

Fast and Accurate Binary Response Mixed Model Analysis via Expectation Propagation

BY P. HALL¹, I.M. JOHNSTONE², J.T. ORMEROD³, M.P. WAND⁴ AND J.C.F. YU⁴

¹University of Melbourne, ²Stanford University, ³University of Sydney
and ⁴University of Technology Sydney

31st May, 2018

Abstract

Expectation propagation is a general prescription for approximation of integrals in statistical inference problems. Its literature is mainly concerned with Bayesian inference scenarios. However, expectation propagation can also be used to approximate integrals arising in *frequentist* statistical inference. We focus on likelihood-based inference for binary response mixed models and show that fast and accurate quadrature-free inference can be realized for the probit link case with multivariate random effects and higher levels of nesting. The approach is supported by asymptotic theory in which expectation propagation is seen to provide consistent estimation of the exact likelihood surface. Numerical studies reveal the availability of fast, highly accurate and scalable methodology for binary mixed model analysis.

Keywords: Best prediction; Generalized linear mixed models; Maximum likelihood; Kullback-Leibler projection; Message passing; Quasi-Newton methods; Scalable statistical methodology.

1 Introduction

Binary response mixed model-based data analysis is ubiquitous in many areas of application, with examples such as analysis of biomedical longitudinal data (e.g. Diggle *et al.*, 2002), social science multilevel data (e.g. Goldstein, 2010), small area survey data (e.g. Rao & Molina, 2015) and economic panel data (e.g. Baltagi, 2013). The standard approach for likelihood-based inference in the presence of multivariate random effects is Laplace approximation, which is well-known to be inconsistent and prone to inferential inaccuracy. Our main contribution is to overcome this problem using expectation propagation. The new approach possesses speed and scalability on par with that of Laplace approximation, but is provably consistent and demonstrably very accurate. Bayesian approaches and Monte Carlo methods offer another route to accurate inference for binary response mixed models (e.g. Gelman & Hill, 2007). However, speed and scalability issues aside, frequentist inference is the dominant approach in many areas in which mixed models are used. Henceforth, we focus on frequentist binary mixed model analysis.

The main obstacle for likelihood-based inference for binary mixed models is the presence of irreducible integrals. For grouped data with one level of nesting, the dimension of the integrals matches the number of random effects. The two most common approaches to dealing with these integrals are (1) quadrature and (2) Laplace approximation. For example, in the R computing environment (R Core Team, 2018) the function `glmer()` in the package `lme4` (Bates *et al.*, 2015) supports both adaptive Gauss-Hermite quadrature and Laplace approximation for univariate random effects. For multivariate random effects only Laplace approximation is supported by `glmer()`, presumably because of the inherent difficulties of higher dimensional quadrature. Laplace approximation eschews multivariate integration via quadratic approximation of the log-integrand. However, the

resultant approximate inference is well-known to be inaccurate, often to an unacceptable degree, in binary mixed models (e.g. McCulloch *et al.*, Section 14.4). An embellishment of Laplace approximation, known as integrated nested Laplace approximation (Rue, Martino & Chopin, 2009), has been successful in various Bayesian inference contexts.

Expectation propagation (e.g. Minka, 2001) is general prescription for approximation of integrals that arise in statistical inference problems. Most of its literature is within the realm of Computer Science and, in particular, geared towards approximate inference for Bayesian graphical models (e.g. Chapter 10, Bishop, 2006). A major contribution of this article is transferral of expectation propagation methodology to frequentist statistical inference. In principle, our approach applies to any generalized linear mixed model situation. However, expectation propagation for binary response mixed model analysis has some especially attractive features and therefore we focus on this class of models. In the special case of probit mixed models, the expectation propagation approximation to the log-likelihood is exact regardless of the dimension of the random effects. This leads to a new practical alternative to multivariate quadrature. Moreover, asymptotic theory reveals that expectation propagation provides consistent approximation of the exact likelihood surface. This implies very good inferential accuracy of expectation propagation, and is supported by our simulation results. We are not aware of any other quadrature-free approaches to generalized mixed model analysis that has such a strong theoretical underpinning.

To facilitate widespread use of the new approach, a new package in the R language (R Core Team, 2018) has been launched. The package, `glmmEP` (Wand & Yu, 2018), uses a low-level language implementation of expectation propagation for speedy approximate likelihood-based inference and scales well to large sample sizes.

Binary response mixed models, and their inherent computational challenges, are summarized in Section 2. The expectation propagation approach to fitting and approximate inference, with special attention given to the quadrature-free probit link situation, is given in Section 3. Section 4 presents the results of numerical studies for both simulated and real data, and shows expectation propagation to be of great practical value as a fast, high quality approximation that scales well to big data and big model situations. Theoretical considerations are summarised in Section 5. Higher level and random effects extensions are touched upon in Section 6. Lastly, we briefly discuss transferral of new approach to other generalized linear mixed model settings in Section 7.

2 Binary Response Mixed Models

Binary mixed models for grouped data with one level of nesting and Gaussian random effects has the general form

$$y_{ij} | \mathbf{u}_i \stackrel{\text{ind.}}{\sim} \text{Bernoulli}(F(\boldsymbol{\beta}^T \mathbf{x}_{ij}^F + \mathbf{u}_i^T \mathbf{x}_{ij}^R)), \quad \mathbf{u}_i \stackrel{\text{ind.}}{\sim} N(\mathbf{0}, \boldsymbol{\Sigma}), \quad 1 \leq i \leq m, \quad 1 \leq j \leq n_i \quad (1)$$

where F , the *inverse link*, is a pre-specified cumulative distribution function and y_{ij} is the j th response for the i th group, where number of groups is m and the number of responses measurements within the i th group is n_i . Also, \mathbf{x}_{ij}^F is a $d^F \times 1$ vector of predictors corresponding to y_{ij} , modeled as having fixed effects with coefficient vector $\boldsymbol{\beta}$. Similarly, \mathbf{x}_{ij}^R is a $d^R \times 1$ vector of predictors modeled as having random effects with coefficient vectors \mathbf{u}_i , $1 \leq i \leq m$. Typically, \mathbf{x}_{ij}^R is a sub-vector of \mathbf{x}_{ij}^F . It is also very common for each of \mathbf{x}_{ij}^F and \mathbf{x}_{ij}^R to have first entry equal to 1, corresponding to fixed and random intercepts. The random effects covariance matrix $\boldsymbol{\Sigma}$ has dimension $d^R \times d^R$.

By far, the most common choices for F are

$$F = \begin{cases} \text{expit} & \text{for logistic mixed models} \\ \Phi & \text{for probit mixed models} \end{cases}$$

where $\text{expit}(x) \equiv 1/(1 + e^{-x})$ and Φ is the cumulative distribution function of the $N(0, 1)$ distribution.

Despite the simple form of (1), likelihood-based inference for the parameters β and Σ and best prediction of the random effects \mathbf{u}_i is very numerically challenging. Assuming that $F(x) + F(-x) = 1$, as is the case for the logistic and probit cases, the log-likelihood is

$$\ell(\beta, \Sigma) = \sum_{i=1}^m \log \int_{\mathbb{R}^{d^R}} \left\{ \prod_{j=1}^{n_i} F((2y_{ij} - 1)(\beta^T \mathbf{x}_{ij}^F + \mathbf{u}^T \mathbf{x}_{ij}^R)) \right\} |2\pi\Sigma|^{-1/2} \exp(-\frac{1}{2}\mathbf{u}^T \Sigma^{-1} \mathbf{u}) d\mathbf{u} \quad (2)$$

and the best predictor of \mathbf{u}_i is

$$\text{BP}(\mathbf{u}_i) = \frac{\int_{\mathbb{R}^{d^R}} \mathbf{u} \left\{ \prod_{j=1}^{n_i} F((2y_{ij} - 1)(\beta^T \mathbf{x}_{ij}^F + \mathbf{u}^T \mathbf{x}_{ij}^R)) \right\} \exp(-\frac{1}{2}\mathbf{u}^T \Sigma^{-1} \mathbf{u}) d\mathbf{u}}{\int_{\mathbb{R}^{d^R}} \left\{ \prod_{j=1}^{n_i} F((2y_{ij} - 1)(\beta^T \mathbf{x}_{ij}^F + \mathbf{u}^T \mathbf{x}_{ij}^R)) \right\} \exp(-\frac{1}{2}\mathbf{u}^T \Sigma^{-1} \mathbf{u}) d\mathbf{u}}, \quad 1 \leq i \leq m.$$

The d^R -dimensional integrals in the $\ell(\beta, \Sigma)$ and $\text{BP}(\mathbf{u}_i)$ expressions cannot be reduced further and multivariate numerical integration must be called upon for their evaluation. In addition, $\ell(\beta, \Sigma)$ has to be maximized over $\{d^F + \frac{1}{2}d^R(d^R + 1)\}$ -dimensional space to obtain maximum likelihood estimates. Lastly, there is the problem of obtaining approximate confidence intervals for the entries of β and Σ and approximate prediction intervals for the entries of \mathbf{u}_i .

Starting around the early 1990s there have been several proposals for likelihood-based estimation and inference for binary response mixed models and their generalized linear mixed model extensions. Section 14.3 of McCulloch, Searle & Neuhaus (2008) provides a summary of the main approaches up until the mid-2000s. Some more recent contributions, loosely related to the approach proposed here, include Jeon, Rijmen & Rabe-Hesketh (2017), Ogden (2015) and Wand & Ormerod (2012). The relative strengths and weakness of the various proposals depend on attributes such as accuracy, ease of implementation, computational speed and theoretical tractability and properties.

3 Expectation Propagation Likelihood Approximation

We will first explain expectation propagation for approximation of the log-likelihood $\ell(\beta, \Sigma)$. Approximation of $\text{BP}(\mathbf{u}_i)$ follows relatively quickly. First note that $\ell(\beta, \Sigma) = \sum_{i=1}^m \ell_i(\beta, \Sigma)$ where

$$\ell_i(\beta, \Sigma) \equiv \log \int_{\mathbb{R}^{d^R}} \left\{ \prod_{j=1}^{n_i} F((2y_{ij} - 1)(\beta^T \mathbf{x}_{ij}^F + \mathbf{u}^T \mathbf{x}_{ij}^R)) \right\} |2\pi\Sigma|^{-1/2} \exp(-\frac{1}{2}\mathbf{u}^T \Sigma^{-1} \mathbf{u}) d\mathbf{u}.$$

Each of the $\ell_i(\beta, \Sigma)$ are approximated individually and then summed to approximate $\ell(\beta, \Sigma)$. The essence is of the approximation of $\ell_i(\beta, \Sigma)$ is replacement of each

$$F((2y_{ij} - 1)(\beta^T \mathbf{x}_{ij}^F + \mathbf{u}^T \mathbf{x}_{ij}^R)), \quad 1 \leq j \leq n_i,$$

by an unnormalized Multivariate Normal density function, chosen according to an appropriate minimum Kullback-Leibler divergence criterion. The resultant integrand is then proportional to a product of Multivariate Normal density functions and admits an explicit form. The number approximating density functions of the same order of magnitude and, together with the properties of minimum Kullback-Leibler divergence, leads to accurate and statistically consistent approximation of $\ell(\beta, \Sigma)$. In probit case, where $F = \Phi$, the minimum Kullback-Leibler divergence steps are explicit. This leads to accurate approximation of $\ell(\beta, \Sigma)$ without the need for any numerical integration – just some fixed-point iteration. The expectation propagation-approximate log-likelihood, which we denote by

$\ell(\boldsymbol{\beta}, \boldsymbol{\Sigma})$, can be evaluated quite rapidly and maximized using established derivative-free methods such as the Nelder-Mead algorithm (Nelder & Mead, 1965) or quasi-Newton optimization methods such as the Broyden-Fletcher-Goldfarb-Shanno approach with numerical derivatives. The latter also facilitates Hessian matrix approximation at the maximum, which can be used to construct approximate confidence intervals.

We now provide the details, with subsections on each of Kullback-Leibler projection onto unnormalized Multivariate Normal density functions, message passing formulation for organizing the required versions of these projections and quasi-Newton-based approximate inference. The upcoming subsections require some specialized matrix notation. If \mathbf{A} is $d \times d$ matrix then $\text{vec}(\mathbf{A})$ is the $d^2 \times 1$ vector obtained by stacking the columns of \mathbf{A} underneath each other in order from left to right. Also, $\text{vech}(\mathbf{A})$ is $\frac{1}{2}d(d+1) + 1$ vector defined similarly to $\text{vec}(\mathbf{A})$ but only involving entries on and below the diagonal. The *duplication matrix of order d* , denoted by \mathbf{D}_d , is the unique $d^2 \times \frac{1}{2}d(d+1)$ matrix of zeros and ones such that

$$\mathbf{D}_d \text{vech}(\mathbf{A}) = \text{vec}(\mathbf{A}) \quad \text{for } \mathbf{A} = \mathbf{A}^T.$$

The Moore-Penrose inverse of \mathbf{D}_d is

$$\mathbf{D}_d^+ \equiv (\mathbf{D}_d^T \mathbf{D}_d)^{-1} \mathbf{D}_d^T.$$

3.1 Projection onto Unnormalized Multivariate Normal Density Functions

Let $L_1(\mathbb{R}^d)$ denote the set of absolutely integrable functions on \mathbb{R}^d . For $f_1, f_2 \in L_1(\mathbb{R}^d)$ such that $f_1, f_2 \geq 0$, the Kullback-Leibler divergence of f_2 from f_1 is

$$\text{KL}(f_1 \| f_2) = \int_{\mathbb{R}^d} [f_1(\mathbf{x}) \log\{f_1(\mathbf{x})/f_2(\mathbf{x})\} + f_2(\mathbf{x}) - f_1(\mathbf{x})] d\mathbf{x} \quad (3)$$

(e.g. Minka, 2005). In the special case where f_1 and f_2 are density functions the right-hand side of (3) reduces to the more common Kullback-Leibler divergence expression. However, we require this more general form that caters for *unnormalized* density functions.

Now consider the family of functions on \mathbb{R}^d of the form

$$f_{\text{UN}}(\mathbf{x}) \equiv \exp \left\{ \begin{bmatrix} 1 \\ \mathbf{x} \\ \text{vech}(\mathbf{x}\mathbf{x}^T) \end{bmatrix}^T \begin{bmatrix} \eta_0 \\ \boldsymbol{\eta}_1 \\ \boldsymbol{\eta}_2 \end{bmatrix} \right\} \quad (4)$$

where $\eta_0 \in \mathbb{R}$, $\boldsymbol{\eta}_1$ is a $d \times 1$ vector and $\boldsymbol{\eta}_2$ is a $\frac{1}{2}d(d+1) \times 1$ vector restricted in such a way that $f_{\text{UN}} \in L_1(\mathbb{R}^d)$. Then (4) is the family of unnormalized Multivariate Normal density functions written in exponential family form with natural parameters η_0 , $\boldsymbol{\eta}_1$ and $\boldsymbol{\eta}_2$.

Expectation propagation for generalized linear mixed models with Gaussian random effects has the following notion at its core:

$$\text{given } f_{\text{input}} \in L_1(\mathbb{R}^d), \text{ determine the } \eta_0, \boldsymbol{\eta}_1 \text{ and } \boldsymbol{\eta}_2 \text{ that minimizes } \text{KL}(f_{\text{input}} \| f_{\text{UN}}). \quad (5)$$

The solution is termed the (Kullback-Leibler) projection onto the family of Multivariate Normal density functions and we write

$$\text{proj}[f_{\text{input}}](\mathbf{x}) \equiv \exp \left\{ \begin{bmatrix} 1 \\ \mathbf{x} \\ \text{vech}(\mathbf{x}\mathbf{x}^T) \end{bmatrix}^T \begin{bmatrix} \eta_0^* \\ \boldsymbol{\eta}_1^* \\ \boldsymbol{\eta}_2^* \end{bmatrix} \right\}$$

where

$$(\eta_0^*, \boldsymbol{\eta}_1^*, \boldsymbol{\eta}_2^*) = \underset{(\eta_0, \boldsymbol{\eta}_1, \boldsymbol{\eta}_2) \in H}{\text{argmin}} \text{KL}(f_{\text{input}} \| f_{\text{UN}}),$$

with H denoting the set of all allowable natural parameters. Note that the special case of Kullback-Leibler projection onto the unnormalized Multivariate Normal family has a simple moment-matching representation, with $(\eta_0^*, \boldsymbol{\eta}_1^*, \boldsymbol{\eta}_2^*)$ being the unique vector such that zeroth-, first- and second-order moments of f_{UN} match those of f_{input} .

For the binary mixed model (1), expectation propagation requires repeated projection of the form

$$f_{\text{input}}(\mathbf{x}) = F(c_0 + \mathbf{c}_1^T \mathbf{x}) \exp \left\{ \left[\begin{array}{c} \mathbf{x} \\ \text{vech}(\mathbf{x}\mathbf{x}^T) \end{array} \right]^T \left[\begin{array}{c} \boldsymbol{\eta}_1^{\text{input}} \\ \boldsymbol{\eta}_2^{\text{input}} \end{array} \right] \right\}$$

onto the unnormalized Multivariate Normal family. An important observation is that case of probit mixed models, $\text{proj}[f_{\text{input}}](\mathbf{x})$ has an exact solution.

Let $\zeta(x) \equiv \log\{2\Phi(x)\}$. It follows that

$$\zeta'(x) = \phi(x)/\Phi(x) \quad \text{and} \quad \zeta''(x) = -\zeta'(x)\{x + \zeta'(x)\}$$

where $\phi(x) \equiv (2\pi)^{-1/2} \exp(-\frac{1}{2}x^2)$ is the $N(0, 1)$ density function. We are now in a position to define two algebraic functions which are fundamental for approximate likelihood-based inference in probit mixed models based on expectation propagation:

Definition 1. For primary arguments \mathbf{a}_1 ($d \times 1$) and \mathbf{a}_2 ($\frac{1}{2}d(d+1) \times 1$) such that $\text{vec}^{-1}(-\mathbf{D}_d^{+T} \mathbf{a}_2)$ is symmetric and positive definite, and auxiliary arguments $c_0 \in \mathbb{R}$ and \mathbf{c}_1 ($d \times 1$) the function K_{probit} is given by

$$K_{\text{probit}} \left(\left[\begin{array}{c} \mathbf{a}_1 \\ \mathbf{a}_2 \end{array} \right]; c_0, \mathbf{c}_1 \right) \equiv \left[\begin{array}{c} \mathbf{R}_5^T (\mathbf{a}_1 + r_3 \mathbf{c}_1) \\ \mathbf{D}_d^T \text{vec}(\mathbf{R}_5^T \mathbf{A}_2) \end{array} \right]$$

with

$$\begin{aligned} \mathbf{A}_2 &\equiv \text{vec}^{-1}(\mathbf{D}_d^{+T} \mathbf{a}_2), \quad r_1 \equiv \sqrt{2(2 - \mathbf{c}_1^T \mathbf{A}_2^{-1} \mathbf{c}_1)}, \quad r_2 \equiv (2c_0 - \mathbf{c}_1^T \mathbf{A}_2^{-1} \mathbf{a}_1)/r_1, \\ r_3 &\equiv 2\zeta'(r_2)/r_1, \quad r_4 \equiv -2\zeta''(r_2)/r_1^2 \quad \text{and} \quad \mathbf{R}_5 \equiv (\mathbf{A}_2 + r_4 \mathbf{c}_1 \mathbf{c}_1^T)^{-1} \mathbf{A}_2 \end{aligned}$$

and the function A_N is given by

$$A_N \left(\left[\begin{array}{c} \mathbf{a}_1 \\ \mathbf{a}_2 \end{array} \right] \right) \equiv -\frac{1}{4} \mathbf{a}_1^T \mathbf{A}_2^{-1} \mathbf{a}_1 - \frac{1}{2} \log | -2\mathbf{A}_2 |.$$

In addition, for primary arguments $\mathbf{a}_1, \mathbf{b}_1$ ($d \times 1$) and $\mathbf{a}_2, \mathbf{b}_2$ ($\frac{1}{2}d(d+1) \times 1$) such that both $\text{vec}^{-1}(-\mathbf{D}_d^{+T} \mathbf{a}_2)$ and $\text{vec}^{-1}(-\mathbf{D}_d^{+T} \mathbf{b}_2)$ are symmetric and positive definite, and auxiliary arguments $c_0 \in \mathbb{R}$ and \mathbf{c}_1 ($d \times 1$), the function C_{probit} is given by

$$C_{\text{probit}} \left(\left[\begin{array}{c} \mathbf{a}_1 \\ \mathbf{a}_2 \end{array} \right], \left[\begin{array}{c} \mathbf{b}_1 \\ \mathbf{b}_2 \end{array} \right]; c_0, \mathbf{c}_1 \right) \equiv \log \Phi(r_2) + \frac{1}{4} \mathbf{b}_1^T \mathbf{B}_2^{-1} \mathbf{b}_1 - \frac{1}{4} \mathbf{a}_1^T \mathbf{A}_2^{-1} \mathbf{a}_1 + \frac{1}{2} \log \{ |\mathbf{B}_2| / |\mathbf{A}_2| \}$$

with $\mathbf{B}_2 \equiv \text{vec}^{-1}(\mathbf{D}_d^{+T} \mathbf{b}_2)$.

Inspection of Definition 1 reveals that the K_{probit} and C_{probit} functions are simple functions up to evaluations of $\log(\Phi)$ and $\zeta' = \phi/\Phi$. Even though software for Φ is widely available, direct computation of $\log(\Phi)$ and ζ' can be unstable and software such as the function `zeta()` in the R package `sn` (Azzalini, 2017) is recommended. Another option is use of continued fraction representation and Lentz's Algorithm (e.g. Wand & Ormerod, 2012).

Expectation propagation for probit mixed models relies heavily upon:

Theorem 1. If

$$f_{\text{input}}(\mathbf{x}) = \Phi(c_0 + \mathbf{c}_1^T \mathbf{x}) \exp \left\{ \left[\begin{array}{c} \mathbf{x} \\ \text{vech}(\mathbf{x}\mathbf{x}^T) \end{array} \right]^T \left[\begin{array}{c} \boldsymbol{\eta}_1^{\text{input}} \\ \boldsymbol{\eta}_2^{\text{input}} \end{array} \right] \right\}$$

then

$$\text{proj}[f_{\text{input}}](\mathbf{x}) = \exp \left\{ \left[\begin{array}{c} 1 \\ \mathbf{x} \\ \text{vech}(\mathbf{x}\mathbf{x}^T) \end{array} \right]^T \left[\begin{array}{c} \eta_0^* \\ \eta_1^* \\ \eta_2^* \end{array} \right] \right\}$$

where

$$\left[\begin{array}{c} \eta_1^* \\ \eta_2^* \end{array} \right] = K_{\text{probit}} \left(\left[\begin{array}{c} \eta_1^{\text{input}} \\ \eta_2^{\text{input}} \end{array} \right]; c_0, \mathbf{c}_1 \right) \quad \text{and} \quad \eta_0^* = C_{\text{probit}} \left(\left[\begin{array}{c} \eta_1^{\text{input}} \\ \eta_2^{\text{input}} \end{array} \right], \left[\begin{array}{c} \eta_1^* \\ \eta_2^* \end{array} \right]; c_0, \mathbf{c}_1 \right).$$

A proof of Theorem 1 is given in Section S.1 of the online supplement.

3.2 Message Passing Formulation

The i th summand of $\ell(\boldsymbol{\beta}, \boldsymbol{\Sigma})$ can be written as

$$\ell_i(\boldsymbol{\beta}, \boldsymbol{\Sigma}) = \log \int_{\mathbb{R}^{d_{\mathbf{R}}}} \left\{ \prod_{j=1}^{n_i} p(y_{ij} | \mathbf{u}_i; \boldsymbol{\beta}) \right\} p(\mathbf{u}_i; \boldsymbol{\Sigma}) d\mathbf{u}_i \quad (6)$$

where, for $1 \leq j \leq n_i$,

$$p(y_{ij} | \mathbf{u}_i; \boldsymbol{\beta}) \equiv F((2y_{ij} - 1)(\boldsymbol{\beta}^T \mathbf{x}_{ij}^F + \mathbf{u}_i^T \mathbf{x}_{ij}^R)) \quad \text{and} \quad p(\mathbf{u}_i; \boldsymbol{\Sigma}) \equiv |2\pi\boldsymbol{\Sigma}|^{-1/2} \exp(-\frac{1}{2}\mathbf{u}_i^T \boldsymbol{\Sigma}^{-1} \mathbf{u}_i)$$

are, respectively, the conditional density functions of each response given its random effect and the density function of that random effect. Note that product structure of the integrand in (6) can be represented using *factor graph* shown in Figure 1. The circle in Figure 1 corresponds to the random vector \mathbf{u}_i and factor graph parlance is a *stochastic variable node*. The solid rectangles correspond to each of the $n_i + 1$ factors in the (6) integrand. Each of these factors depend on \mathbf{u}_i , which is signified by an edges connecting each factor node to the lone stochastic variable node.




Figure 1: Factor graph representation of the product structure of the integrand in (6). The open circle corresponds to the random effect vector \mathbf{u}_i and the solid rectangles indicate factors. Edges indicate dependence of each factor on \mathbf{u}_i .

Expectation propagation approximation of $\ell_i(\boldsymbol{\beta}, \boldsymbol{\Sigma})$ involves projection onto the unnormalized Multivariate Normal family. Suppose that

$$\tilde{p}(y_{ij} | \mathbf{u}_i; \boldsymbol{\beta}) = \exp \left\{ \left[\begin{array}{c} 1 \\ \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{array} \right]^T \boldsymbol{\eta}_{ij} \right\}, \quad 1 \leq j \leq n_i \quad (7)$$

are initialized to be unnormalized Multivariate Normal density functions in \mathbf{u}_i . Then, for each $j = 1, \dots, n_i$, the $\boldsymbol{\eta}_{ij}$ update involves minimization of

$$\text{KL} \left(p(y_{ij} | \mathbf{u}_i; \boldsymbol{\beta}) \left\{ \prod_{j' \neq j} \tilde{p}(y_{ij'} | \mathbf{u}_i; \boldsymbol{\beta}) \right\} p(\mathbf{u}_i; \boldsymbol{\Sigma}) \parallel \left\{ \prod_{j'=1}^{n_i} \tilde{p}(y_{ij'} | \mathbf{u}_i; \boldsymbol{\beta}) \right\} p(\mathbf{u}_i; \boldsymbol{\Sigma}) \right) \quad (8)$$

as functions of \mathbf{u}_i . Noting that this problem has the form (5), Theorem 1 can be used to perform the update explicitly in the case of a probit link. This procedure is then iterated until the $\boldsymbol{\eta}_{ij}$ s converge.

A convenient way to keep track of the updates and compartmentalize the algebra and coding is to call upon the notion of *message passing*. Minka (2005) shows how to express expectation propagation as a message passing algorithm in the Bayesian graphical models context, culminating in his equation (54) and (83) update formulae. Exactly the same formulae arise here, as is made clear in Section S.2 of the online supplement. In particular, in keeping with (83) of Minka (2005), (8) can be expressed as

$$m_{p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) \leftarrow \frac{\text{proj}[m_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta})}(\mathbf{u}_i) p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta})](\mathbf{u}_i)}{m_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta})}(\mathbf{u}_i)}, \quad 1 \leq j \leq n_i, \quad (9)$$

where $m_{p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i)$ is the *message* passed from the factor $p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta})$ to the stochastic node \mathbf{u}_i and $m_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta})}(\mathbf{u}_i)$ is the message passed from \mathbf{u}_i back to $p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta})$. The message passed from $p(\mathbf{u}_i;\boldsymbol{\Sigma})$ to \mathbf{u}_i is

$$m_{p(\mathbf{u}_i;\boldsymbol{\Sigma}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) \leftarrow \frac{\text{proj}[m_{\mathbf{u}_i \rightarrow p(\mathbf{u}_i;\boldsymbol{\Sigma})}(\mathbf{u}_i) p(\mathbf{u}_i;\boldsymbol{\Sigma})](\mathbf{u}_i)}{m_{\mathbf{u}_i \rightarrow p(\mathbf{u}_i;\boldsymbol{\Sigma})}(\mathbf{u}_i)}. \quad (10)$$

In keeping with equation (54) of Minka (2005), the stochastic node to factor messages are updated according to

$$m_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta})}(\mathbf{u}_i) = m_{p(\mathbf{u}_i;\boldsymbol{\Sigma}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) \left\{ \prod_{j' \neq j}^{n_i} m_{p(y_{ij'}|\mathbf{u}_i;\boldsymbol{\beta}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) \right\}, \quad 1 \leq j \leq n_i, \quad (11)$$

and

$$m_{\mathbf{u}_i \rightarrow p(\mathbf{u}_i;\boldsymbol{\Sigma})}(\mathbf{u}_i) = \prod_{j=1}^{n_i} m_{p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i). \quad (12)$$

As laid out at the end of Section 6 of Minka (2005), the expectation message passing protocol is:

Initialize all factor to stochastic node messages.

Cycle until all factor to stochastic node messages converge:

For each factor:

 Compute the messages passed to the factor using (11) or (12).

 Compute the messages passed from the factor using (9) or (10).

Upon convergence, the expectation propagation approximation to $\ell_i(\boldsymbol{\beta}, \boldsymbol{\Sigma})$ is

$$\underline{\ell}_i(\boldsymbol{\beta}, \boldsymbol{\Sigma}) = \log \int_{\mathbb{R}^{d^R}} \left\{ \prod_{j=1}^{n_i} m_{p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) \right\} m_{p(\mathbf{u}_i;\boldsymbol{\Sigma}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) d\mathbf{u}_i. \quad (13)$$

where the integrand is in keeping with the general form given by (44) of Minka & Winn (2008). The success of expectation propagation hinges on the fact that each of the messages in (13) is an unnormalized Multivariate Normal density function and the integral over \mathbb{R}^{d^R}

can be obtained exactly as follows:

$$\begin{aligned}
& \int_{\mathbb{R}^{d^R}} \left\{ \prod_{j=1}^{n_i} m_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) \right\} m_{p(\mathbf{u}_i; \boldsymbol{\Sigma}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) d\mathbf{u}_i \\
&= \int_{\mathbb{R}^{d^R}} \left[\prod_{j=1}^{n_i} \exp \left\{ \begin{bmatrix} 1 \\ \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{bmatrix}^T \eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i} \right\} \right] \\
&\quad \times \exp \left\{ \begin{bmatrix} 1 \\ \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{bmatrix}^T \eta_{p(\mathbf{u}_i; \boldsymbol{\Sigma}) \rightarrow \mathbf{u}_i} \right\} d\mathbf{u}_i \\
&= (2\pi)^{-1/2} \exp \left\{ \left(\eta_{\boldsymbol{\Sigma}} + \text{SUM}\{\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i}\} \right)_0 \right. \\
&\quad \left. + A_N \left(\left(\eta_{\boldsymbol{\Sigma}} + \text{SUM}\{\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i}\} \right)_{-0} \right) \right\}
\end{aligned}$$

where

$$\eta_{\boldsymbol{\Sigma}} \equiv \begin{bmatrix} -\frac{1}{2} \log |2\pi \boldsymbol{\Sigma}| \\ \mathbf{0}_{d^R} \\ -\frac{1}{2} \mathbf{D}_{d^R}^T \text{vec}(\boldsymbol{\Sigma}^{-1}) \end{bmatrix}, \quad \text{SUM}\{\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i}\} \equiv \sum_{j=1}^{n_i} \eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i},$$

A_N is as defined in Definition 1 and, for an unnormalized Multivariate Normal natural parameter vector $\boldsymbol{\eta}$, η_0 denotes the first entry (the zero subscript is indicative of the first entry being the coefficient of 1) and $\boldsymbol{\eta}_{-0}$ denotes the remaining entries.

The full algorithm for expectation propagation approximation of $\ell(\boldsymbol{\beta}, \boldsymbol{\Sigma})$ is summarized as Algorithm 1. The derivational details are given in Section S.2.

We have carried out extensive simulated data tests on Algorithm 1 using the starting values described in Section 3.3 and found convergence to be rapid. Moreover, each of updates in Algorithm 1 involve explicit calculations and low-level language implementation, used in our R package `glmmEP`, affords very fast evaluation of the approximate log-likelihood surface. As explained in (3.4), quasi-Newton methods can be used for maximization of $\ell(\boldsymbol{\beta}, \boldsymbol{\Sigma})$ and approximate likelihood-based inference.

3.3 Recommended Starting Values for Algorithm 1

In Section S.3 we use a Taylor series argument to justify the following starting values for $\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i}$ in Algorithm 1:

$$\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i}^{\text{start}} \equiv \begin{bmatrix} 0 \\ (2y_{ij} - 1)\zeta'(\hat{a}_{ij})\mathbf{x}_{ij}^R - \zeta''(\hat{a}_{ij})\mathbf{x}_{ij}^R(\mathbf{x}_{ij}^R)^T \hat{\mathbf{u}}_i \\ \frac{1}{2}\zeta''(\hat{a}_{ij})\mathbf{D}_{d^R}^T \text{vec}(\mathbf{x}_{ij}^R(\mathbf{x}_{ij}^R)^T) \end{bmatrix}, \quad 1 \leq j \leq n_i, \quad (14)$$

where

$$\hat{a}_{ij} \equiv (2y_{ij} - 1)(\boldsymbol{\beta}^T \mathbf{x}_{ij}^F + \hat{\mathbf{u}}_i^T \mathbf{x}_{ij}^R)$$

and $\hat{\mathbf{u}}_i$ is a prediction of \mathbf{u}_i . A convenient choice for $\hat{\mathbf{u}}_i$ is that based on Laplace approximation. In the R computing environment the function `glmer()` in the package `lme4` (Bates *et al.*, 2015) provides fast Laplace approximation-based predictions for the \mathbf{u}_i . In our numerical experiments, we found convergence of the cycle loop of Algorithm 1 to be quite rapid, with convergents of

$$\left(\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i} \right)_{-0} \quad \text{relatively close to} \quad \left(\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i}^{\text{start}} \right)_{-0}.$$

Inputs: $y_{ij}, \mathbf{x}_{ij}^F, \mathbf{x}_{ij}^R, 1 \leq i \leq m, 1 \leq j \leq n_i$;

$\boldsymbol{\beta} (d^F \times 1), \boldsymbol{\Sigma} (d^R \times d^R, \text{symmetric and positive definite}).$

Set constants: $c_{0,ij} \leftarrow (2y_{ij} - 1)(\boldsymbol{\beta}^T \mathbf{x}_{ij}^F); \mathbf{c}_{1,ij} \leftarrow (2y_{ij} - 1)\mathbf{x}_{ij}^R, \quad 1 \leq i \leq m, 1 \leq j \leq n_i;$

$$\eta_{p(\mathbf{u}_i; \boldsymbol{\Sigma})} \rightarrow \mathbf{u}_i \leftarrow \eta_{\boldsymbol{\Sigma}} \equiv \begin{bmatrix} -\frac{1}{2} \log |2\pi \boldsymbol{\Sigma}| \\ \mathbf{0}_{d^R} \\ -\frac{1}{2} \mathbf{D}_{d^R}^T \text{vec}(\boldsymbol{\Sigma}^{-1}) \end{bmatrix}, \quad 1 \leq i \leq m.$$

For $i = 1, \dots, m$:

Initialize: $\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})} \rightarrow \mathbf{u}_i, \quad 1 \leq j \leq n_i$ (see Section 3.3 for a recommendation)

Cycle:

$$\text{SUM}\{\eta_{p(\mathbf{y}_i|\mathbf{u}_i; \boldsymbol{\beta})} \rightarrow \mathbf{u}_i\} \leftarrow \sum_{j=1}^{n_i} \eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})} \rightarrow \mathbf{u}_i$$

For $j = 1, \dots, n_i$:

$$\begin{aligned} \eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})} &\leftarrow \eta_{p(\mathbf{u}_i; \boldsymbol{\Sigma})} \rightarrow \mathbf{u}_i + \text{SUM}\{\eta_{p(\mathbf{y}_i|\mathbf{u}_i; \boldsymbol{\beta})} \rightarrow \mathbf{u}_i\} - \eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})} \rightarrow \mathbf{u}_i \\ \left(\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})} \rightarrow \mathbf{u}_i\right)_{-0} &\leftarrow K_{\text{probit}}\left(\left(\eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})}\right)_{-0}; c_{0,ij}, \mathbf{c}_{1,ij}\right) \\ &\quad - \left(\eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})}\right)_{-0} \end{aligned}$$

until all natural parameter vectors converge.

For $j = 1, \dots, n_i$:

$$\left(\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})} \rightarrow \mathbf{u}_i\right)_0 \leftarrow C_{\text{probit}}\left(\left(\eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})}\right)_{-0}, \left(\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})} \rightarrow \mathbf{u}_i\right)_{-0} + \left(\eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})}\right)_{-0}; c_{0,ij}, \mathbf{c}_{1,ij}\right)$$

$$\text{SUM}\{\eta_{p(\mathbf{y}_i|\mathbf{u}_i; \boldsymbol{\beta})} \rightarrow \mathbf{u}_i\} \leftarrow \sum_{j=1}^{n_i} \eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})} \rightarrow \mathbf{u}_i$$

Output: The expectation propagation approximate log-likelihood given by

$$\begin{aligned} \underline{\ell}(\boldsymbol{\beta}, \boldsymbol{\Sigma}) &= \frac{1}{2} m \log(2\pi) + \sum_{i=1}^m \left\{ \left(\eta_{\boldsymbol{\Sigma}} + \text{SUM}\{\eta_{p(\mathbf{y}_i|\mathbf{u}_i; \boldsymbol{\beta})} \rightarrow \mathbf{u}_i\} \right)_0 \right. \\ &\quad \left. + A_N \left(\left(\eta_{\boldsymbol{\Sigma}} + \text{SUM}\{\eta_{p(\mathbf{y}_i|\mathbf{u}_i; \boldsymbol{\beta})} \rightarrow \mathbf{u}_i\} \right)_{-0} \right) \right\} \end{aligned}$$

Algorithm 1: Expectation expectation approximation of the log-likelihood for the probit mixed model (1) with $F = \Phi$ via message passing on the Figure 1 factor graph.

Therefore, we strongly recommend the starting values (14).

3.4 Quasi-Newton Optimization and Approximate Inference

Even though Algorithm 1 provides fast approximate evaluation of the probit mixed model likelihood surface, we still need to maximize over $(\boldsymbol{\beta}, \boldsymbol{\Sigma})$ to obtain the expectation propagation-

approximate maximum likelihood estimators $(\hat{\beta}, \hat{\Sigma})$. This is also the issue of approximate inference based on Fisher information theory.

Since $\ell(\beta, \Sigma)$ is defined implicitly via an iterative scheme, differentiation for use in derivative-based optimization techniques is not straightforward. A practical workaround involves the employment of optimization methods such as those of the quasi-Newton variety for which derivatives are approximated numerically. In the R computing environment the function `optim()` supports several derivative-free optimization implementations. The Matlab computing environment (The Mathworks Incorporated, 2018) has similar capabilities via functions such as `fminunc()`. In the `glmmEP` package and the examples in Section 4 we use the Broyden-Fletcher-Goldfarb-Shanno quasi-Newton method (Broyden, 1970; Fletcher 1970; Goldfarb, 1970; Shanno, 1970) with Nelder-Mead starting values. Section 2.2.2.3 of Givens & Hoetig (2005) provides a concise summary of the Broyden-Fletcher-Goldfarb-Shanno method.

Since Σ is constrained to be symmetric and positive definite, we instead perform quasi-Newton optimization over the unconstrained parameter vector (β, θ) where

$$\theta \equiv \text{vech}\left(\frac{1}{2} \log(\Sigma)\right)$$

and $\log(\Sigma)$ is the matrix logarithm of Σ (e.g. Section 2.2 of Pinheiro & Bates, 2000). Note that $\log(\Sigma)$ can be obtained using

$$\log(\Sigma) = U_{\Sigma} \text{diag}\{\log(\lambda_{\Sigma})\} U_{\Sigma}^T \quad \text{where} \quad \Sigma = U_{\Sigma} \text{diag}(\lambda_{\Sigma}) U_{\Sigma}^T$$

is the spectral decomposition of Σ and $\log(\lambda_{\Sigma})$ denotes element-wise evaluation of the logarithm to the entries of λ_{Σ} . If $(\hat{\beta}, \hat{\theta})$ is the maximizer of ℓ then the expectation propagation approximate maximum likelihood estimate of Σ is

$$\hat{\Sigma} = U_{\hat{\theta}} \text{diag}\{\exp(2\lambda_{\hat{\theta}})\} U_{\hat{\theta}}^T \quad \text{where} \quad \text{vech}^{-1}(\hat{\theta}) = U_{\hat{\theta}} \text{diag}(\lambda_{\hat{\theta}}) U_{\hat{\theta}}^T$$

is the spectral decomposition of the $\text{vech}^{-1}(\hat{\theta})$. Note that $\text{vech}^{-1}(a)$ is the symmetric matrix A of appropriate dimension such that $\text{vech}(A) = a$.

The `optim()` function in R and the `fminunc()` function in Matlab each have the option of computing an approximation to the Hessian matrix at the optimum, which can be used for approximate likelihood-based inference. In particular, we can use the approximate Hessian matrix to construct confidence intervals for the entries of β and the standard deviation and correlation parameters of Σ . The full details are given in Section S.4 of the online supplement. Here we sketch the idea for the special case of $d^R = 2$, for which

$$\Sigma = \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix}.$$

For confidence interval construction it is appropriate (e.g. Section 2.4 of Pinheiro & Bates) to work with the parameter vector

$$\omega \equiv \begin{bmatrix} \log(\sigma_1) \\ \log(\sigma_2) \\ \tanh^{-1}(\rho) \end{bmatrix}.$$

Approximate $100(1 - \alpha)\%$ confidence intervals for the entries of $(\beta, \omega)^T$ are

$$\begin{bmatrix} \hat{\beta} \\ \hat{\omega} \end{bmatrix} \pm \Phi^{-1}\left(1 - \frac{1}{2}\alpha\right) \sqrt{-\text{diagonal}\left(\{H_{\ell}(\hat{\beta}, \hat{\omega})\}^{-1}\right)} \quad (15)$$

where $H_{\ell}(\beta, \omega)$ is the Hessian matrix of ℓ with respect to the (β, ω) parameter vector. Confidence intervals for the entries of β , σ_1 , σ_2 and ρ follow from standard inversion manipulations.

Note that (β, θ) is an unconstrained parametrization whilst (β, ω) is a constrained parametrization. Hence, the optimization should be performed with respect to the former parametrization whereas the Hessian matrix in (15) is respect to the latter parametrization. In the examples of Section 4 and the R package `glimmEP` we use the following strategy:

- Obtain $(\hat{\beta}, \hat{\theta})$ using `optim()` with the (β, θ) parametrization in the function being maximized and the `hessian` argument set to `FALSE`.
- Compute $(\hat{\beta}, \hat{\omega})$ and use this as a initial value with a call to `optim()` with the (β, ω) parametrization in the function being maximized and the `hessian` argument set to `TRUE`.

Full details of confidence interval calculations for the general multivariate random effects situation are given in Section S.4 of the online supplement.

In our numerical experiments, we have found Nelder-Mead followed by Broyden-Fletcher-Goldfarb-Shanno optimization of expectation propagation approximate log-likelihood, with confidence intervals based on the approximate Hessian matrix, to be very effective. In Section 4 we present simulation results that show this strategy producing fast and accurate inference for binary mixed models.

3.5 Expectation Propagation Approximate Best Prediction

The best predictors of \mathbf{u}_i are

$$\text{BP}(\mathbf{u}_i) \equiv E(\mathbf{u}_i | \mathbf{y}), \quad 1 \leq i \leq m.$$

We now show that Algorithm 1 provides, as by-products, straightforward empirical best predictions of the \mathbf{u}_i .

Let

$$\hat{\boldsymbol{\eta}}_i \equiv \eta_{\Sigma} + \text{SUM}\{\eta_{p(\mathbf{y}_i | \mathbf{u}_i; \beta)} \rightarrow \mathbf{u}_i\} = \begin{bmatrix} \hat{\boldsymbol{\eta}}_{i1} \\ \hat{\boldsymbol{\eta}}_{i2} \end{bmatrix} \quad (16)$$

where η_{Σ} and $\text{SUM}\{\eta_{p(\mathbf{y}_i | \mathbf{u}_i; \beta)} \rightarrow \mathbf{u}_i\}$ are as in Algorithm 1 with $(\beta, \Sigma) = (\hat{\beta}, \hat{\Sigma})$, $\hat{\boldsymbol{\eta}}_{i1}$ is the sub-vector of $\hat{\boldsymbol{\eta}}_i$ corresponding to the first d^R entries and $\hat{\boldsymbol{\eta}}_{i2}$ contains the remaining entries. Then in Section S.5 of the online supplement we show that a suitable empirical approximation to $\text{BP}(\mathbf{u}_i)$, based on the expectation propagation estimate, is

$$\underline{\text{BP}}(\mathbf{u}_i) = -\frac{1}{2} \left\{ \text{vec}^{-1} \left(\mathbf{D}_d^{+T} \hat{\boldsymbol{\eta}}_{i2} \right) \right\}^{-1} \hat{\boldsymbol{\eta}}_{i1}. \quad (17)$$

The corresponding covariance matrix empirical approximation is

$$\underline{\text{Cov}}(\mathbf{u}_i | \mathbf{y}) = -\frac{1}{2} \left\{ \text{vec}^{-1} \left(\mathbf{D}_d^{+T} \hat{\boldsymbol{\eta}}_{i2} \right) \right\}^{-1}. \quad (18)$$

In view of equation (13.7) of McCulloch, Searle & Neuhaus (2008), $\text{Cov}\{\underline{\text{BP}}(\mathbf{u}_i) - \mathbf{u}_i\}$ is approximated by $E_{\mathbf{y}_i}\{\underline{\text{Cov}}(\mathbf{u}_i | \mathbf{y}_i)\}$. Approximate prediction interval construction is hindered by this expectation over the sampling distribution of the responses. See, for example, Carlin & Gelfand (1991), for discussion and access to some of the relevant literature concerning valid prediction interval construction in the more general empirical Bayes context.

4 Numerical Evaluation and Illustration

We now demonstrate the impressive accuracy and speed of Algorithm 1 combined with quasi-Newton methods for approximate likelihood-based inference for probit mixed models. Firstly, we report the results of some studies involving simulated data. Analysis of actual data is discussed later in this section.

4.1 Simulations

Our simulations involved (1) comparison with exact maximum likelihood for the $d^R = 1$ situation for which quadrature is univariate, and (2) evaluation of inferential accuracy and speed for a larger model involving bivariate random effects.

4.1.1 Comparison with Exact Maximum Likelihood for Univariate Random Effects

Our first simulation study involved simulation of 1,000 datasets according to the $d^R = 1$ version of (1) with true parameter values:

$$\boldsymbol{\beta}_{\text{true}} = [0, 1]^T \quad \text{and} \quad \boldsymbol{\Sigma}_{\text{true}} = \sigma_{\text{true}}^2 = 1. \quad (19)$$

The sample sizes were set to $m = 100$ and $n_i = 2$. The \boldsymbol{x}_{ij}^F and \boldsymbol{x}_{ij}^R vectors were of the form

$$\boldsymbol{x}_{ij}^F = [1, x_{ij}]^T \quad \text{and} \quad \boldsymbol{x}_{ij}^R = 1 \quad (20)$$

where x_{ij} was generated independently from a Uniform distribution on the unit interval.

For each simulated dataset, the probit mixed model defined by (20) was fit using each of the following approaches:

- (1) Exact maximum likelihood with adaptive Gauss-Hermite quadrature used for the univariate intractable integrals. This was achieved using the function `glmer()` in the R package `lme4` (Bates *et al.*, 2015). The number of points for evaluation of the adaptive Gauss-Hermite approximation was fixed at 100.
- (2) The Laplace approximation used by `glmer()`.
- (3) Expectation propagation as described in Section 3.

Of interest is comparison of quadrature-free approximations (2) and (3) against the exact maximum likelihood benchmark. Figure 2 contrasts the point estimates and confidence intervals produced by Laplace approximation and expectation propagation against those produced by exact maximum likelihood. The first row of Figure 2 shows that Laplace approximation results in shoddy statistical inference, with the empirical coverage values falling well below the advertized 95% level. The gray line segments for exact likelihood confidence intervals and black line segments for their Laplace approximations have very noticeable discrepancies. In the second row of Figure 2 we repeat the empirical coverage percentages and gray line segments for exact likelihood inference and, instead, compare these results with those produced by expectation propagation. For the fixed effects, β_0 and β_1 , the empirical coverage of expectation propagation is seen to be very close to 95%. For the standard deviation parameter, σ , expectation propagation delivers slightly more coverage than advertized (97.5% versus 95%). However, the relatively low sample sizes in this study should be kept in mind. The simulation study in the next subsection uses higher sample sizes and expectation propagation is seen to be particularly accurate in terms of confidence interval coverage.

4.1.2 Accuracy and Speed Assessment for Bivariate Random Effects

In this study we simulated 1,000 datasets according to a $d^R = 2$ version of (1) with true parameter values:

$$\boldsymbol{\beta}_{\text{true}} = [0.37, 0.93, -0.46, 0.08, -1.34, 1.09]^T \quad \text{and} \quad \boldsymbol{\Sigma}_{\text{true}} = \begin{bmatrix} 0.53 & -0.36 \\ -0.36 & 0.92 \end{bmatrix}. \quad (21)$$

The number of groups was fixed at $m = 250$ and each n_i value selected randomly from a discrete Uniform distribution on $\{20, 21, \dots, 30\}$. The \boldsymbol{x}_{ij}^F and \boldsymbol{x}_{ij}^R vectors were of the form

$$\boldsymbol{x}_{ij}^F = [1, x_{1,ij}, x_{2,ij}, x_{3,ij}, x_{4,ij}, x_{5,ij}]^T \quad \text{and} \quad \boldsymbol{x}_{ij}^R = [1, x_{1,ij}]^T$$

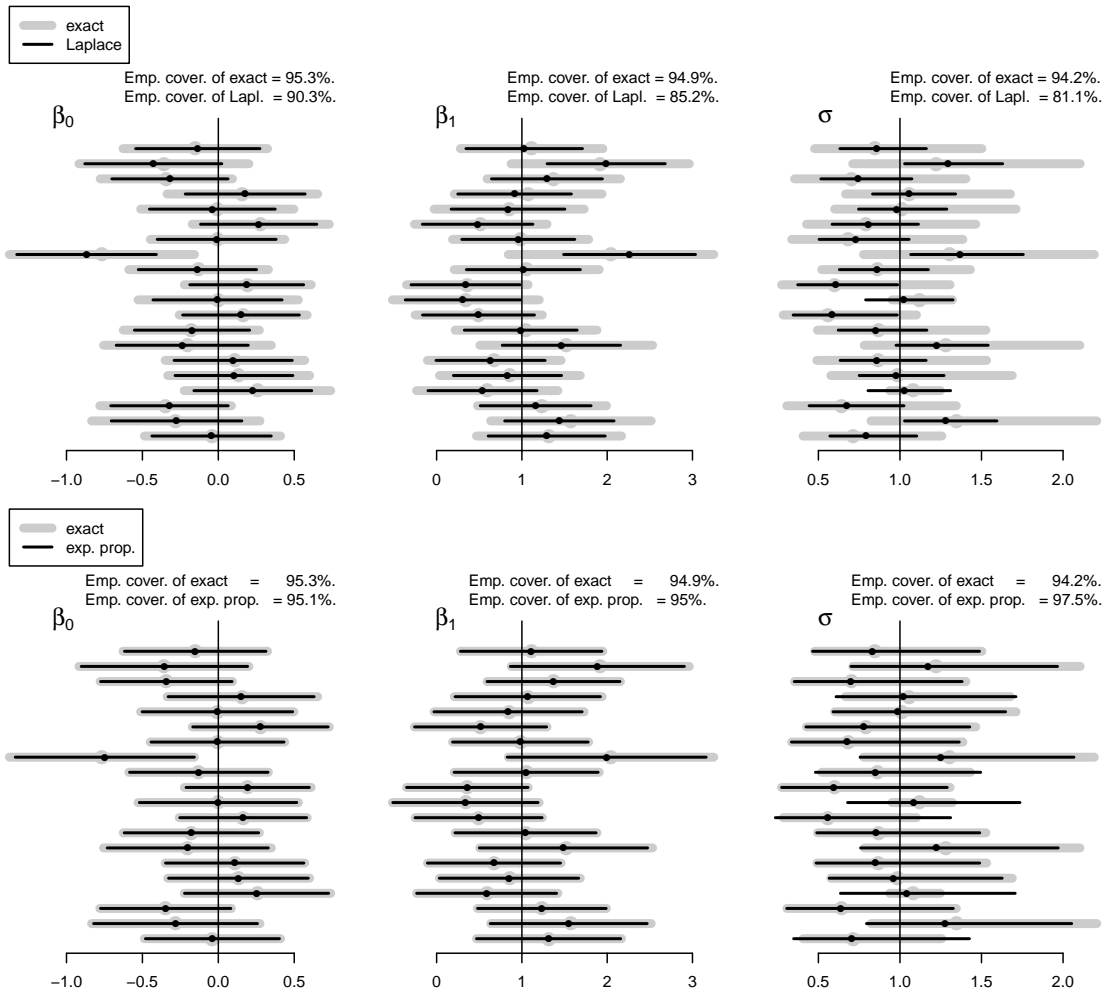


Figure 2: Comparison of point estimation and 95% confidence interval coverage for the first simulation study with true parameter values given by (19). The upper row of panels compares exact maximum likelihood with Laplace approximation. The low row of panels compares exact maximum likelihood with expectation propagation approximation. The horizontal lines indicate expectation propagation-based confidence intervals for 20 randomly chosen replications of the simulation study described in the text. The points indicate the corresponding approximate maximum likelihood estimates. The vertical lines indicate true parameter values. The percentages displayed at the top of each panel are empirical coverages over all 1,000 replications for each method involved in the comparison.

where each $x_{k,ij}$ was generated independently from a Uniform distribution on the unit interval. All relative tolerance values were set to 10^{-5} and the maximum number of iteration values were set to 100, which is relevant for the upcoming speed assessment.

The points and horizontal line segments in Figure 3 are displays of estimates and corresponding 95% confidence intervals for each of the interpretable model parameters, for 50 randomly chosen replications. The numbers in the top right-hand corner of each panel are the empirical coverage values based on all 1,000 replications. For all nine parameters, the empirical coverage values are in keeping with the advertized coverage of 95%, and is an indication of excellent accuracy for this setting.

Despite the higher samples and complexity of the model, we have gotten the fitting times down to tens of seconds in the `gmmEP` package within the R computing environment. This has been achieved by implementation of Algorithm 1 in a low level language so that approximate likelihood evaluations are very rapid. The computing speed depends upon various relative tolerance values and upper bounds on numbers of iterations for the

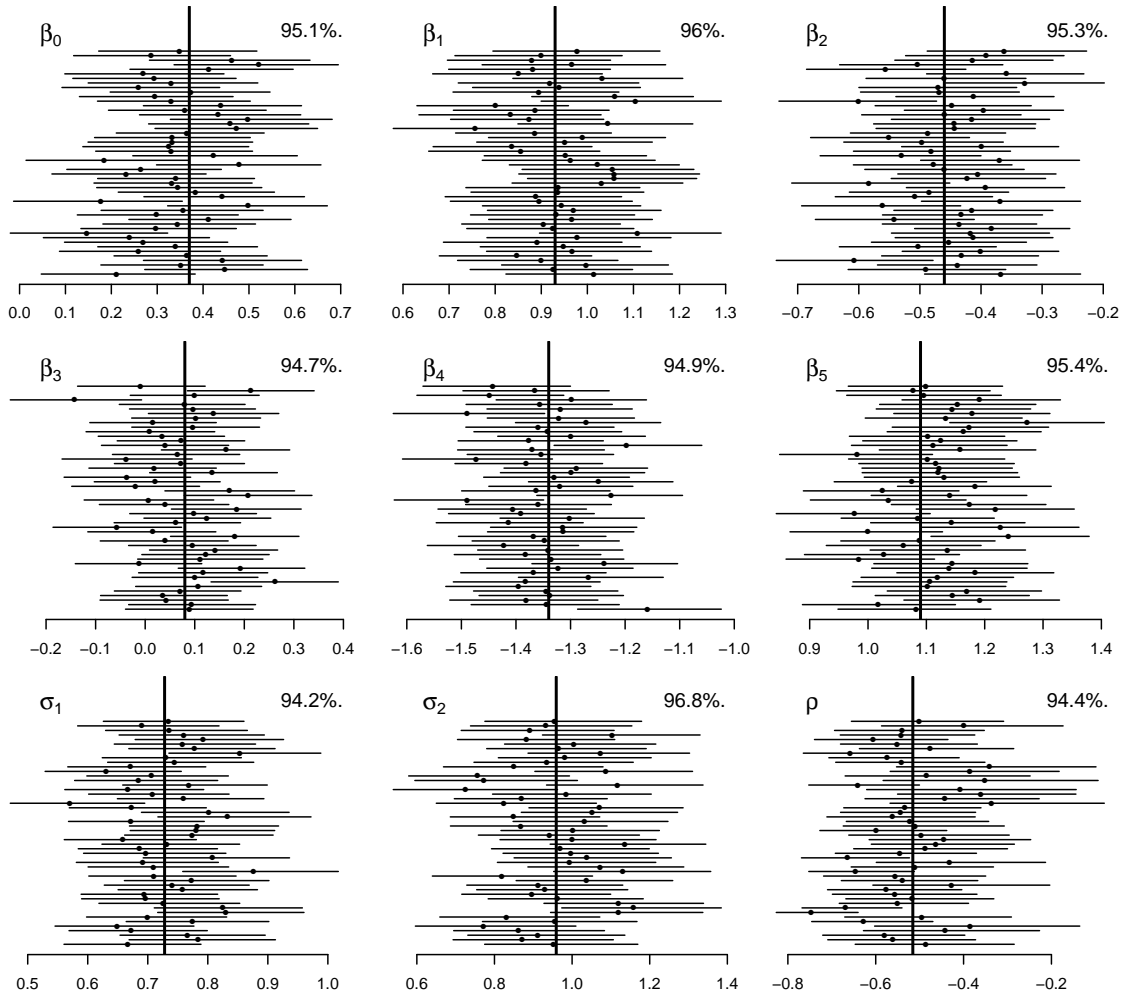


Figure 3: Summary of confidence interval coverage for the second simulation study with true parameter values given by (21). The horizontal lines indicate expectation propagation-based confidence intervals for 50 randomly chosen replications of the simulation study described in the text. The solid circular points indicate the corresponding point estimates. The vertical lines indicate true parameter values. The percentage in the top right-hand corner of each panel is the empirical coverage over all 1,000 replications.

various iterative schemes as well as attributes of the computer. This simulation study was run on a MacBook Air laptop with 8 gigabytes of random access memory and a 2.2 gigahertz processor. The convergence stopping criteria values are given earlier in this section. Over the 1,000 replications the median computing time was 18 seconds, the upper quartile was 20 seconds and the maximum was 34 seconds. Such speed is impressive given that each data set contained tens of thousands of observations and bivariate random effects are accurately handled.

4.2 Application to Data from a Fertility Study

Data from a 1988 Bangladesh fertility study are stored in the data frame `Contraception` within the R package `mlmRev` (Bates, Maechler and Bolker, 2014). Steele, Diamond and Amin (1996) contains details of the study and some multilevel analyses. Variables in the `Contraception` data frame include:

use a two-level factor variable indicating whether a woman is a user of contraception at the time of the survey, with levels `Y` for use and `N` for non-use,

parameter	95% C.I. low.	estimate	95% C.I. upp.
β_0	-1.2185	-1.0418	-0.8651
β_1	0.2956	0.5003	0.7049
β_2	-0.0259	-0.0164	-0.0068
β_3	0.4934	0.6815	0.8698
β_4	0.6223	0.8306	1.0389
β_5	0.6102	0.8244	1.0387
σ_1	0.2748	0.3785	0.5214
σ_2	0.3096	0.4965	0.7962
ρ	-0.9367	-0.7984	-0.4446

Table 1: Expectation propagation approximate maximum likelihood estimates and corresponding 95% confidence interval (C.I.) lower and upper limits for the parameters in model (22) and (23).

age age of the woman in years at the time of the survey, centred about the average age of all women in the study,

district a multi-level factor variable that codes the district, out of 60 districts in total, in which the woman lives,

urban a two-level factor variable indicating whether or not the district in which the woman lives is urban, with levels Y for urban dwelling and N for rural dwelling, and

livch a four-level factor variable that indicates the number of living children of the woman, with levels 0 for no children, 1 for one child, 2 for two children and 3+ for three or more children.

A random intercepts and slopes probit mixed model for these data is

$$I(\text{use}_{ij} = Y) | u_{0i}, u_{1i} \stackrel{\text{ind.}}{\sim} \text{Bernoulli} \left(\Phi \left(\beta_0 + u_{0i} + (\beta_1 + u_{1i}) I(\text{urban}_{ij} = Y) + \beta_2 \text{age}_{ij} + \beta_3 I(\text{livch}_{ij} = 1) + \beta_4 I(\text{livch}_{ij} = 2) + \beta_5 I(\text{livch}_{ij} = 3+) \right) \right) \quad (22)$$

where $I(\mathcal{P}) = 1$ if \mathcal{P} is true and 0 otherwise. Also, use_{ij} denotes the value of use for the j th woman within the i th district, $1 \leq i \leq 60$, with the other variables defined analogously. The bivariate random effects vectors are assumed to satisfy

$$\begin{bmatrix} u_{0i} \\ u_{1i} \end{bmatrix} \stackrel{\text{ind.}}{\sim} N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix} \right). \quad (23)$$

We fitted this model using our expectation propagation approximate likelihood inference scheme. It took about 35 seconds on the fourth author's MacBook Air laptop (2.2 gigahertz processor and 8 gigabytes of random access memory) to produce the inferential summary given in Table 1.

Each of the parameters is seen to be statistically significantly different from zero. As examples, the 95% confidence interval for β_2 of (0.296, 0.705) indicates a higher use of contraception in urban districts and the 95% confidence interval for σ_2 of (0.310, 0.796) shows that there is significant heterogeneity in the urban versus rural effect across the 60 districts.

We also used expectation propagation approximate best prediction to obtain predictions of the u_{0i} and u_{1i} values. The results are plotted in Figure 4 and provide a visualization of between-district heterogeneity.




Figure 4: Scatterplot of the expectation propagation-approximate best predictions of the random slopes and corresponding random intercepts for the fit of model given by (22) and (23) to data from a 1998 Bangladesh fertility study.

5 Theoretical Considerations

We now discuss the question regarding whether the excellent inferential accuracy of the Section 3 methodology is supported by theory. A fuller theoretical analysis is the subject of ongoing work involving the first four authors and, upon completion, will be reported elsewhere. In this section we provide a heuristic explanation for the accuracy of expectation propagation in the binary response mixed model context.

First note that the i th log-likelihood summand is

$$\ell_i(\boldsymbol{\beta}, \boldsymbol{\Sigma}) = \log \int_{\mathbb{R}^{dR}} \left\{ \prod_{j=1}^{n_i} \frac{p(\mathbf{y}_i | \mathbf{u}_i; \boldsymbol{\beta})}{\tilde{p}(\mathbf{y}_i | \mathbf{u}_i; \boldsymbol{\beta})} \right\} \exp \left\{ \left[\begin{array}{c} 1 \\ \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{array} \right]^T \hat{\boldsymbol{\eta}}_i \right\} d\mathbf{u}_i$$

where $\tilde{p}(\mathbf{y}_i | \mathbf{u}_i; \boldsymbol{\beta})$ is given by expression (7) with the $\boldsymbol{\eta}_{ij}$ set to the converged $\eta_{p(y_{ij} | \mathbf{u}_i; \boldsymbol{\beta})} \rightarrow \mathbf{u}_i$ values. We also have

$$\underline{\ell}_i(\boldsymbol{\beta}, \boldsymbol{\Sigma}) = \log \int_{\mathbb{R}^{dR}} \exp \left\{ \left[\begin{array}{c} 1 \\ \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{array} \right]^T \hat{\boldsymbol{\eta}}_i \right\} d\mathbf{u}_i.$$

Now make the change of variables

$$\mathbf{v} = \boldsymbol{\Delta}_i^{-1} \{ \mathbf{u}_i - \text{BP}(\mathbf{u}_i) \} \quad \text{where} \quad \boldsymbol{\Delta}_i \equiv \text{Cov}(\mathbf{u}_i | \mathbf{y})^{1/2}$$

involving the expectation propagation-approximate best predictor quantities given by (17) and (18). Straightforward manipulations then lead to the discrepancy between $\ell(\boldsymbol{\beta}, \boldsymbol{\Sigma})$ and $\underline{\ell}_i(\boldsymbol{\beta}, \boldsymbol{\Sigma})$ equalling

$$\ell(\boldsymbol{\beta}, \boldsymbol{\Sigma}) - \underline{\ell}_i(\boldsymbol{\beta}, \boldsymbol{\Sigma}) = \log \int_{\mathbb{R}^{dR}} \left\{ \prod_{j=1}^{n_i} A_{ij}(\boldsymbol{\Delta}_i \mathbf{v}) \right\} \phi_{\mathbf{I}}(\mathbf{v}) d\mathbf{v} \quad (24)$$

where, for any $\mathbf{x} \in \mathbb{R}^{d^R}$, $\phi_I(\mathbf{x}) \equiv (2\pi)^{-d^R/2} \exp(-\frac{1}{2}\mathbf{x}^T\mathbf{x})$ and

$$A_{ij}(\mathbf{x}) \equiv F\left((2y_{ij} - 1)(\boldsymbol{\beta}^T \mathbf{x}_{ij}^F + (\mathbb{B}\tilde{\mathbf{P}}(\mathbf{u}_i) + \mathbf{x})^T \mathbf{x}_{ij}^R)\right) \\ \times \exp\left\{-\left[\begin{array}{c} 1 \\ \mathbb{B}\tilde{\mathbf{P}}(\mathbf{u}_i) + \mathbf{x} \\ \text{vech}\left((\mathbb{B}\tilde{\mathbf{P}}(\mathbf{u}_i) + \mathbf{x})(\mathbb{B}\tilde{\mathbf{P}}(\mathbf{u}_i) + \mathbf{x})^T\right) \end{array}\right]^T \eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i}\right\}.$$

Using the same change of variables, the moment-matching conditions corresponding to the Kullback-Leibler projection (8) are

$$\int_{\mathbb{R}^{d^R}} \mathbf{v}^{\otimes k} A_{ij}(\boldsymbol{\Delta}_i \mathbf{v}) \phi_I(\mathbf{v}) d\mathbf{v} = \int_{\mathbb{R}^{d^R}} \mathbf{v}^{\otimes k} \phi_I(\mathbf{v}) d\mathbf{v}, \quad k = 0, 1, 2, \quad (25)$$

where $\mathbf{v}^{\otimes 0} \equiv 1$, $\mathbf{v}^{\otimes 1} \equiv \mathbf{v}$ and $\mathbf{v}^{\otimes 2} = \mathbf{v}\mathbf{v}^T$.

To aid intuition, for the remainder of this section we restrict attention to $d^R = 1$ and write δ_i instead of $\boldsymbol{\Delta}_i$ to signify the fact that this quantity is scalar in this special case. Next, we make the

$$\text{working assumption: } \delta_i = O_p(n_i^{-1/2}). \quad (26)$$

This assumption is in keeping with the fact that δ_i is the expectation propagation approximation to the sample standard deviation of $\mathbb{B}\tilde{\mathbf{P}}(\mathbf{u}_i) - \mathbf{u}_i$. Then Taylor series expansion of A_{ij} about zero and substitution into the $d^R = 1$ version of (25) leads to

$$A_{ij}(0) = 1 + O(\delta_i^4), \quad A'_{ij}(0) = O(\delta_i^2) \quad \text{and} \quad A''_{ij}(0) = O(\delta_i^2).$$

Plugging these into (24) and using $\log(1 + \varepsilon) \approx \varepsilon$ for small ε we obtain

$$\ell_i(\boldsymbol{\beta}, \boldsymbol{\Sigma}) - \tilde{\ell}_i(\boldsymbol{\beta}, \boldsymbol{\Sigma}) = O_p(n_i^{-1/2}) \quad \text{under (26)}.$$

These heuristics suggest that expectation propagation provides consistent estimation of the log-likelihood summands as the number of measurements in the i th group increases. The deeper question concerning the asymptotic statistical properties of the expectation propagation-based estimators $(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\Sigma}})$ requires more delicate theoretical analysis. As mentioned earlier in this section, this question is being pursued by authors of this article.

Before closing this section, we mention that there is a small but emerging body of research concerning the large sample behavior of expectation propagation for approximation Bayesian inference. A recent contribution of this type is Dehaene & Barthelmé (2018) which provides Bernstein-von Mises theory for Bayesian expectation propagation.

6 Higher Level and Crossed Random Effects Extensions

The binary mixed model given by (1) is adequate for the common situation of there being only one grouping mechanism. However, more elaborate models are required for situations such as hierarchical and cross-tabulated grouping mechanisms. Goldstein (2010), for example, provides an extensive treatment of mixed models with higher levels of nesting. A major reference for crossed random effects mixed models is Baayen, Davidson & Bates (2008). Here we provide advice regarding extension our expectation propagation approach to these settings.

The *two levels of nesting* extension of (1) is

$$y_{ijk} | \mathbf{u}_i^{L1}, \mathbf{u}_{ij}^{L2} \stackrel{\text{ind.}}{\sim} \text{Bernoulli}\left(F\left(\boldsymbol{\beta}^T \mathbf{x}_{ijk}^F + (\mathbf{u}_i^{L1})^T \mathbf{x}_{ijk}^{R1} + (\mathbf{u}_{ij}^{L2})^T \mathbf{x}_{ijk}^{R2}\right)\right), \\ \mathbf{u}_i^{L1} \stackrel{\text{ind.}}{\sim} N(\mathbf{0}, \boldsymbol{\Sigma}^{L1}) \text{ independently of } \mathbf{u}_{ij}^{L2} \stackrel{\text{ind.}}{\sim} N(\mathbf{0}, \boldsymbol{\Sigma}^{L2}), \quad (27) \\ 1 \leq i \leq m, \quad 1 \leq j \leq n_i, \quad 1 \leq k \leq o_{ij}.$$

The response y_{ijk} and predictor vectors \mathbf{x}_{ijk}^F , \mathbf{x}_{ijk}^{R1} and \mathbf{x}_{ijk}^{R2} correspond to the k th set of measurements within the j th inner group within the i th outer group. The number of outer groups is m and the number of inner groups in the i th outer group is n_i . The sample size of the j th group in the i th outer group is o_{ij} . Also, \mathbf{x}_{ijk}^{R1} is $d^{R1} \times 1$ and \mathbf{x}_{ijk}^{R2} is $d^{R2} \times 1$. The log-likelihood of $(\beta, \Sigma^{L1}, \Sigma^{L2})$ may be written as

$$\ell(\beta, \Sigma^{L1}, \Sigma^{L2}) = \sum_{i=1}^m \log \int_{\mathbb{R}^{d^R}} \prod_{j=1}^{n_i} \prod_{k=1}^{o_{ij}} p \left(y_{ijk} \mid \begin{bmatrix} \mathbf{u}_i^{L1} \\ \mathbf{u}_{ij}^{L2} \end{bmatrix}; \beta \right) p \left(\begin{bmatrix} \mathbf{u}_i^{L1} \\ \mathbf{u}_{ij}^{L2} \end{bmatrix}; \Sigma^{L1}, \Sigma^{L2} \right) d \begin{bmatrix} \mathbf{u}_i^{L1} \\ \mathbf{u}_{ij}^{L2} \end{bmatrix} \quad (28)$$

where $d^R = d^{R1} + d^{R2}$,

$$p \left(y_{ijk} \mid \begin{bmatrix} \mathbf{u}_i^{L1} \\ \mathbf{u}_{ij}^{L2} \end{bmatrix}; \beta \right) \equiv F \left((2y_{ijk} - 1) (\beta^T \mathbf{x}_{ijk}^F + (\mathbf{u}_i^{L1})^T \mathbf{x}_{ijk}^{R1} + (\mathbf{u}_{ij}^{L2})^T \mathbf{x}_{ijk}^{R2}) \right), \quad y_{ijk} = 0, 1,$$

and

$$p \left(\begin{bmatrix} \mathbf{u}_i^{L1} \\ \mathbf{u}_{ij}^{L2} \end{bmatrix}; \Sigma^{L1}, \Sigma^{L2} \right) \equiv |2\pi \Sigma^{L1}|^{-1/2} |2\pi \Sigma^{L2}|^{-1/2} \exp \left\{ -\frac{1}{2} \begin{bmatrix} \mathbf{u}_i^{L1} \\ \mathbf{u}_{ij}^{L2} \end{bmatrix}^T \begin{bmatrix} \Sigma^{L1} & \mathbf{0} \\ \mathbf{0} & \Sigma^{L2} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{u}_i^{L1} \\ \mathbf{u}_{ij}^{L2} \end{bmatrix} \right\}.$$

Expectation propagation approximation of $\ell(\beta, \Sigma^{L1}, \Sigma^{L2})$ then proceeds by message passing on the factor graph displayed in Figure 5. In the probit case Theorem 1 can be called upon to obtain closed form updates for the message natural parameter vectors leading to an algorithm analogous to Algorithm 1.




Figure 5: Factor graph representation of the product structure of the integrand in (6). The open circle corresponds to the random effect vector $[\mathbf{u}_i^{L1} \ \mathbf{u}_{ij}^{L2}]^T$ and the solid rectangles indicate factors in the integrand of (28). Edges indicate dependence of each factor on $[\mathbf{u}_i^{L1} \ \mathbf{u}_{ij}^{L2}]^T$.

A crossed random effects extension of (1) is

$$y_{ii'j} | \mathbf{u}_i, \mathbf{u}'_{i'} \stackrel{\text{ind.}}{\sim} \text{Bernoulli} \left(F \left(\beta^T \mathbf{x}_{ii'j}^F + (\mathbf{u}_i)^T \mathbf{x}_{ii'j}^{R'} + (\mathbf{u}'_{i'})^T \mathbf{x}_{ii'j}^{R''} \right) \right),$$

$$\mathbf{u}_i \stackrel{\text{ind.}}{\sim} N(\mathbf{0}, \Sigma) \text{ independently of } \mathbf{u}'_{i'} \stackrel{\text{ind.}}{\sim} N(\mathbf{0}, \Sigma'),$$

$$1 \leq i \leq m, \quad 1 \leq i' \leq m', \quad 1 \leq j \leq n_{ii'}$$

where the data are cross-tabulated according to membership of two groups of sizes m and m' indexed according to the pair $(i, i') \in \{1, \dots, m\} \times \{1, \dots, m'\}$, with $n_{ii'}$ denoting the sample size within group (i, i') . Note that $n_{ii'} = 0$ is a possibility for some (i, i') . The response $y_{ii'j}$ and the predictor vectors $\mathbf{x}_{ii'j}^F$, $\mathbf{x}_{ii'j}^{R'}$ and $\mathbf{x}_{ii'j}^{R''}$ correspond to the j th set of measurements within group (i, i') . The \mathbf{u}_i , $1 \leq i \leq m$, are $d^R \times 1$ random effects for group-specific departures from the fixed effects for the first group. The $\mathbf{u}'_{i'}$, $1 \leq i' \leq m'$, are $d^{R'}$ random effects for group-specific departures from the fixed effects for the second group. The log-likelihood of (β, Σ, Σ') is

$$\ell(\beta, \Sigma, \Sigma') = \sum_{i=1}^m \sum_{i'=1}^{m'} I(n_{ii'} > 0) \log \int_{\mathbb{R}^{d^R + d^{R'}}} \prod_{j=1}^{n_{ii'}} p \left(y_{ii'j} \mid \begin{bmatrix} \mathbf{u}_i \\ \mathbf{u}'_{i'} \end{bmatrix}; \beta \right) p \left(\begin{bmatrix} \mathbf{u}_i \\ \mathbf{u}'_{i'} \end{bmatrix}; \Sigma, \Sigma' \right) d \begin{bmatrix} \mathbf{u}_i \\ \mathbf{u}'_{i'} \end{bmatrix}$$

where

$$p \left(y_{ii'j} \left| \begin{bmatrix} \mathbf{u}_i \\ \mathbf{u}'_{i'} \end{bmatrix}; \boldsymbol{\beta} \right. \right) \equiv F \left((2y_{ijk} - 1) (\boldsymbol{\beta}^T \mathbf{x}_{ii'j}^F + (\mathbf{u}_i)^T \mathbf{x}_{ii'j}^R + (\mathbf{u}'_{i'})^T \mathbf{x}_{ii'j}^R) \right), \quad y_{ijk} = 0, 1,$$

and

$$p \left(\begin{bmatrix} \mathbf{u}_i \\ \mathbf{u}'_{i'} \end{bmatrix}; \boldsymbol{\Sigma}, \boldsymbol{\Sigma}' \right) \equiv |2\pi\boldsymbol{\Sigma}|^{-1/2} |2\pi\boldsymbol{\Sigma}'|^{-1/2} \exp \left\{ -\frac{1}{2} \begin{bmatrix} \mathbf{u}_i \\ \mathbf{u}'_{i'} \end{bmatrix}^T \begin{bmatrix} \boldsymbol{\Sigma} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}' \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{u}_i \\ \mathbf{u}'_{i'} \end{bmatrix} \right\}.$$

Expectation propagation approximation of $\ell(\boldsymbol{\beta}, \boldsymbol{\Sigma}, \boldsymbol{\Sigma}')$ can be achieved via message passing on a factor graph similar to that shown in Figure 5.

7 Transferral to Other Mixed Models

Until now we have mainly focused on the special case of probit mixed models with Gaussian random effects since the requisite Kullback-Leibler projections have closed form solutions. However, our approach is quite general and, at least in theory, applies to other mixed models. We now briefly describe transferral to other mixed models.

7.1 Logistic Mixed Models

As we mention in Section 2, the probit and logistic cases are distinguished according to whether $F = \Phi$ or $F = \text{expit}$. Therefore, transferral from probit to logistic mixed models involves replacement of f_{input} in Theorem 1 by

$$f_{\text{input}}(\mathbf{x}) = \text{expit}(c_0 + \mathbf{c}_1^T \mathbf{x}) \exp \left\{ \begin{bmatrix} \mathbf{x} \\ \text{vech}(\mathbf{x}\mathbf{x}^T) \end{bmatrix}^T \begin{bmatrix} \boldsymbol{\eta}_1^{\text{input}} \\ \boldsymbol{\eta}_2^{\text{input}} \end{bmatrix} \right\}, \quad \mathbf{x} \in \mathbb{R}^d. \quad (29)$$

In view of Lemma 1 of the online supplement, Kullback-Leibler projection of f_{input} onto the unnormalized Normal family involves univariate integrals of the form

$$\int_{-\infty}^{\infty} x^p \exp\{qx - rx^2 - \log(1 + e^x)\} dx, \quad p = 0, 1, 2, \quad q \in \mathbb{R}, \quad r > 0. \quad (30)$$

In the Bayesian context, Gelman *et al.* (2014; Section 13.8) and Kim & Wand (2017) describe quadrature-based approaches to evaluation of (30), each of which transfers to the frequentist context dealt with here. However, there is a significant speed cost compared with the probit case.

An alternative approach involves use of the family of approximations to expit of the form

$$\text{expit}_k(x) \equiv \sum_{i=1}^k p_{k,i} \Phi(s_{k,i}x)$$

for constants $p_{k,i}$ and $s_{k,i}$, as advocated by Monahan & Stefanski (1989). Since the approximation is a linear combination of scalings of Φ , the function f_{input} in (29) with expit replaced by expit_k admits closed form Kullback-Leibler projections onto the unnormalized Multivariate Normal, leading to fast and accurate inference for logistic mixed models.

Details on the mechanics and performance of expectation propagation for logistic mixed models is to be reported in Yu (2019).

7.2 Other Generalized Linear Mixed Models

Whilst we have focused on the binary response situation in this article, we quickly point out that the principles apply to other generalized linear mixed models such as those based on the Gamma and Poisson families. Note that (2) with $F = \text{expit}$ generalizes to

$$\begin{aligned} \ell(\boldsymbol{\beta}, \boldsymbol{\Sigma}) &= \sum_{i=1}^m \log \int_{d^{\mathbb{R}}} \left[\prod_{j=1}^{n_i} \exp \left\{ y_{ij} (\boldsymbol{\beta}^T \mathbf{x}_{ij}^{\mathbb{F}} + \mathbf{u}^T \mathbf{x}_{ij}^{\mathbb{R}}) - b(\boldsymbol{\beta}^T \mathbf{x}_{ij}^{\mathbb{F}} + \mathbf{u}^T \mathbf{x}_{ij}^{\mathbb{R}}) + c(y_{ij}) \right\} \right] \\ &\quad \times |2\pi\boldsymbol{\Sigma}|^{-1/2} \exp(-\frac{1}{2}\mathbf{u}^T \boldsymbol{\Sigma}^{-1} \mathbf{u}) d\mathbf{u} \end{aligned}$$

where the functions b and c are as given in Table 2.1 of McCullagh & Nelder (1989). Setting $b(x) = \log(1 + e^x)$ and $c(x) = 0$ gives the $F = \text{expit}$ logistic mixed model while putting $b(x) = e^x$ and $c(x) = -\log(x!)$ gives the corresponding Poisson mixed model. The family of integrals

$$\int_{-\infty}^{\infty} x^p \exp\{qx - rx^2 - b(x)\} dx, \quad p = 0, 1, 2, \quad q \in \mathbb{R}, \quad r > 0,$$

is required to facilitate the required Kullback-Leibler projections for the $d^{\mathbb{R}} = 1$ case. With the exception of logistic mixed models, multivariate numerical integration appears to be required when $d^{\mathbb{R}} > 1$. Yu (2019) will contain a detailed account of the practicalities and performance of expectation propagation for this class of models.

Acknowledgments

This research was supported by Australian Research Council Discovery Project DP180100597. We are grateful for advice from Jim Booth, Omar Ghattas and Alan Huang on aspects of this research.

References

- Azzalini, A. (2017). The R package `sn`: The skew-normal and skew-t distributions (version 1.5). <http://azzalini.stat.unipd.it/SN>
- Baayen, R.H., Davidson, D.J. and Bates, D.M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, **59**, 390–412.
- Bates, D., Maechler, M. and Bolker, B. (2014). `lme4`: Examples from multilevel modelling software review. R package version 1.0. <http://cran.r-project.org>.
- Baltagi, B.H. (2013). *Econometric Analysis of Panel Data, Fifth Edition*. Chichester, U.K.: John Wiley & Sons.
- Bates, D., Maechler, M., Bolker, B. and Walker, S. (2015). Fitting linear mixed-effects models using `lme4`. *Journal of Statistical Software*, **67(1)**, 1–48.
- Bishop, C.M. (2006). *Pattern Recognition and Machine Learning*. New York: Springer.
- Broyden, C.G. (1970). The convergence of a class of double-rank minimization algorithms. *Journal of the Institute of Mathematics and Its Applications*, **6**, 76–90.

- Carlin, B.P. and Gelfand, A.E. (1991). A sample reuse method for accurate parametric empirical Bayes confidence intervals. *Journal of the Royal Statistical Society, Series B*, **53**, 189–200.
- Dehaene, G. and Barthelmé, S. (2018). Expectation propagation in the large-data limit. *Journal of the Royal Statistical Society, Series B*, **80**, 199–217.
- Diggle, P., Heagerty, P., Liang, K.-L. and Zeger, S. (2002). *Analysis of Longitudinal Data, Second Edition*. Oxford, U.K.: Oxford University Press.
- Fletcher, R. (1970). A new approach to variable metric algorithms. *Computer Journal*, **13**, 317–322.
- Gelman, A., Carlin, J.B., Stern, H.S., Dunson, D.B., Vehtari, A. and Rubin, D.B. (2014). *Bayesian Data Analysis, Third Edition*, Boca Raton, Florida: CRC Press.
- Gelman, A. and Hill, J. (2007). *Data Analysis using Regression and Multilevel/Hierarchical Models*. New York: Cambridge University Press.
- Givens, G.H. and Hoeting, J.A. (2005). *Computational Statistics*, Hoboken, New Jersey: John Wiley & Sons.
- Goldfarb, D. (1970). A family of variable metric updates derived by variational means. *Mathematics of Computation*, **24**, 23–26.
- Goldstein, H. (2010). *Multilevel Statistical Models, Fourth Edition*. Chichester, U.K.: John Wiley & Sons.
- Harville, D.A. (2008). *Matrix Algebra from a Statistician's Perspective*. New York: Springer.
- Jeon, M., Rijmen, F. and Rabe-Hesketh, S. (2017). A variational maximization-maximization algorithm for generalized linear mixed models with crossed random effects. *Psychometrika*, **82**, 693–716.
- Kim, A.S.I. and Wand, M.P. (2017). On expectation propagation for generalised, linear and mixed models. *Australian and New Zealand Journal of Statistics*, **59**, in press.
- Magnus, J.R. and Neudecker, H. (1999). *Matrix Differential Calculus with Applications in Statistics and Econometrics, Revised Edition*. Chichester U.K.: Wiley.
- Mardia, K.V., Kent, J.T. and Bibby, J.M. (1979). *Multivariate Analysis*. London: Academic Press.
- The Mathworks Incorporated (2018). Natick, Massachusetts, U.S.A.
- McCullagh, P. and Nelder, J.A. (1989). *Generalized Linear Models, Second Edition*. London: Chapman and Hall.
- McCulloch, C.E., Searle, S.R. and Neuhaus, J.M. (2008). *Generalized, Linear, and Mixed Models, Second Edition*. New York: John Wiley & Sons.
- Minka, T.P. (2001). Expectation propagation for approximate Bayesian inference. In J.S. Breese & D. Koller (eds), *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, pp. 362–369. Burlington, Massachusetts: Morgan Kaufmann.

- Minka, T. (2005). Divergence measures and message passing. *Microsoft Research Technical Report Series*, **MSR-TR-2005-173**, 1–17.
- Minka, T. & Winn, J. (2008), Gates: A graphical notation for mixture models. *Microsoft Research Technical Report Series*, **MSR-TR-2008-185**, 1–16.
- Monahan, J.F. & Stefanski, L.A. (1989). Normal scale mixture approximations to $F^*(z)$ and computation of the logistic-normal integral. In Balakrishnan, N. (editor), *Handbook of the Logistic Distribution*. New York: Marcel Dekker, 529–540.
- Nelder, J.A. and Mead, R. (1965). A simplex method for function minimization. *Computer Journal*, **7**, 308–313.
- Ogden, H.E. (2015). A sequential reduction method for inference in generalized linear mixed models. *Electronic Journal of Statistics*, **9**, 135–152.
- Pinheiro, J.C. and Bates, D.M. (2000). *Mixed-Effects Models in S and S-PLUS*. New York: Springer.
- R Core Team (2018). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>.
- Rao, J.N.K. and Molina, I. (2015). *Small Area Estimation, Second Edition*. Hoboken, New Jersey: John Wiley & Sons.
- Rue, H., Martino, S. and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations (with discussion). *Journal of the Royal Statistical Society, Series B*, **71**, 319–392.
- Shanno, D.F. (1970). Conditioning of quasi-Newton methods for function minimization. *Mathematics of Computation*, **24**, 647–656.
- Steele, F., Diamond, I. and Amin, S. (1996). Immunization uptake in rural Bangladesh: a multilevel analysis. *Journal of the Royal Statistical Society, Series A*, **159**, 289–299.
- The Mathworks Incorporated (2018). Natick, Massachusetts, U.S.A.
- Wainwright, M.J. and Jordan, M.I. (2008). Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, **1**, 1–305.
- Wand, M.P. and Ormerod, J.T. (2012). Continued fraction enhancement of Bayesian computing. *Stat*, **1**, 31–41.
- Wand, M.P. and Yu, J.C.F. (2018). glmmEP: Fast and accurate likelihood-based inference in generalized linear mixed models via expectation propagation. R package version 1.0. <http://cran.r-project.org>.
- Yu, J.C.F. (2019). *Fast and Accurate Frequentist Generalized Linear Mixed Model Analysis via Expectation Propagation*. Doctor of Philosophy thesis, University of Technology Sydney.

Supplement for:
**Fast and Accurate Binary Response Mixed Model
 Analysis via Expectation Propagation**

BY P. HALL¹, I.M. JOHNSTONE², J.T. ORMEROD³, M.P. WAND⁴ AND J.C.F. YU⁴

¹University of Melbourne, ²Stanford University, ³University of Sydney
 and ⁴University of Technology Sydney

S.1 Proof of Theorem 1

For $\mathbf{x} \in \mathbb{R}^d$, define

$$\phi_{\Sigma}(\mathbf{x}) \equiv (2\pi)^{-d/2} |\Sigma|^{-1/2} \exp\left(-\frac{1}{2} \mathbf{x}^T \Sigma^{-1} \mathbf{x}\right)$$

so that

$$\phi_{\mathbf{I}}(\mathbf{x}) = (2\pi)^{-d/2} \exp\left(-\frac{1}{2} \mathbf{x}^T \mathbf{x}\right).$$

We continue to use an unadorned ϕ to denote the Univariate Normal density function:

$$\phi(x) = (2\pi)^{-1/2} \exp\left(-\frac{1}{2} x^2\right).$$

The notation $\|\mathbf{v}\| = \sqrt{\mathbf{v}^T \mathbf{v}}$ for a column vector \mathbf{v} is also used.

Lemma 1. For any function $g : \mathbb{R} \rightarrow \mathbb{R}$ and $d \times 1$ vectors α_1, α_2 and α_3 such that the integrals exist:

$$\int_{\mathbb{R}^d} g(\alpha_1^T \mathbf{x}) \phi_{\mathbf{I}}(\mathbf{x}) d\mathbf{x} = \int_{-\infty}^{\infty} g(\|\alpha_1\| z) \phi(z) dz, \quad (\text{S.1})$$

$$\int_{\mathbb{R}^d} g(\alpha_1^T \mathbf{x}) (\alpha_2^T \mathbf{x}) \phi_{\mathbf{I}}(\mathbf{x}) d\mathbf{x} = \{(\alpha_1^T \alpha_2) / \|\alpha_1\|\} \int_{-\infty}^{\infty} z g(\|\alpha_1\| z) \phi(z) dz \quad (\text{S.2})$$

$$\begin{aligned} \text{and } \int_{\mathbb{R}^d} g(\alpha_1^T \mathbf{x}) (\alpha_2^T \mathbf{x}) (\alpha_3^T \mathbf{x}) \phi_{\mathbf{I}}(\mathbf{x}) d\mathbf{x} &= (\alpha_2^T \alpha_3) \int_{-\infty}^{\infty} g(\|\alpha_1\| z) \phi(z) dz \quad (\text{S.3}) \\ &+ \{(\alpha_1^T \alpha_2)(\alpha_1^T \alpha_3) / \|\alpha_1\|^2\} \int_{-\infty}^{\infty} (z^2 - 1) g(\|\alpha_1\| z) \phi(z) dz. \end{aligned}$$

Proof of Lemma 1. Lemma 1 is a consequence of the fact that the integrals on the left-hand side are, respectively,

$$E\{g(\alpha_1^T \mathbf{x})\}, E\{g(\alpha_1^T \mathbf{x})(\alpha_2^T \mathbf{x})\} \quad \text{and} \quad E\{g(\alpha_1^T \mathbf{x})(\alpha_2^T \mathbf{x})(\alpha_3^T \mathbf{x})\}$$

where

$$\mathbf{x} \sim N(\mathbf{0}_d, \mathbf{I}_d).$$

We now focus on simplification of the third integral (S.4). Simplification of the first and second integrals is similar and simpler. Make the change of variables

$$\mathbf{s} \equiv \begin{bmatrix} s_1 \\ s_2 \\ s_3 \end{bmatrix} = \mathbf{A} \mathbf{x} \quad \text{where} \quad \mathbf{A} \equiv \begin{bmatrix} \alpha_1^T \\ \alpha_2^T \\ \alpha_3^T \end{bmatrix}$$

so that

$$E\{g(\boldsymbol{\alpha}_1^T \mathbf{x})(\boldsymbol{\alpha}_2^T \mathbf{x})(\boldsymbol{\alpha}_3^T \mathbf{x})\} = E\{g(s_1)s_2s_3\} \quad \text{where } \mathbf{s} \sim N(\mathbf{0}_3, \mathbf{A}\mathbf{A}^T).$$

We then note that,

$$\begin{aligned} E\{g(s_1)s_2s_3\} &= \int_{-\infty}^{\infty} g(s_1) \left\{ \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} s_2s_3p(s_2s_3|s_1)ds_2ds_3 \right\} p(s_1)ds_1 \\ &= \int_{-\infty}^{\infty} g(s_1) \left\{ \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \{\text{Cov}(s_2, s_3|s_1) + E(s_2|s_1)E(s_3|s_1)\}ds_2ds_3 \right\} p(s_1)ds_1 \end{aligned}$$

and make use of the result (see e.g. Theorem 3.2.4 of Mardia, Kent & Bibby, 1979)

$$\begin{aligned} \begin{bmatrix} s_2 \\ s_3 \end{bmatrix} \Big| s_1 &\sim N\left((s_1/\|\boldsymbol{\alpha}_1\|^2) \begin{bmatrix} \boldsymbol{\alpha}_1^T \boldsymbol{\alpha}_2 \\ \boldsymbol{\alpha}_1^T \boldsymbol{\alpha}_3 \end{bmatrix}, \right. \\ &\quad \left. \begin{bmatrix} \|\boldsymbol{\alpha}_2\|^2 & \boldsymbol{\alpha}_2^T \boldsymbol{\alpha}_3 \\ \boldsymbol{\alpha}_2^T \boldsymbol{\alpha}_3 & \|\boldsymbol{\alpha}_3\|^2 \end{bmatrix} - (1/\|\boldsymbol{\alpha}_1\|^2) \begin{bmatrix} (\boldsymbol{\alpha}_1^T \boldsymbol{\alpha}_2)^2 & (\boldsymbol{\alpha}_1^T \boldsymbol{\alpha}_2)(\boldsymbol{\alpha}_1^T \boldsymbol{\alpha}_3) \\ (\boldsymbol{\alpha}_1^T \boldsymbol{\alpha}_2)(\boldsymbol{\alpha}_1^T \boldsymbol{\alpha}_3) & (\boldsymbol{\alpha}_1^T \boldsymbol{\alpha}_3)^2 \end{bmatrix} \right). \end{aligned}$$

Result (S.4) then follows via simple algebraic manipulations.

Lemma 2. For all $a \in \mathbb{R}$ and $d \times 1$ vectors \mathbf{b}

$$\int_{\mathbb{R}^d} \Phi(a + \mathbf{b}^T \mathbf{x}) \phi_{\mathbf{I}}(\mathbf{x}) d\mathbf{x} = \Phi\left(\frac{a}{\sqrt{\mathbf{b}^T \mathbf{b} + 1}}\right), \quad (\text{S.4})$$

$$\int_{\mathbb{R}^d} \mathbf{x} \Phi(a + \mathbf{b}^T \mathbf{x}) \phi_{\mathbf{I}}(\mathbf{x}) d\mathbf{x} = \frac{\mathbf{b}}{\sqrt{\mathbf{b}^T \mathbf{b} + 1}} \phi\left(\frac{a}{\sqrt{\mathbf{b}^T \mathbf{b} + 1}}\right) \quad \text{and} \quad (\text{S.5})$$

$$\int_{\mathbb{R}^d} \mathbf{x}\mathbf{x}^T \Phi(a + \mathbf{b}^T \mathbf{x}) \phi_{\mathbf{I}}(\mathbf{x}) d\mathbf{x} = \Phi\left(\frac{a}{\sqrt{\mathbf{b}^T \mathbf{b} + 1}}\right) \mathbf{I} - \frac{a\mathbf{b}\mathbf{b}^T}{\sqrt{(\mathbf{b}^T \mathbf{b} + 1)^3}} \phi\left(\frac{a}{\sqrt{\mathbf{b}^T \mathbf{b} + 1}}\right). \quad (\text{S.6})$$

Proof of Lemma 2.

Suppose that Z_1 and Z_2 are independent $N(0, 1)$ random variables. As defined in Section 4.2, for a logical proposition \mathcal{P} , let $I(\mathcal{P}) = 1$ if \mathcal{P} is true and $I(\mathcal{P}) = 0$ if \mathcal{P} is false. Then

$$P(Z_1 \leq a + \|\mathbf{b}\|Z_2) = E\{I(Z_1 \leq a + \|\mathbf{b}\|Z_2)\} = E[E\{I(Z_1 \leq a + \|\mathbf{b}\|Z_2)|Z_2\}] = E\{h(Z_2)\}$$

where $h(z_2) \equiv E\{I(Z_1 \leq a + \|\mathbf{b}\|Z_2)|Z_2 = z_2\}$. But note that

$$h(z_2) = P(Z_1 \leq a + \|\mathbf{b}\|z_2) = \Phi(a + \|\mathbf{b}\|z_2)$$

which implies that

$$P(Z_1 \leq a + \|\mathbf{b}\|Z_2) = E\{\Phi(a + \|\mathbf{b}\|Z_2)\} = \int_{-\infty}^{\infty} \Phi(a + \|\mathbf{b}\|z) \phi(z) dz.$$

Then

$$P(Z_1 \leq a + \|\mathbf{b}\|Z_2) = P(Z_1 - \|\mathbf{b}\|Z_2 \leq a) = P(X_3 \leq a)$$

where $X_3 \equiv Z_1 - \|\mathbf{b}\|Z_2 \sim N(0, 1 + \mathbf{b}^T \mathbf{b})$ by independence of Z_1 and Z_2 . Then (S.4) follows immediately.

Next, let \mathbf{e}_i denote the $d \times 1$ vector with i th entry equal to 1 and with zeroes elsewhere. Then, using Lemma 1, the i th entry of the right-hand side of (S.5) is

$$\begin{aligned} \int_{\mathbb{R}^d} (\mathbf{e}_i^T \mathbf{x}) \Phi(a + \mathbf{b}^T \mathbf{x}) \phi_{\mathbf{I}}(\mathbf{x}) d\mathbf{x} &= \{(\mathbf{e}_i^T \mathbf{b}) / \|\mathbf{b}\|\} \int_{-\infty}^{\infty} z \Phi(a + \|\mathbf{b}\|z) \phi(z) dz \\ &= -\{(\mathbf{e}_i^T \mathbf{b}) / \|\mathbf{b}\|\} \int_{-\infty}^{\infty} \Phi(a + \|\mathbf{b}\|z) \phi'(z) dz \\ &= (\mathbf{e}_i^T \mathbf{b}) \int_{-\infty}^{\infty} \phi(a + \|\mathbf{b}\|z) \phi(z) dz \end{aligned}$$

where the last result follows via integration by parts. The last integrand is

$$(2\pi)^{-1} \exp\{-\frac{1}{2}(a + \|\mathbf{b}\|z)^2 - \frac{1}{2}z^2\} = \phi\left(\frac{a}{\sqrt{\mathbf{b}^T \mathbf{b} + 1}}\right) \phi\left(\frac{z + a\|\mathbf{b}\|/(1 + \mathbf{b}^T \mathbf{b})}{1/\sqrt{\mathbf{b}^T \mathbf{b} + 1}}\right)$$

and (S.5) is an immediate consequence.

Lastly, because of (S.4), the (i, j) entry of the right-hand side of (S.6) is

$$\begin{aligned} \int_{\mathbb{R}^d} (\mathbf{e}_i^T \mathbf{x})(\mathbf{e}_j^T \mathbf{x}) \Phi(a + \mathbf{b}^T \mathbf{x}) \phi_{\mathbf{I}}(\mathbf{x}) d\mathbf{x} \\ &= (\mathbf{e}_i^T \mathbf{e}_j) \int_{-\infty}^{\infty} \Phi(a + \|\mathbf{b}\|z) \phi(z) dz + \frac{(\mathbf{e}_i^T \mathbf{b})(\mathbf{e}_j^T \mathbf{b})}{\|\mathbf{b}\|^2} \int_{-\infty}^{\infty} \Phi(a + \|\mathbf{b}\|z) \phi''(z) dz \\ &= (\mathbf{e}_i^T \mathbf{e}_j) \Phi\left(\frac{a}{\sqrt{\mathbf{b}^T \mathbf{b} + 1}}\right) - \frac{(\mathbf{e}_i^T \mathbf{b})(\mathbf{e}_j^T \mathbf{b})}{\|\mathbf{b}\|} \int_{-\infty}^{\infty} \phi(a + \|\mathbf{b}\|z) \phi'(z) dz \end{aligned}$$

where the last result follows via integration by parts. The last integrand is

$$-(2\pi)^{-1} z \exp\{-\frac{1}{2}(a + \|\mathbf{b}\|z)^2 - \frac{1}{2}z^2\} = -z \phi\left(\frac{a}{\sqrt{\mathbf{b}^T \mathbf{b} + 1}}\right) \phi\left(\frac{z + a\|\mathbf{b}\|/(1 + \mathbf{b}^T \mathbf{b})}{1/\sqrt{\mathbf{b}^T \mathbf{b} + 1}}\right)$$

and (S.6) follows. ■

Next, we note a key connection between Kullback-Leibler projection onto the unnormalized and normalized Multivariate Normal families. For the latter, we introduce the notation

$$\text{proj}_{\mathbf{N}}[p](\mathbf{x}) = q(\mathbf{x})$$

where q is the Multivariate Normal density function minimizes $\text{KL}(p||q)$.

Lemma 3. Let $f \in L_1(\mathbb{R}^d)$ be such that $f \geq 0$ and define $C_f \equiv \int_{\mathbb{R}^d} f(\mathbf{x}) d\mathbf{x}$. Then

$$\text{proj}[f](\mathbf{x}) = C_f \text{proj}_{\mathbf{N}}[f/C_f](\mathbf{x}).$$

Proof Lemma 3.

Let $g(\cdot; \boldsymbol{\eta})$ be a generic unnormalized Multivariate Normal density function with natural parameter vector $\boldsymbol{\eta}$:

$$g(\mathbf{x}; \boldsymbol{\eta}) = \exp \left\{ \left[\begin{array}{c} 1 \\ \mathbf{x} \\ \text{vech}(\mathbf{x}\mathbf{x}^T) \end{array} \right]^T \boldsymbol{\eta} \right\}.$$

Then the Kullback-Leibler divergence of $g(\cdot; \boldsymbol{\eta})$ from f is

$$\text{KL}(f\|g(\cdot; \boldsymbol{\eta})) = \int_{\mathbb{R}^d} [f(\boldsymbol{x}) \log\{f(\boldsymbol{x})/g(\boldsymbol{x}; \boldsymbol{\eta})\} + g(\boldsymbol{x}; \boldsymbol{\eta}) - f(\boldsymbol{x})] d\boldsymbol{x} = \mathcal{K}(\boldsymbol{\eta}) + \text{const}$$

where ‘const’ denotes terms not depending on $\boldsymbol{\eta}$ and

$$\mathcal{K}(\boldsymbol{\eta}) \equiv (2\pi)^{d/2} \exp\{\boldsymbol{\eta}_0 + A_N(\boldsymbol{\eta}_{-0})\} - \begin{bmatrix} \int_{\mathbb{R}^d} f(\boldsymbol{x}) d\boldsymbol{x} \\ \int_{\mathbb{R}^d} \boldsymbol{x} f(\boldsymbol{x}) d\boldsymbol{x} \\ \int_{\mathbb{R}^d} \text{vech}(\boldsymbol{x}\boldsymbol{x}^T) f(\boldsymbol{x}) d\boldsymbol{x} \end{bmatrix}^T \boldsymbol{\eta}.$$

The derivative vector of $\mathcal{K}(\boldsymbol{\eta})$ is

$$D\mathcal{K}(\boldsymbol{\eta}) = (2\pi)^{d/2} \exp\{\boldsymbol{\eta}_0 + A_N(\boldsymbol{\eta}_{-0})\} \begin{bmatrix} 1 \\ D A(\boldsymbol{\eta}_{-0})^T \end{bmatrix}^T - \begin{bmatrix} \int_{\mathbb{R}^d} f(\boldsymbol{x}) d\boldsymbol{x} \\ \int_{\mathbb{R}^d} \boldsymbol{x} f(\boldsymbol{x}) d\boldsymbol{x} \\ \int_{\mathbb{R}^d} \text{vech}(\boldsymbol{x}\boldsymbol{x}^T) f(\boldsymbol{x}) d\boldsymbol{x} \end{bmatrix}^T$$

so the stationary condition, $D\mathcal{K}(\boldsymbol{\eta})^T = \mathbf{0}$, for the minimization of $\text{KL}(f\|g(\cdot; \boldsymbol{\eta}))$ is

$$(2\pi)^{d/2} \exp\{\boldsymbol{\eta}_0 + A_N(\boldsymbol{\eta}_{-0})\} \begin{bmatrix} 1 \\ \nabla A_N(\boldsymbol{\eta}_{-0}) \end{bmatrix} = \begin{bmatrix} \int_{\mathbb{R}^d} f(\boldsymbol{x}) d\boldsymbol{x} \\ \int_{\mathbb{R}^d} \boldsymbol{x} f(\boldsymbol{x}) d\boldsymbol{x} \\ \int_{\mathbb{R}^d} \text{vech}(\boldsymbol{x}\boldsymbol{x}^T) f(\boldsymbol{x}) d\boldsymbol{x} \end{bmatrix}. \quad (\text{S.7})$$

with $\nabla A_N(\boldsymbol{\eta}_{-0}) \equiv D A(\boldsymbol{\eta}_{-0})^T$ denoting the *gradient* vector of $A(\boldsymbol{\eta}_{-0})$. It is easily checked that (S.7) is satisfied by

$$\begin{aligned} (\boldsymbol{\eta}^*)_0 &= \log(C_f) - A_N(\boldsymbol{\eta}_{-0}^*) - \frac{1}{2} d \log(2\pi) \\ \text{where } \boldsymbol{\eta}_{-0}^* &= (\nabla A_N)^{-1} \left(\begin{bmatrix} \int_{\mathbb{R}^d} \boldsymbol{x} \{f(\boldsymbol{x})/C_f\} d\boldsymbol{x} \\ \int_{\mathbb{R}^d} \text{vech}(\boldsymbol{x}\boldsymbol{x}^T) \{f(\boldsymbol{x})/C_f\} d\boldsymbol{x} \end{bmatrix} \right) \end{aligned} \quad (\text{S.8})$$

with existence and uniqueness of $(\nabla A_N)^{-1}$ being guaranteed by Proposition 3.2 of Wainwright & Jordan (2008). The Hessian matrix of $\mathcal{K}(\boldsymbol{\eta})$ is

$$H\mathcal{K}(\boldsymbol{\eta}) = (2\pi)^{d/2} e^{\boldsymbol{\eta}_0 + A_N(\boldsymbol{\eta}_{-0})} \left\{ \begin{bmatrix} 1 \\ \nabla A_N(\boldsymbol{\eta}_{-0}) \end{bmatrix} \begin{bmatrix} 1 \\ \nabla A_N(\boldsymbol{\eta}_{-0}) \end{bmatrix}^T + \begin{bmatrix} 0 & \mathbf{0}^T \\ \mathbf{0} & H A_N(\boldsymbol{\eta}_{-0}) \end{bmatrix} \right\}$$

From Proposition 3.1 of Wainwright & Jordan (2008), A_N is strictly convex on its domain and therefore $H A_N(\boldsymbol{\eta}_{-0})$ is positive definite. Hence $H\mathcal{K}(\boldsymbol{\eta})$ is positive definite for all $\boldsymbol{\eta}$ and so (S.8) is the unique minimizer of $\text{KL}(f\|g(\cdot; \boldsymbol{\eta}))$. Therefore,

$$\text{proj}[f](\boldsymbol{x}) = \exp \left\{ \begin{bmatrix} 1 \\ \boldsymbol{x} \\ \text{vech}(\boldsymbol{x}\boldsymbol{x}^T) \end{bmatrix}^T \boldsymbol{\eta}^* \right\}$$

where $\boldsymbol{\eta}^*$ is as given by (S.8). However, $\boldsymbol{\eta}_{-0}^*$ is the same natural parameter vector that arises via projection of f/C_f onto the family of Multivariate Normal density functions and so

$$\text{proj}_N[f/C_f](\boldsymbol{x}) = \exp \left\{ \begin{bmatrix} \boldsymbol{x} \\ \text{vech}(\boldsymbol{x}\boldsymbol{x}^T) \end{bmatrix}^T \boldsymbol{\eta}_{-0}^* - A_N(\boldsymbol{\eta}_{-0}^*) \right\} (2\pi)^{-d/2}$$

which immediately leads to Lemma 3.

■

The proof of Theorem 1 involves transferral between the common $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ parameters of the d -variate Normal distribution and the natural parameters corresponding to the sufficient statistics \boldsymbol{x} and $\text{vech}(\boldsymbol{x}\boldsymbol{x}^T)$. The transformations in each direction are

$$\begin{cases} \boldsymbol{\eta}_1 = \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu} \\ \boldsymbol{\eta}_2 = -\frac{1}{2} \boldsymbol{D}_d^T \text{vec}(\boldsymbol{\Sigma}^{-1}) \end{cases} \quad \text{and} \quad \begin{cases} \boldsymbol{\mu} = -\frac{1}{2} \{ \text{vec}^{-1}(\boldsymbol{D}_d^{+T} \boldsymbol{\eta}_2) \}^{-1} \boldsymbol{\eta}_1 \\ \boldsymbol{\Sigma} = -\frac{1}{2} \{ \text{vec}^{-1}(\boldsymbol{D}_d^{+T} \boldsymbol{\eta}_2) \}^{-1} \end{cases} \quad (\text{S.9})$$

Recall the notation

$$\boldsymbol{v}^{\otimes k} \equiv \begin{cases} 1 & \text{for } k = 0 \\ \boldsymbol{v} & \text{for } k = 1 \\ \boldsymbol{v}\boldsymbol{v}^T & \text{for } k = 2 \end{cases}$$

and consider Kullback-Leibler projection of $f_{\text{input}}/C_{f_{\text{input}}}$ onto the family of d -variate Normal density functions where

$$f_{\text{input}}(\boldsymbol{x}) \equiv \Phi(c_0 + \boldsymbol{c}_1^T \boldsymbol{x}) \exp \left\{ \begin{bmatrix} \boldsymbol{x} \\ \text{vech}(\boldsymbol{x}\boldsymbol{x}^T) \end{bmatrix}^T \begin{bmatrix} \boldsymbol{\eta}_1^{\text{input}} \\ \boldsymbol{\eta}_2^{\text{input}} \end{bmatrix} \right\},$$

and $C_{f_{\text{input}}} \equiv \int_{\mathbb{R}^d} f_{\text{input}}(\boldsymbol{x}) d\boldsymbol{x}$. Then the projection has mean and covariance matrix

$$\boldsymbol{\mu}^* = \mathcal{M}_1/\mathcal{M}_0 \quad \text{and} \quad \boldsymbol{\Sigma}^* = \mathcal{M}_2/\mathcal{M}_0 - (\mathcal{M}_1/\mathcal{M}_0)(\mathcal{M}_1/\mathcal{M}_0)^T \quad (\text{S.10})$$

where

$$\mathcal{M}_k \equiv \int_{\mathbb{R}^d} \boldsymbol{x}^{\otimes k} \Phi(c_0 + \boldsymbol{c}_1^T \boldsymbol{x}) \exp \left\{ \begin{bmatrix} \boldsymbol{x} \\ \text{vech}(\boldsymbol{x}\boldsymbol{x}^T) \end{bmatrix}^T \begin{bmatrix} \boldsymbol{\eta}_1^{\text{input}} \\ \boldsymbol{\eta}_2^{\text{input}} \end{bmatrix} \right\} d\boldsymbol{x}.$$

Letting

$$\boldsymbol{\Sigma}^{\text{input}} \equiv -\frac{1}{2} \{ \text{vec}^{-1}(\boldsymbol{D}_{d^R}^{+T} \boldsymbol{\eta}_2^{\text{input}}) \}^{-1} \quad \text{and} \quad \boldsymbol{\mu}^{\text{input}} \equiv \boldsymbol{\Sigma}^{\text{input}} \boldsymbol{\eta}_1^{\text{input}}$$

be the common parameters corresponding to $\boldsymbol{\eta}^{\text{input}}$ and making the change of variable $\boldsymbol{z} = (\boldsymbol{\Sigma}^{\text{input}})^{-1/2}(\boldsymbol{x} - \boldsymbol{\mu}^{\text{input}})$ we obtain

$$\mathcal{M}_k = (2\pi)^{d/2} e^{-A_N(\boldsymbol{\eta}^{\text{input}})} \int_{\mathbb{R}^d} (\boldsymbol{\mu}^{\text{input}} + (\boldsymbol{\Sigma}^{\text{input}})^{1/2} \boldsymbol{z})^{\otimes k} \Phi\left((c_0 + \boldsymbol{c}_1^T \boldsymbol{\mu}^{\text{input}}) + \{(\boldsymbol{\Sigma}^{\text{input}})^{1/2} \boldsymbol{c}_1\}^T \boldsymbol{z}\right) \phi_I(\boldsymbol{z}) d\boldsymbol{z}.$$

Lemma 2 and simple algebraic manipulations then give

$$\mathcal{M}_1/\mathcal{M}_0 = \boldsymbol{\mu}^{\text{input}} + \frac{\boldsymbol{\Sigma}^{\text{input}} \boldsymbol{c}_1 \zeta'(r_2)}{\sqrt{\boldsymbol{c}_1^T \boldsymbol{\Sigma}^{\text{input}} \boldsymbol{c}_1 + 1}},$$

and

$$\begin{aligned} \mathcal{M}_2/\mathcal{M}_0 &= \boldsymbol{\mu}^{\text{input}} (\boldsymbol{\mu}^{\text{input}})^T + \frac{\{ \boldsymbol{\Sigma}^{\text{input}} \boldsymbol{c}_1 (\boldsymbol{\mu}^{\text{input}})^T + \boldsymbol{\mu}^{\text{input}} \boldsymbol{c}_1^T \boldsymbol{\Sigma}^{\text{input}} \} \zeta'(r_2)}{\sqrt{\boldsymbol{c}_1^T \boldsymbol{\Sigma}^{\text{input}} \boldsymbol{c}_1 + 1}} \\ &\quad + \boldsymbol{\Sigma}^{\text{input}} - \frac{r \zeta'(r_2) \boldsymbol{\Sigma}^{\text{input}} \boldsymbol{c}_1 \boldsymbol{c}_1^T \boldsymbol{\Sigma}^{\text{input}}}{\boldsymbol{c}_1^T \boldsymbol{\Sigma}^{\text{input}} \boldsymbol{c}_1 + 1}. \end{aligned}$$

where

$$r_2 \equiv \frac{2c_0 - \boldsymbol{c}_1^T \{ \text{vec}^{-1}(\boldsymbol{D}_{d^R}^{+T} \boldsymbol{\eta}_2^{\text{input}}) \}^{-1} \boldsymbol{\eta}_1^{\text{input}}}{\sqrt{2 \left[2 - \boldsymbol{c}_1^T \{ \text{vec}^{-1}(\boldsymbol{D}_{d^R}^{+T} \boldsymbol{\eta}_2^{\text{input}}) \}^{-1} \boldsymbol{c}_1 \right]}}.$$

Combining these last two results and noting (S.10) we obtain the common parameter solutions

$$\begin{aligned}\boldsymbol{\mu}^* &= \boldsymbol{\mu}^{\text{input}} + \frac{\boldsymbol{\Sigma}^{\text{input}} \mathbf{c}_1 \zeta'(r_2)}{\sqrt{\mathbf{c}_1^T \boldsymbol{\Sigma}^{\text{input}} \mathbf{c}_1 + 1}} \\ \boldsymbol{\Sigma}^* &= \boldsymbol{\Sigma}^{\text{input}} + \left\{ \frac{\zeta''(r_2)}{\mathbf{c}_1^T \boldsymbol{\Sigma}^{\text{input}} \mathbf{c}_1 + 1} \right\} \boldsymbol{\Sigma}^{\text{input}} \mathbf{c}_1 \mathbf{c}_1^T \boldsymbol{\Sigma}^{\text{input}}.\end{aligned}$$

Transferral to natural parameters via (S.9) and some simple manipulations then lead to

$$\begin{bmatrix} \boldsymbol{\eta}_1^* \\ \boldsymbol{\eta}_2^* \end{bmatrix} = K_{\text{probit}} \left(\begin{bmatrix} \boldsymbol{\eta}_1^{\text{input}} \\ \boldsymbol{\eta}_2^{\text{input}} \end{bmatrix}; c_0, \mathbf{c}_1 \right).$$

Finally,

$$\begin{aligned}\eta_0^* &= \log(C_{f_{\text{input}}}) - \log \int_{\mathbb{R}^{d^{\mathbb{R}}}} \exp \left\{ \begin{bmatrix} \mathbf{x} \\ \text{vech}(\mathbf{x}\mathbf{x}^T) \end{bmatrix}^T \begin{bmatrix} \boldsymbol{\eta}_1^* \\ \boldsymbol{\eta}_2^* \end{bmatrix} \right\} d\mathbf{x} \\ &= \log(\mathcal{M}_0) - \frac{1}{2} d^{\mathbb{R}} \log(2\pi) - A_N(\boldsymbol{\eta}^*) = \log \Phi(r_2) + A_N(\boldsymbol{\eta}^{\text{input}}) - A_N(\boldsymbol{\eta}^*) \\ &= C_{\text{probit}} \left(\begin{bmatrix} \boldsymbol{\eta}_1^{\text{input}} \\ \boldsymbol{\eta}_2^{\text{input}} \end{bmatrix}, \begin{bmatrix} \boldsymbol{\eta}_1^* \\ \boldsymbol{\eta}_2^* \end{bmatrix}; c_0, \mathbf{c}_1 \right)\end{aligned}$$

S.2 Derivation of Algorithm 1

We now provide full justification of Algorithm 1, starting with a derivation of the message passing representation used in Algorithm 1.

S.2.1 Message Passing Representation Derivation

The derivation of the message passing representation is based on the infrastructure and results laid out in Minka (2005). The treatment given there is for a generalization of Kullback-Leibler divergence, known as α -divergence, and for approximation of (normalized) density functions rather than general non-negative L_1 functions. The Kullback-Leibler divergence minimization problem given by (8) corresponds to $\alpha = 1$ in the notation of Minka (2005). Following Section 4.1 of Minka (2005) we then define the messages passed from the factors neighboring \mathbf{u}_i in Figure 1 to be

$$m_{p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) \equiv \underline{p}(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta}) \quad \text{and} \quad m_{p(\mathbf{u}_i;\boldsymbol{\Sigma}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) \equiv p(\mathbf{u}_i;\boldsymbol{\Sigma}). \quad (\text{S.11})$$

Then, (54) of Minka (2005) invokes the definition

$$m_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta})}(\mathbf{u}_i) \equiv m_{p(\mathbf{u}_i;\boldsymbol{\Sigma}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) \prod_{j' \neq j} m_{p(y_{ij'}|\mathbf{u}_i;\boldsymbol{\beta}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i). \quad (\text{S.12})$$

Result (60) of Minka (2005) with $\alpha = 1$, $s' = 1$ (since we are working with unnormalized rather than normalized Kullback-Leibler divergence) and the simplification that there is only one stochastic node, namely \mathbf{u}_i , provides the main factor to stochastic node message passing updates:

$$m_{p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) \leftarrow \frac{\text{proj}[m_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta})}(\mathbf{u}_i) p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta})](\mathbf{u}_i)}{m_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i;\boldsymbol{\beta})}(\mathbf{u}_i)}, \quad 1 \leq j \leq n_i. \quad (\text{S.13})$$

The other factor to stochastic node message passing update is, trivially from (S.11),

$$m_{p(\mathbf{u}_i; \boldsymbol{\Sigma}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) \leftarrow p(\mathbf{u}_i; \boldsymbol{\Sigma}).$$

The stochastic node to factor updates are, from (S.12),

$$m_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})}(\mathbf{u}_i) \leftarrow m_{p(\mathbf{u}_i; \boldsymbol{\Sigma}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) \prod_{j' \neq j} m_{p(y_{ij'}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i), \quad 1 \leq j \leq n_i.$$

Next, we simplify these message updates to a programmable form.

S.2.2 Simplification of the $m_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i)$ Updates

From (S.12) it is apparent that $m_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i)$ is an unnormalized Multivariate Normal density function and therefore

$$m_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})}(\mathbf{u}_i) = \exp \left\{ \left[\begin{array}{c} 1 \\ \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{array} \right]^T \eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})} \right\}$$

with natural parameter vector $\eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})}$. Introducing the abbreviation:

$$\boldsymbol{\eta}^{\otimes} \equiv \eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})}$$

we have

$$m_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})}(\mathbf{u}_i) = \exp(\eta_0^{\otimes}) \exp \left\{ \left[\begin{array}{c} \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{array} \right]^T \boldsymbol{\eta}_{-0}^{\otimes} \right\}$$

where η_0^{\otimes} denotes the first entry of $\boldsymbol{\eta}^{\otimes}$ and $\boldsymbol{\eta}_{-0}^{\otimes}$ contains the remaining entries. Substitution info (S.13) leads to

$$\begin{aligned} m_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) &\leftarrow \frac{\text{proj} \left[\exp(\eta_0^{\otimes}) \exp \left\{ \left[\begin{array}{c} \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{array} \right]^T \boldsymbol{\eta}_{-0}^{\otimes} \right\} \Phi((2y_{ij} - 1)(\boldsymbol{\beta}^T \mathbf{x}_{ij}^F + \mathbf{u}_i^T \mathbf{x}_{ij}^R)) \right] (\mathbf{u}_i)}{\exp(\eta_0^{\otimes}) \exp \left\{ \left[\begin{array}{c} \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{array} \right]^T \boldsymbol{\eta}_{-0}^{\otimes} \right\}} \\ &= \frac{\text{proj} \left[\Phi(c_{0,ij} + \mathbf{c}_{1,ij}^T \mathbf{u}_i) \exp \left\{ \left[\begin{array}{c} \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{array} \right]^T \boldsymbol{\eta}_{-0}^{\otimes} \right\} \right] (\mathbf{u}_i)}{\exp \left\{ \left[\begin{array}{c} \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{array} \right]^T \boldsymbol{\eta}_{-0}^{\otimes} \right\}} \end{aligned}$$

where

$$c_{0,ij} \equiv (2y_{ij} - 1)(\boldsymbol{\beta}^T \mathbf{x}_{ij}^F) \quad \text{and} \quad \mathbf{c}_{1,ij} \equiv (2y_{ij} - 1)\mathbf{x}_{ij}^R.$$

Using Theorem 1:

$$m_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) \leftarrow \exp \left\{ \left[\begin{array}{c} 1 \\ \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{array} \right]^T \eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i} \right\}$$

where the linear and quadratic coefficient updates are

$$\begin{aligned} (\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i})_{-0} &\leftarrow K_{\text{probit}} \left((\eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})})_{-0}; (2y_{ij} - 1)(\boldsymbol{\beta}^T \mathbf{x}_{ij}^F), (2y_{ij} - 1)\mathbf{x}_{ij}^R \right) \\ &\quad - (\eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})})_{-0} \end{aligned}$$

and the constant coefficient update is

$$\begin{aligned} (\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i})_0 &\leftarrow C_{\text{probit}} \left((\eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})})_{-0}, (\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i})_{-0} \right. \\ &\quad \left. + (\eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})})_{-0}; (2y_{ij} - 1)(\boldsymbol{\beta}^T \mathbf{x}_{ij}^F), (2y_{ij} - 1)\mathbf{x}_{ij}^R \right). \end{aligned}$$

S.2.3 Simplification of the $m_{p(\mathbf{u}_i; \Sigma)} \rightarrow \mathbf{u}_i(\mathbf{u}_i)$ Update

The second definition in (S.11) gives

$$m_{p(\mathbf{u}_i; \Sigma)} \rightarrow \mathbf{u}_i(\mathbf{u}_i) \leftarrow p(\mathbf{u}_i; \Sigma) = \exp \left\{ \left[\begin{array}{c} 1 \\ \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{array} \right]^T \eta_{\Sigma} \right\}.$$

Therefore, if $\eta_{p(\mathbf{u}_i; \Sigma)} \rightarrow \mathbf{u}_i$ denotes the natural parameter vector of $m_{p(\mathbf{u}_i; \Sigma)} \rightarrow \mathbf{u}_i(\mathbf{u}_i)$ then it has the trivial update

$$\eta_{p(\mathbf{u}_i; \Sigma)} \rightarrow \mathbf{u}_i \leftarrow \eta_{\Sigma}.$$

S.2.4 Simplification of the $m_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \beta)}(\mathbf{u}_i)$ Updates

Given the simplified forms of the messages in the two previous subsections we have from (S.12):

$$\begin{aligned} m_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \beta)}(\mathbf{u}_i) &\leftarrow \exp \left\{ \left[\begin{array}{c} 1 \\ \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{array} \right]^T \eta_{p(\mathbf{u}_i; \Sigma)} \rightarrow \mathbf{u}_i \right\} \\ &\times \prod_{j' \neq j} \exp \left\{ \left[\begin{array}{c} 1 \\ \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{array} \right]^T \eta_{p(y_{ij'}|\mathbf{u}_i; \beta)} \rightarrow \mathbf{u}_i \right\} \end{aligned}$$

which leads to

$$\begin{aligned} \eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \beta)} &\leftarrow \eta_{p(\mathbf{u}_i; \Sigma)} \rightarrow \mathbf{u}_i + \sum_{j' \neq j} \eta_{p(y_{ij'}|\mathbf{u}_i; \beta)} \rightarrow \mathbf{u}_i \\ &= \eta_{p(\mathbf{u}_i; \Sigma)} \rightarrow \mathbf{u}_i + \text{SUM}\{\eta_{p(y_i|\mathbf{u}_i; \beta)} \rightarrow \mathbf{u}_i\} - \eta_{p(y_{ij}|\mathbf{u}_i; \beta)} \rightarrow \mathbf{u}_i. \end{aligned}$$

S.2.5 Assembly of All Natural Parameter Updates

We now return to the message passing protocol given in Section 3.2:

Initialize all factor to stochastic node messages.

Cycle until all factor to stochastic node messages converge:

For each factor:

 Compute the messages passed to the factor using (11) or (12).

 Compute the messages passed from the factor using (9) or (10).

For the factors $p(y_{ij}|\mathbf{u}_i; \beta)$:

 computing the messages passed to each of these factors reduces to

$$\eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \beta)} \leftarrow \eta_{p(\mathbf{u}_i; \Sigma)} \rightarrow \mathbf{u}_i + \text{SUM}\{\eta_{p(y_i|\mathbf{u}_i; \beta)} \rightarrow \mathbf{u}_i\} - \eta_{p(y_{ij}|\mathbf{u}_i; \beta)} \rightarrow \mathbf{u}_i$$

 and computing the messages passed from these factors reduces to

$$\left(\eta_{p(y_{ij}|\mathbf{u}_i; \beta)} \rightarrow \mathbf{u}_i \right)_{-0} \leftarrow K_{\text{probit}} \left(\left(\eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \beta)} \right)_{-0}; \mathbf{c}_{0,ij}, \mathbf{c}_{1,ij} \right)$$

$$-(\eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})})_{-0}$$

and

$$\begin{aligned} (\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i})_0 &\leftarrow C_{\text{probit}} \left((\eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})})_{-0}, (\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i})_{-0} \right. \\ &\quad \left. + (\eta_{\mathbf{u}_i \rightarrow p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})})_{-0}; \mathbf{c}_{0,ij}, \mathbf{c}_{1,ij} \right). \end{aligned}$$

For the factors $p(\mathbf{u}_i; \boldsymbol{\Sigma})$:

computing the messages passed from these factors reduces to

$$\eta_{\mathbf{u}_i \rightarrow p(\mathbf{u}_i; \boldsymbol{\Sigma})} \leftarrow \sum_{j=1}^{n_i} \eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i}$$

and computing the messages passed to these factors reduces to

$$\eta_{p(\mathbf{u}_i; \boldsymbol{\Sigma}) \rightarrow \mathbf{u}_i} \rightarrow \eta_{\boldsymbol{\Sigma}}.$$

Algorithm 1 is essentially these natural parameter updates being cycled until convergence. The update for $(\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i})_0$ can be moved outside of the cycle loop without affecting convergence. Also, the $\eta_{\mathbf{u}_i \rightarrow p(\mathbf{u}_i; \boldsymbol{\Sigma})}$ updates are redundant and are omitted from Algorithm 1.

S.3 Derivation of Starting Values Recommendation

We now derive useful starting values for the $\eta_{p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \rightarrow \mathbf{u}_i}$ that have to be initialized in Algorithm 1. Note that

$$\log p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) = \sum_{j=1}^{n_i} \{\zeta(a_{ij}) - \log(2)\} \quad \text{where} \quad a_{ij} \equiv (2y_{ij} - 1)(\boldsymbol{\beta}^T \mathbf{x}_{ij}^F + \mathbf{u}_i^T \mathbf{x}_{ij}^R)$$

and ζ is as defined in Section 3.1. Let $\hat{\mathbf{u}}_i$ be a prediction of \mathbf{u}_i and consider the following expansion of the data-dependent component of $\ell(\boldsymbol{\beta}, \boldsymbol{\Sigma})$:

$$\begin{aligned} \zeta(a_{ij}) &= \zeta(\hat{a}_{ij} + (\mathbf{u}_i - \hat{\mathbf{u}}_i)^T \mathbf{x}_{ij}^R (2y_{ij} - 1)) \\ &= \zeta(\hat{a}_{ij}) + (\mathbf{u}_i - \hat{\mathbf{u}}_i)^T \mathbf{x}_{ij}^R y_{ij}^\dagger \zeta'(\hat{a}_{ij}) + \frac{1}{2} \{(\mathbf{u}_i - \hat{\mathbf{u}}_i)^T \mathbf{x}_{ij}^R (2y_{ij} - 1)\}^2 \zeta''(\hat{a}_{ij}) + \dots \\ &= \begin{bmatrix} 1 \\ \mathbf{u}_i - \hat{\mathbf{u}}_i \\ \text{vech}((\mathbf{u}_i - \hat{\mathbf{u}}_i)(\mathbf{u}_i - \hat{\mathbf{u}}_i)^T) \end{bmatrix}^T \tilde{\boldsymbol{\eta}}_{ij} + \dots \end{aligned}$$

where, as in Section 3.3, $\hat{a}_{ij} \equiv (2y_{ij} - 1)(\boldsymbol{\beta}^T \mathbf{x}_{ij}^F + \mathbf{u}_i^T \mathbf{x}_{ij}^R)$, and

$$\tilde{\boldsymbol{\eta}}_{ij} \equiv \begin{bmatrix} \zeta(\hat{a}_{ij}) \\ \mathbf{x}_{ij}^R (2y_{ij} - 1) \zeta'(\hat{a}_{ij}) \\ \frac{1}{2} \zeta''(\hat{a}_{ij}) \mathbf{D}_{d^R}^T \text{vec}(\mathbf{x}_{ij}^R (\mathbf{x}_{ij}^R)^T) \end{bmatrix}.$$

It follows that the quadratic approximation to $\log p(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})$ based on Taylor expansion about $\hat{\mathbf{u}}_i$ is $\log \check{p}(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta})$ where

$$\check{p}(y_{ij}|\mathbf{u}_i; \boldsymbol{\beta}) \equiv \exp \left\{ \begin{bmatrix} 1 \\ \mathbf{u}_i - \hat{\mathbf{u}}_i \\ \text{vech}((\mathbf{u}_i - \hat{\mathbf{u}}_i)(\mathbf{u}_i - \hat{\mathbf{u}}_i)^T) \end{bmatrix}^T \tilde{\boldsymbol{\eta}}_{ij} \right\}.$$

The starting value recommendation for $\eta_{p(y_{ij}|\mathbf{u}_i; \beta) \rightarrow \mathbf{u}_i}$ is based on replacement of $\check{p}(y_{ij}|\mathbf{u}_i; \beta)$ by $p(y_{ij}|\mathbf{u}_i; \beta)$ in (S.13):

$$m_{\check{p}(y_{ij}|\mathbf{u}_i; \beta) \rightarrow \mathbf{u}_i}(\mathbf{u}_i) \leftarrow \frac{\text{proj}[m_{\mathbf{u}_i \rightarrow \check{p}(y_{ij}|\mathbf{u}_i; \beta)}(\mathbf{u}_i) \check{p}(y_{ij}|\mathbf{u}_i; \beta)](\mathbf{u}_i)}{m_{\mathbf{u}_i \rightarrow \check{p}(y_{ij}|\mathbf{u}_i; \beta)}(\mathbf{u}_i)} = \check{p}(y_{ij}|\mathbf{u}_i; \beta)$$

with the $\text{proj}[\cdot]$ being superfluous in this case due to $\check{p}(y_{ij}|\mathbf{u}_i; \beta)$ being already in the Multivariate Normal family. The starting value for $\eta_{p(y_{ij}|\mathbf{u}_i; \beta) \rightarrow \mathbf{u}_i}$ that arises from this substitution is then given by

$$\exp \left\{ \left[\begin{array}{c} 1 \\ \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{array} \right]^T \eta_{p(y_{ij}|\mathbf{u}_i; \beta) \rightarrow \mathbf{u}_i}^{\text{start}} \right\} = \exp \left\{ \left[\begin{array}{c} 1 \\ \mathbf{u}_i - \hat{\mathbf{u}}_i \\ \text{vech}((\mathbf{u}_i - \hat{\mathbf{u}}_i)(\mathbf{u}_i - \hat{\mathbf{u}}_i)^T) \end{array} \right]^T \check{\eta}_{ij} \right\}.$$

By matching coefficients of like terms we arrive at

$$\eta_{p(y_{ij}|\mathbf{u}_i; \beta) \rightarrow \mathbf{u}_i}^{\text{start}} = \left[\begin{array}{c} \eta_0^{\text{start}} \\ (2y_{ij} - 1)\zeta'(\hat{a}_{ij})\mathbf{x}_{ij}^R - \zeta''(\hat{a}_{ij})\mathbf{x}_{ij}^R(\mathbf{x}_{ij}^R)^T \hat{\mathbf{u}}_i \\ \frac{1}{2}\zeta''(\hat{a}_{ij})\mathbf{D}_{d^R}^T \text{vec}(\mathbf{x}_{ij}^R(\mathbf{x}_{ij}^R)^T) \end{array} \right]$$

where

$$\eta_0^{\text{start}} = \zeta(\hat{a}_{ij}) - (2y_{ij} - 1)\zeta'(\hat{a}_{ij})(\mathbf{x}_{ij}^R)^T \hat{\mathbf{u}}_i + \frac{1}{2}\zeta''(\hat{a}_{ij})\{(\mathbf{x}_{ij}^R)^T \hat{\mathbf{u}}_i\}^2.$$

In Algorithm 1 the cycle loop corresponds to determination of the natural parameter vector

$$\left(\eta_{p(y_{ij}|\mathbf{u}_i; \beta) \rightarrow \mathbf{u}_i} \right)_{-0}$$

implying that the first entry of $\eta_{p(y_{ij}|\mathbf{u}_i; \beta) \rightarrow \mathbf{u}_i}^{\text{start}}$ is not needed for these iterations. Hence, we can instead set $\eta_0^{\text{start}} = 0$ without affecting Algorithm 1. We now have (14).

S.4 Details of Confidence Interval Calculations

Here we provide full details of approximate confidence intervals calculations based on quasi-Newton maximization of $\ell(\beta, \Sigma)$. The calculations depend on the following ingredients:

- some additional convenient matrix notation.
- formulae for transformation from the parameter vector $\theta \equiv \text{vech}(\frac{1}{2} \log(\Sigma))$ to a parameter vector ω that is more appropriate for confidence interval construction.
- formulae for the reverse transformation: from ω to θ .
- a quasi-Newton optimization-based strategy for calculating confidence intervals for the entries of ω , which are then easily transformed to confidence intervals for interpretable covariance matrix parameters, as illustrated in Figures 2 and 3.

S.4.1 Additional Matrix Notation

For a $d \times d$ matrix \mathbf{A} define $\text{diagonal}(\mathbf{A})$ to be the $d \times 1$ vector consisting of the diagonal entries of \mathbf{A} and, provided $d \geq 2$, define $\text{vecbd}(\mathbf{A})$ to be the $\frac{1}{2}d(d-1)$ vector containing the

entries of \mathbf{A} that are below the diagonal of \mathbf{A} in order from left to right and top to bottom. For example,

$$\text{diagonal} \left(\begin{bmatrix} 1 & 5 & 9 & 13 \\ 2 & 6 & 10 & 14 \\ 3 & 7 & 11 & 15 \\ 4 & 8 & 12 & 16 \end{bmatrix} \right) = \begin{bmatrix} 1 \\ 6 \\ 11 \\ 16 \end{bmatrix} \quad \text{and} \quad \text{vecbd} \left(\begin{bmatrix} 1 & 5 & 9 & 13 \\ 2 & 6 & 10 & 14 \\ 3 & 7 & 11 & 15 \\ 4 & 8 & 12 & 16 \end{bmatrix} \right) = \begin{bmatrix} 2 \\ 3 \\ 4 \\ 7 \\ 8 \\ 12 \end{bmatrix}.$$

In addition, if each of \mathbf{a} and \mathbf{b} are $d \times 1$ vectors then $\mathbf{a} \odot \mathbf{b}$ is $d \times 1$ vector of element-wise products and \mathbf{a}/\mathbf{b} is $d \times 1$ vector of element-wise quotients. Similarly, $\log(\mathbf{a})$ and $\tanh(\mathbf{a})$ are obtained in an element-wise fashion.

S.4.2 Transformation from $\boldsymbol{\theta}$ to $\boldsymbol{\omega}$

Given a $\frac{1}{2}d(d+1) \times 1$ vector $\boldsymbol{\theta}$, its corresponding $\boldsymbol{\omega}$ vector of the same length is found via the steps:

1. Obtain the spectral decomposition $\text{vech}^{-1}(\boldsymbol{\theta}) = \mathbf{U}_{\boldsymbol{\theta}} \text{diag}(\boldsymbol{\lambda}_{\boldsymbol{\theta}}) \mathbf{U}_{\boldsymbol{\theta}}^T$.
2. Set $\boldsymbol{\Sigma} = \mathbf{U}_{\boldsymbol{\theta}} \text{diag}\{\exp(2\boldsymbol{\lambda}_{\boldsymbol{\theta}})\} \mathbf{U}_{\boldsymbol{\theta}}^T$.
3. (a) If $d = 1$ then $\boldsymbol{\omega} = \log(\sqrt{\boldsymbol{\Sigma}})$.
(b) If $d > 1$ then

$$\boldsymbol{\omega} = \begin{bmatrix} \log(\sqrt{\text{diagonal}(\boldsymbol{\Sigma})}) \\ \tanh^{-1} \left(\text{vecbd}(\boldsymbol{\Sigma}) / \sqrt{\text{vecbd}(\text{diagonal}(\boldsymbol{\Sigma}) \text{diagonal}(\boldsymbol{\Sigma})^T)} \right) \end{bmatrix}.$$

S.4.3 Transformation from $\boldsymbol{\omega}$ to $\boldsymbol{\theta}$

Given a $\frac{1}{2}d(d+1) \times 1$ vector $\boldsymbol{\omega}$, its corresponding $\boldsymbol{\theta}$ vector of the same length is found via the steps:

1. Form the $d \times d$ symmetric matrix $\boldsymbol{\Sigma}$ as follows:
 - (a) If $d = 1$ then $\boldsymbol{\Sigma} = \exp(2\boldsymbol{\omega})$.
 - (b) If $d > 1$ then let $\boldsymbol{\omega}_1$ denote the first d entries of $\boldsymbol{\omega}$ and $\boldsymbol{\omega}_2$ denote the remaining $\frac{1}{2}d(d-1)$ entries of $\boldsymbol{\omega}$.
 - i. Set $\text{diagonal}(\boldsymbol{\Sigma}) = \exp(2\boldsymbol{\omega}_1)$
 - ii. Obtain the below-diagonal entries of $\boldsymbol{\Sigma}$ so that

$$\text{vecbd}(\boldsymbol{\Sigma}) = \tanh(\boldsymbol{\omega}_2) \odot \text{vecbd}(\exp(\boldsymbol{\omega}_1) \exp(\boldsymbol{\omega}_1)^T)$$

holds. Obtain the above-diagonal entries of $\boldsymbol{\Sigma}$ such that symmetry of $\boldsymbol{\Sigma}$ is enforced.

2. Obtain the spectral decomposition: $\boldsymbol{\Sigma} = \mathbf{U}_{\boldsymbol{\Sigma}} \text{diag}(\boldsymbol{\lambda}_{\boldsymbol{\Sigma}}) \mathbf{U}_{\boldsymbol{\Sigma}}^T$.
3. Obtain $\boldsymbol{\theta} = \text{vech} \left(\frac{1}{2} \mathbf{U}_{\boldsymbol{\Sigma}} \text{diag}\{\log(\boldsymbol{\lambda}_{\boldsymbol{\Sigma}})\} \mathbf{U}_{\boldsymbol{\Sigma}}^T \right)$.

S.4.4 Quasi-Newton Optimization-Based Confidence Interval Calculations

The steps for obtaining confidence intervals for each of the interpretable parameters are:

1. Obtain $(\hat{\beta}, \hat{\theta})$ using a quasi-Newton optimization routine applied the expectation propagation-approximate log-likelihood $\underline{\ell}$ with unconstrained input parameters (β, θ) .
2. Obtain $\hat{\omega}$ corresponding to $\hat{\theta}$ using the steps given in Section S.4.2.
3. Call the quasi-Newton optimization routine with input parameters (β, ω) instead of (β, θ) , and initial value $(\hat{\beta}, \hat{\omega})$. In this call, request that the Hessian matrix $H\ell(\beta, \omega)$ at the maximum $(\hat{\beta}, \hat{\omega})$ be computed. The steps given in Section S.4.3 are used to obtain the corresponding (β, θ) vector for evaluation of $\underline{\ell}$ via the version of $\underline{\ell}$ used in 1. for the optimization.
4. Form $100(1 - \alpha)\%$ confidence intervals for the entries of (β, ω) using

$$\begin{bmatrix} \hat{\beta} \\ \hat{\omega} \end{bmatrix} \pm \Phi^{-1}(1 - \frac{1}{2}\alpha) \sqrt{-\text{diagonal}(\{H\ell(\hat{\beta}, \hat{\omega})\}^{-1})}.$$

5. Transform the confidence intervals limits for the ω component, using the functions \exp and \tanh , to instead correspond to the standard deviation and correlation parameters:

$$\begin{bmatrix} \sqrt{\text{diagonal}(\Sigma)} \\ \text{vecbd}(\Sigma) / \sqrt{\text{vecbd}(\text{diagonal}(\Sigma)\text{diagonal}(\Sigma)^T)} \end{bmatrix}.$$

S.5 Details of Approximate Best Prediction

For the binary mixed model (1), the best prediction of \mathbf{u}_i is

$$\begin{aligned} \text{BP}(\mathbf{u}_i) &= E(\mathbf{u}_i | \mathbf{y}) = E(\mathbf{u}_i | \mathbf{y}_i) = \int_{\mathbb{R}^{dR}} \mathbf{u}_i p(\mathbf{u}_i | \mathbf{y}_i; \beta, \Sigma) d\mathbf{u}_i \\ &= \int_{\mathbb{R}^{dR}} \mathbf{u}_i \left\{ \frac{p(\mathbf{y}_i | \mathbf{u}_i; \beta) p(\mathbf{u}_i; \Sigma)}{\int_{\mathbb{R}^{dR}} p(\mathbf{y}_i | \mathbf{u}_i; \beta) p(\mathbf{u}_i; \Sigma)} \right\} d\mathbf{u}_i \end{aligned}$$

where $\mathbf{y}_i \equiv (y_{i1}, \dots, y_{in_i})$. Now note that Algorithm 1 involves replacement of

$$p(\mathbf{y}_i | \mathbf{u}_i; \beta) p(\mathbf{u}_i; \Sigma) \quad \text{by} \quad \exp \left\{ \begin{bmatrix} 1 \\ \mathbf{u}_i \\ \text{vech}(\mathbf{u}_i \mathbf{u}_i^T) \end{bmatrix}^T \hat{\boldsymbol{\eta}}_i \right\}$$

where $\hat{\boldsymbol{\eta}}_i$ is defined by (16). This leads to the approximation

$$\begin{aligned} \text{BP}(\mathbf{u}_i) &= E(\hat{\mathbf{u}}_i) \quad \text{where } \hat{\mathbf{u}}_i \text{ is Multivariate Normal with natural parameter } \hat{\boldsymbol{\eta}}_i \\ &= -\frac{1}{2} \left\{ \text{vec}^{-1} \left(\mathbf{D}_d^{+T} \hat{\boldsymbol{\eta}}_{i2} \right) \right\}^{-1} \hat{\boldsymbol{\eta}}_{i1}. \end{aligned}$$

Using (13.7) of McCulloch, Searle & Neuhaus (2008), the covariance matrix of $\text{BP}(\mathbf{u}_i) - \mathbf{u}_i$ is

$$\text{Cov}\{\text{BP}(\mathbf{u}_i) - \mathbf{u}_i\} = E_{\mathbf{y}_i}\{\text{Cov}(\mathbf{u}_i|\mathbf{y}_i)\}.$$

The expectation propagation approximation of $\text{Cov}(\mathbf{u}_i|\mathbf{y}_i)$ is

$$\tilde{\text{Cov}}(\mathbf{u}_i|\mathbf{y}) = -\frac{1}{2}\left\{\text{vec}^{-1}\left(\mathbf{D}_d^{+T}\hat{\tilde{\boldsymbol{\eta}}}_{i2}\right)\right\}^{-1}.$$

However, approximation of $\text{Cov}\{\text{BP}(\mathbf{u}_i) - \mathbf{u}_i\}$ is hindered by the expectation over the \mathbf{y}_i vector.