

# Scalable Subset Selection in Linear Mixed Models

Ryan Thompson\*, Matt P. Wand, and Joanna J. J. Wang

*School of Mathematical and Physical Sciences, University of Technology Sydney*

May 14, 2026

## Abstract

Linear mixed models (LMMs), which incorporate fixed and random effects, are key tools for analyzing heterogeneous data, such as in personalized medicine. Nowadays, this type of data is increasingly wide, sometimes containing thousands of candidate predictors, necessitating sparsity for prediction and interpretation. However, existing sparse learning methods for LMMs do not scale well beyond tens or hundreds of predictors, leaving a large gap compared with sparse methods for linear models, which ignore random effects. This paper closes the gap with a new  $\ell_0$  regularized method for LMM subset selection that can run on datasets containing thousands of predictors in seconds to minutes. On the computational front, we develop a coordinate descent algorithm as our main workhorse and provide a guarantee of its convergence. We also develop a local search algorithm to help traverse the nonconvex optimization surface. Both algorithms readily extend to subset selection in generalized LMMs via a penalized quasi-likelihood approximation. On the statistical front, we provide a finite-sample bound on the Kullback–Leibler divergence of the new method. We then demonstrate its excellent performance in experiments involving synthetic and real datasets.

---

\*Corresponding author. Email: [ryan.thompson-1@uts.edu.au](mailto:ryan.thompson-1@uts.edu.au)

*Keywords:* coordinate descent, hierarchical sparsity, mixed modeling, penalized quasi-likelihood, variable selection

# 1 Introduction

Linear mixed models (LMMs) have a long and rich history in statistics, dating back at least to the 1950s (Henderson et al., 1959), and have more recently drawn attention in machine learning. The appeal of LMMs lies in their capacity to flexibly model heterogeneous effects in data containing multiple observations on clusters of related units. Examples include clinical studies with repeated measurements on the same patients at different times to capture patient-specific effects, or educational evaluations with repeated measurements on the same schools across different classes to isolate school-specific effects. The prevalence of data exhibiting this type of structure has motivated an enormous body of research into LMMs. We refer to Jiang and Nguyen (2021) and West et al. (2022) for recent and accessible treatments.

Given  $n_i$  observations on  $\mathbf{y}_i \in \mathbb{R}^{n_i}$ ,  $\mathbf{X}_i \in \mathbb{R}^{n_i \times p}$ , and  $\mathbf{Z}_i \in \mathbb{R}^{n_i \times d}$ , representing a response vector and predictor matrices for cluster  $i = 1, \dots, m$ , an LMM is typically written as

$$\mathbf{y}_i = \mathbf{X}_i \boldsymbol{\beta} + \mathbf{Z}_i \mathbf{u}_i + \boldsymbol{\varepsilon}_i, \quad i = 1, \dots, m,$$

where  $\boldsymbol{\varepsilon}_i \in \mathbb{R}^{n_i}$  is stochastic noise.<sup>1</sup> The coefficients  $\boldsymbol{\beta} \in \mathbb{R}^p$  are known as fixed effects, and the coefficients  $\mathbf{u}_i \in \mathbb{R}^d$  as random effects. The fixed effects are so-called because  $\boldsymbol{\beta}$  is constant across all clusters  $i$ , whereas the random effects  $\mathbf{u}_i$  vary across clusters according to a specified distribution, typically a Gaussian distribution  $\mathbf{u}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Gamma})$ . In modern data-analytic regimes where the predictor dimensionalities  $p$  and  $d$  are large, it is common to operate on the principle that  $\boldsymbol{\beta}$  and  $\mathbf{u}_i$  are sparse. The goal is then to learn if a predictor in  $\mathbf{X}_i$  or  $\mathbf{Z}_i$  has a nonzero effect. More generally, we wish to learn if a predictor has a zero, fixed, or random effect. In that case, which is our paper’s focus, one takes  $\mathbf{Z}_i = \mathbf{X}_i$ , giving

$$\mathbf{y}_i = \mathbf{X}_i (\boldsymbol{\beta} + \mathbf{u}_i) + \boldsymbol{\varepsilon}_i, \quad i = 1, \dots, m. \tag{1.1}$$

---

<sup>1</sup>To simplify presentation, we exclude intercept terms.

Consequently, unless the number of nonzero random effects is small, it is necessary to structure the covariance matrix to prevent the number of estimable parameters from growing quadratically. A standard and widely adopted structure, particularly suitable for modern, predictor-rich settings, is the diagonal matrix  $\mathbf{\Gamma} = \text{diag}(\boldsymbol{\gamma})$ , where  $\boldsymbol{\gamma} \in \mathbb{R}_+^p$  is a vector of random effect variances. Hence, the random effect  $u_{ik} = 0$  whenever its variance  $\gamma_k = 0$ .

Methods for fitting sparse LMMs broadly fit into two categories: (i) those that select either fixed effects or random effects but not both (Schelldorfer et al., 2011; Groll and Tutz, 2014; Pan and Huang, 2014; Schelldorfer et al., 2014; Kowal, 2023), and (ii) those that select both types of effect (Bondell et al., 2010; Ibrahim et al., 2011; Fan and Li, 2012; Peng and Lu, 2012; Lin et al., 2013; Hui et al., 2017*a,b*; Li et al., 2018; Heiling et al., 2023; Sholokhov et al., 2024). While methods in the first category sometimes scale to thousands of predictors (specifically those that select fixed effects), they assume either the set of relevant fixed effects or set of relevant random effects is known a priori—an assumption rarely satisfied in contemporary predictor-rich settings. The second category of methods addresses this limitation by selecting both types of effects simultaneously, but typically do not scale nearly as well and often do not enforce the hierarchical selection principle, whereby a random effect can only be nonzero if the corresponding fixed effect is also nonzero (an exception is Hui et al., 2017*a*). Moreover, apart from Kowal (2023) and Sholokhov et al. (2024), these approaches rely on continuous shrinkage regularizers such as the  $\ell_1$  (lasso) penalty (Tibshirani, 1996), smoothly clipped absolute deviation penalty (Fan and Li, 2001), or minimax concave penalty (Zhang, 2010), which have the side effect of biasing (shrinking) any estimated nonzero coefficients toward zero. Further still, many only support Gaussian likelihoods, restricting their application to continuous responses. Table 1 concisely summarizes these capabilities and limitations.

In parallel to the above line of work, research on sparse linear models (LMs) (i.e., models with fixed effects only) has yielded several remarkable advances over the past few years. In particular, tools now exist that can rapidly solve  $\ell_0$  regularized regression—also known as best subset selection—on datasets with thousands of predictors (Hazimeh and Mazumder, 2020;

	Fixed selection	Random selection	Hier. principle	$\ell_0$ subset selection	GLMMs	Unrestr. covariance	Max predictors
Bondell et al. (2010)	✓	✓	✗	✗	✗	✓	$\approx 10^1$
Ibrahim et al. (2011)	✓	✓	✗	✗	✓	✓	$\approx 10^1$
Schelldorfer et al. (2011)	✓	✗	-	✗	✗	✓	$\approx 10^3$
Fan and Li (2012)	✓	✓	✗	✗	✗	✗	$\approx 10^2$
Peng and Lu (2012)	✓	✓	✗	✗	✗	✓	$\approx 10^1$
Lin et al. (2013)	✓	✓	✗	✗	✗	✓	$\approx 10^1$
Groll and Tutz (2014)	✓	✗	-	✗	✓	✓	$\approx 10^2$
Pan and Huang (2014)	✗	✓	-	✗	✓	✓	$\approx 10^1$
Schelldorfer et al. (2014)	✓	✗	-	✗	✓	✓	$\approx 10^3$
Hui et al. (2017a)	✓	✓	✓	✗	✓	✓	$\approx 10^1$
Hui et al. (2017b)	✓	✓	✗	✗	✓	✓	$\approx 10^1$
Li et al. (2018)	✓	✓	✗	✗	✗	✓	$\approx 10^2$
Heiling et al. (2023)	✓	✓	✗	✗	✓	✓	$\approx 10^1$
Kowal (2023)	✓	✗	-	✓	✗	✓	$\approx 10^2$
Sholokhov et al. (2024)	✓	✓	✗	✓	✗	✗	$\approx 10^3$
Our proposal	✓	✓	✓	✓	✓	✗	$\approx 10^4$

Table 1: Capabilities and limitations of existing tools for fitting sparse LMMs. Hier. principle indicates whether the method can enforce that random effects are selected only when their corresponding fixed effects are selected. Unrestr. covariance indicates whether the method allows unrestricted random effect covariance structures. Max predictors is derived from the largest experiment reported in each paper, with the  $\approx$  symbol meaning “on the order of.”

Dedieu et al., 2021) despite the problem being NP-hard (Natarajan, 1995). Unlike continuous shrinkage regularizers,  $\ell_0$  regularization explicitly penalizes the number of selected predictors, yielding superior performance in selection, interpretation, and, provided the underlying signal is sufficiently strong, prediction (Bertsimas et al., 2016; Hastie et al., 2020). These computational breakthroughs rely on coordinate-wise methods, pathwise optimization, and heuristics inspired by early work in Friedman et al. (2007). Concurrently, these developments have also motivated Bayesian adaptations of  $\ell_0$  regularized regression (Kowal, 2022), which follow a different computational paradigm but similarly leverage discrete selection to yield sparse solutions. Recent work also explores structured variants of  $\ell_0$  subset selection, including grouped and hierarchical selection (Hazimeh et al., 2023; Thompson and Vahid, 2024).

In view of the preceding discussion, we introduce a new estimator that facilitates scalable

$\ell_0$  subset selection in LMMs. The new estimator is defined as a solution to the problem

$$\min_{\substack{\boldsymbol{\beta} \in \mathbb{R}^p, \boldsymbol{\gamma} \in \mathbb{R}_+^p \\ \beta_k = 0 \Rightarrow \gamma_k = 0}} l(\boldsymbol{\beta}, \boldsymbol{\gamma}; \mathbf{y}, \mathbf{X}) + \lambda\alpha\|\boldsymbol{\beta}\|_0 + \lambda(1 - \alpha)\|\boldsymbol{\gamma}\|_0, \quad (1.2)$$

where  $\mathbf{y} = (\mathbf{y}_1^\top, \dots, \mathbf{y}_m^\top)^\top$  and  $\mathbf{X} = (\mathbf{X}_1^\top, \dots, \mathbf{X}_m^\top)^\top$  are the full response vector and predictor matrix, respectively, and  $l$  is a negative log-likelihood. The  $\ell_0$  norm  $\|\cdot\|_0$  counts the number of nonzeros and thus induces sparsity in  $\boldsymbol{\beta}$  and  $\boldsymbol{\gamma}$ , with  $\lambda \geq 0$  and  $\alpha \in (0, 1]$  serving as tuning parameters. The constraint  $\beta_k = 0 \Rightarrow \gamma_k = 0$  means that if a fixed effect’s coefficient  $\beta_k$  is zero, the corresponding random effect’s variance  $\gamma_k$  must also be zero, thereby enforcing a hierarchical selection structure. The estimator (1.2) possesses several favorable properties not simultaneously achieved by any existing alternatives. First, it selects fixed *and* random effects so one need not specify any relevant effects a priori. Indeed, with datasets containing many predictors, it is usually unknown which predictors are relevant and, for those that are, whether their effect is fixed or random.<sup>2</sup> Second, it enforces the hierarchy that a predictor’s random effect can only be active if its fixed effect is also active. From a modeling perspective, this constraint encodes the principle that random effects are deviations from a baseline fixed effect. Third, it allows the fixed and random effects to enter the model with their full effect. In contrast, estimators that use continuous regularizers require post-processing to unwind shrinkage, whereas our approach directly produces unshrunk (unbiased) coefficients.

Inspired by the remarkable computational advances for subset selection in LMs, we develop new algorithmic machinery for subset selection in LMMs on datasets containing at least an order of magnitude more predictors than has previously been feasible. The core components of our algorithmic framework are a coordinate descent method and a local search method, the former for rapidly computing a solution and the latter for refining that solution towards a global minimum. These algorithms readily extend to generalized LMMs (GLMMs) via a penalized quasi-likelihood approximation. On the statistical side, we provide a guarantee for

---

<sup>2</sup>If prior information is available, it is straightforward to enforce certain effect types for specific predictors by introducing parameter-specific penalty factors.

the new estimator in the form of a high-probability bound on the Kullback–Leibler divergence from the true model. The bound does not rely on asymptotic arguments and holds for finite samples. We then evaluate the new estimator on synthetic datasets, stress-testing its scalability across various sample sizes, predictor dimensions, and correlation levels. We further demonstrate its utility on real datasets that require methods capable of handling heterogeneity, enforcing sparsity, and producing interpretable results.

The remaining sections of this paper are arranged as follows. Section 2 introduces the algorithmic framework. Section 3 describes the sparse GLMM extension. Section 4 derives the statistical guarantee. Section 5 presents the synthetic data experiments. Section 6 reports the real data experiments. Finally, Section 7 closes the paper with final remarks.

## 2 Algorithmic Framework

### 2.1 Objective Function

We begin with the Gaussian likelihood and defer discussion of other likelihoods to Section 3. The Gaussian (negative) log-likelihood, ignoring constants, takes the form

$$l(\boldsymbol{\beta}, \boldsymbol{\gamma}, \sigma^2; \mathbf{y}, \mathbf{X}) = \sum_{i=1}^m \log \det\{\sigma^2 \mathbf{V}_i(\boldsymbol{\gamma})\} + \frac{1}{\sigma^2} \sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}), \quad (2.1)$$

where  $\sigma^2 > 0$  is the noise variance and  $\mathbf{V}_i(\boldsymbol{\gamma})$  is a covariance matrix that captures the correlation between observations in the  $i$ th cluster as induced by the random effects:

$$\mathbf{V}_i(\boldsymbol{\gamma}) = \mathbf{I} + \mathbf{X}_i \text{diag}(\boldsymbol{\gamma}) \mathbf{X}_i^\top. \quad (2.2)$$

This restricted covariance structure is a key ingredient in the scalability of the proposed method because changing a single  $\gamma_k$  induces only a rank-one change in  $\mathbf{V}_i(\boldsymbol{\gamma})$ , which, as we will see, enables efficient updates of  $\mathbf{V}_i^{-1}(\boldsymbol{\gamma})$  and  $\log \det\{\mathbf{V}_i(\boldsymbol{\gamma})\}$ . The predictor matrix

$\mathbf{X}$  is assumed to be standardized to have columns with unit  $\ell_2$  norm. Whereas the noise variance in non-mixed LMs can be ignored when estimating regression coefficients, it cannot be ignored in the mixed case since it is not independent of the other parameters. Nonetheless, it is still possible to simplify computation by profiling out  $\sigma^2$  with its maximum-likelihood estimator, which admits a closed-form solution given  $\boldsymbol{\beta}$  and  $\boldsymbol{\gamma}$ :

$$\hat{\sigma}^2(\boldsymbol{\beta}, \boldsymbol{\gamma}) = \frac{1}{n_\bullet} \sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}), \quad (2.3)$$

where  $n_\bullet = \sum_{i=1}^m n_i$ . Substituting (2.3) into (2.1) and removing constants that do not depend on  $\boldsymbol{\beta}$  or  $\boldsymbol{\gamma}$  yields

$$l(\boldsymbol{\beta}, \boldsymbol{\gamma}; \mathbf{y}, \mathbf{X}) = \sum_{i=1}^m \log \det\{\mathbf{V}_i(\boldsymbol{\gamma})\} + n_\bullet \log \left\{ \sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}) \right\}. \quad (2.4)$$

We omit the dependence on  $\mathbf{y}$  and  $\mathbf{X}$  herein and write the negative log-likelihood as  $l(\boldsymbol{\beta}, \boldsymbol{\gamma})$ .

Given  $\boldsymbol{\beta}$  and  $\boldsymbol{\gamma}$ , it is often of interest to predict the (unobservable) random effects  $\mathbf{u}_1, \dots, \mathbf{u}_m$ . The best linear unbiased predictor for the random effects in cluster  $i$  is

$$\hat{\mathbf{u}}_i = \text{diag}(\boldsymbol{\gamma}) \mathbf{X}_i^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}). \quad (2.5)$$

Though not directly part of the optimization process,  $\hat{\mathbf{u}}_i$  are useful for generating predictions.

Returning to the optimization problem, we regularize the negative log-likelihood (2.4) as

$$\min_{\substack{\boldsymbol{\beta} \in \mathbb{R}^p, \boldsymbol{\gamma} \in \mathbb{R}_+^p \\ \beta_k = 0 \Rightarrow \gamma_k = 0}} f(\boldsymbol{\beta}, \boldsymbol{\gamma}) := l(\boldsymbol{\beta}, \boldsymbol{\gamma}) + \lambda \alpha \|\boldsymbol{\beta}\|_0 + \lambda(1 - \alpha) \|\boldsymbol{\gamma}\|_0. \quad (2.6)$$

The role of the parameter  $\alpha$  is to control the composition of sparsity. In particular, having fixed  $\lambda$ , a larger value of  $\alpha$  means that a random effect is more likely to be nonzero when its corresponding fixed effect is nonzero. In the limiting case when  $\alpha = 1$ , there is no penalty for a component in  $\boldsymbol{\gamma}$  being nonzero once the corresponding component in  $\boldsymbol{\beta}$  is nonzero.

The objective function  $f$  in (2.6) is nonsmooth due to the  $\ell_0$  norms. It is also nonconvex with two sources of nonconvexity: the negative log-likelihood  $l$  and the  $\ell_0$  norms. These properties rule out the direct application of off-the-shelf optimization software and necessitate the development of tailored algorithms that scale to real-world problem instances.

## 2.2 Coordinate Descent

Coordinate descent algorithms iteratively minimize multi-dimensional objective functions by optimizing one or more coordinates at a time while keeping all others fixed. This process involves repeated cycles through the coordinates until convergence is achieved. In contrast to simultaneous updating of all coordinates (i.e., gradient descent), coordinate-wise updating is well-suited to sparse, high-dimensional problems where only a small subset of coordinates is important. We refer the interested reader to [Wright \(2015\)](#) for an overview.

The coordinate descent algorithm we develop here optimizes a block of two coordinates simultaneously:  $(\beta_k, \gamma_k)$ , i.e., the parameters corresponding to predictor  $k$ . The motivation for this design is two-fold: (i) these parameters are not easily separable due to the hierarchy constraint  $\beta_k = 0 \Rightarrow \gamma_k = 0$ , and (ii) both parameters need several of the same quantities for their update, so we can reduce overheads by updating both at the same time. A single cycle of our coordinate descent algorithm is thus a single pass over all blocks  $(\beta_1, \gamma_1), \dots, (\beta_p, \gamma_p)$ .

Suppose we are in cycle  $t$  and thus far we have optimized the parameter blocks  $1, \dots, k-1$ . A vanilla coordinate descent approach would proceed by updating the  $k$ th block as

$$(\beta_k^{(t)}, \gamma_k^{(t)}) \leftarrow \arg \min_{\substack{\beta_k \in \mathbb{R}, \gamma_k \in \mathbb{R}_+ \\ \beta_k = 0 \Rightarrow \gamma_k = 0}} f(\beta_1^{(t)}, \dots, \beta_k, \dots, \beta_p^{(t-1)}, \gamma_1^{(t)}, \dots, \gamma_k, \dots, \gamma_p^{(t-1)}). \quad (2.7)$$

The difficulty with this particular update is that no closed-form solution is available to the optimization problem in (2.7). A separate iterative algorithm is required. This difficulty is in contrast to non-mixed LMs, where the coordinate-wise update admits an analytical solution. Hence, rather than fully optimizing in each coordinate block, we partially optimize in each

block by taking a single gradient descent step towards the optimum for that block.

Our partial minimization scheme can be understood in the context of the block descent lemma of [Beck and Tetrushvili \(2013\)](#), with which we can bound the negative log-likelihood:

$$\begin{aligned}
l(\boldsymbol{\beta}, \boldsymbol{\gamma}) \leq \bar{l}_{L_k}(\boldsymbol{\beta}, \boldsymbol{\gamma}; \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}) &:= l(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}) + \nabla_{\beta_k} l(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}})(\beta_k - \tilde{\beta}_k) + \nabla_{\gamma_k} l(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}})(\gamma_k - \tilde{\gamma}_k) \\
&\quad + \frac{\bar{L}_k}{2} \left\{ (\beta_k - \tilde{\beta}_k)^2 + (\gamma_k - \tilde{\gamma}_k)^2 \right\},
\end{aligned} \tag{2.8}$$

which holds for any  $(\boldsymbol{\beta}, \boldsymbol{\gamma})$  and  $(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}})$  differing only in coordinate block  $k$  (for the coordinate descent update below,  $(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}})$  is taken to be the current iterate before block  $k$  is updated). The notation  $\nabla_{\beta_k}$  and  $\nabla_{\gamma_k}$  is shorthand for gradients with respect to  $\beta_k$  and  $\gamma_k$ , and  $\bar{L}_k \geq L_k$  where  $L_k$  is the Lipschitz constant of  $l$  as a function of  $\beta_k$  and  $\gamma_k$ . We may now upper bound the original objective function  $f(\boldsymbol{\beta}, \boldsymbol{\gamma})$  by a new objective function  $\bar{f}_{L_k}(\boldsymbol{\beta}, \boldsymbol{\gamma}; \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}})$ , given by

$$f(\boldsymbol{\beta}, \boldsymbol{\gamma}) \leq \bar{f}_{L_k}(\boldsymbol{\beta}, \boldsymbol{\gamma}; \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}) := \bar{l}_{L_k}(\boldsymbol{\beta}, \boldsymbol{\gamma}; \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}) + \lambda\alpha\|\boldsymbol{\beta}\|_0 + \lambda(1-\alpha)\|\boldsymbol{\gamma}\|_0. \tag{2.9}$$

Taking  $\bar{f}$  as our new objective function, the update to the  $k$ th block of coordinates becomes

$$(\beta_k^{(t)}, \gamma_k^{(t)}) \leftarrow \arg \min_{\substack{\beta_k \in \mathbb{R}, \gamma_k \in \mathbb{R}_+ \\ \beta_k=0 \Rightarrow \gamma_k=0}} \bar{f}_{L_k}(\beta_1^{(t)}, \dots, \beta_k, \dots, \beta_p^{(t-1)}, \gamma_1^{(t)}, \dots, \gamma_k, \dots, \gamma_p^{(t-1)}; \boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)}), \tag{2.10}$$

where the arguments  $(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)})$  after the semicolon denote the current iterate used to evaluate the bound before block  $k$  is updated, i.e.,  $(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}) = (\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)})$  in (2.9), with coordinates  $1, \dots, k-1$  already updated in cycle  $t$  and coordinates  $k, \dots, p$  still at their cycle  $t-1$  values. In contrast to (2.7), the update (2.10) admits a closed-form solution, described by [Proposition 1](#).

**Proposition 1.** Define the thresholding function

$$T_{\bar{L}}(\beta, \gamma; \lambda, \alpha) = \begin{cases} (0, 0) & \text{if } \beta^2 < \frac{2\lambda\alpha}{\bar{L}} \text{ and } \beta^2 + \gamma_+^2 < \frac{2\lambda}{\bar{L}} \\ (\beta, 0) & \text{if } \gamma_+^2 < \frac{2\lambda(1-\alpha)}{\bar{L}} \\ (\beta, \gamma_+) & \text{otherwise,} \end{cases} \quad (2.11)$$

where  $\gamma_+ = \max(\gamma, 0)$ . Furthermore, define the gradients

$$\nabla_{\beta_k} l(\boldsymbol{\beta}, \boldsymbol{\gamma}) = -2n \cdot \frac{\sum_{i=1}^m \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma})(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})}{\sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma})(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})} \quad (2.12)$$

and

$$\nabla_{\gamma_k} l(\boldsymbol{\beta}, \boldsymbol{\gamma}) = \sum_{i=1}^m \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{x}_{ik} - n \cdot \frac{\sum_{i=1}^m \{\mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma})(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})\}^2}{\sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma})(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})}. \quad (2.13)$$

Then the optimization problem (2.10) is solved by

$$(\beta_k^{(t)}, \gamma_k^{(t)}) = T_{\bar{L}_k} \left( \beta_k^{(t-1)} - \frac{1}{\bar{L}_k} \nabla_{\beta_k} l(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)}), \gamma_k^{(t-1)} - \frac{1}{\bar{L}_k} \nabla_{\gamma_k} l(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)}); \lambda, \alpha \right).$$

*Proof.* See Appendix A. □

Proposition 1 gives the update as a gradient descent step with step size equal to the inverse of the Lipschitz constant, followed by a threshold step. The particular form of the threshold that appears here follows from the  $\ell_0$  regularizers and the hierarchy constraint. In practice, the step size scalar  $\bar{L}_k$  is determined via backtracking line search by starting with a large value and iteratively decreasing the estimated constant until inequality (2.8) is satisfied.

We now present Algorithm 1, which summarizes the coordinate descent routine. Our implementation of the algorithm adopts several heuristics including gradient screening, gradient sorting, and active set updates. These heuristics are described in Appendix B.

Each iteration of Algorithm 1 requires computation of the gradients (2.12) and (2.13).

---

**Algorithm 1** Coordinate Descent
 

---

**Require:**  $(\boldsymbol{\beta}^{(0)}, \boldsymbol{\gamma}^{(0)}) \in \mathbb{R}^p \times \mathbb{R}_+^p$  and  $T \in \mathbb{N}$

```

for  $t = 1, \dots, T$  do
   $(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)}) \leftarrow (\boldsymbol{\beta}^{(t-1)}, \boldsymbol{\gamma}^{(t-1)})$ 
  for  $k = 1, \dots, p$  do
     $(\beta_k^{(t)}, \gamma_k^{(t)}) \leftarrow \arg \min_{\substack{\beta_k \in \mathbb{R}, \gamma_k \in \mathbb{R}_+ \\ \beta_k = 0 \Rightarrow \gamma_k = 0}} \bar{f}_{L_k}(\beta_1^{(t)}, \dots, \beta_k, \dots, \beta_p^{(t-1)}, \gamma_1^{(t)}, \dots, \gamma_k, \dots, \gamma_p^{(t-1)}; \boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)})$ 
  end for
  if converged then
    break
  end if
end for
return  $(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)})$ 

```

---

These gradients, in turn, involve the inverses of the matrices  $\mathbf{V}_1(\boldsymbol{\gamma}), \dots, \mathbf{V}_m(\boldsymbol{\gamma})$ , each of dimension  $n_i \times n_i$ . Computing these inverses naïvely takes  $O(n_i^3)$  operations, which is computationally prohibitive if  $n_i$  is large or  $p$  is large (since the inverses need to be computed  $p$  times to do one full cycle over the coordinates). Fortunately, since each coordinate descent iteration only updates a single coordinate of  $\boldsymbol{\gamma}$ , we can perform a rank-one update to get the new inverse. Let  $\tilde{\mathbf{V}}_i^{-1} = \mathbf{V}_i^{-1}(\gamma_1^{(t-1)}, \dots, \gamma_k^{(t-1)}, \dots, \gamma_p^{(t-1)})$  be the inverse from the current iteration and let  $\hat{\mathbf{V}}_i^{-1} = \mathbf{V}_i^{-1}(\gamma_1^{(t)}, \dots, \gamma_k^{(t)}, \gamma_{k+1}^{(t-1)}, \dots, \gamma_p^{(t-1)})$  be the inverse for the next iteration. Then, since  $\boldsymbol{\gamma}$  has changed in the  $k$ th coordinate only, the Sherman-Morrison formula gives

$$\hat{\mathbf{V}}_i^{-1} = \tilde{\mathbf{V}}_i^{-1} - \frac{\tilde{\mathbf{V}}_i^{-1} \mathbf{x}_{ik} \mathbf{x}_{ik}^\top \tilde{\mathbf{V}}_i^{-1}}{(\gamma_k^{(t)} - \gamma_k^{(t-1)})^{-1} + \mathbf{x}_{ik}^\top \tilde{\mathbf{V}}_i^{-1} \mathbf{x}_{ik}}.$$

This update requires only  $O(n_i^2)$  operations in contrast to  $O(n_i^3)$  for the naïve update.

Though the log determinant of  $\mathbf{V}_i(\boldsymbol{\gamma})$  is not required for the gradients and thus not needed by Algorithm 1, it can be useful for other purposes, e.g., monitoring convergence. Similar to the rank-one update strategy for the inverse, we can perform a rank-one update for the log determinant via the matrix determinant lemma (Harville, 1997), which gives

$$\log \det(\hat{\mathbf{V}}_i) = \log \det(\tilde{\mathbf{V}}_i) + \log \left( 1 + (\gamma_k^{(t)} - \gamma_k^{(t-1)}) \mathbf{x}_{ik}^\top \tilde{\mathbf{V}}_i^{-1} \mathbf{x}_{ik} \right).$$

As with the inverse, this update takes  $O(n_i^2)$  operations versus the naïve update of  $O(n_i^3)$ .

We now establish the convergence of our coordinate descent routine in Theorem 1.

**Theorem 1.** *Suppose  $\mathbf{y}$  does not lie in the column space of  $\mathbf{X}$ , and that  $\boldsymbol{\beta} \in \mathbb{R}^p$  and  $\boldsymbol{\gamma} \in \mathbb{R}_+^p$  with  $\|\boldsymbol{\beta}\|_\infty \leq \bar{\beta} \in \mathbb{R}_+$  and  $\|\boldsymbol{\gamma}\|_\infty \leq \bar{\gamma} \in \mathbb{R}_+$ . Let  $\{(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)})\}_{t \in \mathbb{N}}$  be a sequence of iterates generated by Algorithm 1. Then, for any  $\bar{L}_k \geq L_k$  for  $k = 1, \dots, p$ , the sequence of objective values  $\{f(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)})\}_{t \in \mathbb{N}}$  is decreasing, convergent, and satisfies the inequality*

$$f(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)}) - f(\boldsymbol{\beta}^{(t+1)}, \boldsymbol{\gamma}^{(t+1)}) \geq \sum_{k=1}^p \left( \frac{\bar{L}_k - L_k}{2} (\beta_k^{(t)} - \beta_k^{(t+1)})^2 + \frac{\bar{L}_k - L_k}{2} (\gamma_k^{(t)} - \gamma_k^{(t+1)})^2 \right).$$

Furthermore, for any  $\bar{L}_k > L_k$ , the sequence of active sets  $\{\mathcal{A}(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)})\}_{t \in \mathbb{N}}$ , i.e., the set of predictors with nonzero effects, converges to a fixed set in finitely many iterations.

*Proof.* See Appendix C. □

Theorem 1 ensures the sequence of objective values generated by Algorithm 1 is monotonically decreasing and convergent. Moreover, it guarantees the sequence of predictors with nonzero effects (the active set) converges to a fixed set in a finite number of iterations. These results are not direct applications of existing convergence results but rely on proving boundedness and Lipschitz continuity of the objective function, which hold provided  $\boldsymbol{\beta}$  and  $\boldsymbol{\gamma}$  lie in a compact set. While it is straightforward to ensure that Algorithm 1 satisfies this condition by introducing the box constraints  $\|\boldsymbol{\beta}\|_\infty \leq \bar{\beta} \in \mathbb{R}_+$  and  $\|\boldsymbol{\gamma}\|_\infty \leq \bar{\gamma} \in \mathbb{R}_+$ , where  $\|\cdot\|_\infty$  is the maximum absolute value, it appears unnecessary for convergence in practice.

## 2.3 Local Search

While coordinate descent algorithms can guarantee global minimizers for convex objective functions, the nonconvex nature of our objective function means Algorithm 1 cannot provide such a guarantee. This limitation is exacerbated when predictors are highly correlated, since several competing active sets may attain similar objective values, making coordinate descent

more susceptible to converging to a suboptimal solution. To address this issue, we devise a local search routine to refine the solution from coordinate descent. Inspired by similar strategies proposed in [Beck and Eldar \(2013\)](#), [Hazimeh and Mazumder \(2020\)](#), and [Dedieu et al. \(2021\)](#), our routine systematically perturbs the active set of predictors by swapping previously excluded predictors into the model while simultaneously removing active ones. The process is iterated with a configurable neighborhood size  $s \in \mathbb{N}$ , which governs the number of predictors swapped, until no further iterations improve the objective value.

The local search problem described above can be formalized mathematically as

$$\min_{\substack{\boldsymbol{\xi} \in \mathbb{R}^p, \boldsymbol{\zeta} \in \mathbb{R}_+^p, \xi_k=0 \Rightarrow \zeta_k=0 \\ \mathbf{z}_1 \in \{0,1\}^p, \mathbf{z}_2 \in \{0,1\}^p \\ \text{Supp}(\mathbf{z}_1) \subseteq \mathcal{A}, \text{Supp}(\mathbf{z}_2) \subseteq \mathcal{A}^c \\ \|\mathbf{z}_1\|_0 \leq s, \|\mathbf{z}_2\|_0 \leq s}} f(\boldsymbol{\beta} - \mathbf{z}_1 \odot \boldsymbol{\beta} + \mathbf{z}_2 \odot \boldsymbol{\xi}, \boldsymbol{\gamma} - \mathbf{z}_1 \odot \boldsymbol{\gamma} + \mathbf{z}_2 \odot \boldsymbol{\zeta}), \quad (2.14)$$

where  $\text{Supp}(\mathbf{z}) := \{j : z_j \neq 0\}$ . In words, given an incumbent solution  $(\boldsymbol{\beta}, \boldsymbol{\gamma})$ , local search finds a subset of size  $s$  to remove from the current active set  $\mathcal{A}$  and replace with a subset of size  $s$  from the inactive set  $\mathcal{A}^c$  such that the objective function  $f(\boldsymbol{\beta}, \boldsymbol{\gamma})$  is minimized. Here,  $\mathbf{z}_1$  and  $\mathbf{z}_2$  are binary indicator vectors specifying the predictors to be removed and added, respectively, while  $\boldsymbol{\xi}$  and  $\boldsymbol{\zeta}$  contain the parameter values for the newly included predictors. When  $s = 1$ , the problem simplifies to evaluating the objective for all possible single swaps, making it computationally efficient whilst having the capability to escape poor local optima.

After solving the local search problem (2.14), the updated solution serves as the initialization point for a subsequent run of coordinate descent. This iterative interplay between the two methods—described in [Algorithm 2](#)—combines the fast, fine-grained updates from coordinate descent with the broader, exploratory improvements from local search. The algorithm terminates when neither coordinate descent nor local search achieves further improvement. Our implementation of [Algorithm 2](#) solves the local search problem with  $s = 1$ .

---

**Algorithm 2** Coordinate Descent with Local Search

---

**Require:**  $(\hat{\boldsymbol{\beta}}^{(0)}, \hat{\boldsymbol{\gamma}}^{(0)}) \in \mathbb{R}^p \times \mathbb{R}_+^p$  and  $T \in \mathbb{N}$   
**for**  $t = 1, \dots, T$  **do**  
    Take  $(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)})$  as the output of Algorithm 1 initialized at  $(\hat{\boldsymbol{\beta}}^{(t-1)}, \hat{\boldsymbol{\gamma}}^{(t-1)})$   
    Take  $(\hat{\boldsymbol{\beta}}^{(t)}, \hat{\boldsymbol{\gamma}}^{(t)})$  as the solution to the local search problem (2.14) at  $(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)})$   
    **if**  $f(\hat{\boldsymbol{\beta}}^{(t)}, \hat{\boldsymbol{\gamma}}^{(t)}) = f(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)})$  **then**  
        **break**  
    **end if**  
**end for**  
**return**  $(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)})$

---

## 2.4 Pathwise Optimization

Typical practice in sparse modeling is to estimate multiple models with different values of  $\lambda$  so different sparsity levels can be compared (e.g., using cross-validation). It is usually computationally expedient to compute these models in a pathwise fashion by sequentially initializing the optimizer at the previous solution. In our case, the pathwise optimization approach has the benefit of (i) reducing overall computation time since the solution at  $\lambda = \lambda^{(r+1)}$  is likely near the solution at  $\lambda = \lambda^{(r)}$ , and (ii) helping the optimizer navigate the nonconvex optimization surface, thereby improving the quality of the models.

We construct the regularization sequence  $\{\lambda^{(r)}\}_{r=1}^R$  so the model path evolves in a sparse-to-dense manner, beginning with a large amount of regularization and ending with a small amount, i.e.,  $\lambda^{(r)} > \lambda^{(r+1)}$ . Moreover, we configure the sequence such that the model from the solve at  $\lambda = \lambda^{(r+1)}$  has a different set of active predictors to the model from the solve with  $\lambda = \lambda^{(r)}$ . Proposition 2 provides a formula for choosing  $\lambda^{(r+1)}$  to guarantee this behavior.

**Proposition 2.** *Let  $(\boldsymbol{\beta}^{(r)}, \boldsymbol{\gamma}^{(r)})$  be a coordinate descent solution from  $\lambda = \lambda^{(r)}$ . Denote its active set by  $\mathcal{A}^{(r)}$ . Then running coordinate descent initialized at  $(\boldsymbol{\beta}^{(r)}, \boldsymbol{\gamma}^{(r)})$  with  $\lambda$  set to*

$$\lambda^{(r+1)} = c \max_{k \notin \mathcal{A}^{(r)}} \max \left( \frac{\{\nabla_{\beta_k} l(\boldsymbol{\beta}^{(r)}, \boldsymbol{\gamma}^{(r)})\}^2}{2\alpha \bar{L}_k}, \frac{\{\nabla_{\beta_k} l(\boldsymbol{\beta}^{(r)}, \boldsymbol{\gamma}^{(r)})\}^2 + \max(-\nabla_{\gamma_k} l(\boldsymbol{\beta}^{(r)}, \boldsymbol{\gamma}^{(r)}), 0)^2}{2\bar{L}_k} \right)$$

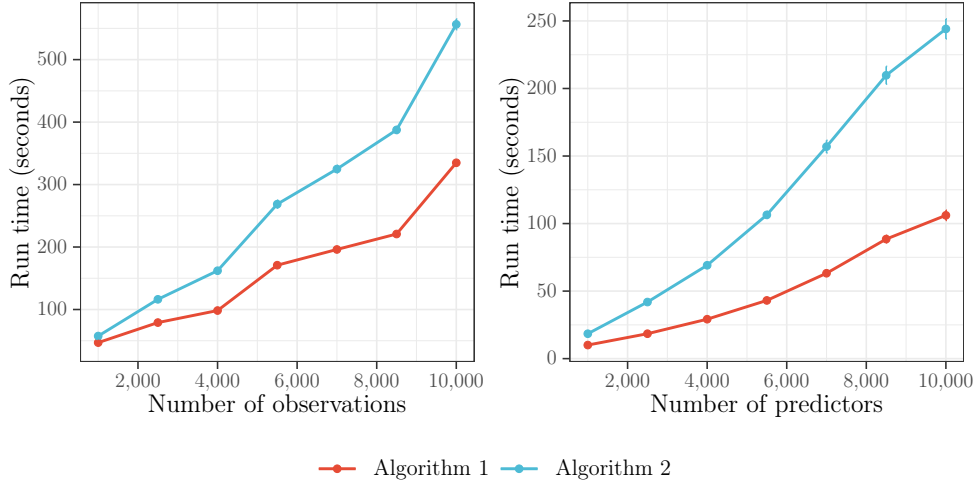


Figure 1: Run time in seconds for Algorithm 1 (coordinate descent) and Algorithm 2 (coordinate descent with local search). The number of predictors  $p = 100$  in the left plot and the number of observations  $n_{\bullet} = 100$  in the right plot. The number of clusters  $m = \lceil n_{\bullet}/10 \rceil$ . The averages (points) and standard errors (error bars) are measured over 100 synthetic datasets.

yields a solution  $(\beta^{(r+1)}, \gamma^{(r+1)})$  whose active set  $\mathcal{A}^{(r+1)}$  differs from  $\mathcal{A}^{(r)}$  for  $c \in [0, 1)$ .

*Proof.* See Appendix D. □

## 2.5 Computational Complexity

Assuming the number of clusters  $m$  is fixed, by leveraging low-rank matrix inverse updates, a single cycle of coordinate descent across all  $p$  coordinates requires  $O(p \sum_{i=1}^m n_i^2)$  operations. If instead  $m$  varies and the cluster size is fixed (say  $n_i = n$  for all  $i = 1, \dots, m$ ), the number of operations required becomes  $O(n_{\bullet} p)$ , where  $n_{\bullet} = mn$ . Running a single round of local search after coordinate descent does not alter the overall complexity since we need only perform  $p$  swaps for each active predictor, with the number of active predictors fixed.

To verify the complexity analysis and obtain a sense of run times, we measure the time to fit a full regularization path over a sequence of 100 values of  $\lambda$ . Figure 1 reports these times as a function of both  $n_{\bullet}$  and  $p$ . As predicted, the curves are mostly linear. Naturally, local search introduces additional computational costs, with the exact cost depending on  $n_{\bullet}$  and  $p$ .

Even so, one can fit a path of 100 models with 10,000 predictors in less than 5 minutes using coordinate descent with local search and 2 minutes using coordinate descent alone.

## 3 Sparse GLMMs

### 3.1 Non-Gaussian Likelihoods

The preceding section focuses on  $\ell_0$  regularized LMMs where the response is continuous and the likelihood is Gaussian. GLMMs are flexible generalizations of LMMs that handle a broad range of response types by extending the likelihood to include the entire exponential family (see, e.g., [Stroup et al., 2024](#)). For instance, similar to non-mixed generalized LMs (GLMs), GLMMs can accommodate binary response data via a Bernoulli likelihood and count response data via a Poisson likelihood. The conditional expectation of a GLMM satisfies

$$g(\mathbb{E}[y_{ij} \mid \mathbf{u}_i]) = \eta_{ij}, \quad \eta_{ij} = \mathbf{x}_{ij}^\top(\boldsymbol{\beta} + \mathbf{u}_i),$$

where  $g$  is a so-called link function (e.g., logit for Bernoulli or log for Poisson). Unfortunately, unlike their non-mixed GLM cousins, GLMMs are not nearly as computationally tractable as the likelihoods involved do not admit analytic forms except in the Gaussian case. As a consequence, it is not possible to fit an  $\ell_0$  regularized GLMM by direct application of [Algorithm 1](#) or [Algorithm 2](#) since these rely on the availability of closed-form gradients.

### 3.2 Penalized Quasi-Likelihood Approximation

To address intractability, we employ a penalized quasi-likelihood (PQL) approximation ([Breslow and Clayton, 1993](#)) to fit  $\ell_0$  regularized GLMMs in the non-Gaussian case. PQL approximations are among the most scalable techniques for fitting GLMMs since they effectively transform the likelihood into a tractable Gaussian form.<sup>3</sup> Briefly, the approximation

---

<sup>3</sup>Although computationally efficient, PQL approximations can lead to biased parameter estimates; see [Capanu et al. \(2013\)](#) for a discussion on the relative merits of various GLMM estimation methods.

replaces the likelihood with its quasi-likelihood and then iteratively linearizes the mean function around the current estimate  $\hat{\mathbf{u}}_i$  of the random effects  $\mathbf{u}_i$ . At each iteration, the Gaussian likelihood is optimized with working responses  $\mathbf{z}_i = (z_{i1}, \dots, z_{in_i})^\top$  and weights  $\mathbf{w}_i = (w_{i1}, \dots, w_{in_i})^\top$ , whose elements are given by

$$z_{ij} = \eta_{ij} + \frac{y_{ij} - \mathbf{E}[y_{ij} \mid \hat{\mathbf{u}}_i]}{\nabla g^{-1}(\eta_{ij})} \quad (3.1)$$

and

$$w_{ij} = \frac{\{\nabla g^{-1}(\eta_{ij})\}^2}{\text{Var}[y_{ij} \mid \hat{\mathbf{u}}_i]}. \quad (3.2)$$

Here,  $\nabla g^{-1}$  is the gradient of the inverse link function, and  $\mathbf{E}[y_{ij} \mid \hat{\mathbf{u}}_i]$  and  $\text{Var}[y_{ij} \mid \hat{\mathbf{u}}_i]$  are the conditional mean and conditional variance of  $y_{ij}$  as prescribed under the distributional family assumed on the response, both evaluated at the current estimate of the random effects. Recall that these estimates are computable using the best linear unbiased predictor in (2.5).

Algorithm 3 implements the above weighting scheme with the Gaussian likelihood optimization performed by coordinate descent and (optionally) local search. To incorporate

---

**Algorithm 3** PQL Approximation

---

**Require:**  $(\boldsymbol{\beta}^{(0)}, \boldsymbol{\gamma}^{(0)}) \in \mathbb{R}^p \times \mathbb{R}_+^p$  and  $T \in \mathbb{N}$

- 1: **for**  $t = 1, \dots, T$  **do**
  - 2:   Plug  $\boldsymbol{\beta}^{(t-1)}$  and  $\boldsymbol{\gamma}^{(t-1)}$  into (2.5) to get  $\hat{\mathbf{u}}_1, \dots, \hat{\mathbf{u}}_m$
  - 3:   Plug  $\boldsymbol{\beta}^{(t-1)}$  and  $\hat{\mathbf{u}}_1, \dots, \hat{\mathbf{u}}_m$  into (3.1) and (3.2) to get  $\mathbf{z}_1, \dots, \mathbf{z}_m$  and  $\mathbf{w}_1, \dots, \mathbf{w}_m$
  - 4:   Run Algorithm 1 or 2 with  $(\boldsymbol{\beta}^{(t-1)}, \boldsymbol{\gamma}^{(t-1)})$ ,  $\mathbf{y}_i = \mathbf{z}_i$ , and  $\mathbf{w}_i$ ,  $i = 1, \dots, m$ , to get  $(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)})$
  - 5:   **if** converged **then**
  - 6:     **break**
  - 7:   **end if**
  - 8: **end for**
  - 9: **return**  $(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)})$
- 

weights into either of these algorithms one need only replace the covariance matrix (2.2) with

$$\mathbf{V}_i(\boldsymbol{\gamma}) = \text{diag}(\mathbf{w}_i^{-1}) + \mathbf{X}_i \text{diag}(\boldsymbol{\gamma}) \mathbf{X}_i^\top,$$

where  $\mathbf{w}_i^{-1}$  is the elementwise reciprocal. The calls to coordinate descent or local search inside Algorithm 3 are initialized at the previous iterates, making the algorithm significantly faster than multiple individual runs of coordinate descent or local search. We remark that the algorithm resembles iteratively weighted least squares, which is used in the well-known and highly scalable `glmnet` package to fit  $\ell_1$  regularized GLMs (Tay et al., 2023).

### 3.3 Toolkit

We implement our method for learning sparse GLMMs in the R package `glmmsel` (generalized linear mixed model selection). At the time of writing, `glmmsel` supports Gaussian and Bernoulli likelihoods for handling regression and classification tasks. The package’s core components are written in C++ to facilitate fast computation, with the R code serving as a user-friendly interface. This modular design also makes it straightforward to develop interfaces for other languages (e.g., Python or Julia) since the algorithms need not be reimplemented. `glmmsel` is available open-source on the R repository CRAN and GitHub at

<https://github.com/ryan-thompson/glmmsel>.

## 4 Statistical Guarantee

### 4.1 Setup

We now develop a statistical guarantee on the finite-sample performance of our estimator. Towards this end, we assume the data-generating process is an LMM of the form

$$\mathbf{y}_i = \mathbf{X}_i \boldsymbol{\beta}^0 + \boldsymbol{\nu}_i, \quad \boldsymbol{\nu}_i \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{V}_i(\boldsymbol{\gamma}^0)), \quad i = 1, \dots, m, \quad (4.1)$$

where the noise variance  $\sigma^2 > 0$  and the covariance matrix  $\mathbf{V}_i(\boldsymbol{\gamma}) = \mathbf{I} + \mathbf{X}_i \text{diag}(\boldsymbol{\gamma}) \mathbf{X}_i^\top$ . We remark that model (4.1) is a standard and equivalent way of writing model (1.1). The predictor matrix  $\mathbf{X}$  is assumed to be deterministic (i.e., non-random) and have columns with

unit  $\ell_2$  norm. The fixed effect coefficients  $\boldsymbol{\beta}^0 \in \mathbb{R}^p$  and random effect variances  $\boldsymbol{\gamma}^0 \in \mathbb{R}_+^p$  are assumed to be  $s$ -sparse with nested support, i.e.,  $\|\boldsymbol{\beta}^0\|_0 \leq s$  and  $\text{Supp}(\boldsymbol{\gamma}^0) \subseteq \text{Supp}(\boldsymbol{\beta}^0)$ .

Given a sample of data  $(\mathbf{y}, \mathbf{X})$ , we fit model (4.1) using the  $\ell_0$  regularized estimator

$$(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}}) \in \arg \min_{(\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \mathcal{C}(s)} l(\boldsymbol{\beta}, \boldsymbol{\gamma}), \quad (4.2)$$

where the negative log-likelihood

$$l(\boldsymbol{\beta}, \boldsymbol{\gamma}) = \sum_{i=1}^m \log \det\{\sigma^2 \mathbf{V}_i(\boldsymbol{\gamma})\} + \frac{1}{\sigma^2} \sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})$$

and the constraint set

$$\mathcal{C}(s) = \{(\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \mathbb{R}^p \times \mathbb{R}_+^p : \|\boldsymbol{\beta}\|_0 \leq s, \|\boldsymbol{\beta}\|_\infty \leq \bar{\beta}, \|\boldsymbol{\gamma}\|_\infty \leq \bar{\gamma}, \text{Supp}(\boldsymbol{\gamma}) \subseteq \text{Supp}(\boldsymbol{\beta})\}.$$

The estimator (4.2) encodes the  $\ell_0$  regularizer as a hard constraint and is hence recognizable as the constrained version of the penalized estimator (2.6) with  $\alpha = 1$  (chosen here to reduce the complexity of the analysis). The correspondence between these two versions of the estimator can be understood by the fact that for any value of the regularization parameter  $\lambda$  there exists a sparsity level  $s$  such that the two estimators produce identical estimates. The box constraints  $\|\boldsymbol{\beta}\|_\infty \leq \bar{\beta}$  and  $\|\boldsymbol{\gamma}\|_\infty \leq \bar{\gamma}$  are included in the constraint set to facilitate the technical analysis. Here, the upper bound parameters  $\bar{\beta}$  and  $\bar{\gamma}$  are taken to satisfy  $\bar{\beta} \geq \|\boldsymbol{\beta}^0\|_\infty$  and  $\bar{\gamma} \geq \|\boldsymbol{\gamma}^0\|_\infty$ . Such bounds are commonly assumed when deriving statistical guarantees on sparse estimators (see, e.g., [Dedieu et al., 2021](#); [Nguyen et al., 2024](#)).

## 4.2 Result

The object of our analysis is the Kullback–Leibler (KL) divergence between the true distribution  $N(\mathbf{X}\boldsymbol{\beta}^0, \sigma^2 \mathbf{V}(\boldsymbol{\gamma}^0))$  and the fitted distribution  $N(\mathbf{X}\hat{\boldsymbol{\beta}}, \sigma^2 \mathbf{V}(\hat{\boldsymbol{\gamma}}))$  produced by the estimator (4.2), where  $\mathbf{V}(\boldsymbol{\gamma}) = \text{bdiag}\{\mathbf{V}_1(\boldsymbol{\gamma}), \dots, \mathbf{V}_m(\boldsymbol{\gamma})\}$  is the covariance matrix over the

full sample. Theorem 2 gives a finite-sample (non-asymptotic) probabilistic upper bound on this divergence, where the notation  $\lesssim$  is used to indicate that the bound holds up to a positive universal multiplicative constant independent of the parameters and data.

**Theorem 2.** Fix  $\delta \in (0, 1]$ . Let  $L(\delta) > 0$  and  $L > 0$  be the upper bounds on the Lipschitz constants of  $l(\boldsymbol{\beta}, \boldsymbol{\gamma})$  and  $\mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})]$  guaranteed by Lemmas 5 and 6 in Appendix E, respectively. Then, with probability at least  $1 - \delta$ , the estimator (4.2) satisfies

$$\begin{aligned} \text{KL} \left[ \mathbb{N}(\mathbf{X}\boldsymbol{\beta}^0, \sigma^2\mathbf{V}(\boldsymbol{\gamma}^0)) \parallel \mathbb{N}(\mathbf{X}\hat{\boldsymbol{\beta}}, \sigma^2\mathbf{V}(\hat{\boldsymbol{\gamma}})) \right] \\ \lesssim \max \left\{ (1 + \bar{\gamma}s) \max(\sqrt{n_{\bullet}\Pi}, \Pi), \frac{s\bar{\beta}}{\sigma} \sqrt{(1 + \bar{\gamma}s)\Pi} \right\}, \end{aligned} \quad (4.3)$$

where

$$\Pi = s \log\left(\frac{p}{s}\right) + \max(L(\delta), L) s\sqrt{s}(\bar{\beta} + \bar{\gamma}) + \log\left(\frac{1}{\delta}\right).$$

*Proof.* See Appendix E. □

The quantities appearing on the right-hand side not previously described can be understood as follows. First, the term  $s \log(p/s)$  represents the cost of searching for the best subset of predictors, indicating that the bound grows logarithmically with the number of candidate predictors. Meanwhile, the term  $\max(L(\delta), L)$  is from the maximal Lipschitz constants  $L(\delta)$  and  $L$  of the negative log-likelihood and its expectation, respectively (exact forms of these terms are detailed in Appendix E). Finally, the term  $\log(1/\delta)$  relates the tightness of the bound to the probability it is satisfied, which is guaranteed to be at least  $1 - \delta$ .

To our knowledge, no comparable guarantees are available for existing sparse LMM estimators, whether  $\ell_0$  regularized or otherwise. Schelldorfer et al. (2011) provided a finite-sample bound on the KL divergence of an  $\ell_1$  regularized LMM estimator, but their estimator does not select random effects. Raskutti et al. (2011) derived a now classic result for the  $\ell_0$  regularized least squares estimator in the non-mixed LM setting. Rather than control the KL divergence, their bounds control the estimation and prediction errors of the estimator but nonetheless feature the same logarithmic scaling factor  $s \log(p/s)$  that is in our bound.

Against this backdrop, Theorem 2 provides a statistical justification for the hierarchical estimator proposed in Section 2, at least in its constrained  $\alpha = 1$  form, which retains the same nested support condition. The bound shows that, under hierarchical sparsity, the dependence on the total number of candidate predictors is only logarithmic through the term  $s \log(p/s)$ . Our estimator should therefore be most effective in high-dimensional settings with a genuinely sparse hierarchical signal. The experiments in the next section examine performance in finite samples and assess robustness beyond the exact assumptions of the theorem.

## 5 Synthetic Datasets

### 5.1 Simulation Design

We generate synthetic datasets with the response distributed as Gaussian or Bernoulli. For the Gaussian case, the response is simulated according to the data-generating process

$$y_{ij} \sim N(\mathbf{x}_{ij}^\top(\boldsymbol{\beta} + \mathbf{u}_i), \sigma^2), \quad j = 1, \dots, n_i, \quad i = 1, \dots, m,$$

where the noise variance  $\sigma^2 > 0$  is set such that the signal-to-noise ratio is one. Similarly, for the Bernoulli case, the response is simulated according to the data-generating process

$$y_{ij} \sim \text{Bern} \left( \frac{1}{1 + \exp\{-\mathbf{x}_{ij}^\top(\boldsymbol{\beta} + \mathbf{u}_i)\}} \right), \quad j = 1, \dots, n_i, \quad i = 1, \dots, m.$$

The total sample size  $n_\bullet = \sum_{i=1}^m n_i$  is swept over a grid of seven values equidistant between  $n_\bullet = 30$  and  $n_\bullet = 1,000$  on the logarithmic scale. Each sample is randomly assigned to one of  $m = \lceil n_\bullet/10 \rceil$  clusters so the cluster sample sizes  $n_1, \dots, n_m$  vary with an average size of 10.

For both types of response, the predictors  $\mathbf{x}_{ij} \in \mathbb{R}^p$  are generated as  $N(\mathbf{0}, \boldsymbol{\Sigma})$ , where  $\boldsymbol{\Sigma}$  has elements  $\Sigma_{ij} = \rho^{|i-j|}$  for  $i, j = 1, \dots, p$ . Since our interest lies in the sparse regime, we only specify a subset of these predictors as relevant. Specifically, for the fixed effects, we randomly

sample a set  $\mathcal{A}_1 \subset \{1, \dots, p\}$  of relevant indices. We then take the fixed effects coefficients corresponding to  $\mathcal{A}_1$  to be one and all other coefficients to be zero, i.e.,  $\boldsymbol{\beta}_{\mathcal{A}_1} = \mathbf{1}$  and  $\boldsymbol{\beta}_{\mathcal{A}_1^c} = \mathbf{0}$ . Following the principle that fixed effects precede random effects, we sample a second subset  $\mathcal{A}_2 \subset \mathcal{A}_1$  indexing the predictors with random effects. The random effects coefficient vector  $\mathbf{u}_i \sim N(\mathbf{0}, \text{diag}(\boldsymbol{\gamma}))$  then has its variances configured as  $\boldsymbol{\gamma}_{\mathcal{A}_2} = \mathbf{1}$  and  $\boldsymbol{\gamma}_{\mathcal{A}_2^c} = \mathbf{0}$ .

## 5.2 Baselines and Metrics

Our implementation `glmmsel` is benchmarked against several state-of-the-art baselines: `pysr3` (Sholokhov et al., 2024), `L0Learn` (Hazimeh and Mazumder, 2020), `glmnet` (Friedman et al., 2010), and `ncvreg` (Brehehy and Huang, 2011). `L0Learn`, `glmnet`, and `ncvreg` are the best available toolkits for learning non-mixed sparse LMs. Meanwhile, `pysr3` is the only existing toolkit for learning sparse LMMs that is capable of scaling beyond tens of predictors. For completeness, Appendix G includes `glmPen` (Heiling et al., 2023) and `rpql` (Hui et al., 2017b) as additional baselines in small-scale experiments. All sparsity parameters are tuned using a separate validation set generated independently and identically to the training set. Appendix F details the regularizer of each method and the candidate tuning parameter values considered.

We evaluate the competing methods according to several metrics that characterize different qualities of the estimated model. As a measure of prediction performance, we report the squared error of the fitted linear predictor relative to the null model:

$$\text{Prediction error} := \frac{\sum_{i=1}^m \|\mathbf{X}_i(\boldsymbol{\beta} + \mathbf{u}_i) - \mathbf{X}_i(\hat{\boldsymbol{\beta}} + \hat{\mathbf{u}}_i)\|_2^2}{\sum_{i=1}^m \|\mathbf{X}_i(\boldsymbol{\beta} + \mathbf{u}_i)\|_2^2},$$

where  $\hat{\boldsymbol{\beta}}$  and  $\hat{\mathbf{u}}_1, \dots, \hat{\mathbf{u}}_m$  are estimates of the fixed and random effect coefficients. Here, the null model is the model with all fixed and random effects equal to zero. Thus, a value of zero indicates perfect prediction accuracy, and a value of one indicates no improvement over the null model. This same linear predictor-based definition is used for both Gaussian and

Bernoulli responses. As a measure of interpretability, we report the sparsity level (i.e., the number of predictors selected to have nonzero fixed effects or nonzero random effects):

$$\text{Sparsity} := \|\hat{\boldsymbol{\beta}} + \hat{\mathbf{u}}_1\|_0.$$

Lower values are desirable from an interpretation standpoint, with a value of five reflecting the correct sparsity level. As a measure of selection performance, we report the F1 score:

$$\text{F1 score} := \frac{2\text{TP}}{2\text{TP} + \text{FP} + \text{FN}},$$

where TP is the number of true positives, FP is the number of false positives, and FN is the number of false negatives. We consider two different types of F1 score. The first is an F1 score of the effect types, under which a predictor is counted as correct only if it is selected with the correct role, fixed or random. The second is an F1 score of the nonzeros, which ignores this distinction and counts a predictor as correct whenever it is selected at all. Hence, the effect type F1 score is the more stringent measure and is better suited to assessing whether a method correctly distinguishes fixed effects from random effects.

### 5.3 Results

Figures 2 and 3 report results for  $p = 1,000$  predictors with Gaussian and Bernoulli responses, respectively, in a sparse setting with 5 nonzero fixed effects and 3 nonzero random effects. In both cases, `glmmsel` demonstrates excellent performance in prediction, interpretation, and selection. At smaller sample sizes, where accurately estimating random effects is inherently challenging, `glmmsel`'s prediction error is comparable to that of the sparse LM methods `L0Learn`, `glmnet`, and `ncvreg`. However, as the sample size increases, `glmmsel` delivers substantially better predictions. Moreover, it accurately recovers the true predictors and, unlike the sparse LM methods, correctly distinguishes between fixed and random effects. In the Gaussian case, when  $n_{\bullet} \approx 1,000$ , `pysr3` emerges as the next best performer, though

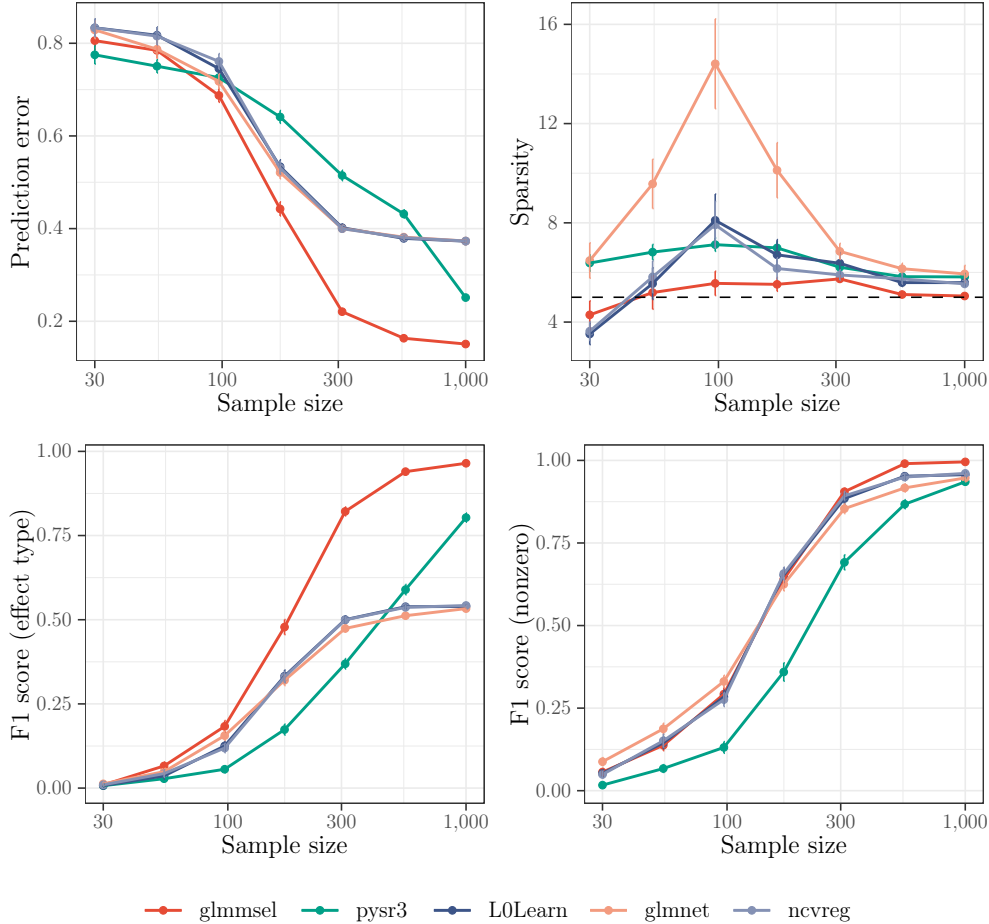


Figure 2: Comparisons on synthetic data with Gaussian response. The number of predictors  $p = 1,000$  with 5 nonzero fixed effects and 3 nonzero random effects. The correlation level  $\rho = 0.5$ . The averages (points) and standard errors (error bars) are measured over 100 datasets. In the upper right panel, the dashed horizontal line indicates the true number of predictors with nonzero effects.

a substantial gap remains in both its prediction and selection performance compared to `glmmsel`. Notably, `pysr3` is incompatible with non-Gaussian responses, making `glmmsel` the first tool capable of learning sparse GLMMs on data of this dimensionality.

Next, we consider a particularly challenging, ultra-high-dimensional setting with  $p = 10,000$  predictors, an order of magnitude more than in the previous experiments. Figure 4 reports the results. Here, `glmmsel` is the only sparse LMM tool that can practically scale to this many predictors. `pysr3`, which was able to handle  $p = 1,000$ , is unable to scale to this experiment. Otherwise, the story mirrors that of the  $p = 1,000$  case: `glmmsel` delivers strong

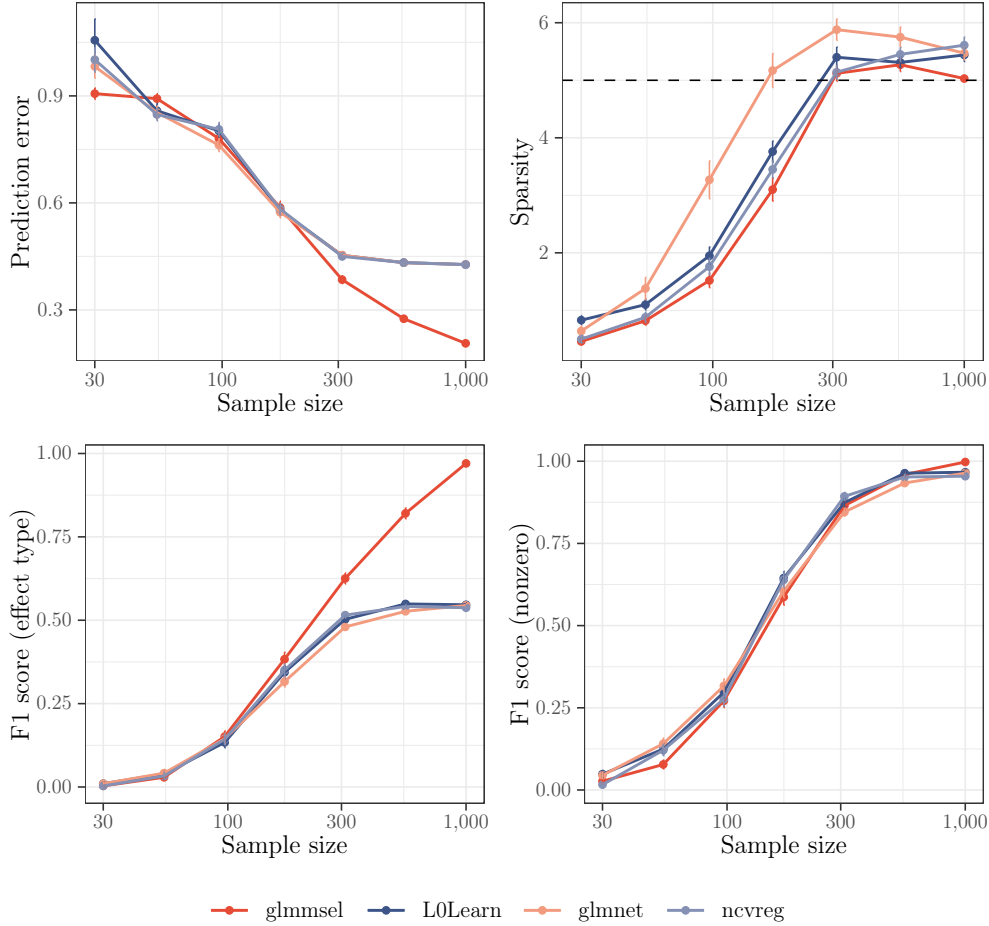


Figure 3: Comparisons on synthetic data with Bernoulli response. The number of predictors  $p = 1,000$  with 5 nonzero fixed effects and 3 nonzero random effects. The correlation level  $\rho = 0.5$ . The averages (points) and standard errors (error bars) are measured over 100 datasets. In the upper right panel, the dashed horizontal line indicates the true number of predictors with nonzero effects.

performance and eventually dominates all competing methods as the sample size increases.

As a second challenging setting, we consider predictors with high correlation levels of  $\rho = 0.9$ . To illustrate the advantage of local search in such settings, we add `glmmsel`'s Algorithm 2 to the experiments. We also run `L0Learn` with its local search algorithm. Figure 5 reports the results. Under Algorithm 2, `glmmsel` demonstrates a clear edge: it maintains superior predictive accuracy relative to all baselines and excels at identifying the true effects, despite the complexity of navigating the highly correlated predictor space.

Additional results are provided in Appendices H and I. Appendix H shows that the main

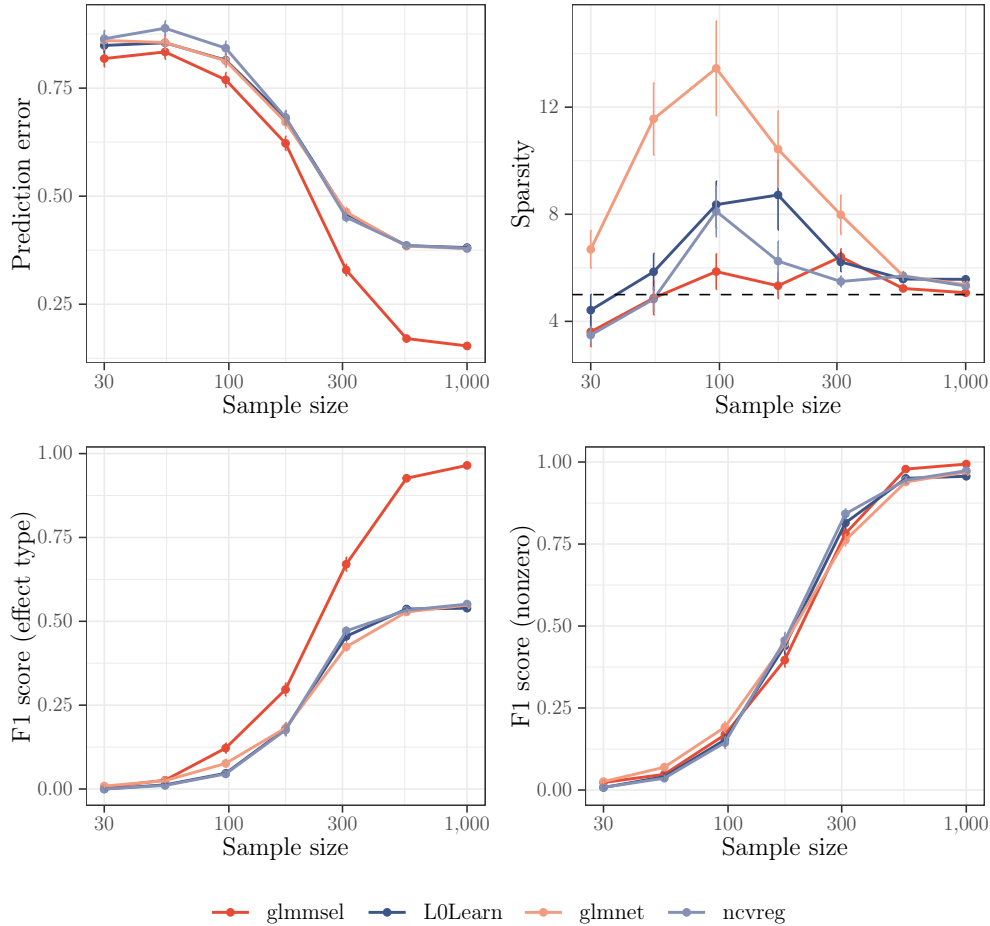


Figure 4: Comparisons on synthetic data with Gaussian response. The number of predictors  $p = 10,000$  with 5 nonzero fixed effects and 3 nonzero random effects. The correlation level  $\rho = 0.5$ . The averages (points) and standard errors (error bars) are measured over 100 datasets. In the upper right panel, the dashed horizontal line indicates the true number of predictors with nonzero effects.

conclusions remain stable in denser settings with more active predictors, while Appendix I shows that additional tuning of the concavity/shrinkage parameter  $\gamma$  for `glmnet` and `ncvreg` can improve prediction error but yields much denser models and weaker selection performance.

## 6 Real Datasets

The `riboflavin` dataset from [Schelldorfer et al. \(2011\)](#) contains data on riboflavin production by *Bacillus subtilis*, a bacterium used in industrial fermentation. Besides a continuous response

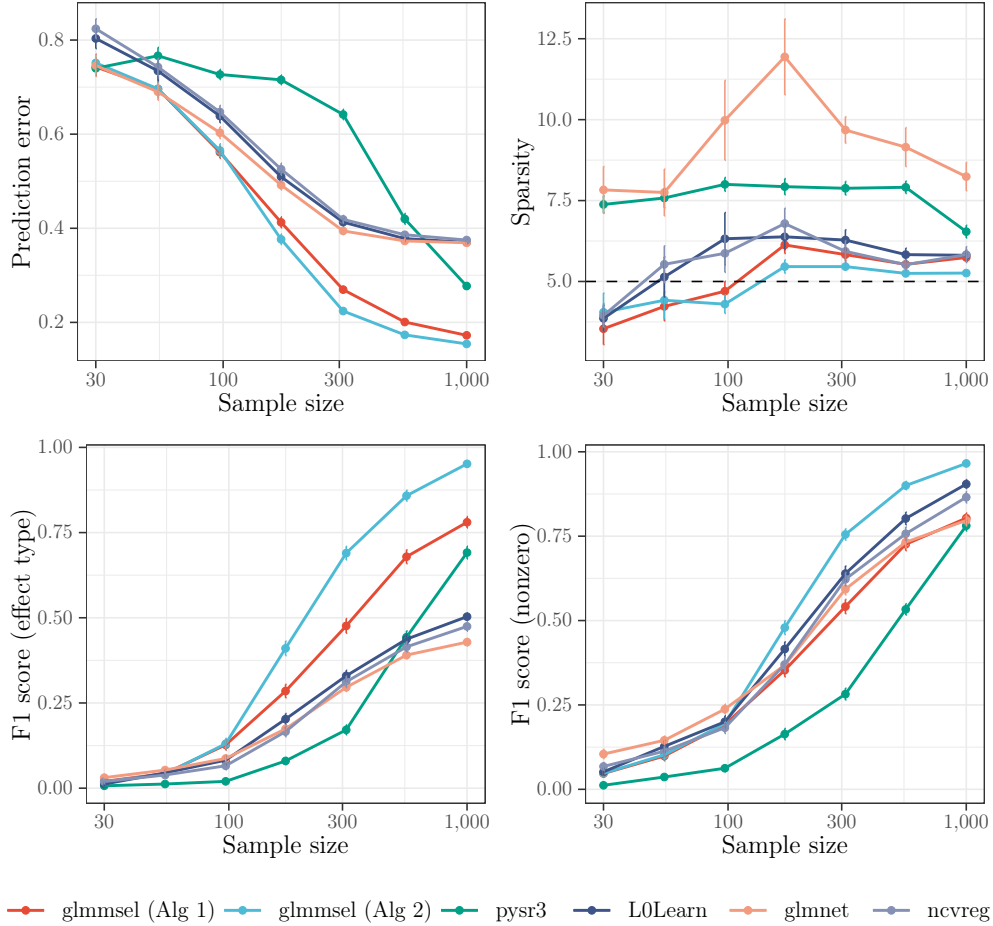


Figure 5: Comparisons on synthetic data with Gaussian response. The number of predictors  $p = 1,000$  with 5 nonzero fixed effects and 3 nonzero random effects. The correlation level  $\rho = 0.9$ . The averages (points) and standard errors (error bars) are measured over 100 datasets. In the upper right panel, the dashed horizontal line indicates the true number of predictors with nonzero effects.

variable measuring the log-transformed production rate in bacterium samples, the dataset comprises  $p = 4,088$  predictor variables in the form of bacterium gene expression levels. The task is to discover the genes relevant to riboflavin production, with the ultimate goal of genetically modifying the bacterium to enhance output in industrial settings. To support this analysis, the dataset includes observations on  $m = 28$  bacterium samples, with each sample measured  $n_i \in \{2, \dots, 6\}$  times, resulting in a total of  $n_{\bullet} = 111$  observations.

We randomly split the dataset into training, validation, and testing sets in 0.70-0.15-0.15 proportions. Table 2 reports the results from 100 such splits. The LMMs fit by `glmmsel`

	Prediction error	Sparsity		
		Total	Fixed	Random
<code>glmmsel</code>	0.304 (0.016)	11.7 (1.1)	10.8 (1.1)	0.9 (0.1)
<code>L0Learn</code>	0.359 (0.015)	7.6 (1.2)	7.6 (1.2)	0.0 (0.0)
<code>glmnet</code>	0.309 (0.015)	23.6 (2.0)	23.6 (2.0)	0.0 (0.0)
<code>ncvreg</code>	0.401 (0.018)	5.3 (0.5)	5.3 (0.5)	0.0 (0.0)

Table 2: Comparisons on the `riboflavin` dataset. The averages and standard errors (parentheses) are measured over 100 splits of the data. Prediction error is the squared error loss of the fitted linear predictor relative to the fixed intercept-only model.

achieve the lowest prediction loss among all methods. `L0Learn` yields sparser LMs, but because they cannot capture heterogeneity across samples, their prediction loss is notably higher. The LMs fit by `glmnet` exhibit statistically comparable prediction loss to the LMMs fit by `glmmsel`, but contain more than twice as many genes on average. Meanwhile, `ncvreg` achieves the highest sparsity, selecting fewer genes on average than all other methods considered, at the cost of substantially worse prediction performance. Taken together, these results show that `glmmsel` achieves the most favorable prediction–sparsity trade-off, attaining the lowest prediction loss using 10.8 fixed effects and only 0.9 random effects on average. The fact that so few random effects are selected indicates that heterogeneity in the `riboflavin` data is present but sparse, with most of the signal explained by genes with shared effects across samples and only a very small number requiring sample-specific deviations. Ultimately, `glmmsel` is the only method capable of capturing heterogeneity while scaling effectively to datasets of this dimension.

As a second real dataset, Appendix J reports the performance of `glmmsel` on the classification of popularity of news articles posted to the online news platform Mashable.

## 7 Concluding Remarks

LMMs are key tools for modeling heterogeneous patterns encountered in contemporary data-analytic problems, such as those that arise in personalized medicine. These problems

increasingly involve datasets with a high dimensionality of predictors, making sparsity a desirable (even necessary) quality from the standpoints of prediction and interpretation. This paper introduces a new  $\ell_0$  regularized estimator for subset selection in LMMs. We demonstrate that coordinate descent, coupled with a low-complexity local search routine, can scale our estimator to wide datasets well beyond the reach of existing methods. These algorithms also readily extend to subset selection in GLMMs via a PQL approximation. Besides scalability, the new estimator is shown to have excellent statistical properties in theory, simulation, and practice. Our R package `glmmsel` is available on CRAN and GitHub.

## Acknowledgments

The authors acknowledge financial support from the Australian Research Council under Discovery Project DP230101179.

## A Proof of Proposition 1

The proof is delivered in three parts. First, we derive the form of the thresholding function. Second, we derive the gradient for  $\beta_k$ . Third, we derive the gradient for  $\gamma_k$ .

**Thresholding Function** Observe that (2.8), the upper bound on the negative log-likelihood, treated as a function of block  $k$ , satisfies

$$\begin{aligned}
\bar{l}_{\bar{L}_k}(\boldsymbol{\beta}, \boldsymbol{\gamma}; \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}) &= l(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}) + \nabla l(\tilde{\boldsymbol{\beta}}_k)^\top (\beta_k - \tilde{\beta}_k) + \nabla l(\tilde{\boldsymbol{\gamma}}_k)^\top (\gamma_k - \tilde{\gamma}_k) \\
&\quad + \frac{\bar{L}_k}{2} \left\{ (\beta_k - \tilde{\beta}_k)^2 + (\gamma_k - \tilde{\gamma}_k)^2 \right\} \\
&= \frac{\bar{L}_k}{2} \left\{ \beta_k - \underbrace{\left( \tilde{\beta}_k - \frac{1}{\bar{L}_k} \nabla l(\tilde{\beta}_k) \right)}_{\hat{\beta}_k} \right\}^2 + \frac{\bar{L}_k}{2} \left\{ \gamma_k - \underbrace{\left( \tilde{\gamma}_k - \frac{1}{\bar{L}_k} \nabla l(\tilde{\gamma}_k) \right)}_{\hat{\gamma}_k} \right\}^2 + \text{const.} \\
&= \frac{\bar{L}_k}{2} (\beta_k - \hat{\beta}_k)^2 + \frac{\bar{L}_k}{2} (\gamma_k - \hat{\gamma}_k)^2 + \text{const.},
\end{aligned}$$

where const. represents terms independent of  $\beta_k$  and  $\gamma_k$ . It immediately follows that (2.9), the upper bound on the objective function, satisfies

$$\bar{f}_{\bar{L}_k}(\boldsymbol{\beta}, \boldsymbol{\gamma}; \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}) = \frac{\bar{L}_k}{2}(\beta_k - \hat{\beta}_k)^2 + \frac{\bar{L}_k}{2}(\gamma_k - \hat{\gamma}_k)^2 + \lambda\alpha 1\{\beta_k \neq 0\} + \lambda(1 - \alpha)1\{\gamma_k \neq 0\} + \text{const.}$$

Hence, to minimize (2.9) as a function of block  $k$ , we need only minimize the right-hand side of the above function:

$$\min_{\substack{\beta \in \mathbb{R}, \gamma \in \mathbb{R}_+ \\ \beta=0 \Rightarrow \gamma=0}} \frac{\bar{L}}{2}(\beta - \hat{\beta})^2 + \frac{\bar{L}}{2}(\gamma - \hat{\gamma})^2 + \lambda\alpha 1\{\beta \neq 0\} + \lambda(1 - \alpha)1\{\gamma \neq 0\}$$

or equivalently

$$\min_{\substack{\beta \in \mathbb{R}, \gamma \in \mathbb{R} \\ \beta=0 \Rightarrow \gamma=0}} \frac{\bar{L}}{2}(\beta - \hat{\beta})^2 + \frac{\bar{L}}{2}(\gamma - \hat{\gamma}_+)^2 + \lambda\alpha 1\{\beta \neq 0\} + \lambda(1 - \alpha)1\{\gamma \neq 0\},$$

where  $\hat{\gamma}_+ = \max(\hat{\gamma}, 0)$  and we drop the subscript  $k$  to simplify notation. The hierarchy constraint gives rise to three possible cases for the optimal solutions  $\beta^*$  and  $\gamma^*$ :

1.  $\beta^* = 0$  and  $\gamma^* = 0$ , so the optimal objective value is  $f_1^* = \bar{L}/2\hat{\beta}^2 + \bar{L}/2\hat{\gamma}_+^2$ ;
2.  $\beta^* = \hat{\beta}$  and  $\gamma^* = 0$ , so the optimal objective value is  $f_2^* = \bar{L}/2\hat{\gamma}_+^2 + \lambda\alpha$ ; and
3.  $\beta^* = \hat{\beta}$  and  $\gamma^* = \hat{\gamma}_+$ , so the optimal objective value is  $f_3^* = \lambda$ .

It is straightforward to check that

1.  $f_1^*$  is less than  $f_2^*$  and  $f_3^*$  if  $\hat{\beta}^2 < 2\lambda\alpha/\bar{L}$  and  $\hat{\beta}^2 + \hat{\gamma}_+^2 < 2\lambda/\bar{L}$ ;
2.  $f_2^*$  is less than  $f_1^*$  and  $f_3^*$  if  $\hat{\beta}^2 > 2\lambda\alpha/\bar{L}$  and  $\hat{\gamma}_+^2 < 2\lambda(1 - \alpha)/\bar{L}$ ; and
3.  $f_3^*$  is less than  $f_1^*$  and  $f_2^*$  if  $\hat{\beta}^2 + \hat{\gamma}_+^2 > 2\lambda/\bar{L}$  and  $\hat{\gamma}_+^2 > 2\lambda(1 - \alpha)/\bar{L}$ .

Combining these conditions, we have

$$(\beta^*, \gamma^*) = \begin{cases} (0, 0) & \text{if } \hat{\beta}^2 < \frac{2\lambda\alpha}{\bar{L}} \text{ and } \hat{\beta}^2 + \hat{\gamma}_+^2 < \frac{2\lambda}{\bar{L}} \\ (\hat{\beta}, 0) & \text{if } \hat{\gamma}_+^2 < \frac{2\lambda(1-\alpha)}{\bar{L}} \\ (\hat{\beta}, \hat{\gamma}_+) & \text{otherwise.} \end{cases}$$

**Gradient for  $\beta_k$**  Consider the gradient of  $l$  with respect to  $\beta_k$ . Since  $\mathbf{V}_i(\boldsymbol{\gamma})$  does not depend on  $\beta_k$ , we have

$$\nabla_{\beta_k} l(\boldsymbol{\beta}, \boldsymbol{\gamma}) = n \cdot \nabla_{\beta_k} \log\{S(\boldsymbol{\beta}, \boldsymbol{\gamma})\},$$

where

$$S(\boldsymbol{\beta}, \boldsymbol{\gamma}) = \sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}).$$

The chain rule gives

$$\nabla_{\beta_k} \log\{S(\boldsymbol{\beta}, \boldsymbol{\gamma})\} = \frac{\nabla_{\beta_k} S(\boldsymbol{\beta}, \boldsymbol{\gamma})}{S(\boldsymbol{\beta}, \boldsymbol{\gamma})},$$

with the gradient on the right-hand side being

$$\nabla_{\beta_k} S(\boldsymbol{\beta}, \boldsymbol{\gamma}) = -2 \sum_{i=1}^m \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}).$$

The gradient claimed by the proposition immediately follows as

$$\nabla_{\beta_k} l(\boldsymbol{\beta}, \boldsymbol{\gamma}) = -2n \cdot \frac{\sum_{i=1}^m \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})}{\sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})}.$$

**Gradient for  $\gamma_k$**  Consider now the gradient of  $l$  with respect to  $\gamma_k$ . It holds

$$\nabla_{\gamma_k} l(\boldsymbol{\beta}, \boldsymbol{\gamma}) = \sum_{i=1}^m \nabla_{\gamma_k} \log \det\{\mathbf{V}_i(\boldsymbol{\gamma})\} + n \cdot \nabla_{\gamma_k} \log\{S(\boldsymbol{\beta}, \boldsymbol{\gamma})\}. \quad (\text{A.1})$$

Using the standard matrix calculus result that  $\nabla \log \det\{\mathbf{A}(x)\} = \text{tr}\{\mathbf{A}^{-1}(x)\nabla\mathbf{A}(x)\}$ , the first term on the right-hand side of (A.1) evaluates to

$$\nabla_{\gamma_k} \log \det\{\mathbf{V}_i(\gamma)\} = \text{tr}(\mathbf{V}_i^{-1}(\gamma)\mathbf{x}_{ik}\mathbf{x}_{ik}^\top) = \text{tr}(\mathbf{x}_{ik}^\top\mathbf{V}_i^{-1}(\gamma)\mathbf{x}_{ik}) = \mathbf{x}_{ik}^\top\mathbf{V}_i^{-1}(\gamma)\mathbf{x}_{ik}, \quad (\text{A.2})$$

which follows from the cyclic property of the trace function and the final term being a scalar.

For the second term on the right-hand side of (A.1), we have

$$\nabla_{\gamma_k} \log\{S(\boldsymbol{\beta}, \gamma)\} = \frac{\nabla_{\gamma_k} S(\boldsymbol{\beta}, \gamma)}{S(\boldsymbol{\beta}, \gamma)}.$$

Applying the standard matrix calculus results that  $\nabla(\mathbf{z}^\top\mathbf{A}^{-1}(x)\mathbf{z}) = \mathbf{z}^\top\nabla\mathbf{A}^{-1}(x)\mathbf{z}$  and  $\nabla\mathbf{A}^{-1}(x) = -\mathbf{A}^{-1}(x)\nabla\mathbf{A}(x)\mathbf{A}^{-1}(x)$  gives

$$\begin{aligned} \nabla_{\gamma_k} S(\boldsymbol{\beta}, \gamma) &= -\sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i\boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\gamma) \nabla_{\gamma_k} \mathbf{V}_i(\gamma) \mathbf{V}_i^{-1}(\gamma) (\mathbf{y}_i - \mathbf{X}_i\boldsymbol{\beta}) \\ &= -\sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i\boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\gamma) \mathbf{x}_{ik} \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\gamma) (\mathbf{y}_i - \mathbf{X}_i\boldsymbol{\beta}) \\ &= -\sum_{i=1}^m \{\mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\gamma) (\mathbf{y}_i - \mathbf{X}_i\boldsymbol{\beta})\}^2, \end{aligned}$$

where the final equality follows from the fact that  $\mathbf{z}^\top(\mathbf{u}\mathbf{u}^\top)\mathbf{z} = (\mathbf{u}^\top\mathbf{z})^2$ . We now have

$$\nabla_{\gamma_k} \log\{S(\boldsymbol{\beta}, \gamma)\} = -\frac{\sum_{i=1}^m \{\mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\gamma) (\mathbf{y}_i - \mathbf{X}_i\boldsymbol{\beta})\}^2}{\sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i\boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\gamma) (\mathbf{y}_i - \mathbf{X}_i\boldsymbol{\beta})}. \quad (\text{A.3})$$

Finally, substituting (A.3) and (A.2) into (A.1) yields the claimed gradient as

$$\nabla_{\gamma_k} l(\boldsymbol{\beta}, \gamma) = \sum_{i=1}^m \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\gamma) \mathbf{x}_{ik} - n \cdot \frac{\sum_{i=1}^m \{\mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\gamma) (\mathbf{y}_i - \mathbf{X}_i\boldsymbol{\beta})\}^2}{\sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i\boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\gamma) (\mathbf{y}_i - \mathbf{X}_i\boldsymbol{\beta})}.$$

## B Coordinate Descent Heuristics

Our implementation of Algorithm 1 adopts three heuristics that accelerate its convergence.

**Gradient Screening** Prior to running coordinate descent, the predictors are designated strong or weak. Strong predictors are those whose gradients are among (say) the 100 largest in magnitude. The remaining predictors are weak predictors. Coordinate descent is then run on the strong set until convergence. At that point, we do one cycle over the weak predictors. If the active set is unchanged, we stop. Otherwise, if any weak predictors become active, they are shifted from the weak set to the strong set, and coordinate descent is run again.

**Gradient Sorting** Before running coordinate descent, the predictors are sorted in descending order according to the magnitude of their gradients. This sorting ensures that predictors with the greatest potential impact are considered first, usually leading to faster convergence.

**Active Set Updates** After several iterations of coordinate descent, the active set usually stabilizes, with subsequent iterations simply polishing the estimates of the nonzero parameters. For this reason, we restrict the coordinate descent updates solely to the active set once it has stabilized (e.g., after three cycles with no changes). After convergence on the active set, one additional cycle through the inactive predictors is performed to confirm overall convergence.

## C Proof of Theorem 1

### C.1 Preliminary Lemmas

The proof requires two technical lemmas which we state and prove in turn. The first lemma ensures the negative log-likelihood (2.4) is bounded below. The second lemma ensures the gradients (2.12) and (2.13) of the negative log-likelihood are Lipschitz continuous.

**Lemma 1.** *Suppose  $\mathbf{y}$  does not lie in the column space of  $\mathbf{X}$ . Then there exists  $c \in \mathbb{R}$  such that the negative log-likelihood  $l(\boldsymbol{\beta}, \boldsymbol{\gamma})$  given in (2.4) satisfies  $l(\boldsymbol{\beta}, \boldsymbol{\gamma}) \geq c$  for all  $\boldsymbol{\beta} \in \mathbb{R}^p$  and  $\boldsymbol{\gamma} \in \mathbb{R}_+^p$  such that  $\|\boldsymbol{\gamma}\|_\infty \leq \bar{\gamma}$ .*

*Proof.* Recall that the negative log-likelihood is

$$l(\boldsymbol{\beta}, \boldsymbol{\gamma}) = \sum_{i=1}^m \log \det\{\mathbf{V}_i(\boldsymbol{\gamma})\} + n_\bullet \log \left\{ \sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}) \right\}, \quad (\text{C.1})$$

where

$$\mathbf{V}_i(\boldsymbol{\gamma}) = \mathbf{I} + \mathbf{X}_i \text{diag}(\boldsymbol{\gamma}) \mathbf{X}_i^\top.$$

Begin by choosing any  $\boldsymbol{\beta} \in \mathbb{R}^p$  and  $\boldsymbol{\gamma} \in \mathbb{R}_+^p$  such that  $\|\boldsymbol{\gamma}\|_\infty \leq \bar{\gamma}$ . Consider the first term  $\sum_{i=1}^m \log \det\{\mathbf{V}_i(\boldsymbol{\gamma})\}$  on the right-hand side of (C.1). Since  $\text{diag}(\boldsymbol{\gamma})$  is positive semi-definite (because  $\gamma_k \geq 0$  for all  $k$ ), the product  $\mathbf{X}_i \text{diag}(\boldsymbol{\gamma}) \mathbf{X}_i^\top$  must also be positive semi-definite. Hence,  $\mathbf{V}_i(\boldsymbol{\gamma})$  must be positive definite and its minimal eigenvalue  $\lambda_{\min}$  must satisfy

$$\lambda_{\min}\{\mathbf{V}_i(\boldsymbol{\gamma})\} = \lambda_{\min}(\mathbf{I} + \mathbf{X}_i \text{diag}(\boldsymbol{\gamma}) \mathbf{X}_i^\top) \geq \lambda_{\min}(\mathbf{I}) = 1.$$

It follows immediately from the property of determinants that  $\det\{\mathbf{V}_i(\boldsymbol{\gamma})\} \geq 1$  and hence

$$\sum_{i=1}^m \log \det\{\mathbf{V}_i(\boldsymbol{\gamma})\} \geq 0. \quad (\text{C.2})$$

Consider now the second term  $n_\bullet \log\{\sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})\}$  on the right-hand side of (C.1). The maximal eigenvalue  $\lambda_{\max}$  of  $\mathbf{V}_i(\boldsymbol{\gamma})$  satisfies

$$\lambda_{\max}\{\mathbf{V}_i(\boldsymbol{\gamma})\} = \lambda_{\max}(\mathbf{I} + \mathbf{X}_i \text{diag}(\boldsymbol{\gamma}) \mathbf{X}_i^\top) \leq 1 + \bar{\gamma} \lambda_{\max}(\mathbf{X}_i \mathbf{X}_i^\top) \leq c_1 := 1 + \bar{\gamma} \lambda_{\max}(\mathbf{X} \mathbf{X}^\top) > 0,$$

and hence the minimal eigenvalue of  $\mathbf{V}_i^{-1}(\boldsymbol{\gamma})$  satisfies

$$\lambda_{\min}\{\mathbf{V}_i^{-1}(\boldsymbol{\gamma})\} = \frac{1}{\lambda_{\max}\{\mathbf{V}_i(\boldsymbol{\gamma})\}} \geq \frac{1}{c_1}.$$

It follows from the property of eigenvalues that

$$\sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma})(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}) \geq \sum_{i=1}^m \frac{1}{c_1} \|\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}\|_2^2 = \frac{1}{c_1} \|\mathbf{y} - \mathbf{X} \boldsymbol{\beta}\|_2^2.$$

By assumption,  $\mathbf{y}$  is not in the column space spanned by  $\mathbf{X}$ , so there is a  $c_2 > 0$  such that

$$\|\mathbf{y} - \mathbf{X} \boldsymbol{\beta}\|_2^2 \geq c_2.$$

Thus, we have

$$\sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma})(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}) \geq \frac{c_2}{c_1},$$

and hence

$$n_\bullet \log \left\{ \sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma})(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}) \right\} \geq c := n_\bullet \log \left( \frac{c_2}{c_1} \right). \quad (\text{C.3})$$

Combining (C.2) and (C.3) yields the result of the lemma.  $\square$

**Lemma 2.** *Suppose  $\mathbf{y}$  does not lie in the column space of  $\mathbf{X}$  and that  $\boldsymbol{\beta} \in \mathbb{R}^p$  and  $\boldsymbol{\gamma} \in \mathbb{R}_+^p$  satisfy  $\|\boldsymbol{\beta}\|_\infty \leq \bar{\beta}$  and  $\|\boldsymbol{\gamma}\|_\infty \leq \bar{\gamma}$ . Then the negative log-likelihood gradients (2.12) and (2.13) are Lipschitz continuous in the  $k$ th block of coordinates.*

*Proof.* To facilitate exposition, define the functions

$$A(\boldsymbol{\beta}, \boldsymbol{\gamma}) = \sum_{i=1}^m \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma})(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})$$

and

$$B(\boldsymbol{\beta}, \boldsymbol{\gamma}) = \sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma})(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}).$$

Consider first the gradient (2.12) of the negative log-likelihood with respect to  $\beta_k$ , which can

be written in terms of  $A(\boldsymbol{\beta}, \boldsymbol{\gamma})$  and  $B(\boldsymbol{\beta}, \boldsymbol{\gamma})$  as

$$\nabla_{\beta_k} l(\boldsymbol{\beta}, \boldsymbol{\gamma}) = -2n_{\bullet} \frac{A(\boldsymbol{\beta}, \boldsymbol{\gamma})}{B(\boldsymbol{\beta}, \boldsymbol{\gamma})}.$$

Our task is to show that  $\nabla_{\beta_k} l(\boldsymbol{\beta}, \boldsymbol{\gamma})$  is Lipschitz continuous in block  $k$ , i.e., that there exists a constant  $L_{\beta_k} \geq 0$  such that

$$|\nabla_{\beta_k} l(\boldsymbol{\beta}^{(1)}, \boldsymbol{\gamma}^{(1)}) - \nabla_{\beta_k} l(\boldsymbol{\beta}^{(2)}, \boldsymbol{\gamma}^{(2)})| \leq L_{\beta_k} \left( |\beta_k^{(1)} - \beta_k^{(2)}| + |\gamma_k^{(1)} - \gamma_k^{(2)}| \right),$$

for any feasible  $\boldsymbol{\beta}^{(1)}$  and  $\boldsymbol{\beta}^{(2)}$  and any feasible  $\boldsymbol{\gamma}^{(1)}$  and  $\boldsymbol{\gamma}^{(2)}$  differing only in coordinate  $k$ . Hereafter, we use the shorthand  $A^{(1)}$  and  $B^{(1)}$  for  $A(\boldsymbol{\beta}^{(1)}, \boldsymbol{\gamma}^{(1)})$  and  $B(\boldsymbol{\beta}^{(1)}, \boldsymbol{\gamma}^{(1)})$ , and similarly for  $A^{(2)}$  and  $B^{(2)}$ . First, observe that

$$\begin{aligned} |\nabla_{\beta_k} l(\boldsymbol{\beta}^{(1)}, \boldsymbol{\gamma}^{(1)}) - \nabla_{\beta_k} l(\boldsymbol{\beta}^{(2)}, \boldsymbol{\gamma}^{(2)})| &= 2n_{\bullet} \left| \frac{A^{(1)}}{B^{(1)}} - \frac{A^{(2)}}{B^{(2)}} \right| \\ &= 2n_{\bullet} \left| \frac{A^{(1)}B^{(2)} - A^{(2)}B^{(1)}}{B^{(1)}B^{(2)}} \right| \\ &= 2n_{\bullet} \left| \frac{B^{(1)}(A^{(1)} - A^{(2)}) - A^{(1)}(B^{(1)} - B^{(2)})}{B^{(1)}B^{(2)}} \right| \\ &\leq \frac{2n_{\bullet}}{|B^{(1)}B^{(2)}|} \left( |B^{(1)}| |A^{(1)} - A^{(2)}| + |A^{(1)}| |B^{(1)} - B^{(2)}| \right). \end{aligned}$$

By the same arguments used in the proof of Lemma 1, there exists a positive constant (call it  $c_1$ ) independent of  $\beta_k$  and  $\gamma_k$  that lower bounds  $B^{(1)}$  and  $B^{(2)}$ , and hence  $1/|B^{(1)}B^{(2)}| \leq 1/c_1^2$ .

To upper bound  $|A^{(1)}|$  and  $|B^{(1)}|$ , we use that

$$\|\mathbf{V}_i^{-1}(\boldsymbol{\gamma})\|_2 = \lambda_{\max}\{\mathbf{V}_i^{-1}(\boldsymbol{\gamma})\} = \frac{1}{\lambda_{\min}\{\mathbf{V}_i(\boldsymbol{\gamma})\}} \leq \frac{1}{1} = 1$$

and

$$\|\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}\|_2 \leq \|\mathbf{y} - \mathbf{X} \boldsymbol{\beta}\|_2 \leq \|\mathbf{y}\|_2 + \|\mathbf{X} \boldsymbol{\beta}\|_2 \leq \|\mathbf{y}\|_2 + \bar{\beta} \sum_{j=1}^p \|\mathbf{x}_j\|_2 \leq c_2 := \|\mathbf{y}\|_2 + \bar{\beta} p,$$

since the columns of  $\mathbf{X}$  have unit  $\ell_2$  norm. It follows from the above two chains of inequalities that

$$|A^{(1)}| \leq \sum_{i=1}^m \|\mathbf{x}_{ik}\|_2 \|\mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(1)})\|_2 \|\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}^{(1)}\|_2 \leq c_2 \sum_{i=1}^m \|\mathbf{x}_{ik}\|_2 \leq c_2 \sqrt{m}$$

and

$$|B^{(1)}| \leq \sum_{i=1}^m \|\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}^{(1)}\|_2 \|\mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(1)})\|_2 \|\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}^{(1)}\|_2 \leq c_2^2 m.$$

Hence, there exists a constant  $c_3 \geq 0$  such that

$$|\nabla_{\beta_k} l(\boldsymbol{\beta}^{(1)}, \boldsymbol{\gamma}^{(1)}) - \nabla_{\beta_k} l(\boldsymbol{\beta}^{(2)}, \boldsymbol{\gamma}^{(2)})| \leq c_3 (|A^{(1)} - A^{(2)}| + |B^{(1)} - B^{(2)}|).$$

We now bound  $|A^{(1)} - A^{(2)}|$  and  $|B^{(1)} - B^{(2)}|$  in terms of  $|\beta_k^{(1)} - \beta_k^{(2)}| + |\gamma_k^{(1)} - \gamma_k^{(2)}|$ . Towards this end, we use that

$$\begin{aligned} \|\mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(1)}) - \mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(2)})\|_2 &= \|\mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(1)}) [\mathbf{V}_i(\boldsymbol{\gamma}^{(2)}) - \mathbf{V}_i(\boldsymbol{\gamma}^{(1)})] \mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(2)})\|_2 \\ &\leq \|\mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(1)})\|_2 \|\mathbf{V}_i(\boldsymbol{\gamma}^{(1)}) - \mathbf{V}_i(\boldsymbol{\gamma}^{(2)})\|_2 \|\mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(2)})\|_2 \\ &\leq \|(\mathbf{I} + \mathbf{X}_i \text{diag}(\boldsymbol{\gamma}^{(1)}) \mathbf{X}_i^\top) - (\mathbf{I} + \mathbf{X}_i \text{diag}(\boldsymbol{\gamma}^{(2)}) \mathbf{X}_i^\top)\|_2 \\ &= \|\mathbf{x}_{ik}\|_2^2 |\gamma_k^{(1)} - \gamma_k^{(2)}| \\ &\leq |\gamma_k^{(1)} - \gamma_k^{(2)}|, \end{aligned}$$

since  $\|\mathbf{x}_{ik}\|_2 \leq 1$ , and

$$\|(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}^{(1)}) - (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}^{(2)})\|_2 = \|\mathbf{X}_i(\boldsymbol{\beta}^{(1)} - \boldsymbol{\beta}^{(2)})\|_2 = \|\mathbf{x}_{ik}\|_2 |\beta_k^{(1)} - \beta_k^{(2)}| \leq |\beta_k^{(1)} - \beta_k^{(2)}|.$$

Using the above two inequalities, we upper bound  $|A^{(1)} - A^{(2)}|$  as

$$\begin{aligned}
|A^{(1)} - A^{(2)}| &= \left| \sum_{i=1}^m \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(1)})(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}^{(1)}) - \sum_{i=1}^m \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(2)})(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}^{(2)}) \right| \\
&\leq \sum_{i=1}^m \|\mathbf{x}_{ik}\|_2 \left\| \mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(1)})(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}^{(1)}) - \mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(2)})(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}^{(2)}) \right\|_2 \\
&= \sum_{i=1}^m \|\mathbf{x}_{ik}\|_2 \left\| \{\mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(1)}) - \mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(2)})\}(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}^{(1)}) \right. \\
&\quad \left. + \mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(2)})\{(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}^{(1)}) - (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}^{(2)})\} \right\|_2 \\
&\leq \sum_{i=1}^m \|\mathbf{x}_{ik}\|_2 \|\mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(1)}) - \mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(2)})\|_2 \|\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}^{(1)}\|_2 \\
&\quad + \sum_{i=1}^m \|\mathbf{x}_{ik}\|_2 \|\mathbf{V}_i^{-1}(\boldsymbol{\gamma}^{(2)})\|_2 \|(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}^{(1)}) - (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}^{(2)})\|_2 \\
&\leq c_2 \sum_{i=1}^m \|\mathbf{x}_{ik}\|_2 |\gamma_k^{(1)} - \gamma_k^{(2)}| + \sum_{i=1}^m \|\mathbf{x}_{ik}\|_2 |\beta_k^{(1)} - \beta_k^{(2)}| \\
&\leq c_2 \sqrt{m} |\gamma_k^{(1)} - \gamma_k^{(2)}| + \sqrt{m} |\beta_k^{(1)} - \beta_k^{(2)}| \\
&\leq c_4 (|\beta_k^{(1)} - \beta_k^{(2)}| + |\gamma_k^{(1)} - \gamma_k^{(2)}|),
\end{aligned}$$

where  $c_4 := (1 + c_2)\sqrt{m}$ . One can upper bound  $|B^{(1)} - B^{(2)}|$  by a similar line of working.

Call the constant in that upper bound  $c_5 \geq 0$ . It follows that

$$\begin{aligned}
|\nabla_{\beta_k} l(\boldsymbol{\beta}^{(1)}, \boldsymbol{\gamma}^{(1)}) - \nabla_{\beta_k} l(\boldsymbol{\beta}^{(2)}, \boldsymbol{\gamma}^{(2)})| &\leq c_3 (|A^{(1)} - A^{(2)}| + |B^{(1)} - B^{(2)}|) \\
&\leq c_3 \max(c_4, c_5) \left( |\beta_k^{(1)} - \beta_k^{(2)}| + |\gamma_k^{(1)} - \gamma_k^{(2)}| \right).
\end{aligned}$$

Hence,  $\nabla_{\beta_k} l(\boldsymbol{\beta}, \boldsymbol{\gamma})$  is Lipschitz with constant  $L_{\beta_k} = c_3 \max(c_4, c_5)$ . The same line of argument, omitted here for brevity, establishes that  $\nabla_{\gamma_k} l(\boldsymbol{\beta}, \boldsymbol{\gamma})$  is also Lipschitz with constant  $L_{\gamma_k} \geq 0$ .

Finally, the proof is completed by taking  $L_k = \max(L_{\beta_k}, L_{\gamma_k})$ .  $\square$

## C.2 Proof of Main Results

With the preliminary lemmas in hand, we can now prove the main results. The proof is presented in two parts: (i) a proof of convergence of the objective value, and (ii) a proof of convergence of the active set. The steps in our proof follow those used in the proof of Theorem 1 in [Dedieu et al. \(2021\)](#).

**Convergence of Objective Value** Let  $(\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*)$  be the result of applying the thresholding function (2.11) to  $(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}})$ . Since  $l(\boldsymbol{\beta}, \boldsymbol{\gamma})$  is Lipschitz on a compact set (Lemma 2), we can invoke inequality (2.8). Substituting  $\boldsymbol{\beta} = \boldsymbol{\beta}^*$  and  $\boldsymbol{\gamma} = \boldsymbol{\gamma}^*$  into (2.9) gives

$$\begin{aligned} f(\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*) &\leq \bar{f}_{L_k}(\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*; \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}) = l(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}) + \nabla l(\tilde{\boldsymbol{\beta}}_k)(\boldsymbol{\beta}_k^* - \tilde{\boldsymbol{\beta}}_k) + \nabla l(\tilde{\boldsymbol{\gamma}}_k)(\boldsymbol{\gamma}_k^* - \tilde{\boldsymbol{\gamma}}_k) \\ &\quad + \frac{L_k}{2} \left\{ (\boldsymbol{\beta}_k^* - \tilde{\boldsymbol{\beta}}_k)^2 + (\boldsymbol{\gamma}_k^* - \tilde{\boldsymbol{\gamma}}_k)^2 \right\} + \lambda\alpha \|\boldsymbol{\beta}^*\|_0 + \lambda(1 - \alpha) \|\boldsymbol{\gamma}^*\|_0. \end{aligned}$$

Next, we add and subtract  $\bar{L}_k/2\{(\boldsymbol{\beta}_k^* - \tilde{\boldsymbol{\beta}}_k)^2 + (\boldsymbol{\gamma}_k^* - \tilde{\boldsymbol{\gamma}}_k)^2\}$  on the right-hand side to get

$$\begin{aligned} f(\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*) &\leq l(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}) + \nabla l(\tilde{\boldsymbol{\beta}}_k)(\boldsymbol{\beta}_k^* - \tilde{\boldsymbol{\beta}}_k) + \nabla l(\tilde{\boldsymbol{\gamma}}_k)(\boldsymbol{\gamma}_k^* - \tilde{\boldsymbol{\gamma}}_k) \\ &\quad + \frac{\bar{L}_k}{2} \left\{ (\boldsymbol{\beta}_k^* - \tilde{\boldsymbol{\beta}}_k)^2 + (\boldsymbol{\gamma}_k^* - \tilde{\boldsymbol{\gamma}}_k)^2 \right\} + \frac{L_k - \bar{L}_k}{2} \left\{ (\boldsymbol{\beta}_k^* - \tilde{\boldsymbol{\beta}}_k)^2 + (\boldsymbol{\gamma}_k^* - \tilde{\boldsymbol{\gamma}}_k)^2 \right\} \\ &\quad + \lambda\alpha \|\boldsymbol{\beta}^*\|_0 + \lambda(1 - \alpha) \|\boldsymbol{\gamma}^*\|_0 \\ &= \bar{f}_{\bar{L}_k}(\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*; \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}) + \frac{L_k - \bar{L}_k}{2} \left\{ (\boldsymbol{\beta}_k^* - \tilde{\boldsymbol{\beta}}_k)^2 + (\boldsymbol{\gamma}_k^* - \tilde{\boldsymbol{\gamma}}_k)^2 \right\}. \end{aligned}$$

Using that  $\bar{f}_{\bar{L}_k}(\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*; \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}) \leq \bar{f}_{\bar{L}_k}(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}; \tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}) = f(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}})$  and rearranging terms, we get

$$f(\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\gamma}}) - f(\boldsymbol{\beta}^*, \boldsymbol{\gamma}^*) \geq \frac{\bar{L}_k - L_k}{2} \left\{ (\tilde{\boldsymbol{\beta}}_k - \boldsymbol{\beta}_k^*)^2 + (\tilde{\boldsymbol{\gamma}}_k - \boldsymbol{\gamma}_k^*)^2 \right\}.$$

Finally, taking  $\boldsymbol{\beta}^* = \boldsymbol{\beta}^{(t+1)}$ ,  $\tilde{\boldsymbol{\beta}} = \boldsymbol{\beta}^{(t)}$ ,  $\boldsymbol{\gamma}^* = \boldsymbol{\gamma}^{(t+1)}$ , and  $\tilde{\boldsymbol{\gamma}} = \boldsymbol{\gamma}^{(t)}$ , and recalling that all  $p$  coordinates are updated in a single cycle of coordinate descent (i.e.,  $p$  block coordinate

updates from iteration  $t$  to  $t + 1$ ), we have

$$f(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)}) - f(\boldsymbol{\beta}^{(t+1)}, \boldsymbol{\gamma}^{(t+1)}) \geq \sum_{k=1}^p \frac{\bar{L}_k - L_k}{2} \left\{ (\beta_k^{(t)} - \beta_k^{(t+1)})^2 + (\gamma_k^{(t)} - \gamma_k^{(t+1)})^2 \right\}.$$

The right-hand side is nonnegative for  $\bar{L}_k \geq L_k$ . Hence, the sequence  $\{f(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)})\}_{t \in \mathbb{N}}$  is decreasing and, because it is bounded below on a compact set (Lemma 1), convergent.

**Convergence of Active Set** We proceed using proof by contradiction. Suppose the active set does not converge in finitely many iterations. Choose  $t$  such that  $\mathcal{A}^{(t)} \neq \mathcal{A}^{(t+1)}$ . Then for some  $k = 1, \dots, p$  either (i)  $\beta_k^{(t)} = 0$  and  $\beta_k^{(t+1)} \neq 0$ , or (ii)  $\beta_k^{(t)} \neq 0$  and  $\beta_k^{(t+1)} = 0$ . Consider case (i). By the previous result we have

$$f(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)}) - f(\boldsymbol{\beta}^{(t+1)}, \boldsymbol{\gamma}^{(t+1)}) \geq \frac{\bar{L}_k - L_k}{2} (\beta_k^{(t+1)})^2.$$

Because  $\beta_k^{(t+1)}$  is the output of the thresholding function, it holds  $|\beta_k^{(t+1)}| \geq \sqrt{2\lambda\alpha/\bar{L}_k}$ , giving

$$f(\boldsymbol{\beta}^{(t)}, \boldsymbol{\gamma}^{(t)}) - f(\boldsymbol{\beta}^{(t+1)}, \boldsymbol{\gamma}^{(t+1)}) \geq \frac{\bar{L}_k - L_k}{\bar{L}_k} \lambda\alpha,$$

which is positive for  $\bar{L}_k > L_k$ . The same line of working establishes a similar inequality for case (ii). The quantity on the right-hand side is positive, indicating the objective value strictly decreases. If the active set does not converge then the objective value must decrease indefinitely, which contradicts the fact that the objective value is bounded below on a compact set (Lemma 1). Therefore, the active set must converge in finitely many iterations.

## D Proof of Proposition 2

Let  $\lambda^{(r)}$  be the current value of the sparsity parameter. Under the conditions of Theorem 1 the active set  $\mathcal{A}^{(r)}$  converges in finitely many iterations, meaning that no active predictors are removed from the model by the thresholding function (2.11) and no inactive predictors

are added. Hence, by the first condition on the inactive predictors in (2.11), it must hold for all  $k \notin \mathcal{A}^{(r)}$  (i.e., all inactive predictors) that

$$\begin{aligned} \left( \frac{-\nabla l(\beta_k^{(r)})}{\bar{L}_k} \right)^2 &< \frac{2\lambda^{(r)}\alpha}{\bar{L}_k} \\ \nabla l(\beta_k^{(r)})^2 &< 2\lambda^{(r)}\alpha\bar{L}_k \\ \frac{\nabla l(\beta_k^{(r)})^2}{2\alpha\bar{L}_k} &< \lambda^{(r)}. \end{aligned}$$

Likewise, by the second condition on the inactive predictors in (2.11), we have

$$\begin{aligned} \left( \frac{-\nabla l(\beta_k^{(r)})}{\bar{L}_k} \right)^2 + \left( \frac{\max(-\nabla l(\gamma_k^{(r)}), 0)}{\bar{L}_k} \right)^2 &< \frac{2\lambda^{(r)}}{\bar{L}_k} \\ \nabla l(\beta_k^{(r)})^2 + \max(-\nabla l(\gamma_k^{(r)}), 0)^2 &< 2\lambda^{(r)}\bar{L}_k \\ \frac{\nabla l(\beta_k^{(r)})^2 + \max(-\nabla l(\gamma_k^{(r)}), 0)^2}{2\bar{L}_k} &< \lambda^{(r)}. \end{aligned}$$

Hence, if we take  $\lambda = \lambda^{(r+1)}$ , where

$$\lambda^{(r+1)} < \max_{k \notin \mathcal{A}^{(r)}} \max \left( \frac{\nabla l(\beta_k^{(r)})^2}{2\alpha\bar{L}_k}, \frac{\nabla l(\beta_k^{(r)})^2 + \max(-\nabla l(\gamma_k^{(r)}), 0)^2}{2\bar{L}_k} \right),$$

then one of the above conditions will be violated and the active set must change.

## E Proof of Theorem 2

### E.1 Preliminary Lemmas

The proof of the theorem requires several technical lemmas, which we now state and prove. The first two lemmas provide probabilistic upper bounds on important quantities. The last two lemmas ensure the negative log-likelihood and its expectation are Lipschitz continuous.

**Lemma 3.** Let  $(\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \mathcal{C}(s)$ . Then, for any  $t > 0$ , it holds

$$\Pr \left[ \left| \frac{1}{\sigma^2} \sum_{i=1}^m [\boldsymbol{\nu}_i^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \boldsymbol{\nu}_i - \sigma^2 \text{tr} \{ \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{V}_i(\boldsymbol{\gamma}^0) \}] \right| \geq t \right] \leq 2 \exp \left\{ -c \min \left( \frac{t^2}{n_\bullet (1 + \bar{\gamma} s)^2}, \frac{t}{1 + \bar{\gamma} s} \right) \right\},$$

where  $c > 0$  is a universal constant.

*Proof.* Let  $\mathbf{h}_i = \sigma^{-1} \mathbf{V}_i^{-\frac{1}{2}}(\boldsymbol{\gamma}^0) \boldsymbol{\nu}_i$  so that  $\mathbf{h}_i \sim \mathbf{N}(\mathbf{0}, \mathbf{I})$ . Additionally, let

$$\mathbf{A}_i = \mathbf{V}_i^{\frac{1}{2}}(\boldsymbol{\gamma}^0) \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{V}_i^{\frac{1}{2}}(\boldsymbol{\gamma}^0).$$

It follows

$$\frac{1}{\sigma^2} \boldsymbol{\nu}_i^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \boldsymbol{\nu}_i = \mathbf{h}_i^\top \mathbf{A}_i \mathbf{h}_i$$

and

$$\text{tr} \{ \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{V}_i(\boldsymbol{\gamma}^0) \} = \text{tr}(\mathbf{A}_i).$$

For any  $\mathbf{h} \sim \mathbf{N}(\mathbf{0}, \mathbf{I})$ , matrix  $\mathbf{A}$ , and  $t > 0$ , the Hanson-Wright inequality ([Rudelson and Vershynin, 2013](#)) gives

$$\Pr [ |\mathbf{h}^\top \mathbf{A} \mathbf{h} - \mathbb{E}[\mathbf{h}^\top \mathbf{A} \mathbf{h}]| \geq t ] \leq 2 \exp \left\{ -c \min \left( \frac{t^2}{\|\mathbf{A}\|_F^2}, \frac{t}{\|\mathbf{A}\|_2} \right) \right\},$$

for some universal constant  $c > 0$ . Hence, using that  $\mathbb{E}[\mathbf{h}^\top \mathbf{A} \mathbf{h}] = \text{tr}(\mathbf{A})$  for any  $\mathbf{h} \sim \mathbf{N}(\mathbf{0}, \mathbf{I})$  and matrix  $\mathbf{A}$ , we have

$$\Pr [ |\mathbf{h}_i^\top \mathbf{A}_i \mathbf{h}_i - \text{tr}(\mathbf{A}_i)| \geq t ] \leq 2 \exp \left\{ -c \min \left( \frac{t^2}{\|\mathbf{A}_i\|_F^2}, \frac{t}{\|\mathbf{A}_i\|_2} \right) \right\}.$$

Now, to take the sum over  $i = 1, \dots, m$ , we use that

$$\sum_{i=1}^m \{\mathbf{h}_i^\top \mathbf{A}_i \mathbf{h}_i - \text{tr}(\mathbf{A}_i)\} = \begin{pmatrix} \mathbf{h}_1 \\ \vdots \\ \mathbf{h}_m \end{pmatrix}^\top \begin{pmatrix} \mathbf{A}_1 & \dots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \dots & \mathbf{A}_m \end{pmatrix} \begin{pmatrix} \mathbf{h}_1 \\ \vdots \\ \mathbf{h}_m \end{pmatrix} - \text{tr} \left\{ \begin{pmatrix} \mathbf{A}_1 & \dots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \dots & \mathbf{A}_m \end{pmatrix} \right\}$$

with

$$\left\| \begin{pmatrix} \mathbf{A}_1 & \dots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \dots & \mathbf{A}_m \end{pmatrix} \right\|_F^2 = \sum_{i=1}^m \|\mathbf{A}_i\|_F^2$$

and

$$\left\| \begin{pmatrix} \mathbf{A}_1 & \dots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \dots & \mathbf{A}_m \end{pmatrix} \right\|_2 = \max_{i=1, \dots, m} \|\mathbf{A}_i\|_2.$$

It immediately follows

$$\Pr \left[ \left| \sum_{i=1}^m \{\mathbf{h}_i^\top \mathbf{A}_i \mathbf{h}_i - \text{tr}(\mathbf{A}_i)\} \right| \geq t \right] \leq 2 \exp \left\{ -c \min \left( \frac{t^2}{\sum_{i=1}^m \|\mathbf{A}_i\|_F^2}, \frac{t}{\max_{i=1, \dots, m} \|\mathbf{A}_i\|_2} \right) \right\}.$$

We now upper bound the right-hand side by upper bounding the terms in the denominator on the right-hand side. First, we have

$$\|\mathbf{A}_i\|_2 = \|\mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{V}_i(\boldsymbol{\gamma}^0)\|_2 \leq \|\mathbf{V}_i^{-1}(\boldsymbol{\gamma})\|_2 \|\mathbf{V}_i(\boldsymbol{\gamma}^0)\|_2 \leq 1 + \bar{\gamma}s,$$

where we use that

$$\|\mathbf{V}_i^{-1}(\boldsymbol{\gamma})\|_2 = \frac{1}{\lambda_{\min}\{\mathbf{V}_i(\boldsymbol{\gamma})\}} \leq \frac{1}{1} = 1,$$

which follows from the lower bound on  $\lambda_{\min}\{\mathbf{V}_i(\boldsymbol{\gamma})\}$  in the proof of Lemma 1, and

$$\|\mathbf{V}_i(\boldsymbol{\gamma}^0)\|_2 \leq 1 + \sum_{k:\gamma_k^0 \neq 0} \gamma_k^0 \|\mathbf{x}_{ik}\|_2^2 \leq 1 + s\bar{\gamma},$$

since the columns of  $\mathbf{X}$  have unit  $\ell_2$  norm and  $\boldsymbol{\gamma}^0$  has  $s$  nonzero components. Next, we use that  $\|\mathbf{A}_i\|_F^2 \leq n_i \|\mathbf{A}_i\|_2^2$  in tandem with the above result to get

$$\sum_{i=1}^m \|\mathbf{A}_i\|_F^2 \leq \sum_{i=1}^m n_i \|\mathbf{A}_i\|_2^2 \leq \sum_{i=1}^m n_i (1 + \bar{\gamma}s)^2 = n_{\bullet} (1 + \bar{\gamma}s)^2.$$

Finally, we conclude

$$\Pr \left[ \left| \sum_{i=1}^m \{\mathbf{h}_i^\top \mathbf{A}_i \mathbf{h}_i - \text{tr}(\mathbf{A}_i)\} \right| \geq t \right] \leq 2 \exp \left\{ -c \min \left( \frac{t^2}{n_{\bullet} (1 + \bar{\gamma}s)^2}, \frac{t}{1 + \bar{\gamma}s} \right) \right\}.$$

□

**Lemma 4.** *Let  $(\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \mathcal{C}(s)$ . Then, for any  $t > 0$ , it holds*

$$\Pr \left[ \left| \frac{1}{\sigma^2} \sum_{i=1}^m 2(\boldsymbol{\beta}^0 - \boldsymbol{\beta})^\top \mathbf{X}_i^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \boldsymbol{\nu}_i \right| \geq t \right] \leq 2 \exp \left( -\frac{t^2}{128\sigma^{-2}(1 + \bar{\gamma}s)s^2\bar{\beta}^2} \right).$$

*Proof.* Let  $\mathbf{h}_i = \mathbf{X}_i^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \boldsymbol{\nu}_i$  so that

$$\frac{1}{\sigma^2} \sum_{i=1}^m 2(\boldsymbol{\beta}^0 - \boldsymbol{\beta})^\top \mathbf{X}_i^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \boldsymbol{\nu}_i = \frac{1}{\sigma^2} \sum_{i=1}^m 2(\boldsymbol{\beta}^0 - \boldsymbol{\beta})^\top \mathbf{h}_i.$$

Since  $\mathbf{X}_i$  and  $\mathbf{V}_i^{-1}(\boldsymbol{\gamma})$  are fixed and  $\boldsymbol{\nu}_i \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{V}_i(\boldsymbol{\gamma}^0))$ , it holds

$$\mathbf{h}_i \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{X}_i^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{V}_i(\boldsymbol{\gamma}^0) \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{X}_i).$$

Pre-multiplying  $\mathbf{h}_i$  by  $2\sigma^{-2}(\boldsymbol{\beta}^0 - \boldsymbol{\beta})^\top$ , which is also a fixed quantity, we get the scalar random

variable

$$\frac{2}{\sigma^2}(\boldsymbol{\beta}^0 - \boldsymbol{\beta})^\top \mathbf{h}_i \sim \text{N}(0, \nu_i^2),$$

where

$$\nu_i^2 = \frac{4}{\sigma^2}(\boldsymbol{\beta}^0 - \boldsymbol{\beta})^\top \mathbf{X}_i^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{V}_i(\boldsymbol{\gamma}^0) \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{X}_i (\boldsymbol{\beta}^0 - \boldsymbol{\beta}).$$

Since  $\nu_i$  and  $\nu_j$  are independent for all  $i \neq j$ , we can sum over  $i = 1, \dots, m$  to get

$$\frac{1}{\sigma^2} \sum_{i=1}^m 2(\boldsymbol{\beta}^0 - \boldsymbol{\beta})^\top \mathbf{h}_i \sim \text{N} \left( 0, \sum_{i=1}^m \nu_i^2 \right).$$

We can now invoke the standard Gaussian tail bound  $\Pr[|X| > t] \leq 2 \exp\{-t^2/(2\nu^2)\}$  for  $X \sim \text{N}(0, \nu^2)$  and  $t > 0$  to get

$$\Pr \left[ \left| \frac{1}{\sigma^2} \sum_{i=1}^m 2(\boldsymbol{\beta}^0 - \boldsymbol{\beta})^\top \mathbf{h}_i \right| \geq t \right] \leq 2 \exp \left( -\frac{t^2}{2 \sum_{i=1}^m \nu_i^2} \right).$$

Now, using that  $\|\mathbf{V}_i^{-1}(\boldsymbol{\gamma})\|_2 \leq 1$  and  $\|\mathbf{V}_i(\boldsymbol{\gamma}^0)\|_2 \leq 1 + \bar{\gamma}s$  as in the proof of Lemma 3, we can upper bound the right-hand side by upper bounding the variance term appearing in the denominator as

$$\begin{aligned} \sum_{i=1}^m \nu_i^2 &= \sum_{i=1}^m \frac{4}{\sigma^2} (\boldsymbol{\beta}^0 - \boldsymbol{\beta})^\top \mathbf{X}_i^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{V}_i(\boldsymbol{\gamma}^0) \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{X}_i (\boldsymbol{\beta}^0 - \boldsymbol{\beta}) \\ &\leq \sum_{i=1}^m \frac{4}{\sigma^2} (1 + \bar{\gamma}s) (\boldsymbol{\beta}^0 - \boldsymbol{\beta})^\top \mathbf{X}_i^\top \mathbf{X}_i (\boldsymbol{\beta}^0 - \boldsymbol{\beta}) \\ &= \sum_{i=1}^m \frac{4}{\sigma^2} (1 + \bar{\gamma}s) \|\mathbf{X}_i (\boldsymbol{\beta}^0 - \boldsymbol{\beta})\|_2^2 \\ &= \frac{4}{\sigma^2} (1 + \bar{\gamma}s) \|\mathbf{X} (\boldsymbol{\beta}^0 - \boldsymbol{\beta})\|_2^2. \end{aligned}$$

Moreover, for any  $\boldsymbol{\Delta}$  satisfying  $\|\boldsymbol{\Delta}\|_0 \leq 2s$  and  $\|\boldsymbol{\Delta}\|_\infty \leq 2\bar{\Delta}$  for some  $\bar{\Delta} > 0$ , we have

$$\|\mathbf{X}\boldsymbol{\Delta}\|_2 = \left\| \sum_{k:\Delta_k \neq 0} \Delta_k \mathbf{x}_k \right\|_2 \leq \sum_{k:\Delta_k \neq 0} |\Delta_k| \|\mathbf{x}_k\|_2 \leq \sum_{k:\Delta_k \neq 0} |\Delta_k| \leq 4s\bar{\Delta}.$$

Taking  $\Delta = \beta^0 - \beta$ , we have  $\|\Delta\|_0 \leq 2s$  and  $\|\Delta\|_\infty \leq 2\bar{\beta}$ , and hence

$$\frac{4}{\sigma^2}(1 + \bar{\gamma}s)\|\mathbf{X}(\beta^0 - \beta)\|_2^2 \leq \frac{64}{\sigma^2}(1 + \bar{\gamma}s)s^2\bar{\beta}^2.$$

Finally, we conclude

$$\Pr \left[ \left| \frac{1}{\sigma^2} \sum_{i=1}^m 2(\beta^0 - \beta)^\top \mathbf{h}_i \right| \geq t \right] \leq 2 \exp \left( -\frac{t^2}{128\sigma^{-2}(1 + \bar{\gamma}s)s^2\bar{\beta}^2} \right).$$

□

**Lemma 5.** Let  $\mathcal{S}_s$  denote the collection of all sets  $\mathcal{A} \subseteq \{1, \dots, p\}$  with cardinality  $|\mathcal{A}| \leq s$ .

Fix any  $\mathcal{A} \in \mathcal{S}_s$ . Define the constraint set

$$\mathcal{B}(\mathcal{A}) := \{(\beta, \gamma) \in \mathbb{R}^p \times \mathbb{R}_+^p : \|\beta\|_\infty \leq \bar{\beta}, \|\gamma\|_\infty \leq \bar{\gamma}, \text{Supp}(\gamma) \subseteq \text{Supp}(\beta) \subseteq \mathcal{A}\}$$

and its convex hull

$$\tilde{\mathcal{B}}(\mathcal{A}) := \{(\beta, \gamma) \in \mathbb{R}^p \times \mathbb{R}_+^p : \|\beta\|_\infty \leq \bar{\beta}, \|\gamma\|_\infty \leq \bar{\gamma}, \text{Supp}(\gamma) \subseteq \mathcal{A}, \text{Supp}(\beta) \subseteq \mathcal{A}\}.$$

Define the random quantity  $L_{\mathcal{A}}(\mathbf{y})$  as

$$L_{\mathcal{A}}(\mathbf{y}) := \sup_{(\beta, \gamma) \in \tilde{\mathcal{B}}(\mathcal{A})} \|\nabla_{\beta, \gamma} l(\beta, \gamma)\|_2.$$

Then almost surely  $L_{\mathcal{A}}(\mathbf{y}) < \infty$  and on that event  $l(\beta, \gamma)$  is Lipschitz continuous on  $\mathcal{B}(\mathcal{A})$ ,

i.e., it holds

$$|l(\beta, \gamma) - l(\beta', \gamma')| \leq L_{\mathcal{A}}(\mathbf{y}) \sqrt{\|\beta - \beta'\|_2^2 + \|\gamma - \gamma'\|_2^2}$$

for any  $(\beta, \gamma) \in \mathcal{B}(\mathcal{A})$  and  $(\beta', \gamma') \in \mathcal{B}(\mathcal{A})$ . Moreover, for any  $\delta \in (0, 1]$ , it holds

$$\Pr \left[ \max_{\mathcal{A} \in \mathcal{S}_s} L_{\mathcal{A}}(\mathbf{y}) \leq L(\delta) \right] \geq 1 - \delta,$$

where  $L(\delta)$  is a fixed quantity defined as

$$L(\delta) := \sqrt{2s} \max \left( \frac{2}{\sigma^2} \{2\bar{\beta}s + g(\delta)\}, 1 + \frac{1}{\sigma^2} \{2\bar{\beta}s + g(\delta)\}^2 \right)$$

with

$$g(\delta) := \sigma \sqrt{1 + \bar{\gamma}s} \sqrt{n_{\bullet} + 2\sqrt{n_{\bullet} \log \left( \frac{1}{\delta} \right) + 2 \log \left( \frac{1}{\delta} \right)}}.$$

*Proof.* We begin by proving the first claim of the lemma that  $l$  is Lipschitz with constant  $L_{\mathcal{A}}(\mathbf{y})$ . Towards this end, we evaluate the components of  $\nabla_{\beta_{\mathcal{A}}, \gamma_{\mathcal{A}}} l(\boldsymbol{\beta}, \boldsymbol{\gamma})$ . First, rewrite  $l(\boldsymbol{\beta}, \boldsymbol{\gamma}) = l_1(\boldsymbol{\gamma}) + l_2(\boldsymbol{\beta}, \boldsymbol{\gamma})$ , where

$$l_1(\boldsymbol{\gamma}) = \sum_{i=1}^m \log \det \{ \sigma^2 \mathbf{V}_i(\boldsymbol{\gamma}) \}$$

and

$$l_2(\boldsymbol{\beta}, \boldsymbol{\gamma}) = \frac{1}{\sigma^2} \sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^{\top} \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}).$$

For  $l_1$ , a routine derivation gives the components of the gradient with respect to  $\boldsymbol{\gamma}$  as

$$\nabla_{\gamma_k} l_1(\boldsymbol{\gamma}) = \sum_{i=1}^m \mathbf{x}_{ik}^{\top} \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{x}_{ik}.$$

Likewise, for  $l_2$ , the components of the gradient with respect to  $\boldsymbol{\beta}$  and  $\boldsymbol{\gamma}$  are respectively given by

$$\nabla_{\beta_k} l_2(\boldsymbol{\beta}, \boldsymbol{\gamma}) = -\frac{1}{\sigma^2} \sum_{i=1}^m 2\mathbf{x}_{ik}^{\top} \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})$$

and

$$\nabla_{\gamma_k} l_2(\boldsymbol{\beta}, \boldsymbol{\gamma}) = -\frac{1}{\sigma^2} \sum_{i=1}^m \{ \mathbf{x}_{ik}^{\top} \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}) \}^2.$$

All three of the above gradients are continuous and  $\tilde{\mathcal{B}}(\mathcal{A})$  is compact. Recall also that  $\mathbf{y} \sim \mathbf{N}(\mathbf{X}\boldsymbol{\beta}^0, \sigma^2 \mathbf{V}(\boldsymbol{\gamma}^0))$ , so all components of  $\mathbf{y}$  are almost surely finite. Hence, the gradient

norm  $\|\nabla_{\beta_{\mathcal{A}}, \gamma_{\mathcal{A}}} l(\beta, \gamma)\|_2$  attains a maximum and thus

$$L_{\mathcal{A}}(\mathbf{y}) = \sup_{(\beta, \gamma) \in \tilde{\mathcal{B}}(\mathcal{A})} \|\nabla_{\beta_{\mathcal{A}}, \gamma_{\mathcal{A}}} l(\beta, \gamma)\|_2 < \infty$$

almost surely. Moreover, since  $\tilde{\mathcal{B}}(\mathcal{A})$  is convex, we can conclude by the mean value theorem that

$$|l(\beta, \gamma) - l(\beta', \gamma')| \leq L_{\mathcal{A}}(\mathbf{y}) \sqrt{\|\beta - \beta'\|_2^2 + \|\gamma - \gamma'\|_2^2}$$

for any  $\beta' \neq \beta$  and any  $\gamma' \neq \gamma$  in  $\tilde{\mathcal{B}}(\mathcal{A})$  almost surely. Furthermore, because  $\mathcal{B}(\mathcal{A}) \subseteq \tilde{\mathcal{B}}(\mathcal{A})$ , the above inequality also holds for any  $\beta' \neq \beta$  and any  $\gamma' \neq \gamma$  in  $\mathcal{B}(\mathcal{A})$ , thereby establishing the first claim of the lemma that  $l(\beta, \gamma)$  is almost surely  $L_{\mathcal{A}}(\mathbf{y})$  Lipschitz on  $\mathcal{B}(\mathcal{A})$ . To establish the second claim of the lemma, we bound the components of  $L_{\mathcal{A}}(\mathbf{y})$ . Observe that

$$|\nabla_{\gamma_k} l_1(\gamma)| = \left| \sum_{i=1}^m \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\gamma) \mathbf{x}_{ik} \right| \leq \sum_{i=1}^m \|\mathbf{x}_{ik}\|_2 \|\mathbf{V}_i^{-1}(\gamma)\|_2 \leq \sum_{i=1}^m \|\mathbf{x}_{ik}\|_2^2 = 1,$$

where we use that  $\|\mathbf{V}_i^{-1}(\gamma)\|_2 \leq 1$  as in the proof of Lemma 3 and that the columns of  $\mathbf{X}$  have unit  $\ell_2$  norm. Next, see that

$$\begin{aligned} |\nabla_{\beta_k} l_2(\beta, \gamma)| &= \left| -\frac{1}{\sigma^2} \sum_{i=1}^m 2\mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\gamma) (\mathbf{y}_i - \mathbf{X}_i \beta) \right| \\ &\leq \frac{1}{\sigma^2} \sum_{i=1}^m 2\|\mathbf{x}_{ik}\|_2 \|\mathbf{y}_i - \mathbf{X}_i \beta\|_2 \\ &\leq \frac{2}{\sigma^2} \sqrt{\sum_{i=1}^m \|\mathbf{x}_{ik}\|_2^2} \sqrt{\sum_{i=1}^m \|\mathbf{y}_i - \mathbf{X}_i \beta\|_2^2} \\ &= \frac{2}{\sigma^2} \|\mathbf{y} - \mathbf{X} \beta\|_2 \\ &\leq \frac{2}{\sigma^2} (\|\mathbf{y}\|_2 + \|\mathbf{X} \beta\|_2) \\ &\leq \frac{2}{\sigma^2} (\|\mathbf{y}\|_2 + \bar{\beta} s), \end{aligned}$$

where the second inequality follows from the Cauchy-Schwarz inequality and we use that

$$\|\mathbf{X}\boldsymbol{\beta}\|_2 = \left\| \sum_{k:\beta_k \neq 0} \beta_k \mathbf{x}_k \right\|_2 \leq \sum_{k:\beta_k \neq 0} |\beta_k| \|\mathbf{x}_k\|_2 \leq \sum_{k:\beta_k \neq 0} |\beta_k| \leq \bar{\beta}s.$$

Finally, observe that

$$\begin{aligned} |\nabla_{\gamma_k} l_2(\boldsymbol{\beta}, \boldsymbol{\gamma})| &= \left| -\frac{1}{\sigma^2} \sum_{i=1}^m [\mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma})(\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})]^2 \right| \\ &\leq \frac{1}{\sigma^2} \sum_{i=1}^m \|\mathbf{x}_{ik}\|_2^2 \|\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}\|_2^2 \\ &= \frac{1}{\sigma^2} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 \\ &\leq \frac{1}{\sigma^2} (\|\mathbf{y}\|_2 + \|\mathbf{X}\boldsymbol{\beta}\|_2)^2 \\ &\leq \frac{1}{\sigma^2} (\|\mathbf{y}\|_2 + \bar{\beta}s)^2. \end{aligned}$$

Now, to bound  $\|\mathbf{y}\|_2$ , we use the inequality

$$\|\mathbf{y}\|_2 \leq \|\mathbf{X}\boldsymbol{\beta}^0\|_2 + \|\boldsymbol{\nu}\|_2 = \|\mathbf{X}\boldsymbol{\beta}^0\|_2 + \sigma \|\mathbf{V}^{\frac{1}{2}}(\boldsymbol{\gamma}^0)\mathbf{z}\|_2 \leq \bar{\beta}s + \sigma \sqrt{\lambda_{\max}\{\mathbf{V}(\boldsymbol{\gamma}^0)\}} \|\mathbf{z}\|_2,$$

where  $\mathbf{z} \sim \mathbf{N}(\mathbf{0}, \mathbf{I})$ . As in the proof of Lemma 3, it holds

$$\lambda_{\max}\{\mathbf{V}(\boldsymbol{\gamma}^0)\} = \max_{i=1, \dots, m} \lambda_{\max}\{\mathbf{V}_i(\boldsymbol{\gamma}^0)\} \leq 1 + \bar{\gamma}s.$$

Since  $\mathbf{z}$  has standard normal components, a Laurent-Massart  $\chi^2$  tail bound (Laurent and Massart, 2000) gives

$$\Pr \left[ \|\mathbf{z}\|_2^2 \geq n_\bullet + 2\sqrt{n_\bullet t} + 2t \right] \leq \exp(-t).$$

Substituting  $t = \log(1/\delta)$  and taking square roots gives

$$\Pr \left[ \|\mathbf{z}\|_2 \geq \sqrt{n_\bullet + 2\sqrt{n_\bullet \log\left(\frac{1}{\delta}\right)} + 2\log\left(\frac{1}{\delta}\right)} \right] \leq \delta$$

and hence it holds with probability at least  $1 - \delta$  that

$$\|\mathbf{y}\|_2 \leq \bar{\beta}s + g(\delta),$$

where

$$g(\delta) = \sigma\sqrt{1 + \bar{\gamma}s} \sqrt{n_{\bullet} + 2\sqrt{n_{\bullet} \log\left(\frac{1}{\delta}\right)} + 2\log\left(\frac{1}{\delta}\right)}.$$

Substituting the above result into the gradient bounds and taking the maximum yields

$$L_{\mathcal{A}}(\mathbf{y}) \leq \sqrt{2s} \max\left(\frac{2}{\sigma^2}\{2\bar{\beta}s + g(\delta)\}, 1 + \frac{1}{\sigma^2}\{2\bar{\beta}s + g(\delta)\}^2\right) = L(\delta)$$

with probability at least  $1 - \delta$ . Finally, since  $L(\delta)$  does not depend on  $\mathcal{A}$  it holds

$$\Pr\left[\max_{\mathcal{A} \in \mathcal{S}_s} L_{\mathcal{A}}(\mathbf{y}) \leq L(\delta)\right] \geq 1 - \delta.$$

□

**Lemma 6.** *Let  $\mathcal{S}_s$  denote the collection of all sets  $\mathcal{A} \subseteq \{1, \dots, p\}$  with cardinality  $|\mathcal{A}| \leq s$ .*

*Fix any  $\mathcal{A} \in \mathcal{S}_s$ . Define the constraint set*

$$\mathcal{B}(\mathcal{A}) := \{(\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \mathbb{R}^p \times \mathbb{R}_+^p : \|\boldsymbol{\beta}\|_{\infty} \leq \bar{\beta}, \|\boldsymbol{\gamma}\|_{\infty} \leq \bar{\gamma}, \text{Supp}(\boldsymbol{\gamma}) \subseteq \text{Supp}(\boldsymbol{\beta}) \subseteq \mathcal{A}\}$$

*and its convex hull*

$$\tilde{\mathcal{B}}(\mathcal{A}) := \{(\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \mathbb{R}^p \times \mathbb{R}_+^p : \|\boldsymbol{\beta}\|_{\infty} \leq \bar{\beta}, \|\boldsymbol{\gamma}\|_{\infty} \leq \bar{\gamma}, \text{Supp}(\boldsymbol{\gamma}) \subseteq \mathcal{A}, \text{Supp}(\boldsymbol{\beta}) \subseteq \mathcal{A}\}.$$

*Define the quantity*

$$L_{\mathcal{A}} := \sup_{(\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \tilde{\mathcal{B}}(\mathcal{A})} \|\nabla_{\boldsymbol{\beta}, \boldsymbol{\gamma}} \mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})]\|_2.$$

Then  $L_{\mathcal{A}} < \infty$  and  $\mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})]$  is Lipschitz continuous on  $\mathcal{B}(\mathcal{A})$ , i.e., it holds

$$|\mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})] - \mathbb{E}[l(\boldsymbol{\beta}', \boldsymbol{\gamma}')]| \leq L_{\mathcal{A}} \sqrt{\|\boldsymbol{\beta} - \boldsymbol{\beta}'\|_2^2 + \|\boldsymbol{\gamma} - \boldsymbol{\gamma}'\|_2^2}$$

for any  $(\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \mathcal{B}(\mathcal{A})$  and  $(\boldsymbol{\beta}', \boldsymbol{\gamma}') \in \mathcal{B}(\mathcal{A})$ . Moreover, it holds

$$\max_{\mathcal{A} \in \mathcal{S}_s} L_{\mathcal{A}} \leq L := \sqrt{2s} \max \left( \frac{8}{\sigma^2} \bar{\beta} s, 2 + \frac{16}{\sigma^2} \bar{\beta}^2 s^2 + \bar{\gamma} s \right).$$

*Proof.* We start by establishing the first claim of the lemma that the expectation of  $l$  is Lipschitz with constant  $L_{\mathcal{A}}$ . Towards that end, we consider the components of  $\nabla_{\boldsymbol{\beta}, \boldsymbol{\gamma}, \mathcal{A}} \mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})]$ .

First, rewrite  $\mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})] = \mathbb{E}[l_1(\boldsymbol{\gamma})] + \mathbb{E}[l_2(\boldsymbol{\beta}, \boldsymbol{\gamma})]$ , where

$$\mathbb{E}[l_1(\boldsymbol{\gamma})] = \mathbb{E} \left[ \sum_{i=1}^m \log \det \{ \sigma^2 \mathbf{V}_i(\boldsymbol{\gamma}) \} \right] = \sum_{i=1}^m \log \det \{ \sigma^2 \mathbf{V}_i(\boldsymbol{\gamma}) \}$$

and

$$\begin{aligned} \mathbb{E}[l_2(\boldsymbol{\beta}, \boldsymbol{\gamma})] &= \mathbb{E} \left[ \frac{1}{\sigma^2} \sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}) \right] \\ &= \frac{1}{\sigma^2} \sum_{i=1}^m \left[ (\boldsymbol{\beta}^0 - \boldsymbol{\beta})^\top \mathbf{X}_i^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{X}_i (\boldsymbol{\beta}^0 - \boldsymbol{\beta}) + \sigma^2 \text{tr} \{ \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{V}_i(\boldsymbol{\gamma}^0) \} \right]. \end{aligned}$$

For the expectation of  $l_1$ , some standard calculations give the components of the gradient with respect to  $\boldsymbol{\gamma}$  as

$$\nabla_{\gamma_k} \mathbb{E}[l_1(\boldsymbol{\gamma})] = \sum_{i=1}^m \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{x}_{ik}.$$

Similarly, for the expectation of  $l_2$ , the components of the gradient with respect to  $\boldsymbol{\beta}$  and  $\boldsymbol{\gamma}$  are respectively given as

$$\nabla_{\beta_k} \mathbb{E}[l_2(\boldsymbol{\beta}, \boldsymbol{\gamma})] = -\frac{1}{\sigma^2} \sum_{i=1}^m 2 \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{X}_i (\boldsymbol{\beta}^0 - \boldsymbol{\beta})$$

and

$$\nabla_{\gamma_k} \mathbb{E}[l_2(\boldsymbol{\beta}, \boldsymbol{\gamma})] = -\frac{1}{\sigma^2} \sum_{i=1}^m \{ \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{X}_i (\boldsymbol{\beta}^0 - \boldsymbol{\beta}) \}^2 - \sum_{i=1}^m \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{V}_i(\boldsymbol{\gamma}^0) \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{x}_{ik}.$$

The above three gradients are continuous and  $\tilde{\mathcal{B}}(\mathcal{A})$  is compact. Thus, the gradient norm  $\|\nabla_{\boldsymbol{\beta}, \boldsymbol{\gamma}, \mathcal{A}} \mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})]\|_2$  attains a maximum and so it holds

$$L_{\mathcal{A}} = \sup_{(\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \tilde{\mathcal{B}}(\mathcal{A})} \|\nabla_{\boldsymbol{\beta}, \boldsymbol{\gamma}, \mathcal{A}} \mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})]\|_2 < \infty.$$

Furthermore, since  $\tilde{\mathcal{B}}(\mathcal{A})$  is convex, we can apply the mean value theorem to get

$$|\mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})] - \mathbb{E}[l(\boldsymbol{\beta}', \boldsymbol{\gamma}')]| \leq L_{\mathcal{A}} \sqrt{\|\boldsymbol{\beta} - \boldsymbol{\beta}'\|_2^2 + \|\boldsymbol{\gamma} - \boldsymbol{\gamma}'\|_2^2}$$

for any  $\boldsymbol{\beta}' \neq \boldsymbol{\beta}$  and any  $\boldsymbol{\gamma}' \neq \boldsymbol{\gamma}$  in  $\tilde{\mathcal{B}}(\mathcal{A})$ . Moreover, because  $\mathcal{B}(\mathcal{A}) \subseteq \tilde{\mathcal{B}}(\mathcal{A})$ , the above inequality also holds for any  $\boldsymbol{\beta}' \neq \boldsymbol{\beta}$  and any  $\boldsymbol{\gamma}' \neq \boldsymbol{\gamma}$  in  $\mathcal{B}(\mathcal{A})$ , thereby proving the first claim of the lemma that  $\mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})]$  is  $L_{\mathcal{A}}$  Lipschitz on  $\mathcal{B}(\mathcal{A})$ . To prove the second claim of the lemma, we bound the components of  $L_{\mathcal{A}}$ . See that

$$|\nabla_{\gamma_k} \mathbb{E}[l_1(\boldsymbol{\gamma})]| = \left| \sum_{i=1}^m \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{x}_{ik} \right| \leq \sum_{i=1}^m \|\mathbf{x}_{ik}\|_2^2 \|\mathbf{V}_i^{-1}(\boldsymbol{\gamma})\|_2 \leq \sum_{i=1}^m \|\mathbf{x}_{ik}\|_2^2 = 1.$$

Next, observe that

$$\begin{aligned}
|\nabla_{\beta_k} \mathbb{E}[l_2(\boldsymbol{\beta}, \boldsymbol{\gamma})]| &= \left| -\frac{1}{\sigma^2} \sum_{i=1}^m 2\mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{X}_i(\boldsymbol{\beta}^0 - \boldsymbol{\beta}) \right| \\
&\leq \frac{1}{\sigma^2} \sum_{i=1}^m 2\|\mathbf{x}_{ik}\|_2 \|\mathbf{X}_i(\boldsymbol{\beta}^0 - \boldsymbol{\beta})\|_2 \\
&\leq \frac{2}{\sigma^2} \sqrt{\sum_{i=1}^m \|\mathbf{x}_{ik}\|_2^2} \sqrt{\sum_{i=1}^m \|\mathbf{X}_i(\boldsymbol{\beta}^0 - \boldsymbol{\beta})\|_2^2} \\
&= \frac{2}{\sigma^2} \|\mathbf{X}(\boldsymbol{\beta}^0 - \boldsymbol{\beta})\|_2 \\
&\leq \frac{8}{\sigma^2} \bar{\beta} s,
\end{aligned}$$

where we employ the same tricks used in the previous lemma. Finally, see that

$$\begin{aligned}
|\nabla_{\gamma_k} \mathbb{E}[l_2(\boldsymbol{\beta}, \boldsymbol{\gamma})]| &= \left| -\frac{1}{\sigma^2} \sum_{i=1}^m \left\{ \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{X}_i(\boldsymbol{\beta}^0 - \boldsymbol{\beta}) \right\}^2 \right. \\
&\quad \left. - \sum_{i=1}^m \mathbf{x}_{ik}^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{V}_i(\boldsymbol{\gamma}^0) \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{x}_{ik} \right| \\
&\leq \frac{1}{\sigma^2} \sum_{i=1}^m \|\mathbf{x}_{ik}\|_2^2 \|\mathbf{X}_i(\boldsymbol{\beta}^0 - \boldsymbol{\beta})\|_2^2 + \sum_{i=1}^m \|\mathbf{x}_{ik}\|_2^2 \|\mathbf{V}_i(\boldsymbol{\gamma}^0)\|_2 \\
&\leq \frac{1}{\sigma^2} \|\mathbf{X}(\boldsymbol{\beta}^0 - \boldsymbol{\beta})\|_2^2 + 1 + \bar{\gamma} s \\
&\leq \frac{16}{\sigma^2} \bar{\beta}^2 s^2 + 1 + \bar{\gamma} s.
\end{aligned}$$

Taking the maximum over the gradient bounds yields

$$L_{\mathcal{A}} \leq \sqrt{2s} \max \left( \frac{8}{\sigma^2} \bar{\beta} s, 2 + \frac{16}{\sigma^2} \bar{\beta}^2 s^2 + \bar{\gamma} s \right) = L,$$

and, since  $L$  does not depend on  $\mathcal{A}$ , it holds

$$\max_{\mathcal{A} \in \mathcal{S}_s} L_{\mathcal{A}} \leq L.$$

□

## E.2 Proof of Main Result

Since  $(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})$  minimizes the negative log-likelihood, we necessarily have

$$l(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}}) - l(\boldsymbol{\beta}^0, \boldsymbol{\gamma}^0) \leq 0.$$

Subtracting and adding  $\mathbb{E}[l(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})]$  and  $\mathbb{E}[l(\boldsymbol{\beta}^0, \boldsymbol{\gamma}^0)]$  on the left-hand side gives

$$l(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}}) - \mathbb{E}[l(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})] + \mathbb{E}[l(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})] - \mathbb{E}[l(\boldsymbol{\beta}^0, \boldsymbol{\gamma}^0)] + \mathbb{E}[l(\boldsymbol{\beta}^0, \boldsymbol{\gamma}^0)] - l(\boldsymbol{\beta}^0, \boldsymbol{\gamma}^0) \leq 0.$$

Now, isolating  $\mathbb{E}[l(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})] - \mathbb{E}[l(\boldsymbol{\beta}^0, \boldsymbol{\gamma}^0)]$  on the left-hand side gives the upper bound

$$\begin{aligned} \mathbb{E}[l(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})] - \mathbb{E}[l(\boldsymbol{\beta}^0, \boldsymbol{\gamma}^0)] &\leq \mathbb{E}[l(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})] - l(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}}) + l(\boldsymbol{\beta}^0, \boldsymbol{\gamma}^0) - \mathbb{E}[l(\boldsymbol{\beta}^0, \boldsymbol{\gamma}^0)] \\ &\leq \left| \mathbb{E}[l(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})] - l(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}}) \right| + \left| l(\boldsymbol{\beta}^0, \boldsymbol{\gamma}^0) - \mathbb{E}[l(\boldsymbol{\beta}^0, \boldsymbol{\gamma}^0)] \right| \quad (\text{E.1}) \\ &\leq 2 \sup_{(\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \mathcal{C}(s)} |l(\boldsymbol{\beta}, \boldsymbol{\gamma}) - \mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})]|. \end{aligned}$$

We now seek a high-probability upper bound for the supremum on the right-hand side.

Towards this end, consider the difference inside the supremum, given by

$$\begin{aligned} l(\boldsymbol{\beta}, \boldsymbol{\gamma}) - \mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})] &= \sum_{i=1}^m \log \det\{\sigma^2 \mathbf{V}_i(\boldsymbol{\gamma})\} + \frac{1}{\sigma^2} \sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}) \\ &\quad - \mathbb{E} \left[ \sum_{i=1}^m \log \det\{\sigma^2 \mathbf{V}_i(\boldsymbol{\gamma})\} + \frac{1}{\sigma^2} \sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}) \right]. \end{aligned}$$

Since the log determinant terms do not depend on  $\boldsymbol{\nu}_i$ , which is the only source of randomness, we can simplify the above expression to

$$\begin{aligned} l(\boldsymbol{\beta}, \boldsymbol{\gamma}) - \mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})] &= \frac{1}{\sigma^2} \sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}) \\ &\quad - \mathbb{E} \left[ \frac{1}{\sigma^2} \sum_{i=1}^m (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta})^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}) \right]. \end{aligned}$$

Expanding the quadratic terms and evaluating the expectations yields

$$\begin{aligned} l(\boldsymbol{\beta}, \boldsymbol{\gamma}) - \mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})] &= \frac{1}{\sigma^2} \sum_{i=1}^m [\boldsymbol{\nu}_i^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \boldsymbol{\nu}_i - \sigma^2 \text{tr} \{ \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{V}_i(\boldsymbol{\gamma}^0) \}] \\ &\quad + \frac{1}{\sigma^2} \sum_{i=1}^m 2(\boldsymbol{\beta}^0 - \boldsymbol{\beta})^\top \mathbf{X}_i^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \boldsymbol{\nu}_i. \end{aligned}$$

For the first term on the right-hand side, Lemma 3 gives the probabilistic upper bound

$$\begin{aligned} \Pr \left[ \left| \frac{1}{\sigma^2} \sum_{i=1}^m [\boldsymbol{\nu}_i^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \boldsymbol{\nu}_i - \sigma^2 \text{tr} \{ \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \mathbf{V}_i(\boldsymbol{\gamma}^0) \}] \right| \geq t \right] \\ \leq 2 \exp \left\{ -c \min \left( \frac{t^2}{n_\bullet (1 + \bar{\gamma}s)^2}, \frac{t}{1 + \bar{\gamma}s} \right) \right\}, \end{aligned}$$

where  $c > 0$  is a universal constant and  $t > 0$ . Similarly, for the second term, Lemma 4 gives the probabilistic upper bound

$$\Pr \left[ \left| \frac{1}{\sigma^2} \sum_{i=1}^m 2(\boldsymbol{\beta}^0 - \boldsymbol{\beta})^\top \mathbf{X}_i^\top \mathbf{V}_i^{-1}(\boldsymbol{\gamma}) \boldsymbol{\nu}_i \right| \geq t \right] \leq 2 \exp \left( -\frac{t^2}{128\sigma^{-2}(1 + \bar{\gamma}s)s^2\beta^2} \right).$$

These bounds hold for a single fixed  $(\boldsymbol{\beta}, \boldsymbol{\gamma})$ . We now extend them to hold uniformly over  $\mathcal{C}(s)$  via an  $\epsilon$ -net argument. Towards this end, define the set  $\mathcal{B}(\mathcal{A})$  as

$$\mathcal{B}(\mathcal{A}) := \{ (\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \mathbb{R}^p \times \mathbb{R}_+^p : \|\boldsymbol{\beta}\|_\infty \leq \bar{\beta}, \|\boldsymbol{\gamma}\|_\infty \leq \bar{\gamma}, \text{Supp}(\boldsymbol{\gamma}) \subseteq \text{Supp}(\boldsymbol{\beta}) \subseteq \mathcal{A} \}.$$

Then, letting  $\mathcal{S}_s$  be the set of all possible active sets  $\mathcal{A}$  of size at most  $s$ , we have

$$\sup_{(\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \mathcal{C}(s)} |l(\boldsymbol{\beta}, \boldsymbol{\gamma}) - \mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})]| = \max_{\mathcal{A} \in \mathcal{S}_s} \sup_{(\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \mathcal{B}(\mathcal{A})} |l(\boldsymbol{\beta}, \boldsymbol{\gamma}) - \mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})]|.$$

Next, we require a useful inequality. For any  $\beta' \neq \beta$  and any  $\gamma' \neq \gamma$  in  $\mathcal{B}(\mathcal{A})$  it holds

$$\begin{aligned}
& |l(\beta, \gamma) - \mathbb{E}[l(\beta, \gamma)]| \\
&= |l(\beta, \gamma) - l(\beta', \gamma') + l(\beta', \gamma') - \mathbb{E}[l(\beta', \gamma')] + \mathbb{E}[l(\beta', \gamma')] - \mathbb{E}[l(\beta, \gamma)]| \\
&\leq |l(\beta', \gamma') - \mathbb{E}[l(\beta', \gamma')]| + |l(\beta, \gamma) - l(\beta', \gamma')| + |\mathbb{E}[l(\beta', \gamma')] - \mathbb{E}[l(\beta, \gamma)]| \\
&\leq |l(\beta', \gamma') - \mathbb{E}[l(\beta', \gamma')]| + (L_{\mathcal{A}}(\mathbf{y}) + L_{\mathcal{A}})\sqrt{\|\beta - \beta'\|_2^2 + \|\gamma - \gamma'\|_2^2},
\end{aligned}$$

where the last inequality follows from the first results of Lemmas 5 and 6, which give that  $l(\beta, \gamma)$  and  $\mathbb{E}[l(\beta, \gamma)]$  are Lipschitz on  $\mathcal{B}(\mathcal{A})$  with constants  $L_{\mathcal{A}}(\mathbf{y})$  (a random quantity) and  $L_{\mathcal{A}}$  (a fixed quantity), respectively. Now, let  $\mathcal{E}(\mathcal{A})$  be an  $\epsilon$ -net of  $\mathcal{B}(\mathcal{A})$  such that for any  $(\beta, \gamma) \in \mathcal{B}(\mathcal{A})$  there exists a  $(\beta', \gamma') \in \mathcal{E}(\mathcal{A})$  with  $\sqrt{\|\beta - \beta'\|_2^2 + \|\gamma - \gamma'\|_2^2} \leq \epsilon$ . We then have

$$(L_{\mathcal{A}}(\mathbf{y}) + L_{\mathcal{A}})\sqrt{\|\beta - \beta'\|_2^2 + \|\gamma - \gamma'\|_2^2} \leq (L_{\mathcal{A}}(\mathbf{y}) + L_{\mathcal{A}})\epsilon \leq 2 \max(L_{\mathcal{A}}(\mathbf{y}), L_{\mathcal{A}})\epsilon,$$

from which it follows

$$\begin{aligned}
& \max_{\mathcal{A} \in \mathcal{S}_s} \sup_{(\beta, \gamma) \in \mathcal{B}(\mathcal{A})} |l(\beta, \gamma) - \mathbb{E}[l(\beta, \gamma)]| \\
&\leq \max_{\mathcal{A} \in \mathcal{S}_s} \max_{(\beta', \gamma') \in \mathcal{E}(\mathcal{A})} (|l(\beta', \gamma') - \mathbb{E}[l(\beta', \gamma')]| + 2 \max(L_{\mathcal{A}}(\mathbf{y}), L_{\mathcal{A}})\epsilon).
\end{aligned}$$

Invoking the second result of Lemma 5 gives the high-probability bound

$$\Pr \left[ \max_{\mathcal{A} \in \mathcal{S}_s} L_{\mathcal{A}}(\mathbf{y}) \leq L(\delta_1) \right] \geq 1 - \delta_1 \tag{E.2}$$

for  $\delta_1 \in (0, 1]$ . Similarly, by the second result of Lemma 6, we have the deterministic bound

$$\max_{\mathcal{A} \in \mathcal{S}_s} L_{\mathcal{A}} \leq L.$$

Both  $L(\delta_1)$  and  $L$  are deterministic quantities. Define  $L^*(\delta_1)$  to be

$$L^*(\delta_1) := \max(L(\delta_1), L).$$

Then, with probability  $1 - \delta_1$ , it holds

$$\begin{aligned} & \max_{\mathcal{A} \in \mathcal{S}_s} \max_{(\boldsymbol{\beta}', \boldsymbol{\gamma}') \in \mathcal{E}(\mathcal{A})} (|l(\boldsymbol{\beta}', \boldsymbol{\gamma}') - \mathbb{E}[l(\boldsymbol{\beta}', \boldsymbol{\gamma}')]| + 2 \max(L_{\mathcal{A}}(\mathbf{y}), L_{\mathcal{A}})\epsilon) \\ & \leq \max_{\mathcal{A} \in \mathcal{S}_s} \max_{(\boldsymbol{\beta}', \boldsymbol{\gamma}') \in \mathcal{E}(\mathcal{A})} |l(\boldsymbol{\beta}', \boldsymbol{\gamma}') - \mathbb{E}[l(\boldsymbol{\beta}', \boldsymbol{\gamma}')]| + 2L^*(\delta_1)\epsilon. \end{aligned}$$

In what follows, every inequality that contains  $L^*(\delta_1)$  is stated on the event in (E.2). We now have

$$\begin{aligned} & \Pr \left[ \sup_{(\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \mathcal{C}(s)} |l(\boldsymbol{\beta}, \boldsymbol{\gamma}) - \mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})]| \geq t \right] \\ & \leq \Pr \left[ \max_{\mathcal{A} \in \mathcal{S}_s} \max_{(\boldsymbol{\beta}', \boldsymbol{\gamma}') \in \mathcal{E}(\mathcal{A})} |l(\boldsymbol{\beta}', \boldsymbol{\gamma}') - \mathbb{E}[l(\boldsymbol{\beta}', \boldsymbol{\gamma}')]| + 2L^*(\delta_1)\epsilon \geq t \right]. \end{aligned}$$

Setting  $\epsilon = t/[4L^*(\delta_1)]$  and applying a union bound on the right-hand side gives

$$\begin{aligned} & \Pr \left[ \max_{\mathcal{A} \in \mathcal{S}_s} \max_{(\boldsymbol{\beta}', \boldsymbol{\gamma}') \in \mathcal{E}(\mathcal{A})} |l(\boldsymbol{\beta}', \boldsymbol{\gamma}') - \mathbb{E}[l(\boldsymbol{\beta}', \boldsymbol{\gamma}')]| \geq \frac{1}{2}t \right] \\ & \leq \sum_{\mathcal{A} \in \mathcal{S}_s} \sum_{(\boldsymbol{\beta}', \boldsymbol{\gamma}') \in \mathcal{E}(\mathcal{A})} \Pr \left[ |l(\boldsymbol{\beta}', \boldsymbol{\gamma}') - \mathbb{E}[l(\boldsymbol{\beta}', \boldsymbol{\gamma}')]| \geq \frac{1}{2}t \right]. \end{aligned}$$

A standard binomial coefficient bound (see, e.g., [Knuth, 1997](#), Section 1.2.6, Exercise 67) gives

$$|\mathcal{S}_s| = \binom{p}{s} \leq \left(\frac{ep}{s}\right)^s.$$

Next, since  $\|\boldsymbol{\beta}'\|_2 \leq \sqrt{s}\bar{\beta}$  and  $\|\boldsymbol{\gamma}'\|_2 \leq \sqrt{s}\bar{\gamma}$ , a standard  $\ell_2$  norm  $\epsilon$ -net cardinality bound (see, e.g., [Vershynin, 2018](#), Corollary 4.2.13) gives

$$|\mathcal{E}(\mathcal{A})| \leq \left(1 + \frac{2\sqrt{s}\bar{\beta}}{\epsilon}\right)^s \left(1 + \frac{2\sqrt{s}\bar{\gamma}}{\epsilon}\right)^s \leq \exp\left(2s\sqrt{s}\frac{\bar{\beta} + \bar{\gamma}}{\epsilon}\right) = \exp\left(8L^*(\delta_1)s\sqrt{s}\frac{\bar{\beta} + \bar{\gamma}}{t}\right).$$

It now follows that

$$\begin{aligned} \sum_{A \in \mathcal{S}_s} \sum_{(\beta', \gamma') \in \mathcal{E}(A)} \Pr \left[ |l(\beta', \gamma') - \mathbb{E}[l(\beta', \gamma')]| \geq \frac{1}{2}t \right] &\leq \left(\frac{ep}{s}\right)^s \exp \left( 8L^*(\delta_1)s\sqrt{s} \frac{\bar{\beta} + \bar{\gamma}}{t} \right) \\ &\times 2 \left[ \exp \left\{ -c \min \left( \frac{t^2}{4n_\bullet(1 + \bar{\gamma}s)^2}, \frac{t}{2(1 + \bar{\gamma}s)} \right) \right\} + \exp \left( -\frac{t^2}{512\sigma^{-2}(1 + \bar{\gamma}s)s^2\bar{\beta}^2} \right) \right]. \end{aligned} \quad (\text{E.3})$$

We wish to bound the right-hand side by a user chosen  $\delta_2 \in (0, 1]$ . To this end, we solve for a  $t$  that simultaneously satisfies

$$2 \left(\frac{ep}{s}\right)^s \exp \left( 8L^*(\delta_1)s\sqrt{s} \frac{\bar{\beta} + \bar{\gamma}}{t} \right) \exp \left\{ -c \min \left( \frac{t^2}{4n_\bullet(1 + \bar{\gamma}s)^2}, \frac{t}{2(1 + \bar{\gamma}s)} \right) \right\} \leq \frac{\delta_2}{2} \quad (\text{E.4})$$

and

$$2 \left(\frac{ep}{s}\right)^s \exp \left( 8L^*(\delta_1)s\sqrt{s} \frac{\bar{\beta} + \bar{\gamma}}{t} \right) \exp \left( -\frac{t^2}{512\sigma^{-2}(1 + \bar{\gamma}s)s^2\bar{\beta}^2} \right) \leq \frac{\delta_2}{2}. \quad (\text{E.5})$$

Some routine algebra shows that (E.4) is satisfied whenever

$$t \geq t_1 = (1 + \bar{\gamma}s) \max \left( \sqrt{\frac{4n_\bullet}{c}\Pi}, \frac{2}{c}\Pi \right),$$

where

$$\Pi = s \log \left( \frac{ep}{s} \right) + 8L^*(\delta_1)s\sqrt{s}(\bar{\beta} + \bar{\gamma}) + \log \left( \frac{4}{\delta_2} \right).$$

The above holds provided  $t \geq 1$ , which is true for our final choice of  $t$ . Similarly, some standard derivations give that (E.5) is satisfied if

$$\begin{aligned} t \geq t_2 &= \sqrt{512\sigma^{-2}(1 + \bar{\gamma}s)s^2\bar{\beta}^2\Pi} \\ &= \frac{s\bar{\beta}}{\sigma} \sqrt{512(1 + \bar{\gamma}s)\Pi}. \end{aligned}$$

Hence, (E.3) is guaranteed less than  $\delta_2$  provided

$$t \geq t_0 = \max(t_1, t_2).$$

Now, setting  $\delta_1 = \delta/2$  and  $\delta_2 = \delta/2$ , it follows from the previous inequalities that

$$2 \sup_{(\boldsymbol{\beta}, \boldsymbol{\gamma}) \in \mathcal{C}(s)} |l(\boldsymbol{\beta}, \boldsymbol{\gamma}) - \mathbb{E}[l(\boldsymbol{\beta}, \boldsymbol{\gamma})]| \leq 2t_0 \quad (\text{E.6})$$

with probability at least  $1 - \delta$ . Finally, by the definition of KL divergence, we have

$$\begin{aligned} 2 \text{KL} \left[ \text{N}(\mathbf{X}\boldsymbol{\beta}^0, \sigma^2 \mathbf{V}(\boldsymbol{\gamma}^0)) \parallel \text{N}(\mathbf{X}\hat{\boldsymbol{\beta}}, \sigma^2 \mathbf{V}(\hat{\boldsymbol{\gamma}})) \right] &= \mathbb{E}[-l(\boldsymbol{\beta}^0, \boldsymbol{\gamma}^0)] - \mathbb{E}[-l(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})] \\ &= \mathbb{E}[l(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})] - \mathbb{E}[l(\boldsymbol{\beta}^0, \boldsymbol{\gamma}^0)], \end{aligned}$$

from which we can combine (E.1) and (E.6) to complete the proof.

## F Implementation Details

The configuration of each method is as follows.

- `glmmsel` is configured with a sequence of 100 values of  $\lambda$  chosen using Proposition 2, where the first  $\lambda$  sets all coefficients to zero. For each value of  $\lambda$ , we use a sequence of 10 values of  $\alpha$  equispaced between 0.1 and 1.
- `pysr3` is configured with the  $\ell_0$  regularizer and a grid of  $11 \times 11$  values of  $(k_1, k_2)$  equispaced between 0 and 10, where  $k_1$  and  $k_2$  are the maximum allowable number of nonzeros in the fixed effects and random effects, respectively.
- `L0Learn` is configured with the  $\ell_0$  regularizer and a sequence of 100 values of  $\lambda$  chosen using the default method of the package.
- `glmnet` is configured with the relaxed  $\ell_1$  regularizer and a sequence of 100 values of  $\lambda$  chosen using the default method of the package. The relaxation parameter  $\gamma = 0$ .
- `ncvreg` is configured with the minimax concave penalty and a sequence of 100 values of  $\lambda$  chosen using the default method of the package. The concavity parameter  $\gamma = 1 + 10^{-4}$ .

- `glmmPen` is configured with the minimax concave penalty and a grid of  $10 \times 10$  values of  $(\lambda_1, \lambda_2)$  chosen using the default method of the package, where  $\lambda_1$  and  $\lambda_2$  are the penalty parameters for the fixed effects and random effects, respectively.<sup>4</sup> The concavity parameter  $\gamma = 1 + 10^{-4}$ .
- `rpq1` is configured with the minimax concave penalty and a grid of  $30 \times 30$  values of  $(\lambda_1, \lambda_2)$  chosen using the default method of `glmmPen`, where  $\lambda_1$  and  $\lambda_2$  are the penalty parameters for the fixed effects and random effects, respectively. The package uses `glmmPen`'s grid construction method as it does not provide a method for computing the  $\lambda$  sequence. The concavity parameter  $\gamma = 1 + 10^{-4}$ .

The specific choices of  $\gamma = 0$  in `glmnet` and  $\gamma = 1 + 10^{-4}$  in `ncvreg`, `glmmPen`, and `rpq1` are made to eliminate continuous shrinkage effects as far as possible. These parameter choices bring the behavior of these methods as close as possible to an  $\ell_0$  regularizer, ensuring that the comparison across methods focuses on sparsity rather than bias.

## G Small-Scale Experiments

To accommodate a broader set of baselines, namely `glmmPen` and `rpq1`, we run several small-scale experiments with just  $p = 10$  predictors. Figures 6 and 7 report the results for Gaussian and Bernoulli responses, respectively. In this low-dimensional setting, `glmmPen` emerges as the next best performer behind `glmmse1`, eventually delivering predictions on par with `glmmse1` when the sample size reaches  $n_{\bullet} = 1,000$ . Nonetheless, its ability to accurately distinguish zeros from nonzeros and fixed effects from random effects is typically weaker. The underperformance of `rpq1` is likely due to it modeling the random effect covariance matrix as unstructured, whereas the other methods use a more efficient diagonal representation.

---

<sup>4</sup>`glmmPen` uses only 10 values each of  $\lambda_1$  and  $\lambda_2$  because its run time is significantly longer than the other methods. However, it only runs in the  $p = 10$  setting where `glmmse1` also uses at most 10 values of  $\lambda$ . Indeed, by Proposition 2, the number of active predictors changes between consecutive values of  $\lambda$ , and hence there can be at most 10 values when  $p = 10$ .

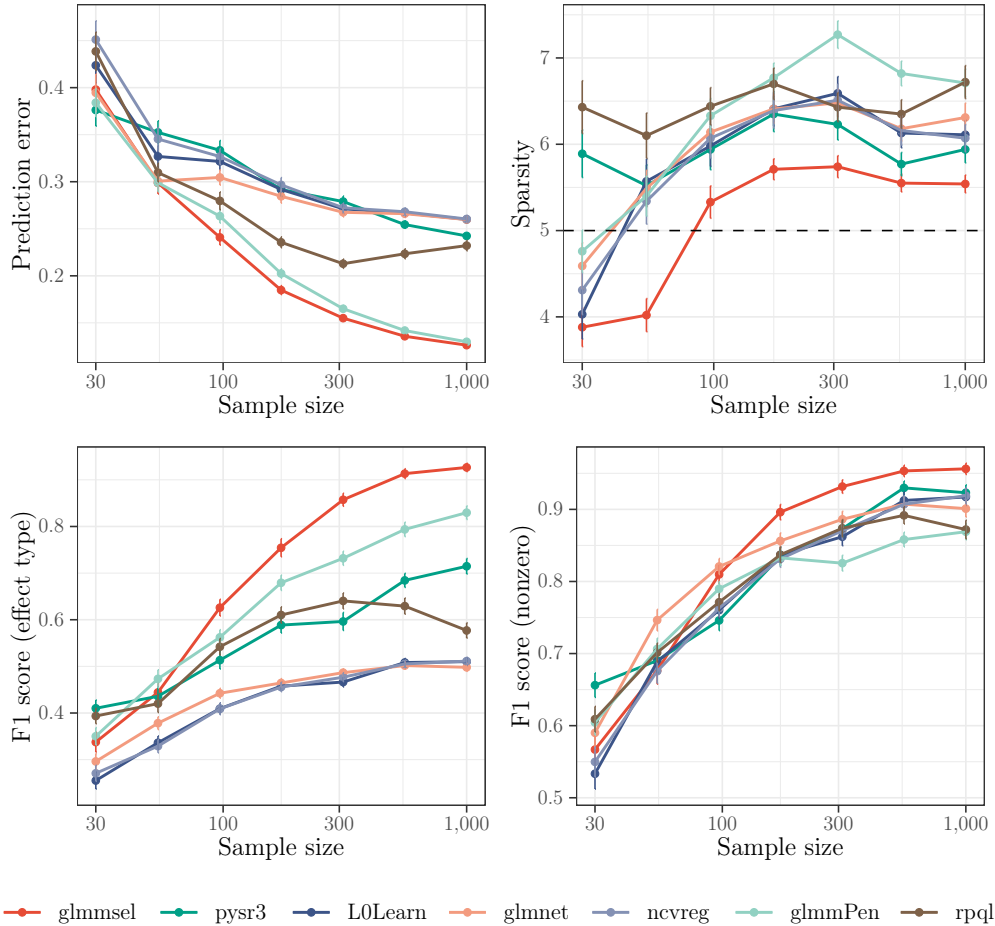


Figure 6: Comparisons on synthetic data with Gaussian response. The number of predictors  $p = 10$  with 5 nonzero fixed effects and 3 nonzero random effects. The correlation level  $\rho = 0.5$ . The averages (points) and standard errors (error bars) are measured over 100 datasets. In the upper right panel, the dashed horizontal line indicates the true number of predictors with nonzero effects.

## H Denser Experiments

To assess performance beyond the highly sparse regime, we consider denser settings with 15 nonzero fixed effects and 5 nonzero random effects. Figures 8 and 9 report the results for  $p = 1,000$  and  $p = 10,000$ , respectively. As expected, these settings are more challenging for all methods, especially at smaller sample sizes. Even so, the main conclusions are unchanged: `glmmsel` continues to perform strongly in prediction, achieves competitive recovery of effect type, particularly at larger sample sizes, and selects models with sparsity close to the truth.

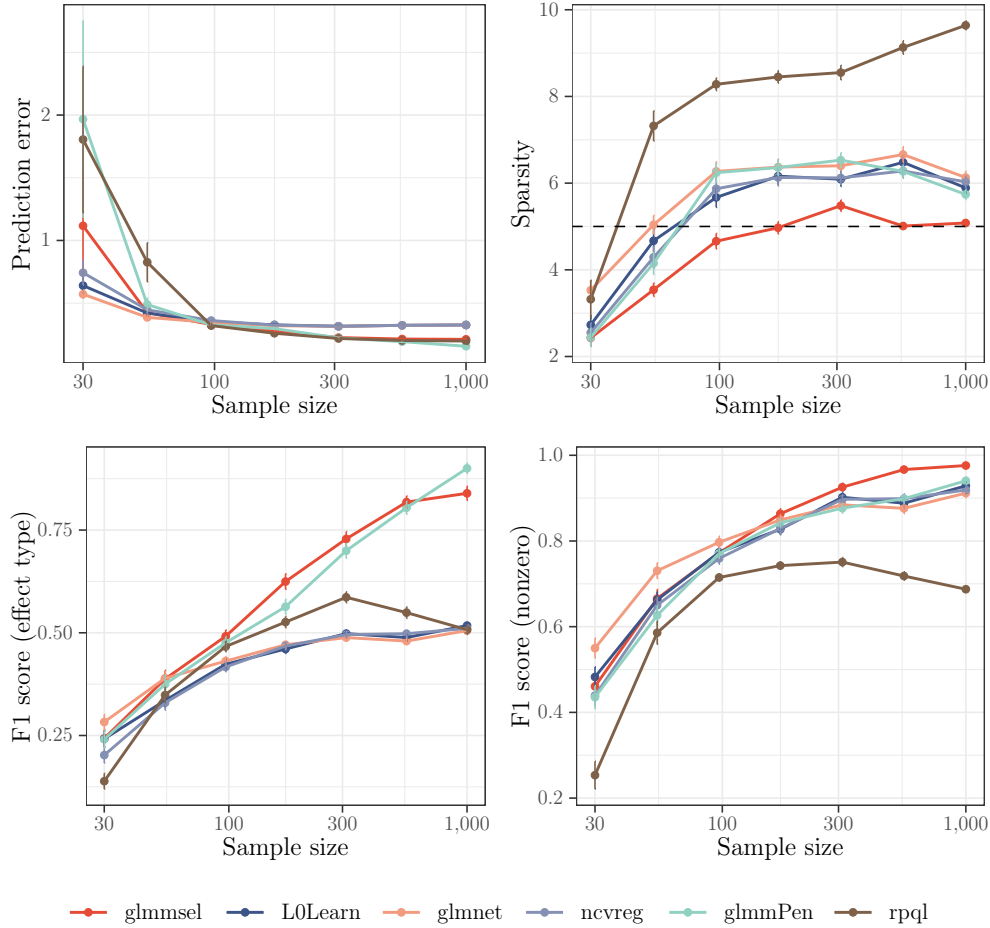


Figure 7: Comparisons on synthetic data with Bernoulli response. The number of predictors  $p = 10$  with 5 nonzero fixed effects and 3 nonzero random effects. The correlation level  $\rho = 0.5$ . The averages (points) and standard errors (error bars) are measured over 100 datasets. In the upper right panel, the dashed horizontal line indicates the true number of predictors with nonzero effects.

Overall, the results suggest that `glmmsel` remains stable beyond the highly sparse regime, including in higher-dimensional settings.

## I Sensitivity to $\gamma$ Tuning

To examine the sensitivity of the competing continuous shrinkage methods to their additional tuning parameter  $\gamma$ , we run additional experiments for `glmnet` and `ncvreg` in which  $\gamma$  is selected by validation over a 10-point grid. Figure 10 reports the results. Relative to the main

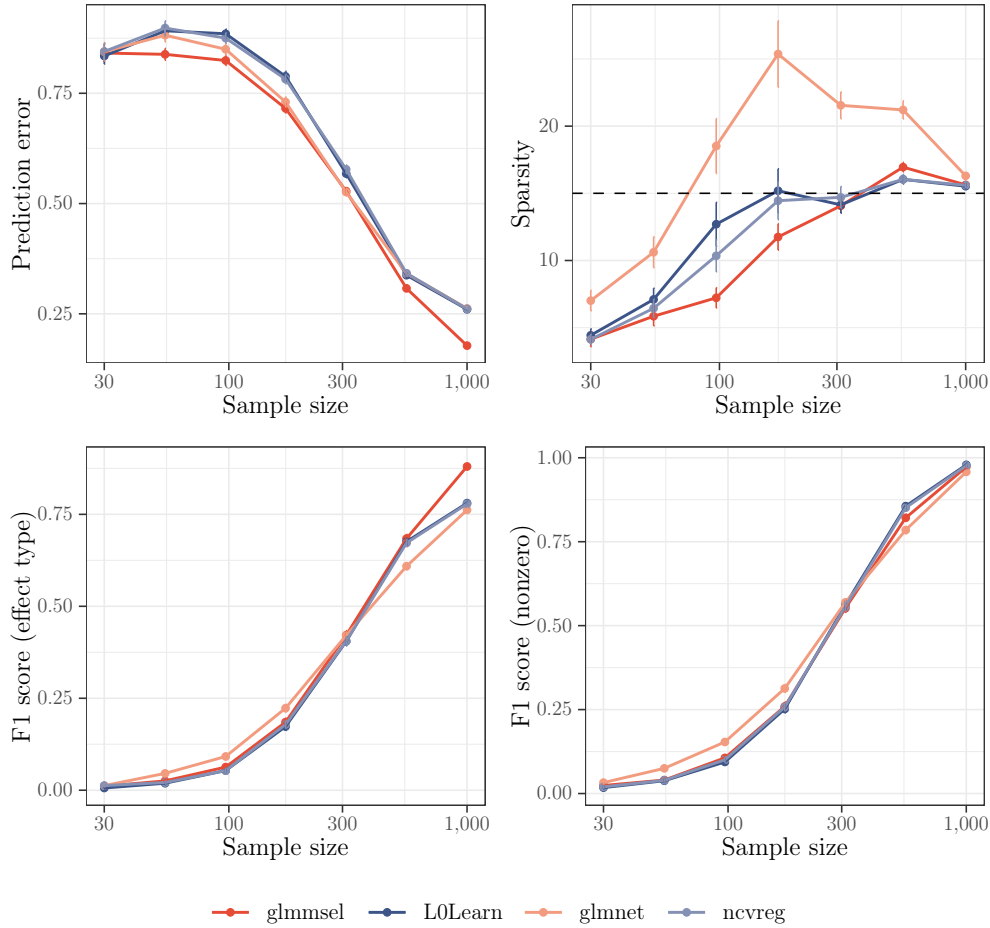


Figure 8: Comparisons on synthetic data with Gaussian response. The number of predictors  $p = 1,000$  with 15 nonzero fixed effects and 5 nonzero random effects. The correlation level  $\rho = 0.5$ . The averages (points) and standard errors (error bars) are measured over 100 datasets. In the upper right panel, the dashed horizontal line indicates the true number of predictors with nonzero effects.

experiments, tuning  $\gamma$  improves prediction for `glmnet` and `ncvreg`, but substantially worsens selection performance. The fitted models become much denser, and both F1 measures decline. Figure 10 also shows that the number of selected predictors tends to increase with sample size. This increase in selected predictors appears to reflect misspecification when random effects are present in the data-generating process, since methods without random effects can improve their fit by selecting more fixed effects to absorb heterogeneity. This pattern is not observed in unreported experiments under a classical data-generating process with fixed effects only.

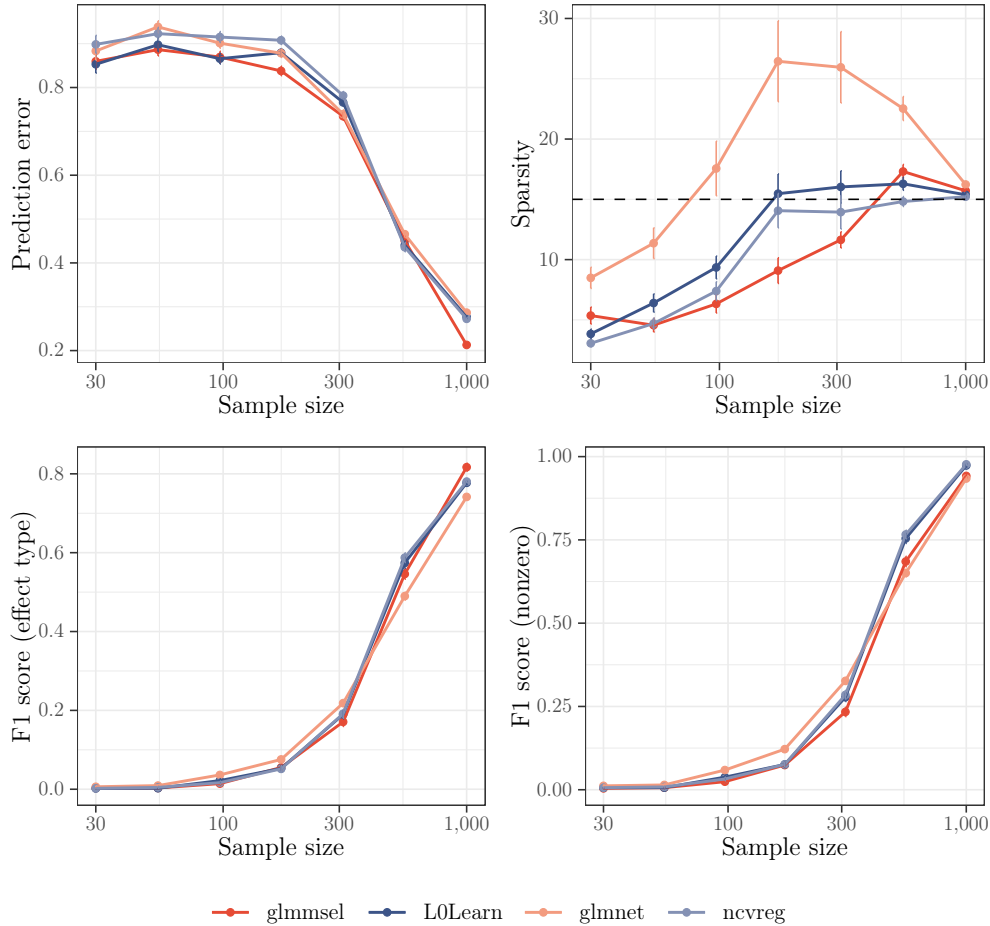


Figure 9: Comparisons on synthetic data with Gaussian response. The number of predictors  $p = 10,000$  with 15 nonzero fixed effects and 5 nonzero random effects. The correlation level  $\rho = 0.5$ . The averages (points) and standard errors (error bars) are measured over 100 datasets. In the upper right panel, the dashed horizontal line indicates the true number of predictors with nonzero effects.

## J Mashable Dataset

The `mashable` dataset from [Fernandes et al. \(2015\)](#) contains data on the popularity of news articles posted to the online news platform Mashable. In addition to a binary response variable indicating whether an article is popular (classified as achieving more than the median number of social media shares), the dataset comprises  $p = 53$  predictor variables that characterize the article (e.g., number of total words, positive words, and images). The task is to identify the article characteristics that are reliable determinants of popularity. Towards that end, the

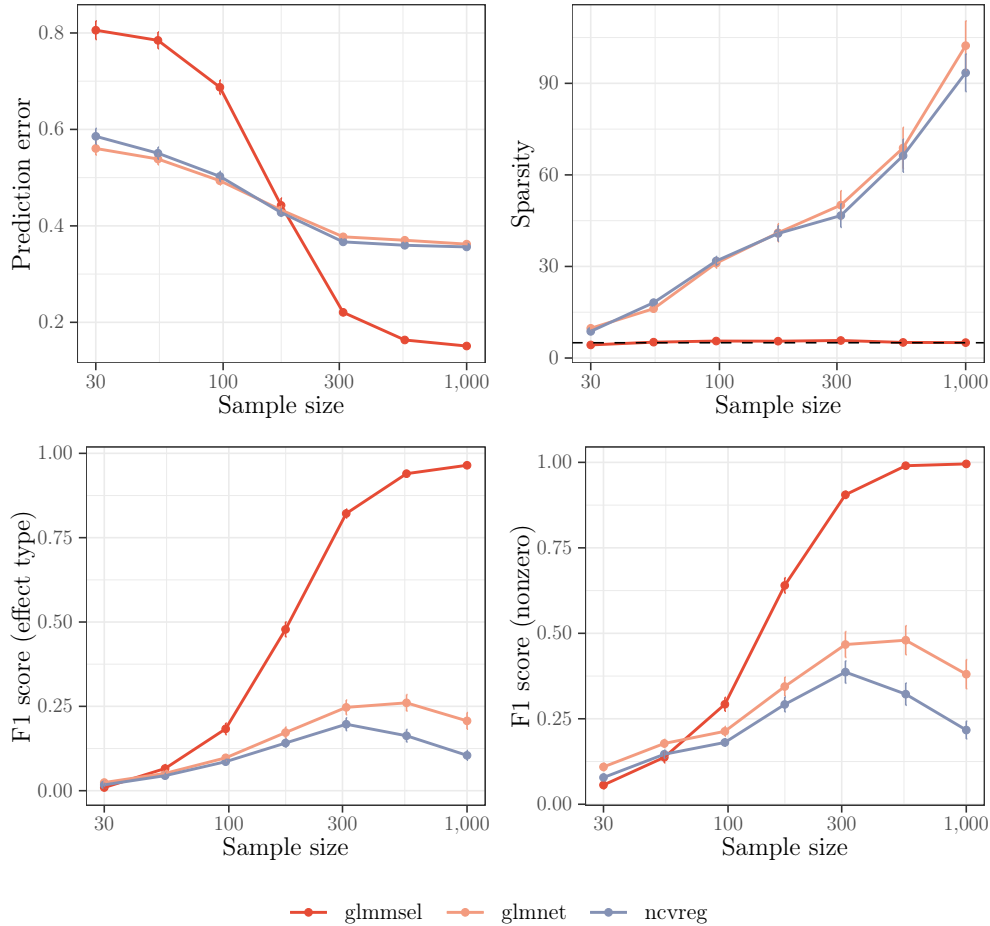


Figure 10: Comparisons on synthetic data with Gaussian response when the  $\gamma$  parameter is tuned for `glmnet` and `ncvreg`. The number of predictors  $p = 1,000$  with 5 nonzero fixed effects and 3 nonzero random effects. The correlation level  $\rho = 0.5$ . The averages (points) and standard errors (error bars) are measured over 100 datasets. In the upper right panel, the dashed horizontal line indicates the true number of predictors with nonzero effects.

dataset includes observations over  $m = 27$  days in December 2014, with  $n_i \in \{20, \dots, 100\}$  articles posted each day, giving a total of  $n_{\bullet} = 1,700$  observations.

We randomly split the dataset into training, validation, and testing sets in 0.70-0.15-0.15 proportions. Table 3 reports the results from 100 of these splits. `glmmsel` again clearly stands out by producing GLMMs that attain the lowest prediction loss and contain the fewest predictors (approximately 14 article characteristics on average). In contrast, the GLMs from `L0Learn` and `ncvreg`, which rely solely on fixed effects and thus cannot capture heterogeneity across days, select roughly 20–30% more article characteristics, resulting in substantially

	Prediction error	Sparsity		
		Total	Fixed	Random
<code>glmmsel</code>	0.908 (0.003)	14.4 (0.9)	13.4 (0.8)	1.0 (0.2)
<code>LOLearn</code>	0.916 (0.004)	19.0 (1.5)	19.0 (1.5)	0.0 (0.0)
<code>glmnet</code>	0.918 (0.004)	21.9 (1.4)	21.9 (1.4)	0.0 (0.0)
<code>ncvreg</code>	0.916 (0.004)	17.3 (1.4)	17.3 (1.4)	0.0 (0.0)

Table 3: Comparisons on the `mashable` dataset. The averages and standard errors (parentheses) are measured over 100 splits of the data. Prediction error is the cross entropy loss of the fitted linear predictor relative to the fixed intercept-only model.

worse prediction performance. `glmnet` is even less competitive, averaging 50% more predictors and achieving notably higher prediction loss. Consistent with the results reported on the `riboflavin` dataset, `glmmsel` again emerges as the most performant method and the only method capable of successfully fitting sparse GLMMs on the `mashable` dataset.

## References

- Beck, A. and Eldar, Y. C. (2013), ‘Sparsity constrained nonlinear optimization: Optimality conditions and algorithms’, *SIAM Journal on Optimization* **23**(3), 1480–1509.
- Beck, A. and Tetrushvili, L. (2013), ‘On the convergence of block coordinate descent type methods’, *SIAM Journal on Optimization* **23**(4), 2037–2060.
- Bertsimas, D., King, A. and Mazumder, R. (2016), ‘Best subset selection via a modern optimization lens’, *Annals of Statistics* **44**(2), 813–852.
- 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**(4), 1069–1077.
- Breheny, P. and Huang, J. (2011), ‘Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection’, *Annals of Applied Statistics* **5**(1), 232–253.

- Breslow, N. E. and Clayton, D. G. (1993), ‘Approximate inference in generalized linear mixed models’, *Journal of the American Statistical Association* **88**(421), 9–25.
- Capanu, M., Gönen, M. and Begg, C. B. (2013), ‘An assessment of estimation methods for generalized linear mixed models with binary outcomes’, *Statistics in Medicine* **32**(26), 4550–4566.
- Dedieu, A., Hazimeh, H. and Mazumder, R. (2021), ‘Learning sparse classifiers: Continuous and mixed integer optimization perspectives’, *Journal of Machine Learning Research* **22**(135), 1–47.
- Fan, J. and Li, R. (2001), ‘Variable selection via nonconcave penalized likelihood and its oracle properties’, *Journal of the American Statistical Association* **96**(456), 1348–1360.
- Fan, Y. and Li, R. (2012), ‘Variable selection in linear mixed effects models’, *Annals of Statistics* **40**(4), 2043–2068.
- Fernandes, K., Vinagre, P. and Cortez, P. (2015), A proactive intelligent decision support system for predicting the popularity of online news, in ‘Progress in Artificial Intelligence’, Vol. 9273, pp. 535–546.
- Friedman, J., Hastie, T., Höfling, H. and Tibshirani, R. (2007), ‘Pathwise coordinate optimization’, *Annals of Applied Statistics* **1**(2), 302–332.
- Friedman, J., Hastie, T. and Tibshirani, R. (2010), ‘Regularization paths for generalized linear models via coordinate descent’, *Journal of Statistical Software* **33**(1), 1–22.
- Groll, A. and Tutz, G. (2014), ‘Variable selection for generalized linear mixed models by  $l_1$ -penalized estimation’, *Statistics and Computing* **24**(2), 137–154.
- Harville, D. A. (1997), *Matrix Algebra From a Statistician’s Perspective*, Springer, New York, NY, USA.

- Hastie, T., Tibshirani, R. and Tibshirani, R. (2020), ‘Best subset, forward stepwise or lasso? Analysis and recommendations based on extensive comparisons’, *Statistical Science* **35**(4), 579–592.
- Hazimeh, H. and Mazumder, R. (2020), ‘Fast best subset selection: Coordinate descent and local combinatorial optimization algorithms’, *Operations Research* **68**(5), 1517–1537.
- Hazimeh, H., Mazumder, R. and Radchenko, P. (2023), ‘Grouped variable selection with discrete optimization: Computational and statistical perspectives’, *Annals of Statistics* **51**(1), 1–32.
- Heiling, H. M., Rashid, N. U., Li, Q. and Ibrahim, J. G. (2023), ‘glmmPen: High dimensional penalized generalized linear mixed models’, *R Journal* **15**(4), 106–128.
- Henderson, C. R., Kempthorne, O., Searle, S. R. and von Krosigk, C. M. (1959), ‘The estimation of environmental and genetic trends from records subject to culling’, *Biometrics* **15**(2), 192–218.
- Hui, F. K. C., Müller, S. and Welsh, A. H. (2017*a*), ‘Hierarchical selection of fixed and random effects in generalized linear mixed models’, *Statistica Sinica* **27**(2), 501–518.
- Hui, F. K. C., Müller, S. and Welsh, A. H. (2017*b*), ‘Joint selection in mixed models using regularized PQL’, *Journal of the American Statistical Association* **112**(519), 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**(2), 495–503.
- Jiang, J. and Nguyen, T. (2021), *Linear and Generalized Linear Mixed Models and Their Applications*, 2nd edn, Springer, New York, NY, USA.
- Knuth, D. E. (1997), *The Art of Computer Programming, Volume 1: Fundamental Algorithms*, Vol. 1, 3rd edn, Addison-Wesley Professional, Reading, MA, USA.

- Kowal, D. R. (2022), ‘Bayesian subset selection and variable importance for interpretable prediction and classification’, *Journal of Machine Learning Research* **23**(108), 1–38.
- Kowal, D. R. (2023), ‘Subset selection for linear mixed models’, *Biometrics* **79**(3), 1853–1867.
- Laurent, B. and Massart, P. (2000), ‘Adaptive estimation of a quadratic functional by model selection’, *Annals of Statistics* **28**(5), 1302–1338.
- Li, Y., Wang, S., Song, P. X.-K., Wang, N., Zhou, L. and Zhu, J. (2018), ‘Doubly regularized estimation and selection in linear mixed-effects models for high-dimensional longitudinal data’, *Statistics and Its Interface* **11**(4), 721–737.
- Lin, B., Pang, Z. and Jiang, J. (2013), ‘Fixed and random effects selection by REML and pathwise coordinate optimization’, *Journal of Computational and Graphical Statistics* **22**(2), 341–355.
- Natarajan, B. K. (1995), ‘Sparse approximate solutions to linear systems’, *SIAM Journal on Computing* **24**(2), 227–234.
- Nguyen, T., Nguyen, H. D., Chamroukhi, F. and McLachlan, G. J. (2024), ‘Non-asymptotic oracle inequalities for the lasso in high-dimensional mixture of experts’. arXiv:2009.10622.
- Pan, J. and Huang, C. (2014), ‘Random effects selection in generalized linear mixed models via shrinkage penalty function’, *Statistics and Computing* **24**(5), 725–738.
- Peng, H. and Lu, Y. (2012), ‘Model selection in linear mixed effect models’, *Journal of Multivariate Analysis* **109**, 109–129.
- Raskutti, G., Wainwright, M. J. and Yu, B. (2011), ‘Minimax rates of estimation for high-dimensional linear regression over  $\ell_q$ -balls’, *IEEE Transactions on Information Theory* **57**(10), 6976–6994.
- Rudelson, M. and Vershynin, R. (2013), ‘Hanson-Wright inequality and sub-Gaussian concentration’, *Electronic Communications in Probability* **18**(82), 1–9.

- Schelldorfer, J., Bühlmann, P. and van de Geer, S. (2011), ‘Estimation for high-dimensional linear mixed-effects models using  $\ell_1$ -penalization’, *Scandinavian Journal of Statistics* **38**(2), 197–214.
- Schelldorfer, J., Meier, L. and Bühlmann, P. (2014), ‘GLMMLasso: An algorithm for high-dimensional generalized linear mixed models using  $\ell_1$ -penalization’, *Journal of Computational and Graphical Statistics* **23**(2), 460–477.
- Sholokhov, A., Burke, J. V., Santomauro, D. F., Zheng, P. and Aravkin, A. (2024), ‘A relaxation approach to feature selection for linear mixed effects models’, *Journal of Computational and Graphical Statistics* **33**(1), 261–275.
- Stroup, W. W., Ptukhina, M. and Garai, J. (2024), *Generalized Linear Mixed Models: Modern Concepts, Methods, and Applications*, 2nd edn, Chapman and Hall/CRC, New York, NY, USA.
- Tay, J. K., Narasimhan, B. and Hastie, T. (2023), ‘Elastic net regularization paths for all generalized linear models’, *Journal of Statistical Software* **106**(1), 1–31.
- Thompson, R. and Vahid, F. (2024), ‘Group selection and shrinkage: Structured sparsity for semiparametric additive models’, *Journal of Computational and Graphical Statistics* **33**(4), 1286–1297.
- Tibshirani, R. (1996), ‘Regression shrinkage and selection via the lasso’, *Journal of the Royal Statistical Society: Series B (Methodological)* **58**(1), 267–288.
- Vershynin, R. (2018), *High-Dimensional Probability: An Introduction with Applications in Data Science*, Cambridge University Press, Cambridge, UK.
- West, B. T., Welch, K. B. and Galecki, A. T. (2022), *Linear Mixed Models: A Practical Guide Using Statistical Software*, 3rd edn, CRC Press, Boca Raton, FL, USA.

Wright, S. J. (2015), ‘Coordinate descent algorithms’, *Mathematical Programming: Series B* **151**(1), 3–34.

Zhang, C.-H. (2010), ‘Nearly unbiased variable selection under minimax concave penalty’, *Annals of Statistics* **38**(2), 894–942.