

Collegio Carlo Alberto



A note on MCMC for nested multilevel regression models via belief propagation

Omiros Papaspiliopoulos
Giacomo Zanella

No. 540
December 2017

Carlo Alberto Notebooks

www.carloalberto.org/research/working-papers

A note on MCMC for nested multilevel regression models via belief propagation

Omiros Papaspiliopoulos* and Giacomo Zanella†

September 15, 2017

Abstract

In the quest for scalable Bayesian computational algorithms we need to exploit the full potential of existing methodologies. In this note we show that the posterior distribution of the regression parameters conditionally on covariance hyperparameters in a large family of multilevel regression models is a high-dimensional Gaussian that can be sampled exactly (as well as marginalized) using belief propagation at a cost that scales linearly in the number of parameters and data. We derive an algorithm that works efficiently even for conditionally singular Gaussian distributions, e.g., when there are linear constraints between the parameters at different levels. We provide a synthesis of similar ideas that have been suggested in the literature before.

1 Belief propagation

Graphical models are now mainstream for modelling high-dimensional vectors of dependent random variables and for doing efficient computations with the joint density thereof, such as marginalisations and maximisations. Among others, two algorithmic paradigms have been established for scalable computations in graphs. More popular within statistics is the junction tree algorithm, see for example Chapter 6 of Cowell et al. (1999), and more popular within machine learning is belief propagation, see for example Section 8.4 in Bishop (2006). The connections between the two paradigms are understood, see for example Chapter 2 of Wainwright and Jordan (2008). For the class of models we consider in this article the two paradigms are essentially equivalent, so we will use the belief propagation formulation, which we find more intuitive in our context and we describe below. The starting point of belief propagation is a factorisation of the joint density of a vector of random variables. (In Section 2 we work with a broader framework that starts with a factorisation of the joint law of the variables.) The set of variables and factors are then represented by a bipartite graph known as the factor graph. There are two sets of nodes: variable nodes (one for each of the random variables) and factor nodes (one for each of the factors in the density factorisation). The edges in the graph connect each factor node to the variable nodes that correspond to those variables involved in the corresponding factor in the factorisation. Details on this construction can be found for example in Section 8.4 of Bishop (2006).

*Institució Catalana de Recerca i Estudis Avançats, Ramon Trias Fargas 25-27, Barcelona 08005 omiros.papaspiopoulos@upf.edu

†Department of Decision Sciences, BIDSa and IGIER, Bocconi University, via Roentgen 1, 20136 Milan, Italy. giacomo.zanella@unibocconi.it. GZ supported by the European Research Council (ERC) through StG “N-BNP” 306406.

An example is the two-level hierarchical model

$$\begin{aligned}\beta &\sim \mathcal{N}(\mu_0, V_0) \\ \beta_i | \beta &\sim \mathcal{N}(\beta, \Sigma) \\ y_i | \beta_i &\sim \mathcal{N}(X_i^y \beta_i, \sigma_i^2 I)\end{aligned}\tag{1}$$

where I is the identity matrix whose dimensions vary depending on the context; $\mathcal{N}(\mu, \Sigma)$ denotes the Gaussian distribution, later $\mathcal{N}(x; \mu, \Sigma)$ will denote the corresponding Gaussian density with argument x ; X_i^y is a matrix of covariates, one for each i , where superscription by y is explained in Section 2 in the context of a generalisation of (1). The last line of (1) defines a linear regression model, $y_{ij} \sim \mathcal{N}((X_{ij}^y)^T \beta_i, \sigma_i^2)$, where $(X_{ij}^y)^T$ is the j th row of matrix X_i^y . The second line of the model pools the local regression coefficients for each i towards a global parameter vector β . The directed acyclic graph of this model for given covariance and design matrices is shown in Figure 1(left). Each of the conditional densities in the model specification corresponds to a factor and the associated factor graph is shown in Figure 1(right).

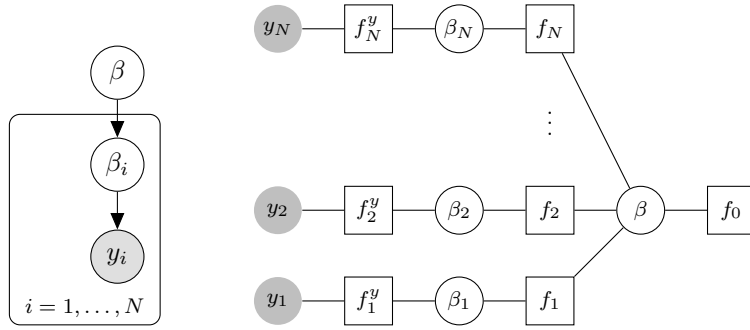


Figure 1: Two-level hierarchical model as directed acyclic graph (left) and as a factor graph (right). The notation used for the factor nodes is explained in Section 2.

Belief propagation involves exchange of messages between variable and factor nodes. Let t denote a variable node in the graph, β_t the corresponding variable, s a neighbouring factor node, f_s the corresponding factor in the density factorisation, and $ne(\cdot)$ be a function that takes as input a node index and returns the set of neighbouring nodes in the graph. By construction factor nodes neighbour variable nodes only and variable nodes neighbour factor nodes only. Hence, f_s is a function of β_t and $\beta_{ne(s)\setminus t}$. Messages exchanged between variable and factor nodes are functions of the variable attached to the variable node and are defined by the following equations:

$$\begin{aligned}m_{\beta_t \rightarrow f_s}(x) &= \prod_{k \in ne(t) \setminus s} m_{f_k \rightarrow \beta_t}(x) \\ m_{f_s \rightarrow \beta_t}(z_t) &= \int f_s(z_{ne(s)}) \prod_{j \in ne(s) \setminus t} m_{\beta_j \rightarrow f_s}(z_j) dz_j.\end{aligned}\tag{2}$$

In terms of notation used in this paper, for sets A and B , $A \setminus B$ denotes the set difference, we identify the variable index t with the one-element set that contains it, for a set of indices A , β_A denotes the set of variables indexed by elements in A , x and z denote generic function arguments. If the factor graph is a tree, as for example in Figure 1(right), and the integrals in (2) can be computed, then the system of equations defined by (2) can

be efficiently solved with a forward-backward algorithm. After choosing an arbitrary variable node as the root of the tree, the direction from the leaves to the root is defined as forward and the reverse as backward. For a given node t , its neighbour on its unique path towards the root will be called its parent and denoted by $pa(t)$. The algorithm is initialised by setting the messages at the variables nodes at the leaves (if any) equal to 1. The forward pass of belief propagation computes messages going from the leaves to the root and the backward pass computes messages from the root to the leaves. After two sweeps we obtain the set of messages solving (2). The resulting messages can then be used to compute the marginal density at any variable node t as $p_{\beta_t}(x) = \prod_{s \in ne(t)} m_{f_s \rightarrow \beta_t}(x)$. When some variables in the definition of the graphical model have been observed, say y_A , the message passing steps that involve these variables keep them fixed at their observed values and do not integrate them out. Then, after the forward and backward step the product of the messages at a variable node t is proportional to the conditional density $p_{\beta_t|y_A}(x)$. Although belief propagation is typically developed for computing marginal distributions and normalising constants, it can also be used to simulate from conditional distributions. This is exploited in the following section, where in the backward pass simulation replaces message computation.

2 Bayesian nested multilevel models

2.1 Model formulation The methods we develop are appropriate for $(k + 1)$ -level models with the following nested structure

$$\begin{aligned}
 \beta &\sim \mathcal{N}(\mu_0, V_0) \quad \text{or} \quad \beta \sim p(\beta) \propto 1 \\
 \beta_i &| \beta \sim \mathcal{N}(X_i \beta, \Sigma) \\
 \beta_{i_1 \dots i_d} &| \beta_{i_1 \dots i_{d-1}} \sim \mathcal{N}(X_{i_1 \dots i_d} \beta_{i_1 \dots i_{d-1}}, \Sigma_{i_1 \dots i_{d-1}}), \quad d = 2, \dots, k \\
 y_{i_1 \dots i_k} &| \beta_{i_1 \dots i_k} \sim \mathcal{N}(X_{i_1 \dots i_k}^y \beta_{i_1 \dots i_k}, \sigma_{i_1 \dots i_k}^2 I).
 \end{aligned} \tag{3}$$

The level closest to the data will be understood as the deepest one. The last line of (3) is a linear regression model where we have concatenated observations at the deepest level into the vector $y_{i_1 \dots i_k}$, like we did in (1). We allow for level-specific covariates that explain the variation of regression coefficients at that level. Throughout the following presentation there will be two sets of quantities associated to the deepest level and that are subscripted in the same way. To distinguish between them, we add the superscript y to those that link to the data; therefore, $X_{i_1 \dots i_k}^y$ is the matrix of covariates used to explain the dataset $y_{i_1 \dots i_k}$, whereas $X_{i_1 \dots i_k}$ is used to explain $\beta_{i_1 \dots i_k}$. We also allow for flat prior on the hyperparameters, which yields a proper posterior under a condition on the matrix P_0 defined in (10) below. The set of all observations $y_{i_1 \dots i_k}$ will be referred to as data, that of all regression parameters, i.e. β and all $\beta_{i_1 \dots i_d}$, as coefficients and that of all covariance hyperparameters as covariances, e.g. we will use $p(\text{data} | \text{covariances})$ to refer to marginal likelihood of observations with regression parameters integrated out.

By allowing rank deficient covariance matrices we can fit in this framework mixed-effects models (also known as linear mixed models) as discussed for example in Section 13.3 of Gelman and Hill (2007). To clarify the concept, an example more complex than (1) is helpful. Suppose we have data on individuals j that belong to different countries i , that we model as a regression $y_{ij} \sim \mathcal{N}(\alpha + \theta_i z_{ij} + \gamma_i w_{ij}, \sigma_i^2)$, where z_{ij} and w_{ij} are individual-level covariates. We want to allow variation of the corresponding coefficients that we partially explain using a country-level covariate s_i , $\theta_i \sim \mathcal{N}(a + b s_i, \tau)$ and $\gamma_i \sim \mathcal{N}(c + d s_i, \lambda)$ with θ_i and γ_i independent. In order to obtain a factor graph that is a tree, we need to clump coefficients at the deepest level into one vector $\beta_i = (\alpha_i, \theta_i, \gamma_i)^T$,

and the covariates and a vector of 1's into the matrix X_i^y . The α_i 's are artificial copies of α . Even though we model the variation on θ_i independently of that on γ_i , having them clumped together forces us to work with them as a single vector. Hence, we write $\beta_i \sim \mathcal{N}(X_i\beta, \Sigma)$ where $\beta = (\alpha, a, b, c, d)^T$, $\Sigma = \text{diag}(0, \tau, \lambda)$, and X_i is a 3×5 matrix with rows $(1, 0, 0, 0, 0)$, $(0, 1, s_i, 0, 0)$, and $(0, 0, 0, 1, s_i)$. Hence, we have obtained the nested multilevel structure of (3) with $k = 1$. An example of a structure that does not fit the framework of (3) is the crossed-effects model, $y_{ij} \sim \mathcal{N}(\alpha_i + \theta_j, \sigma^2)$, to which the belief propagation algorithm we describe below does not apply.

2.2 Message passing for posterior simulation We describe the messages for belief propagation in nested multilevel models as defined in (3), and how they can be used for computing $p(\text{data} \mid \text{covariances})$ and simulating from $p(\text{coefficients} \mid \text{data}, \text{covariances})$. The factor graph associated with (3) is a tree; throughout our presentation we will take β to be its root. The factor nodes will be denoted by $f_{i_1 \dots i_d}$; as already discussed the factor node that links $y_{i_1 \dots i_k}$ to $\beta_{i_1 \dots i_k}$ will be denoted by $f_{i_1 \dots i_k}^y$ and that linking $\beta_{i_1 \dots i_k}$ to $\beta_{i_1 \dots i_{k-1}}$ will be denoted by $f_{i_1 \dots i_k}$; see Figure 1. Messages are functions that can be parameterised in terms of a triplet (c, C, u) , where c is a positive scalar, C a positive semidefinite matrix and u a vector, and for brevity we write

$$m = (c, C, u) \quad \text{to imply} \quad m(x) = c \exp \left(- (1/2)x^T C x + u^T x \right). \quad (4)$$

First we show how to update this set of parameters during the forward step, which works from leaves to the root, i.e., from the deepest level in the hierarchy to the most shallow. With every factor node, except for those at the deepest level linked to the observations, there are associated two messages, one that arrives from the variable node deeper and one that is sent to the variable node one level higher up. We distinguish between the two corresponding triplets by using tildes for the latter. In this context, the message that arrives at a variable node from below coincides with the density of the data on the leaves that originate from the given branch of the tree, conditional on the variable at this node but marginal with respect to all other variables in-between. The messages from factor to variable nodes at the leaves are

$$\begin{aligned} m_{f_{i_1 \dots i_k}^y \rightarrow \beta_{i_1 \dots i_k}} &= (\tilde{c}_{i_1 \dots i_k}^y, \tilde{C}_{i_1 \dots i_k}^y, \tilde{u}_{i_1 \dots i_k}^y) \\ \tilde{C}_{i_1 \dots i_k}^y &= (1/\sigma_{i_1 \dots i_k}^2)(X_{i_1 \dots i_k}^y)^T X_{i_1 \dots i_k}^y \\ \tilde{c}_{i_1 \dots i_k}^y &= (2\pi\sigma_{i_1 \dots i_k}^2)^{-d_{i_1 \dots i_k}/2} \exp \left(- (1/2\sigma_{i_1 \dots i_k}^2) y_{i_1 \dots i_k}^T y_{i_1 \dots i_k} \right) \\ \tilde{u}_{i_1 \dots i_k}^y &= (1/\sigma_{i_1 \dots i_k}^2)(X_{i_1 \dots i_k}^y)^T y_{i_1 \dots i_k} \end{aligned} \quad (5)$$

where $d_{i_1 \dots i_k}$ is the size of $y_{i_1 \dots i_k}$. At higher levels, the factor to variable messages are

$$\begin{aligned} m_{f_{i_1 \dots i_d} \rightarrow \beta_{i_1 \dots i_{d-1}}} &= (\tilde{c}_{i_1 \dots i_d}, \tilde{C}_{i_1 \dots i_d}, \tilde{u}_{i_1 \dots i_d}) \\ \tilde{c}_{i_1 \dots i_d} &= c_{i_1 \dots i_d} |G|^{1/2} \exp \left(- (1/2) u_{i_1 \dots i_d}^T \Gamma^T G \Gamma u_{i_1 \dots i_d} \right) \\ \tilde{C}_{i_1 \dots i_d} &= X_{i_1 \dots i_d}^T B^T (B \Sigma_{i_1 \dots i_{d-1}} B^T + I)^{-1} B X_{i_1 \dots i_d} \\ \tilde{u}_{i_1 \dots i_d} &= X_{i_1 \dots i_d}^T (\Sigma_{i_1 \dots i_{d-1}} C_{i_1 \dots i_d} + I)^{-1} u_{i_1 \dots i_d} \end{aligned} \quad (6)$$

where B , Γ and G implicitly depend on $(i_1 \dots i_d)$ and are defined by

$$C_{i_1 \dots i_d} = B^T B, \quad \Sigma_{i_1 \dots i_{d-1}} = \Gamma^T \Gamma, \quad G = (\Gamma C_{i_1 \dots i_d} \Gamma^T + I)^{-1}. \quad (7)$$

The triplet $(c_{i_1 \dots i_d}, C_{i_1 \dots i_d}, u_{i_1 \dots i_d})$ in (6) corresponds to the message $m_{\beta_{i_1 \dots i_d} \rightarrow f_{i_1 \dots i_d}}$. The matrix decompositions in (7) are used with Schur's complement to obtain formulae

above that are valid without any assumptions on invertibility of either $C_{i_1 \dots i_d}$ or $\Sigma_{i_1 \dots i_{d-1}}$. $\Sigma_{i_1 \dots i_{d-1}} C_{i_1 \dots i_d} + I$ is invertible by construction because $\Sigma_{i_1 \dots i_{d-1}}$ is positive semidefinite.

The messages from variable to factor nodes are as follows: at the deepest level

$$m_{\beta_{i_1 \dots i_k} \rightarrow f_{i_1 \dots i_k}} = (c_{i_1 \dots i_k}, C_{i_1 \dots i_k}, u_{i_1 \dots i_k}) = \left(\tilde{c}_{i_1 \dots i_k}^y, \tilde{C}_{i_1 \dots i_k}^y, \tilde{u}_{i_1 \dots i_k}^y \right) \quad (8)$$

while for any $d \in \{1, \dots, k-1\}$

$$m_{\beta_{i_1 \dots i_d} \rightarrow f_{i_1 \dots i_d}} = (c_{i_1 \dots i_d}, C_{i_1 \dots i_d}, u_{i_1 \dots i_d}) = \left(\prod_j \tilde{c}_{i_1 \dots i_d j}, \sum_j \tilde{C}_{i_1 \dots i_d j}, \sum_j \tilde{u}_{i_1 \dots i_d j} \right) \quad (9)$$

where j runs over the appropriate index set, which depends on the numbers of offsprings $\beta_{i_1 \dots i_d}$ has on the Bayesian network.

At the root we collect messages from f_i for $i \geq 1$, which come from the level below, and from f_0 , which comes from the prior. The normalised message is the posterior density at the root:

$$\begin{aligned} \beta \mid \text{data, covariances} &\sim \mathcal{N}(P^{-1}m, P^{-1}) \\ P &= P_0 + V_0^{-1} \quad \text{or} \quad P = P_0, \quad P_0 = \sum_i \tilde{C}_i \\ m &= \sum_i \tilde{u}_i + V_0^{-1}\mu_0 \quad \text{or} \quad m = \sum_i \tilde{u}_i \end{aligned} \quad (10)$$

where the alternative expressions depend on whether a Gaussian or flat prior is used for β . Flat prior leads to proper posterior provided P_0 defined above is positive definite. Additionally,

$$p(\text{data} \mid \text{covariances}) = \frac{\mathcal{N}(x; \mu_0, V_0) \prod_i \tilde{c}_i \exp \left\{ -(1/2)x^T \tilde{C}_i x + \tilde{u}_i^T x \right\}}{\mathcal{N}(x; P^{-1}m, P^{-1})}, \quad (11)$$

where the first term in nominator is omitted for flat prior, and any value x can be used. Then, for any intermediate level regression coefficients, we have

$$\begin{aligned} \beta_{i_1 \dots i_d} \mid \beta_{i_1 \dots i_{d-1}}, \text{data, covariances} &\sim \\ \mathcal{N} \left\{ G(X_{i_1 \dots i_d} \beta_{i_1 \dots i_{d-1}} + \Sigma_{i_1 \dots i_{d-1}} u_{i_1 \dots i_d}), \Gamma^T (\Gamma C_{i_1 \dots i_d} \Gamma^T + I)^{-1} \Gamma \right\} & \quad (12) \\ G = (\Sigma_{i_1 \dots i_{d-1}} C_{i_1 \dots i_d} + I)^{-1}, \quad \Sigma_{i_1 \dots i_{d-1}} = \Gamma^T \Gamma. & \end{aligned}$$

Simulating backwards according to the distributions described in (10) and (12) we obtain a draw from $p(\text{coefficients} \mid \text{data, covariances})$.

We have not given the details of the calculations that produce the formulae for the messages and the conditional distributions. We have worked under a framework mathematically richer than that of (2), where the definition of a factor graph is extended to correspond to a factorisation of the joint probability measure in terms of regular conditional densities, along the lines of the disintegration theorem as in Theorem 5.4 of Kallenberg (1997). This type of construction is suitable for the factor graph representation of Bayesian networks with conditional Gaussian distributions with semi-definite covariance matrices. In the context of a multilevel model, the messages are Radon-Nikodym derivatives between Gaussian measures. It can be checked (but we have omitted the details here) that indeed the posterior Gaussian laws (12) are absolutely continuous with respect to the prior Gaussian laws in the definition of the multilevel model in (3) with Radon-Nikodym derivative proportional to the message $m_{f_{i_1 \dots i_{d+1}} \rightarrow \beta_{i_1 \dots i_d}}$.

2.3 Complexity considerations Both computation of $p(\text{data} \mid \text{covariances})$ and simulation from $p(\text{coefficients} \mid \text{data}, \text{covariances})$ via belief propagation as described above involve a computational cost that scales linearly with the total number of regression coefficient vectors. The forward and backward steps at each level involve computations that without further structural assumptions scale cubically with the dimension of the level-specific coefficients. On the other hand, if we wish to draw several samples for the coefficients for given values of the covariances, the cost per step can be made quadratic in the dimension of the coefficient vector, since matrix decompositions do not have to be redone. In other words, whereas the computational cost of the algorithm has the same dependence on the characteristics of the model as that of a Gibbs sampling algorithm that simulates each regression coefficient conditionally on the rest, it achieves exact draws from the high-dimensional distribution $p(\text{coefficients} \mid \text{data}, \text{covariances})$. On the other hand, to obtain a factor tree graph we might have to clump together variables, hence increasing the dimension of the variable node dimensions, as we did for the mixed-effects model example in Section 2.1. In such cases the computational cost per iteration of the Gibbs sampler will be smaller. However, such differences will typically be small relative to the overall cost of the algorithms. Additionally, one can use belief propagation to update a subset of the variables conditionally on the rest, as we discuss in the next section, hence obtaining again the same cost per iteration as the Gibbs sampler. Sampling using belief propagation lends itself to parallelisation, but we do not develop this idea further here, although it should be considered for software development.

3 Blocked Gibbs sampling and marginal algorithms

The algorithm of Section 2 can be used in the popular context where the covariance components are given a prior distribution and Bayesian inferences are performed on the joint distribution $p(\text{coefficients}, \text{covariances} \mid \text{data})$. Posterior computations are typically performed using a Gibbs sampling scheme that alternates sampling from $p(\text{covariances} \mid \text{data}, \text{coefficients})$ and $p(\text{coefficients} \mid \text{covariances}, \text{data})$. If conditionally conjugate priors (e.g. Wishart) are used, sampling from $p(\text{covariances} \mid \text{data}, \text{coefficients})$ can be efficiently accomplished exploiting the conditional independence across covariance components given the coefficients and the data. Belief propagation can then be used to sample from $p(\text{coefficients} \mid \text{data}, \text{covariances})$ efficiently.

An alternative approach is to perform inferences on the marginal space $p(\text{covariances} \mid \text{data})$. Even in this case, the belief propagation algorithm of Section 2 is crucial to allow efficient point-wise evaluation of the marginal likelihood $p(\text{data} \mid \text{covariances})$, and thus of $p(\text{covariances} \mid \text{data})$ up to proportionality. The relative merits of working with the extended posterior $p(\text{coefficients}, \text{covariances} \mid \text{data})$ or with the marginal one $p(\text{covariances} \mid \text{data})$ are case-specific. In general one would expect the marginal approach to be preferable when the distribution $p(\text{covariances} \mid \text{data})$ is small or moderate dimensional, so that numerical or Monte Carlo integration can be efficiently employed in the marginal space. When instead the distribution $p(\text{covariances} \mid \text{data})$ is high-dimensional it may be challenging to perform inferences in the marginal space because, having integrated out the regression coefficients, the covariance components are not anymore conditionally independent and one ends up with a potentially complex high-dimensional correlated distribution. In this case the blocked Gibbs Sampling approach may be more efficient.

The methodology we presented so far is well-suited to contexts where the number of regression coefficients is large (i.e. tree with many nodes) and the size of each coefficient is small, since the cost of operations grows as the cube of that size but only linearly with the number of nodes in the tree. When the size of variable nodes is large, we can still use

the message passing algorithm of Section 2 to update a subset of regression coefficients (i.e. some components of each $\beta_{i_1 \dots i_d}$) jointly across all nodes of the tree conditionally on the remaining coefficients. The resulting blocked Gibbs sampling scheme can still exhibit large improvements over the single-site updating. A more detailed study of this case is left to future work.

4 Synthesis of the related literature

The forward-backward iteration we describe for sampling efficiently the high-dimensional Gaussian distribution, $p(\text{coefficients}|\text{data}, \text{covariances})$, is a generalisation of the so-called forward filtering backward sampling algorithm, which is used in the time series community for Bayesian inference for state space models, see for example Algorithm 13.4 in Frühwirth-Schnatter (2006). State space models have factor graphs with single-branch tree structure. The extension of Kalman filtering recursion to tree structures has been long-known, especially in the context of multiscale systems (e.g. Chou et al., 1994). Similar ideas have been exploited in spatial statistics contexts, see for example Huang and Cressie (2001) where algorithms for spatial Gaussian models with tree-structured dependence are developed. Zhang and Agarwal (2008) exploit Kalman filter recursions to perform posterior maximization for some multilevel models that are subset of the framework we consider here. These previous works typically focus on computing marginal distributions and do not develop sampling algorithms, and they do not make a clear connection with belief propagation. The connection between belief propagation and Kalman recursions is recognised in the early technical report Dempster (1990) and using a message-passing formulation in Normand and Tritchler (1992), who provide references to works even older than Dempster (1990).

The role of belief propagation within Bayesian computation for multilevel models is fully recognised in Wilkinson and Yeung (2002) who work along the same lines we have followed in this note, in particular their Section 2.4 on message passing for sampling posteriors that arise in Gaussian tree models (the approach we use in Section 2.2 corresponds to what they call the canonical parameterisation of the multivariate Gaussian). Relative to that work the main novelty in this note is that we have worked out messages with no assumptions on invertibility of prior covariance matrices. This extension has allowed us to also cast mixed effect models as nested multilevel models and use belief propagation for those too. For the two-level mixed effect models Chib and Carlin (1999) derive an efficient algorithm for sampling the posterior distribution of the regression coefficients, which is an instance of the generic algorithm of Section 2, where the structure of 0's in prior covariances that results from clumping is explicitly exploited to simplify some of the matrix computations.

Seen as latent Gaussian models, the nested multilevel models we consider give rise to large sparse precision matrices. Wilkinson and Yeung (2002) and Wilkinson and Yeung (2004) show how to compute the canonical parameters of the Gaussian prior and posterior using sparse linear algebra computations and exploiting the graphical model structure to identify the zeros in the precision matrices. The potential of sparse linear algebra computations for inference and simulation in latent Gaussian models has been recognised at least since Rue (2001); this work, and inter alia the follow-ups Knorr-Held and Rue (2002) and Rue and Held (2005) have illustrated how the precision matrices that arise in spatiotemporal latent Gaussian models can be treated using sparse linear algebra algorithms, such as algorithms for banded matrices and algorithms that attempt to produce Cholesky factors nearly as sparse as the posterior precision. This work has also recognised the connections between their fast linear algebra approach and Kalman

filters/smoothers; see for example the overview in Section 1.2.1 of Rue and Held (2005). The way numerical analysis techniques for sparse matrices can be used for computations with Gaussian graphical models is exposed in Section 2.4 of Rue and Held (2005), and Section 2.5 of the same book carries out a simulation study that tries two black-box algorithms for sparse matrices in the context of precision matrices that arise in spatiotemporal graphical models. Wilkinson and Yeung (2004) suggests the use of such sparse linear algebra methods for inference in Gaussian hierarchical models; first compile the sparse precision using computations that exploit the known positions of the zeros and then feed the precision into a sparse linear algebra algorithm that returns its Cholesky factor. The article does not study whether the structure of zeros in the precision from common families of hierarchical models, such as for example the nested hierarchical models considered in this note, is such that the extraction of the Cholesky factor is actually fast or whether the resultant factor is sparse, or the complexity of the associated operations. These type of questions have been studied in more depth in the case of spatial models in Chapter 2 of Rue and Held (2005).

In the opposite direction, the linear algebra community has explored the use of belief propagation as an efficient iterative method for solving linear systems. An example of this line of work is Shental et al. (2008), who use belief propagation for solving linear systems with positive definite matrices by linking the system solution to the marginal means (or the mode) of a Gaussian distribution with precision matrix given by the system matrix. They discuss connections of this approach to direct and iterative methods for solving systems and find it competitive.

Acknowledgements

The authors are grateful to Darren Wilkinson for his input.

References

- Bishop, C. M. (2006). *Pattern recognition and machine learning*. Information Science and Statistics. Springer, New York.
- Chib, S. and Carlin, B. P. (1999). On mcmc sampling in hierarchical longitudinal models. *Statistics and Computing*, (9):17–26.
- Chou, K. C., Willsky, A. S., and Nikoukhah, R. (1994). Multiscale systems, Kalman filters, and Riccati equations. *IEEE Transactions on Automatic Control*, 39(3):479–492.
- Cowell, R. G., Dawid, A. P., Lauritzen, S. L., and Spiegelhalter, D. J. (1999). *Probabilistic networks and expert systems*. Statistics for Engineering and Information Science. Springer-Verlag, New York.
- Dempster, A. (1990). Normal belief functions and the kalman filter.
- Frühwirth-Schnatter, S. (2006). *Finite mixture and Markov switching models*. Springer Series in Statistics. Springer, New York.
- Gelman, A. and Hill, J. (2007). *Data analysis using regression and multilevel/hierarchical models*, volume 3. Cambridge University Press New York, New York, USA.

- Huang, H.-C. and Cressie, N. (2001). Multiscale graphical modeling in space: applications to command and control. In *Spatial statistics: methodological aspects and applications*, volume 159 of *Lect. Notes Stat.*, pages 83–113. Springer, New York.
- Kallenberg, O. (1997). *Foundations of modern probability*. Probability and its Applications (New York). Springer-Verlag, New York.
- Knorr-Held, L. and Rue, H. v. (2002). On block updating in Markov random field models for disease mapping. *Scand. J. Statist.*, 29(4):597–614.
- Normand, S.-L. and Trichtler, D. (1992). Parameter updating in a Bayes network. *Journal of the American Statistical Association*, 87(420):1109–1115.
- Rue, H. and Held, L. (2005). *Gaussian Markov random fields: theory and applications*. Chapman & Hall.
- Rue, H. v. (2001). Fast sampling of Gaussian Markov random fields. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 63(2):325–338.
- Shental, O., Bickson, D., Siegel, P. H., Wolf, J. K., and Dolev, D. (2008). Gaussian belief propagation solver for systems of linear equations. In *EEE Int. Symp. on Inform. Theory (ISIT)*.
- Wainwright, M. J. and Jordan, M. I. (2008). Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305.
- Wilkinson, D. J. and Yeung, S. K. H. (2002). Conditional simulation from highly structured Gaussian systems, with application to blocking-MCMC for the Bayesian analysis of very large linear models. *Stat. Comput.*, 12(3):287–300.
- Wilkinson, D. J. and Yeung, S. K. H. (2004). A sparse matrix approach to Bayesian computation in large linear models. *Comput. Statist. Data Anal.*, 44(3):493–516.
- Zhang, L. and Agarwal, D. (2008). Fast computation of posterior mode in multi-level hierarchical models. In Koller, D., Schuurmans, D., Bengio, Y., and Bottou, L., editors, *Advances in Neural Information Processing Systems 21*, pages 1913–1920.