots in Figure 1 (a) and (b) give the quadratic losses in the estimation of α and $\mathbf{D}$ respectively arising from the standard and the shrinkage models. As the dimension of the problem increases and the underlying model becomes sparser, the advantage of the shrinkage model is highly pronounced.

5. Real Data Application

We apply the proposed method to US Collaborative Perinatal Project (CPP) data on maternal pregnancy smoking in relation to child growth (Chen et al., 2005). Chen et al. (2005) examine the relationship between maternal smoking habits during pregnancy and childhood obesity within $n = 34866$ children in the CPP using generalized estimating equations (GEE) (Liang and Zeger, 1986). The size of the data, in particular the number of subjects, hinders a random effects analysis (Pennell and Dunson, 2007).

Having removed the missing observations we were left with 28211 subjects and 115811 observations. We set aside 211 observations across the subjects as a hold-out sample to test the performance of our procedure which leaves us with 115600 observations to train our model with. The main predictors considered in the model are, child's age, recruitment center, race, whether ever breast-fed, number of prior live births by mother, socioeconomic index, mothers BMI, number of cigarettes smoked per day, mother's age, mother's education, marital status of mother at recruitment and trimester at recruitment. Dummy variables were created for the categorical variables. We also include a quadratic term for age and interaction terms between all the predictors (except for center) and age. Our design matrix, $\mathbf{X} = \mathbf{Z}$, (with a column of 1s for the intercept term) has 72 columns. Each column of $\mathbf{X} = \mathbf{Z}$ is scaled to have unit length. The response was continuous, the weight of the children in kilograms. A detailed description of the data is available in Chen et al. (2005).

We apply our shrinkage model to the data. Figure 2 give the 99% credible sets for the fixed effect coefficients and for the diagonals of the random effects covariance matrix respectively. We can see for both fixed effects coefficients and the diagonals of the random effects covariance, many credible sets are concentrated around 0. Figure 3 (a) also gives the point estimates and 99% credible sets for the hold-out sample. Here the R^2 on the test set was found to be 94.8%.

The computational advantage of the procedure is undeniable. For the shrinkage model, the algorithm was implemented in MATLAB on a computer with a 2.8 GHz processor and 12 GB RAM. Figure 3(d) tracks the lower-bound and the relative error between two subsequent lower-bound values, $\psi = |\mathbf{L}^{(t)} - \mathbf{L}^{(t-1)}|/|\mathbf{L}^{(t)}|$, for convergence where $\mathbf{L}^{(t)}$ denotes the lower-bound evaluated at iteration t . The preset value of $\psi = 10^{-6}$ is reached after 2485 iterations which takes 2.4 days. It should be noted that the computational intensity for one iteration is almost identical to a Gibbs sampling scenario, which suggests, if Gibbs sampling procedure were to be used, only 2485 samples would have been drawn. Considering the burn-in period required for convergence and the thinning of the chain to obtain less correlated draws, this number is far from sufficient. Thus, with data sets this large or larger, MCMC is not a computationally feasible option.

6. Conclusion

Here we provided a fast approximate solution to fully Bayes inference to be used in the analysis of *large* longitudinal data sets. The proposed parameterization also allows for

identifying the predictors that contribute as fixed and/or random effects. Although this parameterization leads to an unidentifiable likelihood, and would also cause the so-called label-switching problem with the application of Gibbs sampling, the variational approach allows us to converge to one of many solutions which lead to identical inferences. The utility of the new parameterization is justified through a simulation study. The application to a large epidemiological data set also demonstrates computational advantages obtained through the proposed method over conventional sampling techniques.

The computational advantages are contingent upon the availability of the required moments in analytical form which are easily obtained with the conjugate priors we considered. When the likelihood is no longer normal, these moments may not be available in analytical form and we may have to resort in some numerical solution. In such cases computational advantages may not be as pronounced. As a special case, the same computational advantages should be observed in linear mixed models with a binary outcome through latent variable modeling of Albert and Chib (1993). A similar approach for shrinkage estimation in probit models was considered in Armagan (2009).

References

- Albert JH, Chib S. Bayesian Analysis of Binary and Polychotomous Response Data. *Journal of the American Statistical Association*. 1993;88.
- Armagan A. Variational Bridge Regression. *JMLR: W&CP*. 2009; 5:17–24.
- Bishop, CM.; Tipping, ME. Variational Relevance Vector Machines; UAI '00: Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence; San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.; 2000. p. 46–53.
- Chen A, Pennell M, Klebanoff MA, Rogan WJ, Longnecker MP. Maternal smoking during pregnancy in relation to child overweight: follow-up to age 8 years. *International Journal of Epidemiology*. 2005
- Chen Z, Dunson DB. Random Effects Selection in Linear Mixed Models. *Biometrics*. 2003; 59(4): 762–769. [PubMed: 14969453]
- Figueiredo MAT. Adaptive sparseness for supervised learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2003; 25:1150–1159.
- Foster, Scott D, Verbyla, Arunas P, Pitchford, Wayne S. *Journal of Agricultural, Biological and Environmental Statistics*. 2007; 12(2):300–314.
- Frühwirth-Schnatter S, Tüchler R. Bayesian parsimonious covariance estimation for hierarchical linear mixed models. *Statistics and Computing*. 2008; 18(1):1–13.
- George E, McCulloch R. Approaches for Bayesian variable selection. *Statistica Sinica*. 1997; 7:339–373.
- Huang Z, Gelman A. Sampling for Bayesian computation with large datasets. Technical Report, Department of Statistics, Columbia University. 2008; 7:339–373.
- Jaakkola TS, Jordan MI. Bayesian parameter estimation via variational methods. *Statistics and Computing*. 2000; 10(1):25–37.
- Jordan MI, Ghahramani Z, Jaakkola TS, Saul LK. An introduction to variational methods for graphical models. *Machine Learning*. 1999; 37:183–233.
- Kinney SK, Dunson DB. Fixed and Random Effects Selection in Linear and Logistic Models. *Biometrics*. 2007; 63(3):690–698. [PubMed: 17403104]
- Laird NM, Ware JH. Random-Effects Models for Longitudinal Data. *Biometrics*. 1982; 38(4):963–974. [PubMed: 7168798]
- Liang K-Y, Zeger SL. Longitudinal data analysis using generalized linear models. *Biometrika*. 1986; 73(1):13–22.
- Park T, Casella G. The Bayesian Lasso. *Journal of the American Statistical Association*. 2008; 103(6): 681–686.

- Pennell ML, Dunson DB. Fitting semiparametric random effects models to large data sets. *Biostat.* 2007; 8(4):821–834.
- Smith M, Kohn R. Parsimonious Covariance Matrix Estimation for Longitudinal Data. *Journal of the American Statistical Association.* 2002; 97:1141–1153.
- Tibshirani R. Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological).* 1996; 58(1):267–288.
- Tierney L, Kadane JB. Accurate Approximations for Posterior Moments and Marginal Densities. *Journal of the American Statistical Association.* 1986; 81(393):82–86.
- Tipping ME. Sparse Bayesian Learning and the Relevance Vector Machine. *Journal of Machine Learning Research.* 2001; 1
- West M. On Scale Mixtures of Normal Distributions. *Biometrika.* 1987; 74(3):646–648.

Figure 1.
Quadratic loss for (a) α and (b) D . The vertical axis is given in log-scale.

Figure 2.

Vertical lines represent 99% credible sets for (a) the fixed effect coefficients for 72 predictors used in the model and (b) the diagonal elements of the random effects covariance matrix. (c) and (d) provide a closer look at those credible sets that are closer to zero.

Figure 3.

(a) gives the predicted vs. observed response (weight in kg) plot where black circles represent the point estimates for the shrinkage model, dashed line is the 45° line and the solid line is the linear fit between the predicted and observed values. Shaded area gives the 99% point-wise credible region for the mean response. (b) tracks the lower-bound (upper) and $\log \psi$ for convergence (lower) over time. Vertical dashed line marks the iteration/time the pre-specified convergence criterion is reached ($\psi = 10^{-6}$).