

Efficient computation of high-dimensional penalized generalized linear mixed models by latent factor modeling of the random effects

Hillary M. Heiling^{1,*}, Naim U. Rashid¹, Quefeng Li¹, Xianlu L. Peng², Jen Jen Yeh^{2,3,4}, Joseph G. Ibrahim¹

¹Department of Biostatistics, University of North Carolina Chapel Hill, Chapel Hill, NC 27599, United States, ²Lineberger Comprehensive Cancer Center, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599, United States, ³Department of Surgery, University of North Carolina Chapel Hill, Chapel Hill, NC 27599, United States, ⁴Department of Pharmacology, University of North Carolina Chapel Hill, Chapel Hill, NC 27599, United States

*Corresponding author: Hillary M. Heiling, Department of Biostatistics, University of North Carolina Chapel Hill, Chapel Hill, NC 27599, USA (hmheiling@gmail.com).

ABSTRACT

Modern biomedical datasets are increasingly high-dimensional and exhibit complex correlation structures. Generalized linear mixed models (GLMMs) have long been employed to account for such dependencies. However, proper specification of the fixed and random effects in GLMMs is increasingly difficult in high dimensions, and computational complexity grows with increasing dimension of the random effects. We present a novel reformulation of the GLMM using a factor model decomposition of the random effects, enabling scalable computation of GLMMs in high dimensions by reducing the latent space from a large number of random effects to a smaller set of latent factors. We also extend our prior work to estimate model parameters using a modified Monte Carlo Expectation Conditional Minimization algorithm, allowing us to perform variable selection on both the fixed and random effects simultaneously. We show through simulation that through this factor model decomposition, our method can fit high-dimensional penalized GLMMs faster than comparable methods and more easily scale to larger dimensions not previously seen in existing approaches.

KEYWORDS: factor model decomposition; generalized linear mixed models; variable selection.

1 INTRODUCTION

Modern biomedical datasets are increasingly high-dimensional and exhibit complex correlation structures. Generalized Linear Mixed Models (GLMMs) have long been employed to account for such dependencies. However, proper specification of the fixed and random effects is a critical step in estimation of GLMMs. For instance, omitting important random effects can lead to bias in the estimated variance of the fixed effects (Gurka et al., 2011; Bondell et al., 2010), whereas including unnecessary random effects could increase the computational difficulty of fitting the GLMM. Despite the importance of properly specifying the set of fixed and random effects in such models, it is often unknown *a priori* which variables should be specified as fixed or random in the model, particularly in high-dimensional settings in which the feature space of both the fixed and random effects is generally assumed to be sparse.

There are several existing methods used to select fixed and/or random effects within mixed models. One such class of methods are what we term here as candidate model selection. In these methods, researchers manually specify a set of candidate models, use existing software such as R packages `lme4` (Bates et al., 2015) or `MCMCglmm` (Hadfield, 2010) to fit the candidate

models, and then select the best one using appropriate mixed effects model selection criteria such as the BIC-ICQ criterion (Ibrahim et al., 2011), the hybrid Bayesian information criterion BIC_h (Delattre et al., 2014), or fence methods (Jiang, 2014). However, candidate model selection approaches are not feasible for moderate or large dimensions as there are 2^{2^p} possible combinations of fixed and random effects for p predictors of interest.

Penalized likelihood techniques, such as LASSO, have been expanded to apply to mixed models. However, most of these existing approaches have limitations in their applicability. Some methods only select fixed effects, such as the R packages `glmmLasso` (Groll and Tutz, 2014) and `glmmixedLASSO` (Schelldorfer et al., 2014), or only select random effects (Pan and Huang, 2014). Other methods that do select both fixed and random effects are generally limited in their scalability due to their computational burden in high dimensions. Examples include the adaptive LASSO method proposed by Ibrahim et al. (2011), which utilizes a computationally intensive expectation step in their Monte Carlo Expectation Minimization (MCEM) algorithm, and the regularized PQL method proposed by Hui et al. (2017), which has prohibitive initialization requirements in high dimensions and requires the calculation of an inverse

matrix with dimensions equal to the random effects remaining in the model.

Rashid et al. (2020) developed a penalization method that selects both fixed and random effects in larger dimensions not seen in previous approaches. The method developed by Rashid et al. (2020), which uses a Monte Carlo Expectation Conditional Minimization (MCECM) algorithm, was applied to simulations and a case study comprised of 50 covariates and has since been incorporated into the `glmmPen` R package (Heiling et al., 2024). Although this `glmmPen` framework extends the feasible dimensionality of performing variable selection within GLMMs relative to existing methods, new methodology is needed to alleviate the computational burden as the dimension increases even further.

We present a novel reformulation of the GLMM, which we call `glmmPen_FA`, using a factor model decomposition of the random effects. This factor model is used as a dimension reduction tool to represent a large number of latent random effects as a function of a smaller set of latent factors. By reducing the latent space of the random effects, this new model formulation enables us to extend the feasible dimensionality of performing variable selection in GLMMs to hundreds of predictors. We estimate model parameters and perform simultaneous selection of fixed and random effects using an MCECM algorithm. We show through simulations that through this factor model decomposition, our method can fit high-dimensional penalized GLMMs (pGLMMs) faster than comparable methods and more easily scale to larger dimensions not previously seen in existing approaches.

We illustrate the utility of our method by applying our method to a case study that we present in Section 4. The case study data combines gene expression data from pancreatic ductal adenocarcinoma patients across 5 separate studies. We aim to select important features that predict the pancreatic cancer subtypes basal or classical (Moffitt et al., 2015) (ie, identify nonzero fixed effects) as well as identify features that have varied effects across the groups (ie, identify nonzero random effects). Due to the large number of features that we consider, it is difficult to have *a priori* knowledge of which features have nonzero fixed or random effects. Therefore, we will use our method to fit a penalized logistic mixed effects model to select important fixed and random effects.

The remainder of this paper is organized as follows: Section 2 reviews the statistical models and algorithm used to estimate pGLMMs in our new factor model decomposition framework, termed `glmmPen_FA`. In Section 3, simulations are conducted to assess the performance of the new `glmmPen_FA` method. Section 4 describes the motivating case study, where we aim to utilize our method to identify gene expression features important in the prediction of pancreatic cancer subtypes. We close the article with some discussion in Section 5.

2 METHODS

2.1 Model formulation

In this section, we review the notation and model formulation of our approach. We consider the case where we want to analyze

data from K independent groups of observations. For each group $k = 1, \dots, K$, there are n_k observations for a total sample size of $N = \sum_{k=1}^K n_k$. For group k , let $\mathbf{y}_k = (y_{k1}, \dots, y_{kn_k})^T$ be the vector of n_k independent responses, $\mathbf{x}_{ki} = (x_{ki,1}, \dots, x_{ki,p})^T$ be the p -dimensional vector of predictors, and $\mathbf{X}_k = (\mathbf{x}_{k1}, \dots, \mathbf{x}_{kn_k})^T$. For simplification of notation, we will set $n_1 = \dots = n_K = n$ without loss of generality. In GLMMs, we assume that the conditional distribution of \mathbf{y}_k given \mathbf{X}_k belongs to the exponential family and has the following density:

$$f(\mathbf{y}_k | \mathbf{X}_k, \boldsymbol{\alpha}_k; \theta) = \prod_{i=1}^n c(y_{ki}) \exp[\tau^{-1}\{y_{ki}\eta_{ki} - b(\eta_{ki})\}], \quad (1)$$

where $c(y_{ki})$ is a constant that only depends on y_{ki} , τ is the dispersion parameter, $b(\cdot)$ is a known link function, η_{ki} is the linear predictor, $\boldsymbol{\alpha}_k$ are group-specific latent variables, and $\boldsymbol{\theta}$ represents all model coefficients (see more detailed descriptions later in this section). We standardize the fixed effects covariates matrix $\mathbf{X} = (\mathbf{X}_1^T, \dots, \mathbf{X}_K^T)^T$ such that

$$\sum_{k=1}^K \sum_{i=1}^{n_k} x_{ki,j} = 0 \text{ and } N^{-1} \sum_{k=1}^K \sum_{i=1}^{n_k} x_{ki,j}^2 = 1 \text{ for } j = 1, \dots, p.$$

As outlined in Rashid et al. (2020), the traditional GLMM formulation of the linear predictor can be represented as:

$$\eta_{ki} = \mathbf{x}_{ki}^T \boldsymbol{\beta} + \mathbf{z}_{ki}^T \boldsymbol{\gamma}_k = \mathbf{x}_{ki}^T \boldsymbol{\beta} + \mathbf{z}_{ki}^T \boldsymbol{\Gamma} \boldsymbol{\delta}_k, \quad (2)$$

where $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T$ is a p -dimensional vector for the fixed effects coefficients (including the intercept), \mathbf{z}_{ki} is a q -dimensional subvector of \mathbf{x}_{ki} representing the random effect predictors (including the random intercept), $\boldsymbol{\Gamma}$ is the Cholesky decomposition of the random effects covariance matrix $\boldsymbol{\Sigma}$ such that $\boldsymbol{\Gamma} \boldsymbol{\Gamma}^T = \boldsymbol{\Sigma}$, and $\boldsymbol{\gamma}_k = \boldsymbol{\Gamma} \boldsymbol{\delta}_k$ is a q -dimensional vector of unobservable random effects for group k where $\boldsymbol{\delta}_k \sim N_q(\mathbf{0}, \mathbf{I})$.

In its current representation, the model assumes a latent space of dimension q , the number of random effect predictors. When q is large, the estimation of the covariance matrix $\boldsymbol{\Sigma} = \text{Var}(\boldsymbol{\gamma}_k)$ can be computationally burdensome to compute due to both the number of parameters needed to estimate this matrix ($q(q+1)/2$ parameters are needed for an unstructured covariance matrix) as well as the need to approximate a q -dimensional integral (see Section 2.3 for details). Prior work such as Fan et al. (2013) and Tran et al. (2020) have assumed a factor model structure in order to estimate high-dimensional covariance matrices in other settings, such as the estimation of sample covariance matrices for time series data and the covariance matrix in variational inference used to approximate the posterior distribution, respectively. Here, we introduce a novel formulation of the GLMM where we decompose the random effects $\boldsymbol{\gamma}_k$ into a factor model with r latent common factors ($r \ll q$) such that $\boldsymbol{\gamma}_k = \mathbf{B} \boldsymbol{\alpha}_k$, where \mathbf{B} is the $q \times r$ loading matrix and $\boldsymbol{\alpha}_k$ represents the r latent common factors. We assume the latent factors $\boldsymbol{\alpha}_k$ are uncorrelated and follow a $N_r(\mathbf{0}, \mathbf{I})$ distribution. We re-write the linear predictor as:

$$\eta_{ki} = \mathbf{x}_{ki}^T \boldsymbol{\beta} + \mathbf{z}_{ki}^T \mathbf{B} \boldsymbol{\alpha}_k. \quad (3)$$

In this representation, the random component of the linear predictor has variance $\text{Var}(\mathbf{B} \boldsymbol{\alpha}_k) = \mathbf{B} \mathbf{B}^T = \boldsymbol{\Sigma}$. By assuming that $\boldsymbol{\Sigma}$ is low rank, we also reduce the dimension of the latent space from q to r , which reduces the dimension of the integral in the likeli-

hood and thereby reduces the computational complexity of the Expectation step (E-step) in the EM algorithm. Further details are given in Section 2.3.

In order to estimate \mathbf{B} , let $\mathbf{b}_t \in \mathbb{R}^r$ be the t -th row of \mathbf{B} and $\mathbf{b} = (\mathbf{b}_1^T, \dots, \mathbf{b}_q^T)^T$. We can then reparameterize the linear predictor as:

$$\eta_{ki} = \mathbf{x}_{ki}^T \boldsymbol{\beta} + \mathbf{z}_{ki}^T \mathbf{B} \boldsymbol{\alpha}_k = (\mathbf{x}_{ki}^T \quad (\boldsymbol{\alpha}_k \otimes \mathbf{z}_{ki})^T \mathbf{J}) \begin{pmatrix} \boldsymbol{\beta} \\ \mathbf{b} \end{pmatrix} \quad (4)$$

in a manner similar to Chen and Dunson (2003) and Ibrahim et al. (2011), where \mathbf{J} is a matrix that transforms \mathbf{b} to $\text{vec}(\mathbf{B})$ such that $\text{vec}(\mathbf{B}) = \mathbf{J} \mathbf{b}$ and \mathbf{J} is of dimension $(qr) \times (qr)$. The vector of parameters $\boldsymbol{\theta} = (\boldsymbol{\beta}^T, \mathbf{b}^T, \boldsymbol{\tau})^T$ are the main parameters of interest. We denote the true value of $\boldsymbol{\theta}$ as $\boldsymbol{\theta}^* = (\boldsymbol{\beta}^{*T}, \mathbf{b}^{*T}, \boldsymbol{\tau}^*)^T = \text{argmin}_{\boldsymbol{\theta}} E_{\boldsymbol{\theta}}[-\ell(\boldsymbol{\theta})]$ where $\ell(\boldsymbol{\theta})$ is the observed log-likelihood across all K groups such that $\ell(\boldsymbol{\theta}) = \sum_{k=1}^K \ell_k(\boldsymbol{\theta})$, where $\ell_k(\boldsymbol{\theta}) = (1/n) \log \int f(\mathbf{y}_k | \mathbf{X}_k, \boldsymbol{\alpha}_k; \boldsymbol{\theta}) \phi(\boldsymbol{\alpha}_k) d\boldsymbol{\alpha}_k$.

Our main interest lies in selecting the true nonzero fixed and random effects. In other words, we aim to identify the set $S = S_1 \cup S_2 = \{j : \beta_j^* \neq 0\} \cup \{t : \|\mathbf{b}_t^*\|_2 \neq 0\}$, where S_1 and S_2 represent the true fixed and random effects, respectively. When $\mathbf{b}_t = \mathbf{0}$, this indicates that the effect of covariate t is fixed across the K groups (ie, the corresponding t -th row and column of $\boldsymbol{\Sigma}$ is set to $\mathbf{0}$).

We aim to solve the following penalized likelihood problem:

$$\hat{\boldsymbol{\theta}} = \text{argmin}_{\boldsymbol{\theta}} -\ell(\boldsymbol{\theta}) + \lambda_0 \sum_{j=1}^p \rho_0(\beta_j) + \lambda_1 \sum_{t=1}^q \rho_1(\|\mathbf{b}_t\|_2), \quad (5)$$

where $\ell(\boldsymbol{\theta})$ is the observed log-likelihood for all K groups defined earlier, $\rho_0(t)$ and $\rho_1(t)$ are general folded-concave penalty functions, and λ_0 and λ_1 are positive tuning parameters. The $\rho_0(t)$ penalty functions could include the L_1 penalty (aka the Least Absolute Shrinkage and Selection Operator (LASSO) penalty), the Minimax Concave Penalty (MCP), and the Smoothly Clipped Absolute Deviation (SCAD) penalty (Friedman et al., 2010; Breheny and Huang, 2011). For the $\rho_1(t)$ penalty, we treat the elements of \mathbf{b}_t as a group and penalize them in a groupwise manner using the group LASSO, group MCP, or group SCAD penalties presented by Breheny and Huang (2015)). These groups of \mathbf{b}_t are then estimated to be either all zero or all nonzero. In this way, we select covariates to have varying effects ($\hat{\mathbf{b}}_t \neq \mathbf{0}$) or fixed effects ($\hat{\mathbf{b}}_t = \mathbf{0}$) across the K groups.

2.2 MCECM algorithm

We solve (5) for some specific penalty combination (λ_0, λ_1) using a MCECM algorithm (Garcia et al., 2010).

In the s^{th} iteration of the MCECM algorithm, we aim to evaluate the expectation of (E-step) and minimize (M-step) the fol-

lowing penalized Q-function:

$$\begin{aligned} Q_{\lambda}(\boldsymbol{\theta} | \boldsymbol{\theta}^{(s)}) &= \sum_{k=1}^K E \left\{ -\log(f(\mathbf{y}_k, \mathbf{X}_k, \boldsymbol{\alpha}_k; \boldsymbol{\theta} | \mathbf{d}_o; \boldsymbol{\theta}^{(s)})) \right\} \\ &\quad + \lambda_0 \sum_{j=1}^p \rho_0(\beta_j) + \lambda_1 \sum_{t=1}^q \rho_1(\|\mathbf{b}_t\|_2) \\ &= Q_1(\boldsymbol{\theta} | \boldsymbol{\theta}^{(s)}) + Q_2(\boldsymbol{\theta}^{(s)}) + \lambda_0 \sum_{j=1}^p \rho_0(\beta_j) \\ &\quad + \lambda_1 \sum_{t=1}^q \rho_1(\|\mathbf{b}_t\|_2), \end{aligned} \quad (6)$$

where $(\mathbf{y}_k, \mathbf{X}_k, \boldsymbol{\alpha}_k)$ gives the complete data for group k , $\mathbf{d}_{k,o} = (\mathbf{y}_k, \mathbf{X}_k)$ gives the observed data for group k , \mathbf{d}_o represents the entirety of the observed data, and

$$Q_1(\boldsymbol{\theta} | \boldsymbol{\theta}^{(s)}) = - \sum_{k=1}^K \int \log[f(\mathbf{y}_k | \mathbf{X}_k, \boldsymbol{\alpha}_k; \boldsymbol{\theta})] \phi(\boldsymbol{\alpha}_k | \mathbf{d}_{k,o}; \boldsymbol{\theta}^{(s)}) d\boldsymbol{\alpha}_k, \quad (7)$$

$$Q_2(\boldsymbol{\theta}^{(s)}) = - \sum_{k=1}^K \int \log[\phi(\boldsymbol{\alpha}_k)] \phi(\boldsymbol{\alpha}_k | \mathbf{d}_{k,o}; \boldsymbol{\theta}^{(s)}) d\boldsymbol{\alpha}_k. \quad (8)$$

2.2.1 Monte-Carlo E-step

The E-step of the algorithm aims to approximate the r -dimensional integral expressed in (7). The integrals in the Q-function do not have closed forms when $f(\mathbf{y}_k | \mathbf{X}_k, \boldsymbol{\alpha}_k^{(s,m)}; \boldsymbol{\theta})$ is assumed to be non-Gaussian. We approximate these integrals using a Markov Chain Monte Carlo (MCMC) sample of size M from the posterior density $\phi(\boldsymbol{\alpha}_k | \mathbf{d}_{k,o}; \boldsymbol{\theta}^{(s)})$. Let $\boldsymbol{\alpha}_k^{(s,m)}$ be the m^{th} simulated r -dimensional vector from the posterior of the latent common factors, $m = 1, \dots, M$, at the s^{th} iteration of the algorithm for group k . The integral in (7) can be approximated as:

$$Q_1(\boldsymbol{\theta} | \boldsymbol{\theta}^{(s)}) \approx - \frac{1}{M} \sum_{m=1}^M \sum_{k=1}^K \log f(\mathbf{y}_k | \mathbf{X}_k, \boldsymbol{\alpha}_k^{(s,m)}; \boldsymbol{\theta}). \quad (9)$$

We use the fast and efficient No-U-Turn Hamiltonian Monte Carlo (NUTS HMC) sampling procedure from the Stan software (Carpenter et al. 2017) in order to perform the E-step efficiently.

2.2.2 M-step

In the M-step of the algorithm, we aim to minimize

$$\begin{aligned} Q_{1,\lambda}(\boldsymbol{\theta} | \boldsymbol{\theta}^{(s)}) &= Q_1(\boldsymbol{\theta} | \boldsymbol{\theta}^{(s)}) + \lambda_0 \sum_{j=1}^p \rho_0(\beta_j) \\ &\quad + \lambda_1 \sum_{t=1}^q \rho_1(\|\mathbf{b}_t\|_2), \end{aligned} \quad (10)$$

with respect to $\boldsymbol{\theta} = (\boldsymbol{\beta}^T, \mathbf{b}^T, \boldsymbol{\tau})^T$. We do this by using a Majorization-Minimization algorithm with penalties applied to the fixed effects and the rows of \mathbf{B} . The M-step of the s^{th} iteration of the MCECM algorithm proceeds as described in Algorithm 1 given in Web Appendix Section 1.1.

2.2.3 MCECM algorithm

Algorithm 2 in [Web Appendix Section 1.1](#) describes the full MCECM algorithm for estimating the parameters with a particular penalty combination (λ_0, λ_1) . The process of model selection and finding optimal tuning parameters are described in full in the [Web Appendix](#) (see Sections 1.2 and 1.3). Briefly, we identify a maximum penalty that penalizes all fixed effects to zero using techniques from the `ncvreg` R package (Breheny and Huang, 2011). We calculate a sequence of penalties from a small proportion of the max penalty (the minimum penalty) to the max penalty that are equidistant on the log scale, and apply this sequence to both the fixed and random effects in a two-stage model selection approach. For further details on initialization and convergence, see [Web Appendix Section 1.4](#).

2.3 Advantages of `glmmPen_FA` model formulation

There are several advantages to our proposed factor model decomposition of the random effects. By representing the random effects with a factor model, we reduce the latent space from a high dimension of q (the number of candidate random effect predictors) to r . In the more traditional GLMM model formulation, $Q_1(\boldsymbol{\theta}|\boldsymbol{\theta}^{(s)})$ would be represented as:

$$Q_1(\boldsymbol{\theta}|\boldsymbol{\theta}^{(s)}) = - \sum_{k=1}^K \int \log[f(y_k|\mathbf{X}_k, \boldsymbol{\delta}_k; \boldsymbol{\theta})] \phi(\boldsymbol{\delta}_k|\mathbf{d}_{k,0}; \boldsymbol{\theta}^{(s)}) d\boldsymbol{\delta}_k, \quad (11)$$

where $\boldsymbol{\theta}$ includes the fixed effects $\boldsymbol{\beta}$, the nonzero elements of $\boldsymbol{\Gamma}$ given in (2), and τ , and the $\boldsymbol{\delta}_k$ are q -dimensional latent variables. However, by using the novel model formulation given in (3), this changes the integral of interest such that now $Q_1(\boldsymbol{\theta}|\boldsymbol{\theta}^{(s)})$ expressed in (7) is of dimension $r \ll q$. This significantly reduces the computational complexity of estimating this integral in the E-step of the algorithm since we only have to estimate a latent space of dimension r . Consequently, this reduces the computational time. This also enables us to scale our method to hundreds of predictors since the practical dimension of the latent space will be much smaller than the total number of candidate random effects predictors. See [Web Appendix Section 1.6](#) for further discussion about the values of p , q , and r used in this paper.

Furthermore, this proposed formulation allows for more complex correlation structures in higher dimensions. In Rashid et al. (2020), the authors approximated the random effect covariance matrix $\boldsymbol{\Sigma}$ as a diagonal matrix when the dimensions are large as recommended by Fan and Li (2012). This approximation reduced the computational complexity of the algorithm and therefore increased the speed of the model fit. However, in our new formulation, we do not need to assume $\boldsymbol{\Sigma}$ is a diagonal matrix when the dimension is high.

2.4 Estimation of the number of latent factors

Performing our proposed `glmmPen_FA` method requires specifying the number of latent factors r . Since r is typically unknown *a priori*, this value needs to be estimated. We estimate r at the very beginning of the algorithm, and then use this estimated value in all later model estimations during the variable selection procedure.

There have been several proposed methods of estimating r for the approximate factor model. We tried the Eigenvalue Ratio method and Growth Ratio (GR) method developed by Ahn and

Horenstein (2013) as well as the method proposed in Bai and Ng (2002). We found that the GR method gave the most accurate estimates of r within our application. Therefore, in this section, we will describe how we implement the GR method to estimate r . The GR method is used in all of our numerical works.

To apply the GR method to our problem, we need a $q \times K$ matrix of observed random effects. Since we can never observe the random effects, we instead calculate pseudo random effects by first fitting a penalized generalized linear model with a small penalty to each group separately. We then take these group-specific estimates and center them so that all features have a mean of 0. Let these q -dimensional group-specific estimates be denoted as $\hat{\boldsymbol{y}}_k$ for each group $k = 1, \dots, K$. We then define $\mathbf{G} = (\hat{\boldsymbol{y}}_1, \dots, \hat{\boldsymbol{y}}_K)$ as the final $q \times K$ matrix of pseudo random effects.

Let $\psi_j(A)$ be the j -th largest eigenvalue of the positive semidefinite matrix A , and let $\tilde{\mu}_{qK,j} \equiv \psi_j(\mathbf{G}\mathbf{G}^T/(qK)) = \psi_j(\mathbf{G}^T\mathbf{G}/(qK))$.

To find the GR estimator, we first order the eigenvalues of $\mathbf{G}\mathbf{G}^T/(qK)$ from largest to smallest. Then, we calculate the following ratios:

$$GR(j) \equiv \frac{\log[V(j-1)/V(j)]}{\log[V(j)/V(j+1)]} = \frac{\log(1 + \tilde{\mu}_{qK,j}^*)}{\log(1 + \tilde{\mu}_{qK,j+1}^*)}, \quad (12)$$

for $j = 1, 2, \dots, U$, where $V(j) = \sum_{l=j+1}^{\min(q,K)} \tilde{\mu}_{qK,l}$, $\tilde{\mu}_{qK,j}^* = \tilde{\mu}_{qK,j}/V(j)$, and U is a pre-defined constant. Then, we estimate r by

$$\hat{r}_{GR} = \max_{1 \leq j \leq U} GR(j). \quad (13)$$

3 SIMULATIONS: VARIABLE SELECTION IN BINOMIAL DATA WITH 100 PREDICTORS

We examine the variable selection performance of the `glmmPen_FA` algorithm in high dimensions of $p = 100$ total predictors under several different conditions of the underlying data. In all of these simulations, we use a prescreening step to remove some random effects at the start of the algorithm in order to help improve the speed of the algorithm (see [Web Appendix Section 1.2](#) for details), the BIC-ICQ (Ibrahim et al., 2011) criterion for tuning parameter selection, the MCP penalty (MCP penalty for the fixed effects, group MCP penalty for the rows of the \mathbf{B} matrix), and the abbreviated two-stage grid search as described in the [Web Appendix Section 1.2](#) using the penalty sequences described in [Web Appendix Section 1.3](#). In order to determine the robustness of our variable selection procedure based on the assumed value of r , we fit models in one of 2 ways: we estimated the number of common factors r using the GR estimation procedure discussed in Section 2.4, or we input the true value of r for the algorithm to use. All simulation conditions used 100 replicates.

We simulated binary responses from a logistic mixed effects model with $p = 100$ predictors. Of p total predictors, we assume that the first 10 predictors have truly nonzero fixed and random effects, and the other $p - 10$ predictors have zero-valued fixed and random effects. We specified a full model as input for the algorithm such that the candidate random effect predictors equalled the candidate fixed effect predictors (ie, assumed $q =$

TABLE 1 Variable selection results for the $p = 100$ logistic mixed effects simulations, including TP percentages for fixed and random effects, FP percentages for fixed and random effects, the median time in hours for the algorithm to complete (T^{med}), and the average of the mean absolute deviation (Abs. Dev. (Mean)) between the fixed effect coefficient estimates β and the true β values across all simulation replicates. Column **B** describes the general size of both the variances and eigenvalues of the resulting $\Sigma = \mathbf{B}\mathbf{B}^T$ random effects covariance matrix (moderate vs large). Column “ r Est.” refers to the method used to specify r in the algorithm: the GR estimate or the true value of r . Column $\|\mathbf{D}\|_F$ represents the average across simulation replicates of the Frobenius norm of the difference (\mathbf{D}) between the estimated random effects covariance matrix $\hat{\Sigma}$ and the true random effects covariance matrix Σ ; the Frobenius norm was standardized by the number of true random effects selected in the model.

True r	β	B	r Est.	TP %		FP %		T^{med}	Abs. Dev. (Mean)	$\ \mathbf{D}\ _F$
				Fixef	FP % Fixef	Ranef	Ranef			
3	1	Mod.	GR	98.50	2.00	97.20	0.22	2.05	0.26	0.93
			True	99.00	2.14	98.40	0.16	2.36	0.26	0.93
	Large	GR	95.50	2.19	98.60	0.18	2.52	0.33	1.94	
		True	95.50	2.31	98.90	0.17	2.42	0.33	1.96	
	2	Mod.	GR	100.00	2.46	89.00	0.53	1.45	0.37	0.89
			True	100.00	2.78	90.10	0.50	2.07	0.31	0.92
Large	GR	100.00	3.39	94.60	0.80	2.39	0.43	1.78		
	True	100.00	3.40	96.20	0.49	2.41	0.41	1.60		
5	1	Mod.	GR	96.80	2.02	96.20	0.04	3.56	0.35	1.54
			True	96.70	1.86	96.80	0.03	3.60	0.35	1.59
	Large	GR	90.40	2.22	96.80	0.08	4.39	0.44	2.73	
		True	90.50	1.97	96.90	0.07	4.44	0.44	2.83	
	2	Mod.	GR	100.00	2.11	89.00	0.18	2.29	0.52	1.22
			True	100.00	2.42	88.40	0.24	2.99	0.44	1.33
Large	GR	99.90	3.28	93.10	0.50	3.03	0.57	2.26		
	True	99.90	3.36	93.40	0.47	3.98	0.55	2.29		

p), and our aim was to select the set of true predictors and random effects.

We set the sample size to $N = 2500$ and number of groups to $K = 25$, with an equal number of subjects per group. We set up the random effects covariance matrix by specifying a \mathbf{B} matrix with dimensions $(p + 1) \times r$, where $p + 1$ represents the p predictors plus the random intercept, and r represents the number of latent common factors with $r = \{3, 5\}$. Eleven of these $p + 1$ rows—corresponding to the true 10 predictors plus the intercept—had nonzero elements, while the remaining $p - 10$ rows were set to zero. For each value of r , we considered \mathbf{B} matrices that produced covariance matrices $\Sigma = \mathbf{B}\mathbf{B}^T$ with moderate variances and eigenvalues and large variances and eigenvalues (see [Web Appendix Section 2.1](#) for further details). These \mathbf{B} matrices are referred to as the “moderate” and “large” \mathbf{B} matrices, respectively. We use both moderate predictor effects and strong predictor effects, where all 10 of the true fixed effects have coefficient values of 1 or 2, respectively.

For group k , we generated the binary response y_{ki} , $i = 1, \dots, n_k$ such that $y_{ki} \sim \text{Bernoulli}(p_{ki})$ where $p_{ki} = P(y_{ki} = 1 | \mathbf{x}_{ki}, \mathbf{z}_{ki}, \boldsymbol{\gamma}_k, \boldsymbol{\theta}) = \exp(\mathbf{x}_{ki}^T \boldsymbol{\beta} + \mathbf{z}_{ki}^T \boldsymbol{\gamma}_k) / \{1 + \exp(\mathbf{x}_{ki}^T \boldsymbol{\beta} + \mathbf{z}_{ki}^T \boldsymbol{\gamma}_k)\}$, and $\boldsymbol{\gamma}_k \sim N_{11}(0, \mathbf{B}\mathbf{B}^T)$. Each condition was evaluated using 100 total simulated datasets.

For individual i in group k , the vector of predictors for the fixed effects was $\mathbf{x}_{ki} = (1, x_{ki,1}, \dots, x_{ki,p})^T$, and we set the random effects $\mathbf{z}_{ki} = \mathbf{x}_{ki}$, where $x_{ki,j} \sim N(0, 1)$ for $j = 1, \dots, p$, and each \mathbf{x}_j was standardized as described in [Section 2.1](#).

The results for these simulations are presented in [Tables 1 and 2](#). [Table 1](#) provides the average true positive (TP) rates (percent of true predictors selected) and false positive (FP) rates (percent of false predictors selected) for both the fixed and random effects variable selection, the median time in hours to complete the

variable selection procedure, the average of the mean absolute deviation between the fixed effect coefficient estimates and the true coefficients across all simulation replicates, and the average of the Frobenius norm of the difference between the estimated random effect covariance matrix $\hat{\Sigma} = \hat{\mathbf{B}}\hat{\mathbf{B}}^T$ and the true covariance matrix $\Sigma = \mathbf{B}\mathbf{B}^T$ (the Frobenius norm was standardized by the number of true random effects selected in the best model). [Table 2](#) gives the GR r estimation procedure results, including the average estimate of r and the proportion of times that the GR estimate of r was underestimated, correct, or overestimated. All simulations were completed on a Longleaf computing cluster (CPU Intel processors between 2.3Ghz and 2.5GHz).

We see from [Table 1](#) that the glmmPen_FA method is able to accurately select both the fixed and random effects across a variety of conditions, which is supported by the TPs generally being above 90% for both the fixed and random effects and the FPs generally being small: across all conditions, less than 3.5% for fixed effects and less than 1% for random effects.

We can see from [Table 2](#) that the GR estimation procedure applied to the pseudo random effect estimates described in [Section 2.4](#) has varying levels of accuracy depending on the structure of the underlying data. Generally, the GR estimation procedure becomes more accurate as the eigenvalues of the covariance matrix increase and the true predictor effects are moderate. We have found that the estimation of r generally improves when the sample size per group increases (simulations not shown) or when the total number of predictors used in the GR estimation procedure decreases (compare [Table 2](#) with [Web Appendix Sections 2.2 and 2.3](#)). Under conditions that reduce the accuracy of the GR procedure, the GR procedure underestimates r on average. However, when we compare the true and FP rates for the fixed and random effects given in [Table 1](#) between scenarios us-

TABLE 2 Results of the GR r estimation procedure for $p = 100$ logistic mixed effects simulation results, including the average estimate of r across simulations and percent of times that the estimation procedure underestimated r , gave the true r , or overestimated r . Column **B** describes the general size of both the variances and eigenvalues of the resulting $\Sigma = \mathbf{B}\mathbf{B}^T$ random effects covariance matrix (moderate vs large).

True r	β	B	Avg. r	r Underestimated %	r Correct %	r Overestimated %
3	1	Mod.	2.79	21	79	0
		Large	2.95	5	95	0
	2	Mod.	2.21	80	19	1
		Large	2.51	49	51	0
5	1	Mod.	4.60	26	72	2
		Large	4.83	15	83	2
	2	Mod.	3.83	70	28	2
		Large	4.43	46	50	4

ing the true r and those using the estimated r , we see very similar results, even in situations when the GR procedure tended to underestimate r . The misspecification of r does not significantly impact the estimation of the fixed effects coefficients (see the mean absolute deviation values) nor does it significantly impact the estimation of the random effect covariance matrix coefficients (see the Frobenius norm values).

We also used glmmPen to perform variable selection on the simulations where the true number of latent factors was $r = 3$. We let the glmmPen variable selection procedure proceed for 100 hours. In that time, glmmPen was able to complete the following number of replicates out of the 100 total replicates: 83 for ($\beta = 1$, $\mathbf{B} = \text{Moderate}$), 71 for ($\beta = 1$, $\mathbf{B} = \text{Large}$), 100 for ($\beta = 2$, $\mathbf{B} = \text{Moderate}$), and 96 for ($\beta = 2$, $\mathbf{B} = \text{Large}$). The minimum times needed to complete the glmmPen variable selection procedures were 39.91, 57.60, 23.63, and 42.79 hours, respectively.

The [Web Appendix Section 2](#) contains additional simulation results not included in the main paper due to space considerations. These additional simulations include simulations with $p = 500$ predictors, a comparison with glmmPen in moderate dimensions, and alternative data set-ups (eg, sample size, effect magnitudes, correlated predictors).

4 CASE STUDY: PANCREATIC DUCTAL ADENOCARCINOMA

Patients diagnosed with pancreatic ductal adenocarcinoma (PDAC) generally face a very poor prognosis, where the 5-year survival rate is 6% (Khorana et al. 2016). The study by Moffitt et al. (2015) identified genes that are expressed exclusively in pancreatic tumor cells. Using these tumor-specific genes, Moffitt et al. (2015) were able to identify and validate 2 novel tumor subtypes, termed “basal-like” and “classical”. It was found that patients diagnosed with basal-like tumors had significantly worse median survival than those diagnosed with the classical tumors. Consequently, it is of clinical interest to robustly predict this basal-like subtype in order to make and improve tailored treatment recommendations.

In order to improve replicability in the prediction of subtypes in PDAC, we combine PDAC gene expression data from 5 different studies (Aguirre et al. 2018; Cao et al. 2021; Dijk et al. 2020; Hayashi et al. 2020; Raphael et al. 2017) with a total sam-

ple size of 360 subjects; see [Web Appendix Table 9](#) for further details. In order to account and adjust for between-study heterogeneity, we apply our new method glmmPen_FA to fit a penalized logistic mixed effects model to our data to select predictors with study-replicable effects, where we assume that predictor effects may vary between studies.

The basal or classical subtype outcome was calculated using the clustering algorithm specified in Moffitt et al. (2015). Further details are provided in [Web Appendix Section 3.1](#), and the code for this procedure is provided in a GitHub repository, see [Supplementary Materials](#) for more details.

Moffitt et al. (2015) identified a list of 500 genes that were likely to be expressed solely in PDAC tumor cells. The 5 studies had RNA-seq gene expression data for 432 of these 500 genes. There were some significant correlations between some of these 432 genes, as evaluated by Spearman correlations applied to the subject-level rank-transformed gene expression; therefore, we combined highly correlated genes together into meta-genes. The clustering process used to create these meta-genes is described in [Web Appendix Section 3.1](#). The final dataset included 117 meta-genes. The raw gene expression values of each meta-gene, calculated as the sum of the gene expression across all of the genes comprising the meta-gene, were converted from their raw values to their ranks for each subject.

Due to the presence of several pairwise Spearman correlation values greater than 0.5 in this final dataset, we used the Elastic Net penalization procedure (Friedman et al. 2010) to balance between ridge regression and the MCP penalty. We let π represent the balance between ridge regression and the MCP penalty, where $\pi = 0$ represents ridge regression and $\pi = 1$ represents the MCP penalty. We restricted our consideration to $\pi = \{0.6, 0.7, 0.8, 0.9\}$ based on Elastic Net simulation results given in [Web Appendix Section 2.4](#); we then used sensitivity analyses (see [Web Appendix Section 3.2](#)) to choose the optimal π used in this case study variable selection analysis. Within each value of π , we considered the GR estimated value of r (evaluated at 2 for all values of π) and a manually set larger value of $r = 3$. We found that within values of π , the selection results and coefficient values were consistent for the different values of r considered; we therefore restricted our consideration to the GR estimate of r . We also found that the selected meta-genes within the final model were consistent across the values of π , with only small deviations between values of π . For the results

TABLE 3 Covariate meta-gene label within the case study dataset of the meta-genes that had nonzero fixed effects in the final best model, the sign of the fixed effect coefficient (ie, the sign of the log odds ratio) associated with the meta-gene, and the gene symbols of the genes within the meta-gene.

Meta-gene No.	Log Odds Ratio Sign	Gene Component Symbols for Meta-Gene
5	+	ADORA2B, C16orf74, HES2, ULBP2
7	+	AHNAK2, FAM83A, GJB6, ITGA3, IVL, KRT6A, MUC16, PPL, SCEL, SLC2A1, ZNF185
28	+	COL17A1, DHRS9, SPRR1B, SPRR3
52	+	KRT23, S100A4, SPINT2
81	+	BACE2, CEACAM3, RNF43
85	-	BTNL8, CDH17, MYO1A, MYO7B, PDZD3, VIL1
104	-	GATA6, PAQR8, PIP5K1B, TOX3
117	-	TFF1, VSIG2

reported here, we let $\pi = 0.8$, which provided the results that best reflected the conclusions from the overall sensitivity analyses. The same value of π was used for both the fixed effects and random effects penalization. The sequence of λ penalties used in the variable selection procedure was the same as those used in the Binomial variable selection simulations for $p = 100$ (see [Web Appendix Section 1.3](#)).

In the final results, 8 of the 117 total meta-gene covariates had nonzero fixed effect values in the best model selected by the BIC-ICQ criteria, implying these covariates were important for the prediction of the basal outcome. These 8 meta-gene covariates represented 37 genes in total. Table 3 includes the label for these 8 meta-genes, the sign of the associated fixed effect coefficient (ie, the log odds ratio estimate), and the gene symbols of the genes that make up the meta-gene. Meta-genes with positive log odds ratios indicate that having greater relative expression of these meta-genes increases the odds of a subject being in the basal subtype, and vice versa for negative log odds ratios. The best model contained a random intercept (variance value 0.54) and no other random slopes.

We also applied the glmmPen variable selection procedure to this data. Using $\pi = 0.8$, the 8 meta-gene covariates selected by glmmPen_FA were also selected by glmmPen. The glmmPen method selected 2 additional meta-genes (meta-gene 59, genes PKIB, DNAJC15, with negative log odds ratio; meta-gene 71, genes AKRIC3, CA2, MGST2, with positive log odds ratio) and selected meta-gene 117 to have a nonzero random effect (variance 0.99). The main difference in these variable selection procedures was the time to needed to complete the procedure, where glmmPen_FA finished within 0.8 hours and glmmPen finished within 49.4 hours. More details about glmmPen sensitivity analyses (ie, results for different values of π) are provided in [Web Appendix Section 3.2](#).

5 DISCUSSION

By adopting a factor model structure to estimate the high-dimensional random effect covariance matrix in the generalized linear mixed model setting, we are assuming that a small number of underlying latent variables (ie, latent factors) can fully describe the high-dimensional set of candidate random effects we consider in the model. The main benefit of this assumption is that we are able to reduce the latent space from a large number

of random effects to a smaller set of latent factors, thereby greatly simplifying the E-step of the algorithm. We have shown through simulations (both in Section 3 and in the [Supplementary simulations in Web Appendix Section 2](#)) that by reducing the complexity of the integral in the E-step, we can significantly improve the overall time needed to perform variable selection in high-dimensional GLMMs.

The simulations also show how reducing the latent space increases the feasible dimensionality of performing variable selection in GLMMs. By using our novel formulation of the random effects, we can perform variable selection on mixed models with hundreds of predictors within a reasonable time-frame without any *a priori* knowledge of which predictors are relevant for the model, either in terms of fixed or random effects. From the simulation results, we see that the glmmPen_FA method results in accurate selection of the fixed and random effects across several conditions.

One limitation of this assumption is that there may be applications where the random effects cannot be represented as a function of a relatively small set of latent factors. Our method provides the greatest computational benefit (ie, the greatest improvement in time) when the true value of the number of latent factors r is much smaller than the number of random effects considered by the variable selection procedure, q . If the value of r from the estimation procedure is large, or not much smaller than the number of random effects q , then our method has little computational advantage over the existing method of glmmPen.

Additionally, our method is limited by the need to provide an estimate for the number of latent common factors. However, the simulation results show that our data-driven estimation of the number of latent factors, based on the GR estimation procedure by Ahn and Horenstein (2013), provides reasonable estimates. Even when it was estimated incorrectly by this procedure, this misspecification had very little impact on the general variable selection performance or the coefficient estimates. Therefore, our method is not sensitive to the estimation of the number of latent factors.

SUPPLEMENTARY MATERIALS

Supplementary material is available at [Biometrics](#) online.

Web Appendices and Tables referenced in Sections 2, 3, and 4, Section 3 code and simulation output, and Section 4 data and code are available with this paper at the Biometrics website

on Oxford Academic and the GitHub repository https://github.com/hheiling/paper_glmPen_FA. The glmPen R package containing both the original glmPen formulation and the new glmPen_FA method is available for download through CRAN at <https://cran.r-project.org/package=glmPen>.

FUNDING

Research reported in this publication was supported the National Institutes of Health under the following award numbers: RO1 AG073259, U01 CA274298, P50 CA257911, P50 CA058223, and T32 CA106209.

CONFLICT OF INTEREST

None declared.

DATA AVAILABILITY

The data that support the findings in this paper are provided in the Supplementary Materials for this paper.

REFERENCES

- Aguirre, A. J., Nowak, J. A., Camarda, N. D., Moffitt, R. A., Ghazani, A. A., Hazar-Rethinam, M. et al. (2018). Real-time genomic characterization of advanced pancreatic cancer to enable precision medicine. *Cancer Discovery*, 8, 1096–1111.
- Ahn, S. C. and Horenstein, A. R. (2013). Eigenvalue ratio test for the number of factors. *Econometrica*, 81, 1203–1227.
- Bai, J. and Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, 70, 191–221.
- Bates, D., Mächler, M., Bolker, B. and Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67, 1–48.
- Bondell, H. D., Krishna, A. and Ghosh, S. K. (2010). Joint variable selection for fixed and random effects in linear mixed-effects models. *Biometrics*, 66, 1069–1077.
- Breheny, P. and Huang, J. (2011). Coordinate descent algorithms for non-convex penalized regression, with applications to biological feature selection. *Annals of Applied Statistics*, 5, 232–253.
- Breheny, P. and Huang, J. (2015). Group descent algorithms for nonconvex penalized linear and logistic regression models with grouped predictors. *Statistics and Computing*, 25, 173–187.
- Cao, L., Huang, C., Zhou, D. C., Hu, Y., Lih, T. M., Savage, S. R. et al. (2021). Proteogenomic characterization of pancreatic ductal adenocarcinoma. *Cell*, 184, 5031–5052.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M. et al. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76, 1–32.
- Chen, Z. and Dunson, D. B. (2003). Random effects selection in linear mixed models. *Biometrics*, 59, 762–769.
- Delattre, M., Lavielle, M. and Poursat, M.-A. (2014). A note on bic in mixed-effects models. *Electronic Journal of Statistics*, 8, 456–475.
- Dijk, F., Veenstra, V. L., Soer, E. C., Dings, M. P., Zhao, L., Halfwerk, J. B. et al. (2020). Unsupervised class discovery in pancreatic ductal adenocarcinoma reveals cell-intrinsic mesenchymal features and high concordance between existing classification systems. *Scientific Reports*, 10, 337.
- Fan, J., Liao, Y. and Mincheva, M. (2013). Large covariance estimation by thresholding principal orthogonal complements. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 75, 603–680.
- Fan, Y. and Li, R. (2012). Variable selection in linear mixed effects models. *Annals of Statistics*, 40, 2043.
- Friedman, J., Hastie, T. and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33, 1–22.
- Garcia, R. I., Ibrahim, J. G. and Zhu, H. (2010). Variable selection for regression models with missing data. *Statistica Sinica*, 20, 149.
- Groll, A. and Tutz, G. (2014). Variable selection for generalized linear mixed models by l1-penalized estimation. *Statistics and Computing*, 24, 137–154.
- Gurka, M. J., Edwards, L. J. and Muller, K. E. (2011). Avoiding bias in mixed model inference for fixed effects. *Statistics in Medicine*, 30, 2696–2707.
- Hadfield, J. D. (2010). Mcmc methods for multi-response generalized linear mixed models: The MCMCglmm R package. *Journal of Statistical Software*, 33, 1–22.
- Hayashi, A., Fan, J., Chen, R., Ho, Y.-J., Makohon-Moore, A. P., Lecomte, N. et al. (2020). A unifying paradigm for transcriptional heterogeneity and squamous features in pancreatic ductal adenocarcinoma. *Nature Cancer*, 1, 59–74.
- Heiling, H., Rashid, N., Li, Q. and Ibrahim, J. (2024). *glmPen: High Dimensional Penalized Generalized Linear Mixed Models (pGLMM)*, R package version 1.5.4.4.
- Hui, F. K., Müller, S. and Welsh, A. (2017). Joint selection in mixed models using regularized pql. *Journal of the American Statistical Association*, 112, 1323–1333.
- Ibrahim, J. G., Zhu, H., Garcia, R. I. and Guo, R. (2011). Fixed and random effects selection in mixed effects models. *Biometrics*, 67, 495–503.
- Jiang, J. (2014). The fence methods. *Advances in Statistics*, 2014, 1–14.
- Khorana, A. A., Mangu, P. B., Berlin, J., Engebretson, A., Hong, T. S., Maitra, A. et al. (2016). Potentially curable pancreatic cancer: American society of clinical oncology clinical practice guideline. *Journal of Clinical Oncology*, 34, 2541–2556.
- Moffitt, R. A., Marayati, R., Flate, E. L., Volmar, K. E., Herrera Loeza, S. G., Hoadley, K. A. et al. (2015). Virtual microdissection identifies distinct tumor-and stroma-specific subtypes of pancreatic ductal adenocarcinoma. *Nature Genetics*, 47, 1168.
- Pan, J. and Huang, C. (2014). Random effects selection in generalized linear mixed models via shrinkage penalty function. *Statistics and Computing*, 24, 725–738.
- Raphael, B. J., Hruban, R. H., Aguirre, A. J., Moffitt, R. A., Yeh, J. J., Stewart, C. et al. (2017). Integrated genomic characterization of pancreatic ductal adenocarcinoma. *Cancer Cell*, 32, 185–203.
- Rashid, N. U., Li, Q., Yeh, J. J. and Ibrahim, J. G. (2020). Modeling between-study heterogeneity for improved replicability in gene signature selection and clinical prediction. *Journal of the American Statistical Association*, 115, 1125–1138.
- Schelldorfer, J., Meier, L. and Bühlmann, P. (2014). Gmmlasso: an algorithm for high-dimensional generalized linear mixed models using l1-penalization. *Journal of Computational and Graphical Statistics*, 23, 460–477.
- Tran, M.-N., Nguyen, N., Nott, D. and Kohn, R. (2020). Bayesian deep net glm and glmm. *Journal of Computational and Graphical Statistics*, 29, 97–113.