

The Function-on-Scalar LASSO with Applications to Longitudinal GWAS

Rina Foygel Barber¹, Matthew Reimherr^{*2}, and Thomas Schill²

¹University of Chicago

²Pennsylvania State University

Abstract

We present a new methodology for simultaneous variable selection and parameter estimation in function-on-scalar regression with an ultra-high dimensional predictor vector. We extend the LASSO to functional data in both the *dense* functional setting and the *sparse* functional setting. We provide theoretical guarantees which allow for an exponential number of predictor variables. Simulations are carried out which illustrate the methodology and compare the sparse/functional methods. Using the Framingham Heart Study, we demonstrate how our tools can be used in genome-wide association studies, finding a number of genetic mutations which affect blood pressure and are therefore important for cardiovascular health.

Keywords: Functional Data Analysis, High-Dimensional Regression, Variable Selection, Functional Regression

1 Introduction

Over the last several decades, technological advances have supported and necessitated the rapid growth of high-dimensional statistical methods. One of the most important applications for these methods are genome-wide association studies (GWAS). In these studies, researchers search through hundreds of thousands or millions of genetic mutations known as SNPs, single nucleotide polymorphisms, finding those which significantly impact an outcome or phenotype of interest, e.g. blood pressure, diabetes, asthma, etc. GWAS have been hugely successful at finding gene/disease associations as evidenced by the massive repositories of genetic studies

^{*}Corresponding Author: Department of Statistics, Pennsylvania State University, University Park, PA, 16802

and findings, such as dbGaP (<http://www.ncbi.nlm.nih.gov/gap>). Next-generation or high-throughput sequencing technologies are capable of sequencing entire genomes, producing ever-larger genetic datasets to explore. Thus, to further our understanding of the genetic architecture of complex human diseases, there is a substantial and continuing need for powerful, high-dimensional association techniques.

The vast majority of GWAS are cross-sectional, examining only one point in time. The aim of this paper is to present a new framework which combines functional data analysis, FDA, and machine learning for finding and estimating significant effects on longitudinally measured outcomes. We refer to this methodology as *Function-on-Scalar LASSO*, or FS-LASSO. As in Reiss et al. (2010), we use the term *Function-on-Scalar* to help distinguish our setting from the wide array of regression problems now in existence. In GWAS, the “function” is the time-varying phenotype of the individual, and it is being regressed on the “scalar”, which is the individual’s constant (non-time-varying) genotype.

Our aim is to simultaneously exploit the sparse effect of the SNPs and the smooth nature of the longitudinal outcomes. While we are strongly motivated by genetic studies, our methods are general and allow for any setting with a longitudinal/functional outcome and high-dimensional scalar predictors. Our primary goal is to select and estimate the effect of predictors in the following functional linear model

$$(1) \quad Y_n(t) = \mu(t) + \sum_{i=1}^I X_{ni}\beta_i(t) + \varepsilon_n(t).$$

Here $Y_n(t)$ is the value of a quantitative outcome for subject $n \in \{1, \dots, N\}$, at time $t \in \mathcal{T} \subset \mathbb{R}$. The scalars X_{ni} are real-valued, though in our application they take values in $\{0, 1, 2\}$ indicating the minor allele count for the SNP. The number of predictors, I , is allowed to be much greater than the sample size N . This is known as a “scalar predictor/functional response” model. Functional data analysis (FDA) now consists of two main branches: (1) sparse FDA, where the outcomes, Y_n , are observed at a relatively small number of time points and are contaminated with noise, and (2) dense FDA, where outcomes are observed at a large number of time points (possibly with a small level of noise). In practice, some data clearly fall into one of these two categories, say with hundreds or thousands of observations per subject, or with only a handful. There are also many settings where it is not clear which scenario one is in. While establishing a clear cutoff still a relatively open problem, we refer the interested reader to Li and Hsing (2010) and Zhang and Wang (2015+) who suggest the line is when the number of points per curve is greater/less than the quad-root of the sample size. We emphasize that this is just a basic rule of thumb, and care must be taken with each application. We present our methodology for each scenario and we later explore how their performances differ via simulations. In the functional case, as we will see, our methods and theory will actually include a wide range of settings beyond just (1).

At the heart of our methodology is the now classic idea of combining basis expansions for estimating the functions β_i , with proper penalties which induce sparse estimates. While the

methods and theory are new, we can phrase the problem as a type of group LASSO, which allows us to utilize existing computational tools.

The contributions of this paper include the following. First, we provide new asymptotic results for the sparse setting. Our results can be viewed as an extension of the varying coefficient methods discussed in Wei et al. (2011). In particular, we present a new restricted eigenvalue result which is one of the cornerstones for developing convergence rates in high dimensions, as well as the accompanying asymptotic convergence rates for our estimates. Second, we provide a new methodology and accompanying asymptotic theory for a broad class of dense settings, which, to the best of our knowledge, has not been explored before. Interestingly, we show that the FS-LASSO applied to outcomes in any separable Hilbert space achieves the exact same rates of convergence as in the scalar setting. This could have broad implications for not only traditional functional outcomes, but also spatial processes, functional panels, and imaging data such as fMRI. Finally, we demonstrate via simulations some surprising results concerning the choice between the sparse and dense tool sets. In particular, we find that a dense approach is comparable to the sparse in terms of variable selection even in traditionally sparse settings. Furthermore, the dense methods can be carried out at a fraction of the computational cost (both in terms of power and memory). However, the sparse methods produce estimates which are more accurate in traditionally sparse settings. This opens the doors to interesting two-stage procedures where variable screening is done using dense tools, and final estimation is done via sparse ones.

Related literature For foundations on functional data analysis we refer to Ramsay and Silverman (2005) and Horváth and Kokoszka (2012), while an overview of machine learning can be found in Hastie et al. (2001) and James et al. (2013). The literature on functional regression is now quite large, but we attempt to provide several key methodological papers, which help outline the field. Functional regression methods for sparse data include the following. Hoover et al. (1998) examine spline based methods for estimating a simpler form of the functional predictor/functional response model. Fan and Zhang (2000) provide a two step method based on local polynomial smoothing to estimate a functional predictor/scalar response model. Yao et al. (2005) present what has likely become the most common method based on scatter plot smoothing (i.e. local polynomial smoothing) and functional principal components for estimating a full functional predictor/functional response model. Zhu et al. (2012) extend local linear smoothing methods for estimating scalar predictor/multivariate functional response models. Methods for high dense FDA include the following. Cardot et al. (2003) explore spline based methods for a functional predictor/scalar response model. Kokoszka et al. (2008) use a PCA based approach for estimating a full functional predictor/functional response model. James et al. (2009) incorporate shrinkage methods in estimating a functional predictor/scalar response model (for the purposes of estimating the "zero" parts of the regression function). Reiss et al. (2010) examine a B-spline approach for scalar predictor/functional response models, while Reimherr and Nicolae (2014) explore more direct least squares (of a functional norm) estimators for scalar predictor/functional re-

sponse models. However, all of these methods assume a fixed number of covariates. Related work which incorporates an increasing number of predictors mainly comes from the literature on varying coefficient models, for example Wang et al. (2008), Zhao and Xue (2010) and Wei et al. (2011). The only related work in the FDA literature concerns scalar-on-function regression, the reverse of our focus. In particular, Matsui and Konishi (2011) combine basis expansions with a SCAD style penalty, but give no theoretical guarantees. Lian (2013) also uses a SCAD penalty, but combined with FCPA instead of a general basis function. Theory is provided, but only in the case where the number of predictors is fixed. Gertheiss et al. (2013) developed a procedure for a generalized scalar-on-function linear model again using basis expansions, but combined with a LASSO style penalty which simultaneously shrinks and smooths parameter estimates; no theoretical guarantees are given. Fan et al. (2014, 2015+) consider functional additive regression models with a large number of predictors. Their methods allow for nonlinear relationships and theoretical properties are established. Lastly, Kong et al. (2016) consider the problem of variable selection with a scalar outcome and scalar predictors, but with a small set of functional predictors that must be accounted for as well. We emphasize that all of these methods are for the scalar-on-function case; for the function-on-scalar case the only work we are aware of is by Chen et al. (2016) who consider a basis expansion approach with a MCP style penalty and fixed number of covariates, but their method cannot be applied to settings where $I \gg N$.

Outline The remainder of the paper is organized as follows. In Section 2 we provide a framework for scalar predictor/functional response regression with sparsely observed outcomes. We combine basis expansions with a group LASSO style penalty to carry out variable selection and parameter estimation simultaneously. In Section 3 we present the high dense setting. We provide a very general framework where the response functions are allowed to take values from a separable Hilbert space. This allows the functions to be from say $L^2(\mathcal{T})$, which is commonly used, as well as more general spaces such as Sobolev spaces or product spaces of function spaces (i.e. vectors of functions). In Section 4 we provide computational details, including a pre-screening rule which allows one to reduce I substantially before fitting the FS-LASSO, and details for making use of established computational machinery for the group LASSO. We present a simulation study in Section 5 where we explore our procedure in terms of variable selection, parameter estimation, and computational time. In Section 6 we apply our methods to the Framingham Heart Study (FHS), to identify genetic markers that influence blood pressure, in hopes of gaining a better understanding of cardiovascular disease. Genome-wide SNP data and phenotype information were downloaded from dbGaP (<http://www.ncbi.nlm.nih.gov/gap>) study accession phs000007.v25.p9. Concluding remarks are given in Section 7, while all theoretical results are proven in the Appendix.

2 The Sparse Setting

Sparse FDA occurs quite often in longitudinal studies where one has a relatively small number of observed time points per subject or unit. In this case, it is common to work with the raw measurements as opposed to the functional embeddings discussed in Section 3. A brief introduction to sparse FDA can be found in Müller (2008).

We begin with the underlying model.

Assumption 1. Assume, for $1 \leq n \leq N$ and $1 \leq m \leq M_n \leq M < \infty$, that

$$Y_{nm} = \sum_{i=1}^I X_{ni} \beta_i^*(t_{nm}) + \varepsilon_n(t_{nm}), \quad t_{nm} \in \mathcal{T}.$$

The error processes $\{\varepsilon_n(t) : t \in \mathcal{T}\}$ are iid Gaussian and we define the $M_n \times M_n$ covariance matrix $\Sigma_n := \text{Cov}((\varepsilon_n(t_{n1}), \dots, \varepsilon_n(t_{nM_n})))$. The design matrix, \mathbf{X} , is deterministic.

Throughout, we will use a \star to denote true parameter values. For the moment, we do not include assumptions on $\{t_{nm}\}$ or coefficient functions β_i^* . Note that the Gaussian assumption is not crucial—what is really needed is that the processes have subgaussian tails, though we make the Gaussian assumption to simplify the arguments.

Let $\{e_j(\cdot)\}$ be a basis in $L_2(\mathcal{T})$. Approximating the functions β_i^* using this basis, we have that

$$Y_{nm} = \sum_{i=1}^I \sum_{j=1}^J X_{ni} B_{ij}^* e_j(t_{nm}) + T_{nm} + \varepsilon_n(t_{nm}),$$

where T_{nm} is the truncation error obtained after cutting off the basis expansion at J . The following notation will be used repeatedly throughout this section and in the Appendix:

$$\mathbf{T} = (T_{11}, T_{12}, \dots, T_{NM_N})^\top \in \mathbb{R}^{\sum_n M_n},$$

$$\mathbf{E}_{nm} = (e_1(t_{nm}), \dots, e_J(t_{nm}))^\top \in \mathbb{R}^J,$$

$$\mathbf{E}_n = (E_{n1}, \dots, E_{nM_n})^\top \in \mathbb{R}^{M_n \times J},$$

$$\mathbf{F} = \frac{1}{N} \sum_{n=1}^N \mathbf{E}_n^\top \mathbf{E}_n \in \mathbb{R}^{J \times J}.$$

For a matrix $\mathbf{B} \in \mathbb{R}^{I \times J}$, define the *sparse target function* as

$$L(\mathbf{B}) = \frac{1}{2} \sum_{n=1}^N \sum_{m=1}^{M_n} \left(Y_{nm} - \sum_{i=1}^I \sum_{j=1}^J X_{ni} B_{ij} e_j(t_{nm}) \right)^2 + \lambda \|\mathbf{B}\|_{\ell_1/\ell_2},$$

where $\|\mathbf{B}\|_{\ell_1/\ell_2} = \sum_i \|B_{i*}\|_2$ promotes row-wise sparsity (here B_{i*} denotes the i th row of the matrix \mathbf{B}). The estimate $\hat{\mathbf{B}}$ is the minimizer of the above expression. The estimated

coefficient functions are then given by $\hat{\beta}_i(\cdot) = \sum_{j=1}^J \hat{B}_{ij} e_j(\cdot)$. The target function can be rephrased so that traditional group LASSO machinery can be invoked. Notice that

$$\sum_{i=1}^I \sum_{j=1}^J X_{ni} B_{ij} e_j(t_{nm}) = X_n^\top \mathbf{B} E_{nm} = E_{nm}^\top \mathbf{B}^\top X_n = (X_n^\top \otimes E_{nm}^\top) \text{vec}(\mathbf{B}^\top).$$

One can then stack the $(X_n^\top \otimes E_{nm}^\top)$ vectors into a matrix, \mathbf{A} ,

$$\mathbf{A} = \begin{pmatrix} X_1^\top \otimes E_{11}^\top \\ X_1^\top \otimes E_{12}^\top \\ \vdots \\ X_N^\top \otimes E_{NM_N}^\top \end{pmatrix} = \begin{pmatrix} X_1^\top \otimes \mathbf{E}_1 \\ X_2^\top \otimes \mathbf{E}_2 \\ \vdots \\ X_N^\top \otimes \mathbf{E}_N \end{pmatrix} \in \mathbb{R}^{(\sum_n M_n) \times IJ}.$$

The target function can now be expressed as

$$(2) \quad L(\mathbf{B}) = \frac{1}{2} \|Y - \mathbf{A} \text{vec}(\mathbf{B}^\top)\|_2^2 + \lambda \|\mathbf{B}\|_{\ell_1/\ell_2},$$

where

$$Y = (Y_{11}, \dots, Y_{1M_1}, \dots, Y_{N,M_N})^\top.$$

Group LASSO computational tools can then be used to find $\hat{\mathbf{B}}$.

Since the matrix \mathbf{A} will generally have more columns than rows, the linear system $Y \approx \mathbf{A} \text{vec}(\mathbf{B}^\top)$ is underdetermined; in particular the (right) null space of \mathbf{A} is large. However, the grouped sparsity structure in \mathbf{B} allows us to resolve this difficulty. We first recall the notion of restricted eigenvalues Bickel et al. (2009), used for sparse regression (without group structure):

Definition 2.1 (Restricted eigenvalue condition). *A matrix $\mathbf{A} \in \mathbb{R}^{N \times I}$ satisfies the $\text{RE}(I_0, \alpha)$ condition if, for all subsets $S \subset \{1, \dots, I\}$ with $|S| \leq I_0$,*

$$\|\mathbf{A}w\|_2^2 \geq \alpha N \|w\|_2^2 \text{ for all } w \in \mathbb{R}^I \text{ with } \|w_{S^c}\|_{\ell_1} \leq 3 \|w_S\|_{\ell_1}.$$

The constant 3 here is somewhat arbitrary, but it is standard in the literature and we use it here for convenience. Lounici et al. (2011) extend this definition to the group-sparse setting:

Definition 2.2 (Grouped restricted eigenvalue condition). *A matrix $\mathbf{A} \in \mathbb{R}^{N \times IJ}$ satisfies the $\text{RE}_{\text{group}}(I_0, \alpha)$ condition if, for all subsets $S \subset \{1, \dots, I\}$ with $|S| \leq I_0$,*

$$\|\mathbf{A} \text{vec}(\mathbf{W}^\top)\|_2^2 \geq \alpha N \|\mathbf{W}\|_{\text{F}}^2 \text{ for all } \mathbf{W} \in \mathbb{R}^{I \times J} \text{ with } \|\mathbf{W}_{S^c}\|_{\ell_1/\ell_2} \leq 3 \|\mathbf{W}_S\|_{\ell_1/\ell_2}.$$

Here \mathbf{W}_S denotes the submatrix of \mathbf{W} obtained by extracting the rows indexed by S , while \mathbf{W}_{S^c} contains only the rows in S^c , the complement of S .

2.1 Bounding the error

With these definitions in place, our first result is a modified version of Lounici et al. (2011)'s analysis of the group LASSO. It proves resulting bounds on the error $\hat{\mathbf{B}} - \mathbf{B}^*$, where \mathbf{B}^* is the (truncated) true parameter matrix, as long as λ is sufficiently large and \mathbf{A} satisfies the grouped restricted eigenvalue condition.

Theorem 1. *Suppose that Assumption 1 holds and that for some $\delta \in (0, 1)$*

$$(3) \quad \lambda \geq 2\sqrt{N \|\mathbf{F}\|_{\text{op}}} \left(\|T\|_2 + \sqrt{\max_n \|\Sigma_n\|_{\text{op}} (2J + 3 \log(2I/\delta))} \right).$$

Assume also that \mathbf{B}^ has at most I_0 nonzero rows, and that \mathbf{A} satisfies the $\text{RE}_{\text{group}}(I_0, \alpha)$ condition for some $\alpha > 0$. Then with probability at least $1 - \delta$, any minimizer $\hat{\mathbf{B}}$ of (2) satisfies*

$$\|\hat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}} \leq \frac{3\lambda\sqrt{I_0}}{\alpha N} \text{ and } \|\hat{\mathbf{B}} - \mathbf{B}^*\|_{\ell_1/\ell_2} \leq \frac{12\lambda I_0}{\alpha N}.$$

Remark 2.1. *Results for sparse regression in a non-grouped setting commonly give error bounds in both an ℓ_2 norm and an ℓ_1 norm; the former typically gives a more favorable scaling of sample size with respect to sparsity, while the latter ensures that for a larger sample size the error itself is nearly sparse, i.e. does not have a “long tail” of small errors on many coordinates. Similarly, here we give results in both the Frobenius norm (with more favorable sample size scaling) and the ℓ_1/ℓ_2 norm.*

Remark 2.2. *As is standard in the LASSO literature, the consistency results given in this theorem can yield a selection consistency result as well: if we assume that $\min_{i \in I_0} \|\mathbf{B}_i^*\|_2 > \frac{6\lambda\sqrt{I_0}}{\alpha N}$, then the bound on $\|\hat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}}$ implies that the following estimated support,*

$$\hat{S} = \left\{ i \in \{1, \dots, I\} : \|\hat{\mathbf{B}}_{i*}\|_2 > \frac{3\lambda\sqrt{I_0}}{\alpha N} \right\},$$

is correct (i.e. \hat{S} is equal to the row support of \mathbf{B}^) with high probability.*

As is widely appreciated by researchers working on such results, Theorem 1 essentially hinges on concentration inequalities and restricted eigenvalue type conditions. Neither of these tools can be readily applied to our setting. Thus, much of our theoretical work is focussed on showing how to extend these ideas to the functional setting. Below we discuss the restricted eigenvalue condition at length, while discussion of the concentration inequalities involved can be found at the end of Appendix B.1. Concentration inequalities allow us to control the stochastic error in the model and largely dictate the rates of convergence, while the restricted eigenvalue conditions allow us to bound the error $\|\hat{\mathbf{B}} - \mathbf{B}^*\|_{\text{F}}$ from above by a term involving $\|\mathbf{A} \text{vec}((\hat{\mathbf{B}} - \mathbf{B}^*)^\top)\|_2$, and therefore to obtain convergence rates for the parameters of interest.

Note that the phrasing of Theorem 1 allows for the possibility of multiple minimizers of (2), guaranteeing only that any minimizer must achieve the specified converge rates. However, in practice it rare is for LASSO style estimates to not be unique. We direct the interested reader to Tibshirani (2013) for a thorough discussion of this issue.

Example: We next give an example to illustrate the type of scaling that is obtained in Theorem 1. Let $W^{\tau,2}[0,1]$ denote the Sobelev space consisting of functions over $[0,1]$, which have square integrable derivatives of up to order τ . Suppose that the coefficient functions, β_i , take values in a finite ball in $W^{\tau,2}$. Let the basis $e_1(\cdot), e_2(\cdot), \dots$ denote the Fourier basis and assume that the time points t_{nm} are iid $U[0,1]$. Examining λ , there are now a number to terms which can be made more explicit. First, since the basis is orthogonal we have $\mathbf{F} \approx \mathbf{I}_{J \times J}$, so that $\|\mathbf{F}\|_{\text{op}}$ is behaves like a constant. Next, examining the truncation term T , there are $\sum_n M_n \sim N$ coordinates, and each coordinate consists of a sum of I_0 functions each of which lies in a Sobelev ball, thus we have that $\|T\| \approx \sqrt{NI_0}J^{-\tau}$. Examining the second term in λ , we have that

$$\sqrt{\max_n \|\Sigma_n\|_{\text{op}} (2J + 3 \log(2I/\delta))} \sim \sqrt{J + \log(I)} \leq \sqrt{J \log(I)}.$$

We therefore want to choose J to balance the two errors, which means taking J such that

$$\sqrt{NI_0}J^{-\tau} = \sqrt{J \log(I)} \implies J = \left(\frac{NI_0}{\log(I)} \right)^{\frac{1}{1+2\tau}}.$$

Plugging this into our expression for the convergence rates, and assuming that the restricted eigenvalue condition $\mathbf{RE}_{\text{group}}(I_0, \alpha)$ holds for the matrix \mathbf{A} with some $\alpha > 0$ that we treat as a constant, we have that

$$\|\Delta\|_{\mathbf{F}} = O_P(1) \cdot \frac{N^{1/2} \sqrt{J \log(I)} I_0^{1/2}}{N} = O_P(1) \cdot I_0 \left(\frac{\log(I)}{NI_0} \right)^{\frac{\tau}{1+2\tau}},$$

where $\Delta := \mathbf{B} - \widehat{\mathbf{B}}$. We see that the standard nonparametric rate of convergence appears which relates the convergence rates to the smoothness of the underlying parameter functions. This rate also applies to the difference $\|\widehat{\beta} - \beta\|_{L^2}$. Notice that by Parceval's identity and the Triangle inequality

$$\|\widehat{\beta} - \beta\|_{L^2} = \left(\sum_{i=1}^I \int (\widehat{\beta}_i(t) - \beta_i(t))^2 dt \right)^{1/2} \leq \|\Delta\|_{\mathbf{F}} + \|\Pi_J^\perp \beta\|_{L^2},$$

where Π_J^\perp is the projection onto the remaining basis functions $e_{J+1}(\cdot), \dots$. By the assumed smoothness of the β , we have that

$$\|\Pi_J^\perp \beta\|_{L^2} \sim \sqrt{I_0}J^{-\tau} = \sqrt{I_0} \left(\frac{NI_0}{\log(I)} \right)^{\frac{-\tau}{1+2\tau}} \leq I_0 \left(\frac{\log(I)}{NI_0} \right)^{\frac{\tau}{1+2\tau}},$$

Thus the rate for $\|\widehat{\beta} - \beta\|_{L^2}$ is the same as for $\|\Delta\|_{\mathbf{F}}$.

For the ℓ_1/ℓ_2 norm we then have

$$\|\Delta\|_{\ell_1/\ell_2} = O_P(1) \cdot \frac{\lambda I_0}{N} = O_P(1) \cdot I_0^{3/2} \left(\frac{\log(I)}{N I_0} \right)^{\frac{\tau}{1+2\tau}}.$$

By the same arguments as before, the above rate also applies to the sum of the normed differences between the functions, i.e. $\|\widehat{\beta} - \beta\|_{\ell_1}$. As we will see, when $\tau \rightarrow \infty$, the rate above converges to the rate found for the dense setting.

2.2 Restricted eigenvalue condition

Our next result concerns the restricted eigenvalue condition. It proves that, if the covariates X_n are drawn from a Gaussian or subgaussian distribution with well-conditioned covariance structure, then the grouped restricted eigenvalue condition will hold for \mathbf{A} .

Theorem 2. *Suppose that $X_n \stackrel{iid}{\sim} N(0, \Sigma)$ with $\|\Sigma\| \leq \nu^2$, or more generally, the X_n 's are iid where each X_n has mean 0, covariance Σ , and is ν -subgaussian meaning that $\mathbb{E}[e^{v^\top X_n}] \leq e^{\nu^2 \|v\|_2^2/2}$ for any fixed vector v . Suppose that $\lambda_{\min}(\Sigma) > 0$. Consider any fixed sequence of matrices $\mathbf{E}_1 \in \mathbb{R}^{M_1 \times J}$, \dots , $\mathbf{E}_N \in \mathbb{R}^{M_N \times J}$ satisfying*

$$(4) \quad \min_{w \in \mathbf{R}^J \setminus \{0\}} \frac{\sum_n \|\mathbf{E}_n w\|_2}{N \|w\|_2} \geq \gamma_0 > 0, \quad \max_{w \in \mathbf{R}^J \setminus \{0\}} \sqrt{\frac{\sum_n \|\mathbf{E}_n w\|_2^2}{N \|w\|_2^2}} \leq \gamma_1,$$

and define $\mathbf{A} \in \mathbb{R}^{\sum_n M_n \times IJ}$ as before. Then there exist $c_0, c_1, c_2 > 0$ depending only on $\nu, \lambda_{\min}(\Sigma), \gamma_0, \gamma_1$, such that, if

$$(5) \quad N \geq c_0 \cdot k \cdot \log(IJ) \cdot (J + \log(I)),$$

then with probability at least $1 - e^{-c_1 N}$, \mathbf{A} satisfies the $\text{RE}_{\text{group}}(k, c_2)$ condition.

Remark 2.3. *This theorem is a primary distinction between our work discussed thus far and that of Wei et al. (2011). Their work cites existing results which imply that the matrix \mathbf{A} behaves appropriately for each approximately-row-sparse \mathbf{W} in expectation, that is, $\mathbb{E}[\|\mathbf{A} \text{vec}(\mathbf{W}^\top)\|_2^2]$ is lower-bounded for each \mathbf{W} . In contrast, our theorem above proves the much stronger statement that $\|\mathbf{A} \text{vec}(\mathbf{W}^\top)\|_2^2$ can be lower-bounded simultaneously with high probability for all approximately-row-sparse \mathbf{W} , which is necessary for estimation in high dimensions.*

Remark 2.4. *The conditions (4) on the basis matrices \mathbf{E}_n are of course dependent on the choice of basis; this type of basis-dependent assumption is standard in the nonparametric literature. As a simple example of a choice of basis where (4) is satisfied, suppose that there are J evenly spaced time points and our basis is given by J orthonormal functions (e.g. a*

Fourier basis). Suppose measurements for each individual are taken at a random subset of J' time points from the original J . Then each \mathbf{E}_n is a $J' \times J$ submatrix chosen at random from some larger $J \times J$ orthogonal matrix (corresponding to e.g. the Fourier basis), in which case $\gamma_1 = 1$ and $\gamma_0 \sim \sqrt{J'/J}$.

The assumption given in (4) is a condition number type constraint. The upper bound is simply a bound on the largest eigenvalue of $\mathbf{F} = \frac{1}{N} \sum_n \mathbf{E}_n^\top \mathbf{E}_n$, while the lower bound is a bit stronger than a corresponding bound on the smallest eigenvalue. It is, in essence, saying that the norms $\|\mathbf{E}_n w\|_2^2$ should be on the order of $\|w\|_2^2$ for many of the indices n . (In contrast, upper and lower eigenvalue conditions would only ensure that $\|\mathbf{E}_n w\|_2^2$ is on the order of $\|w\|_2^2$ *on average*, but would not prevent degenerate scenarios such as $\|\mathbf{E}_1 w\|_2^2 = N \|w\|_2^2$ and $\|\mathbf{E}_2 w\|_2^2 = \dots = \|\mathbf{E}_N w\|_2^2 = 0$.)

3 The Dense Setting

When the underlying functions are observed at a relatively large number of time points, a different approach than the one described in Section 2 is commonly employed. In particular, for each subject $n \in \{1, \dots, N\}$, the observations $Y_{nm} = Y_n(t_{nm})$ are embedded into a function space, which are then treated as though they were fully observed functions. Implicitly, one is assuming that the error from the embedding is negligible compared to other sources of variability. More details and background can be found in Horváth and Kokoszka (2012). We begin by defining the underlying model.

Assumption 2. Let Y_1, \dots, Y_N be independent random elements of a real separable Hilbert space, \mathcal{H} , satisfying the functional linear model

$$Y_n = \sum_{i=1}^I X_{ni} \beta_i + \varepsilon_n.$$

Assume the $N \times I$ design matrix $\mathbf{X} = \{X_{ni}\}$ is deterministic and has standardized columns, the ε_n are iid square-integrable Gaussian random elements of \mathcal{H} with mean 0 and covariance operator C , and β_i are deterministic elements of \mathcal{H} . Let $\Lambda = (\lambda_1, \lambda_2, \dots)$ denote the vector of eigenvalues of C .

Remark 3.1. We emphasize that, throughout this section, the responses Y_n , coefficients β_i , and errors ε_n are all elements of the Hilbert space (e.g. functions), rather than scalars as in the usual high-dimensional regression setting. Our work in this section can be viewed as a functional analogue of existing high-dimensional sparse signal recovery results for the scalar setting (i.e. the ordinary LASSO).

The most common choice for the function space \mathcal{H} is $L^2(\mathcal{T})$ where \mathcal{T} is a closed and bounded interval. However, the advantage of our phrasing is that we allow for a number of other spaces, including Sobolev spaces (if one wants to better utilize the smoothness of the data), product spaces (if the response is actually a vector of functions), or multidimensional domains (which could arise in areas such as spatial statistics). Throughout this section, when we write $\|\cdot\|_{\mathcal{H}}$ we will always mean the inner product norm on \mathcal{H} . Norms which deviate from this inner product norms will include alternative subscripts.

As before, when applying the FS-LASSO, we are assuming that the true underlying model is sparse, with I_0 denoting the number of true predictors and S_0 denoting the true set of predictors. The FS-LASSO estimate is then the solution to the following minimization problem

$$\hat{\beta} = \operatorname{argmin}_{\beta \in \mathcal{H}^I} L_F(\beta),$$

where

$$L_F(\beta) = \frac{1}{2} \sum_{n=1}^N \|Y_n - X_n^\top \beta\|_{\mathcal{H}}^2 + \lambda \|\beta\|_{\ell_1/\mathcal{H}} \quad \text{for} \quad \|\beta\|_{\ell_1/\mathcal{H}} = \sum_{i=1}^I \|\beta_i\|_{\mathcal{H}}.$$

The norm $\|\cdot\|_{\ell_1/\mathcal{H}}$ is a type of ℓ_1 norm on the product space \mathcal{H}^I , which encourages sparsity among the list of functions β_1, \dots, β_I (that is, the function β_i will be uniformly zero for many indices i). The target function $L_F(\cdot)$ is a direct Hilbert space generalization of the LASSO target function.

Next we introduce a functional restricted eigenvalue assumption, which is a direct analogue of similar assumptions in the LASSO literature (e.g. Bickel et al. (2009)).

Definition 3.1. *We say a matrix $\mathbf{A} \in \mathbb{R}^{N \times I}$ satisfies a functional restricted eigenvalue condition, $\operatorname{RE}_F(I_0, \alpha)$, if for all subsets $S \subset \{1, \dots, I\}$ with $|S| \leq I_0$, we have*

$$\|\mathbf{A}x\|_{\mathcal{H}^N}^2 \geq \alpha N \|x\|_{\mathcal{H}^I}^2 \quad \text{for all } x \in \mathcal{H}^I \text{ that satisfy } \|x_{S^c}\|_{\ell_1/\mathcal{H}} \leq 3\|x_S\|_{\ell_1/\mathcal{H}}.$$

In fact, this functional restricted eigenvalue assumption is no stronger than the usual (scalar) restricted eigenvalue assumption—in the following theorem we show that any matrix \mathbf{A} satisfying the usual (scalar) restricted eigenvalue assumption, will also satisfy the functional version given in Definition 3.1. We do this by showing that the common inequality used for proving that a matrix satisfies a scalar restricted eigenvalue condition immediately implies the same inequality in the functional setting, and leads to the restricted eigenvalue condition.

Theorem 3. *For a fixed matrix $\mathbf{X} \in \mathbb{R}^{N \times I}$, suppose that for some $c_1, c_2 > 0$, \mathbf{X} satisfies*

$$\|\mathbf{X}z\|_2 \geq c_1 \sqrt{N} \cdot \|z\|_2 - c_2 \sqrt{\log(I)} \cdot \|z\|_{\ell_1},$$

for all $z \in \mathbb{R}^I$. Then the same inequality holds for all $x \in \mathcal{H}^I$, that is,

$$(6) \quad \|\mathbf{X}x\|_{\mathcal{H}} \geq c_1 \sqrt{N} \cdot \|x\|_{\mathcal{H}} - c_2 \sqrt{\log(I)} \cdot \|x\|_{\ell_1/\mathcal{H}}.$$

Furthermore, if $c_1 > 4c_2\sqrt{\frac{I_0 \log(I)}{N}}$, then \mathbf{X} satisfies the $\text{RE}_F(I_0, \alpha)$ property with

$$\alpha = \left(c_1 - 4c_2\sqrt{\frac{I_0 \log(I)}{N}} \right)^2.$$

We are now ready to present the main result of the section.

Theorem 4. *If Assumption 2 holds, \mathbf{X} satisfies $\text{RE}_F(I_0, \alpha)$, and*

$$\lambda \geq 2\sqrt{N}\sqrt{\|\Lambda\|_1 + 2\|\Lambda\|_2\sqrt{\log(I/\delta)} + 2\|\Lambda\|_\infty(\log(I/\delta))},$$

then with probability at least $1 - \delta$, any minimizer $\hat{\beta}$ of $L_F(\beta)$ satisfies

$$\|\mathbf{X}(\hat{\beta} - \beta^*)\|_{\mathcal{H}} \leq \frac{4\lambda\sqrt{I_0}}{\sqrt{\alpha N}} \quad \text{and} \quad \|\hat{\beta} - \beta^*\|_{\ell_1/\mathcal{H}} \leq \frac{16\lambda I_0}{\alpha N}.$$

We mention that since Λ is vector of the eigenvalues of C in decreasing order, then one has that $\|\Lambda\|_2^2 = \sum_{i=1}^{\infty} \lambda_i^2$, $\|\Lambda\|_1 = \sum_{i=1}^{\infty} \lambda_i$, and $\|\Lambda\|_\infty = \lambda_1$. The target function $L_F(\beta)$ is convex and coercive, and thus has a least one solution. As in the scalar case, the function may have multiple solutions, but this still depends heavily on the design \mathbf{X} and does not happen often in practice. In the next section we will discuss how this estimate can be computed.

Interestingly, the rates of convergence for \mathcal{H} -valued response variables are exactly the same as those for a scalar response. This is due to our ability to extend scalar concentration inequalities to general Hilbert spaces, see Lemma 9. If we take $\lambda \sim \sqrt{N \log(I)}$ then the prediction error becomes $\|\mathbf{X}(\hat{\beta} - \beta^*)\|_{\mathcal{H}} \sim \sqrt{I_0 \log(I)}$ while the estimation error becomes $\|\hat{\beta} - \beta^*\|_{\ell_1/\mathcal{H}} \sim I_0 \sqrt{\frac{\log(I)}{N}}$. Since this result applies to any separable Hilbert space, our methodology is applicable to a wide range of applications beyond just the GWAS we consider here. From random fields in spatial statistics to brain imaging in fMRI studies, all fall under this umbrella if they are embedded into a Hilbert space. As a final note, we mention that the λ given in Theorem 4 gives a very tight control of the probability that the two inequalities hold, however, for the convergence rates all that matters is that λ has the right order, as is typical in LASSO style results.

4 Computational Details

Here we present several computational tools which we utilize and can be found in Matlab code available through the corresponding authors's website. We begin by providing computational details for the methods in Section 3. Traditionally, when handling dense functional data, one constructs the functional objects by utilizing basis expansions. This is the cornerstone

of the FDA package in R. If $e_1(\cdot), e_2(\cdot), \dots$ is a basis of \mathcal{H} (e.g. B-splines or Fourier), then we can approximate L_F as

$$L_F(\beta) \approx \frac{1}{2} \sum_{n=1}^N \sum_{j=1}^J \langle Y_n - X_n^\top \beta, e_j \rangle^2 + \sum_{i=1}^I \sqrt{\sum_{j=1}^J \langle \beta_i, e_j \rangle^2},$$

for J large (often over one hundred). Letting $B_{i,j} = \langle \beta_i, e_j \rangle$ and $Y_{nj} = \langle Y_n, e_j \rangle$, we have that the above can be expressed as

$$\begin{aligned} & \frac{1}{2} \sum_{n=1}^N \sum_{j=1}^J (Y_{nj} - X_n^\top B_{*j})^2 + \sum_{i=1}^I \|B_{i*}\| \\ &= \frac{1}{2} \|Y - \mathbf{A}_F \text{vec}(\mathbf{B}^\top)\|^2 + \sum_{i=1}^I \|B_{i*}\|, \end{aligned}$$

where $\mathbf{A}_F = \mathbf{X} \otimes \mathbf{I}_{J \times J}$. For large values of J , this can quickly become a substantial computational burden. However, one can use a data driven basis, such as FPCA, so that J can be taken relatively small. However, we stress that dimension reduction is not our intent, and thus we can choose the number of FPCs to explain nearly all of the variability of the processes. Using such an approach, it is common to move from 200 B-spline basis functions down to 5-10 functional principal components, with nearly no information loss. This phrasing now allows us to use the same group LASSO computational tools as in the sparse setting. However, it is now possible (even in sparse data settings) that J is less than the number observed time points, resulting in \mathbf{A}_F being smaller than \mathbf{A} and with a far simpler form.

Since I , the number of groups (SNPs), will generally be extremely large, we implement a screening rule that allows us to substantially reduce the potential number of groups (SNPs) we consider when we minimize $L(\cdot)$ or $L_F(\cdot)$ for some fixed λ . Using Wang et al. (2013, Theorem 4), we do the following:

- (In parallel.) For each SNP i , define $A_i \in \mathbb{R}^{\sum_n M_n \times J}$ when using the sparse algorithm with entries

$$(A_i)_{nm,j} = X_{ni} e_j(t_{nm}).$$

For the dense algorithm define $A_i \in \mathbb{R}^{NJ \times J}$ as $A_i = X^{(i)} \otimes \mathbf{I}_{J \times J}$, and compute

$$\|A_i^\top Y\|_2 \text{ and } \|A_i\|_F.$$

- Find any $\lambda_0 \geq \max_i \|A_i^\top Y\|_2$. At penalty parameter λ_0 , the group LASSO solution will be $\hat{B}_{\lambda_0} = 0$.
- For any $\lambda < \lambda_0$, according to Wang et al. (2013),

$$\frac{1}{\lambda_0} \|A_i^\top Y\|_2 + \left(\frac{1}{\lambda} - \frac{1}{\lambda_0} \right) \|A_i\|_F \|Y\|_2 < 1 \implies (\hat{B}_\lambda)_{i*} = 0.$$

Therefore, if our aim is to apply a convex optimizer to solve the group LASSO problem with at most s SNPs, then after finding $\lambda_0 = \max_i \|A_i^\top Y\|_2$ we can choose any λ sufficiently large so that no more than s SNPs violate the inequality above. SNPs which do violate the above are then dropped since they will not enter the solution path. This allows us to make a substantial reduction in the number of predictors. In our simulations and application we could handle on the order of 10000 SNPs jointly when fitting the FS-LASSO using the ADMM procedure, with the sparse tools requiring around 30gb of RAM, and the smooth requiring around 4gb. By using this screening rule, we can reduce millions of predictors to tens of thousands, and then do a final fit using any number of convex optimization routines.

Finally, we mention the choice of the smoothing parameter λ . In our simulations we utilize the BIC, though we include several other options in the application section. We calculate the BIC as

$$\log(\hat{\sigma}^2) \sum M_n + JI_{active} \log(N) \quad \text{or} \quad \log(\hat{\sigma}^2) NM_{pc} + M_{pc} I_{active} \log(N),$$

for the sparse and dense methods respectively. Here I_{active} is the number of predictors in the current model. The error $\hat{\sigma}^2$ is calculated as

$$\frac{1}{\sum M_n} \sum_n \sum_m (Y_n(t_{nm}) - \hat{Y}_n(t_{nm}))^2 \quad \text{or} \quad \frac{1}{NM_{pc}} \sum_n \sum_j (Y_{nj} - \hat{Y}_{nj})^2,$$

for the sparse and dense methods respectively. The predicted values are computed by recomputing the corresponding least squares estimates, to eliminate the effect of the bias. Given the dependence in the data, the BIC is an adequate, though not optimal choice. Further work is needed on tuning parameter selection for functional models, but we leave this for future research. Fitting the model for a grid of values for λ can be done relatively quickly by utilizing a warm start, i.e. using the previous solution for a particular λ value as the starting point for finding the next solution.

5 Simulations

In this section we present a simulation study to compare the performance of the discussed methods. The predictors, X_{ni} , are generated from a normally distribution with $\text{Cov}(X_{ni}, X_{n'j}) = 1_{n=n'} \rho^{|i-j|}$, which is an autoregressive covariance with the rate of decay controlled by ρ (with independence across subjects, i.e. X_{n*} and $X_{n'*}$ are independent for $n \neq n'$). We mention that in our genetic application the predictors take values 0/1/2 and are thus not normal. Additional simulations using Binomial predictors can be found in Appendix A, and while the performance of all of the methods presented here decrease, the relative conclusions from comparing the methods stay the same, we thus focus on the normal case here. We take $I = 1000$ and $I_0 = 10$. The nonzero functions $\{\beta_i(t) : t \in [0, 1]\}$ are randomly generated from the Matérn process with parameters $(0, I_0^{-1}, 1/4, 0, 5/2)$ and the errors $\{\varepsilon_n(t) : t \in [0, 1]\}$ are generated in the same way but with parameters $(0, 1, 1/4, 0, 3/2)$, which results in errors that

are less smooth than the parameter functions. The Matérn process is very flexible family of stationary processes which produce more realistic structures for biological applications, as compared to Brownian motion or simple low dimensional structures. Each subject is observed at 10 uniformly distributed locations which differ by curve. We consider $N = 50, 100, 200$ and $\rho = 0.5$ and 0.75 , and use 1000 repetitions of each scenario. For the functional method, FPCA is carried out using the PACE package in Matlab. For the sparse method we use cubic Bsplines with $J = 30$.

To compare the methods without having to worry about tuning parameter selection, we examine smoothed ROC curves which give the proportion of true positives found as a function of the false positives. The curves for all scenarios are given in Figure 1. Surprisingly, the ROC curves for the two method are about the same, meaning that in terms of variable selection the sparse and functional methods have nearly equivalent performance on average, though, as we will see in the application section, for a single iteration they can still disagree on the selected subset. Next we examine the average prediction error in Figure 2, with the penalty parameters chosen by BIC, AIC, and 2-fold cross validation. There we see that the sparse method has an advantage, though this decreases for large sample sizes. The BIC and AIC perform about the same, while the CV criteria seems to perform the worst. We note that, since the prediction error compares linear combinations of the parameter estimates, we use it as a single number proxy summarizing the estimation error as well, though clearly if there were particular patterns for $\beta(t)$ one was interested in examining (linear, sinusoidal, etc.), it would be interesting to include them in the simulations and examine their estimation error, however we don't pursue this further here.

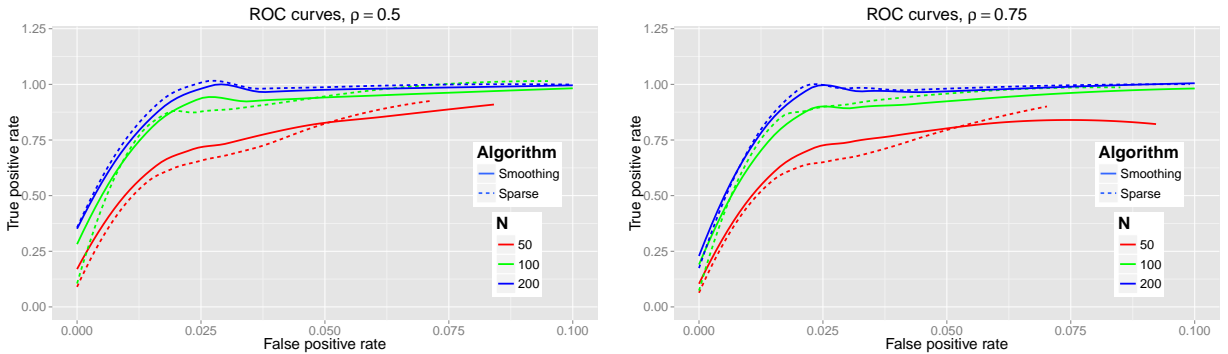


Figure 1: ROC curves comparing the smooth and sparse methods.

Next we compare to two forms of scalar LASSO, one which “time-corrects” by incorporating a nonparametric estimate of the mean (TC-LASSO) and one which does not. Using the nonparametric smoother is an attempt to account for the non-stationarity of the processes. The resulting ROC curves are given in Figure 3 for $\rho = 0.5$ and $\rho = 0.75$, with the solid lines indicating scalar LASSO and the dashed lines the time-corrected version. Examining the plots we see a potentially interesting pattern. For smaller false positive rates, LASSO is actually doing better than FS-LASSO (Figure 1). However, as one moves along the x-

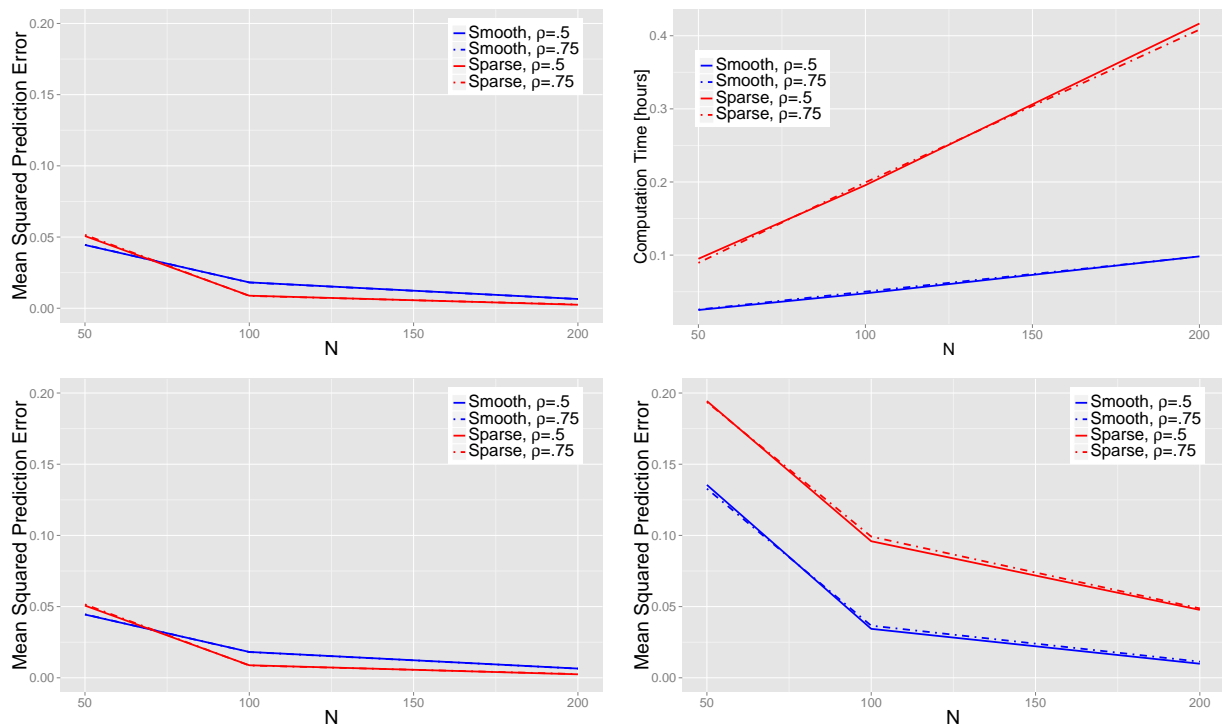


Figure 2: Prediction error and computation times for smooth and sparse methods. The top row is the MSE and computation time for the BIC, while the first panel on the second row is the MSE for AIC, and the second panel on the second row is the MSE for cross-validation.

axis (the false positive rate), the neither version of LASSO changes much, while FS-LASSO quickly tops out, finding all significant predictors. This suggests that LASSO is able to find a subset of predictors rather quickly, likely those which do not vary much over time, but is unable to capture the more complex signals. Surprisingly, incorporating a nonparametric estimate of the mean (TC LASSO) actually results in a large decrease in power. This further supports the need for FS-LASSO, if one believes there is nonlinear structure in the data, then patching LASSO is substantially worse than moving to a functional framework.

As a final comparison of statistical performance, we examine the sensitivity of the sparse FS-LASSO on the choice of J . We consider $J = 27$ and $J = 33$, which constitutes a 10% change in the number of basis functions. The resulting MSE curves are given in Figure 4. Unfortunately, there does appear to be some sensitivity in the performance of FS-LASSO relative to the number of basis functions. While $J = 33$ and $J = 27$ perform similarly, both are slightly worse than $J = 30$, and in fact, worse than the smoothing algorithm. This can be remedied by also choosing the number of basis functions by BIC, AIC, or CV, or by using the smoothing method to select the predictors. Namely, one would ideally like to use the penalty to also control the level of smoothing (not just the selection of the predictors). FS-LASSO behaves very similarly to a ridge regression on the selected predictors, but a penalty which acts more like a smoother would likely perform better (at least in terms of MSE). However,

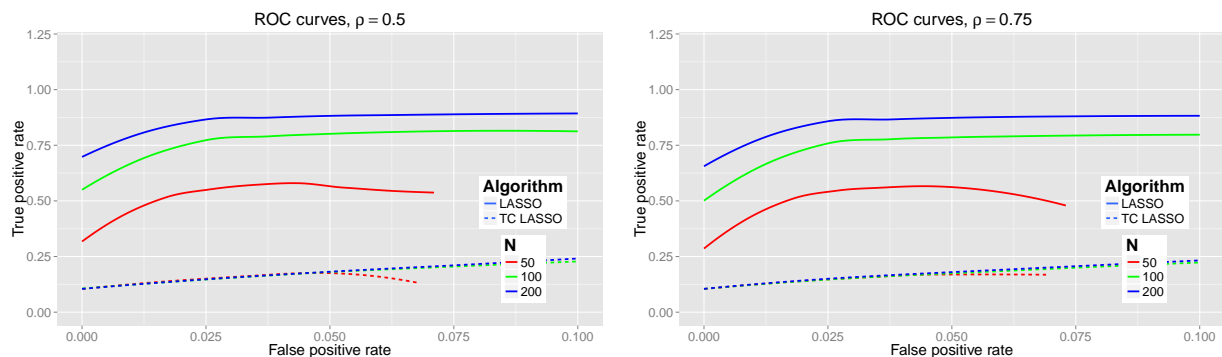


Figure 3: ROC curves for scalar LASSO with (TC LASSO) and without (LASSO) a non-parametric estimate of the mean.

we do not explore this further here as this is an important topic of further research.

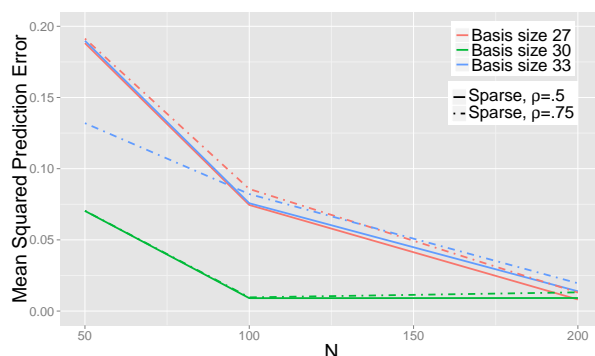


Figure 4: Mean squared prediction error for sparse FS-LASSO with differing numbers of basis functions.

Lastly, we mention computation time and memory. Methods which scale well are critical for applications such as genetic association studies. There, one needs to work with millions of predictors and both computation time and memory become major issues. On a desktop computer, the sparse method presented here can be applied with predictors of the order of 10,000 or so, but not much higher. The smoothing method, which allows for the inclusion of an FPCA, can be applied with predictors of the order of 20,000. On clusters these numbers can be increased depending on the available memory. In Figure 2 we plot average computation times for the two methods when taking $I = 10000$, but keeping everything else the same. We see that the smooth method has a substantial edge in computation time, resulting in shorter times which also scale better with sample size.

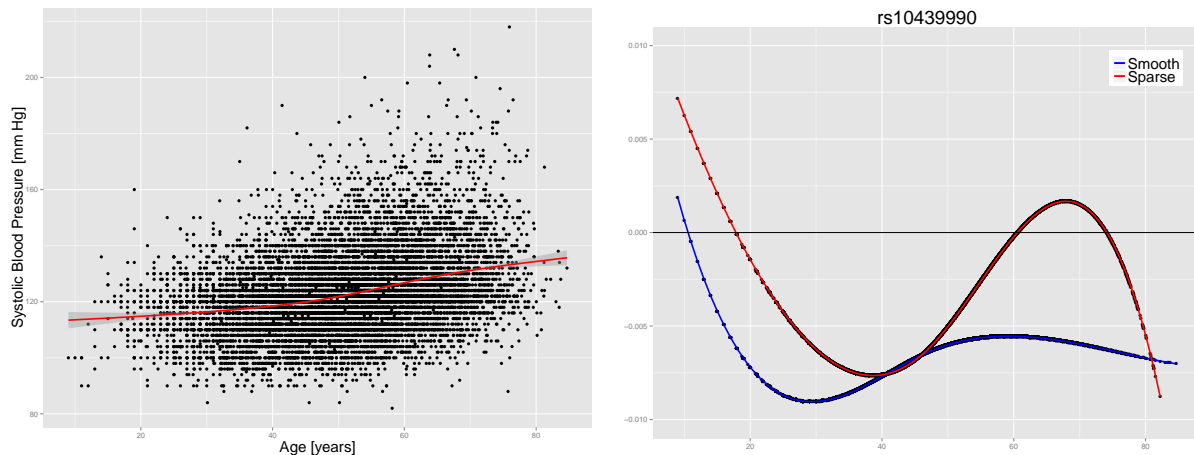


Figure 5: Left Panel: Systolic blood pressure as a function of age from the FHS. The red line is a local linear smoother estimating the mean. Right Panel: Plot of estimated effect for SNP rs10439990 using the ‘smooth’ algorithm (blue) and the ‘sparse’ algorithm (red).

6 Framingham Heart Study

The Framingham Heart Study, FHS, is a long term longitudinal study with the goal of better understanding the risk factors for heart disease. Genome-wide SNP data and phenotype information were downloaded from dbGaP (<http://www.ncbi.nlm.nih.gov/gap>) study accession phs000007.v25.p9. The study consists of three cohorts of subjects with the first cohort recruited in 1958. FHS has had a tremendous impact on our understanding of health and risk factors for heart disease. Since a full account would take far too much space, we refer the interested reader to O’Donnell and Elosua (2008); Mahmood et al. (2014); Chen and Levy (2016) for details on different findings and impacts due to FHS. Here we examine the second cohort which consists of 1924 subjects. Each subject contributed up to 7 clinical exams over 29 years, though some subjects passed away during the study and contributed less. The subjects were genotyped using the Illumina Omni 5M platform resulting in approximately 4.3 million SNPs genotyped. Our goal here is to find a subset of SNPs impacting systolic blood pressure, a primary risk factor for heart disease. As an illustration, in Figure 5 we plot all of the blood pressure measurements versus age and include a local linear smoother. We see that (as is well known) blood pressure increase monotonically, though nonlinearly, with age. When we combine this fact with the high within-subject correlation of blood pressure, we see that scalar procedures such as LASSO are inappropriate. When applying our procedures, we first remove the effects of gender, height, and HDL cholesterol levels nonparametrically using a local linear smoother.

Due to the size of the data, we cannot apply the methods to the entire set jointly. We thus first use our method to screen the number of predictors down to a computationally manageable size. One could take the top performers from marginal regression, but instead we use our procedure to rank the SNPs according to when they enter the model (as one

varies the tuning parameter), and take the top ones. First, we divide up the genome into segments of 10K SNPs for the sparse methods and 20K for the smooth. Using our algorithm, we rank the SNPs in each segment according to the order in which they enter the model (as one decreases the tuning parameter). We take the top ranking SNPs so that, when pooled, we have screened down to approximately 10K SNPs for the sparse and 20K for the smooth. A final application of our methods are applied to these SNPs with the smoothing parameter chosen in four different ways: AIC, BIC, extended BIC with parameter 0.2, and 2-fold cross-validation. The top SNPs for each method are presented in Table 6. As we can see, the smoothing method has found substantially more significant SNPs than the sparse. However, there were 4 candidates which appeared in both lists. Given how similar the sparse functional methods compared in simulations, it might at first be surprising that they only overlap on four SNPs, however, given the huge number of predictors, it seems reasonable that there would be substantial differences between the two. As an illustration, in Figure 5, we provide a plot of one of the estimated SNPs, rs10439990, which was selected using both algorithms. There we see that the effect is negative, meaning that a subject with one or two copies of the mutation has a lower chance of having high blood pressure. However, the effect is not constant over time as it starts lower a younger ages, 20-40, and rises as subjects age, peaking around 70 or so. While the two methods produce similar plots, they disagree on the magnitude of the effect later in life, with the sparse algorithm indicating that the protection is essentially gone, while the smooth algorithm indicates the protection is still present, though diminished. However, we stress that the two submodels selected were different, and thus the difference in estimates could be due in part to having different predictors. This plot also illustrates a cautionary note about interpreting endpoints. With so few observations at the ends, it is safer to focus interpretations on sections with the most observations.

Quantifying the uncertainty of the output from these types of high-dimensional procedures is an ongoing area of research even in the scalar setting. However, to test if either method is dramatically over fitting, we also applied a 10-fold cross-validation to compare the predictive performance between the sparse and smooth method. The resulting values were 0.00961 for the sparse approach and 0.00972 for the smooth approach, thus showing that the sparse approach is doing a slightly better job in terms of prediction. Given the inherently sparse nature of the data, it is maybe surprising that the smooth approach was still so close in terms of prediction. Regardless, as is now common in genetic studies, a follow up on a separate data set would be required to validate the results. Note that for serious follow up studies, it would be useful to extend techniques such as Shi (2015) to the functional setting, which attempt to recover some of the relatively small effects that LASSO style procedures shrink to zero. This is especially useful in genetic studies as the effect sizes are relatively small. Lastly, an extension incorporating dependence between curves would be useful as the FHS data included some related individuals, though we do not pursue this further here.

As a final form of validation, we examined the different association results found in the literature for our selected SNPs. This was accomplished using GWAS Central, <http://www.gwascentral.org/>.

gwascentral.org/index, which provides a searchable database of GWAS results. There were at least a few results that validate some of our findings. In particular, SNP rs10439990 is located the gene ZBTB20, which has been associated with Triglyceride levels (Kathiresan et al., 2007), a common risk factor for heart disease. SNP kgp29965466 is located on the gene ATRNL1, which is associated with several negative health outcomes and indicators including nicotine dependence (Bierut et al., 2007). SNP rs10497371 is on gene MYO3B, which has been associated with Diabetes (Scott et al., 2007). Lastly SNP rs7692492 is on gene FAM190A, which has been associated with coronary heart disease (Samani et al., 2007; Larson et al., 2007).

Top SNPs selected by sparse algorithm					Top SNPs selected by smooth algorithm				
Chromosome, Name	AIC	BIC	EBIC	CV	Chromosome, Name	AIC	BIC	EBIC	CV
2, rs3845756					1, kgp5933154				
2, rs6414023					1, rs17107710				
2, rs10497371					2, kgp5982336				
3, rs10439990					3, kgp26868438				
4, rs7692492					3, rs10439990				
5, kgp30202888					3, kgp11928513				
7, kgp3808198					5, kgp30202888				
7, kgp4510449					7, kgp3808198				
8, kgp8137960					7, kgp4510449				
9, rs1702645					10, kgp29965466				
12, rs978561					11, kgp10123049				
13, rs1924783					12, kgp6953877				
15, kgp6228266					12, rs10859106				

Figure 6: Top SNPs selected by the Sparse Algorithm and by the Smooth Algorithm. Gray boxes indicated the SNP is selected when using the corresponding variable selection criteria. Red indicates SNPs which were chosen by both the smooth and sparse methods.

7 Conclusion

We have provided powerful new tools for analyzing functional or longitudinal data with a large number of predictors. In the sparse case we provided new theory in the form of a restricted eigenvalue condition and accompanying asymptotic theory. While phrased as a functional data method, the sparse case is closely related to varying coefficient models, and thus this work can be viewed as an extension of the work in Wei et al. (2011). In the dense case, we provide a completely new methodology and accompanying asymptotic theory which allows the response functions to take values from any separable Hilbert space. Such generality means the methods can be applied to a variety of settings including traditional

functional outcomes, functional panels, spatial processes, and image data such as fMRI. We also provide accompanying computer code which takes advantage of the structure of group LASSO so that our methods can be applied efficiently.

Our simulations suggest that the choice between sparse and smooth tool sets is not straightforward. The simulations were done in a traditionally sparse setting, but the two methods were nearly equivalent in terms of variable selection and the functional method required far fewer computational resources. However, estimation accuracy was better for the sparse methods. This opens the door to future work which could involve multistage methods where selection is done via smooth methods and estimation via sparse methods.

While our methods perform well, they also highlight the need for better tuning parameter selection methods. Methods such as BIC and cross-validation can be applied in this setting with good results, but such methods are not tailored to the dependence inherent in functional data, and we hope to investigate possible adaptations to these selection criteria for the functional setting in future work.

References

- R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- P. Bickel, Y. Ritov, and A. Tsybakov. Simultaneous analysis of lasso and dantzig selector. *The Annals of Statistics*, 37(4):1705–1732, 2009.
- L. J. Bierut, P. A. Madden, N. Breslau, E. O. Johnson, D. Hatsukami, O. F. Pomerleau, G. E. Swan, J. Rutter, S. Bertelsen, and L. Fox. Novel genes identified in a high-density genome wide association study for nicotine dependence. *Human molecular genetics*, 16(1):24–35, 2007.
- P. Bühlmann and S. Van de Geer. *Statistics for high-dimensional data: methods, theory and applications*. Springer, 2011.
- H. Cardot, F. Ferraty, and P. Sarda. Spline estimators for the functional linear model. *Statistica Sinica*, 13:571–591, 2003.
- G. Chen and D. Levy. Contributions of the framingham heart study to the epidemiology of coronary heart disease. *JAMA cardiology*, 2016.
- Y. Chen, J. Goldsmith, and R. T. Ogden. Variable selection in function-on-scalar regression. *Stat*, 2016.
- J. Fan and J. Zhang. Two-step estimation of functional linear models with applications to longitudinal data. *Journal of the Royal Statistical Society (B)*, 62:303–322, 2000.
- Y. Fan, N. Foutz, G. M. James, W. Jank, et al. Functional response additive model estimation with online virtual stock markets. *The Annals of Applied Statistics*, 8(4):2435–2460, 2014.

- Y. Fan, G. James, and P. Radchenko. Functional Additive Regression. *Annals of Statistics*, 00(000):0000–0000, 2015+.
- J. Gertheiss, A. Maity, and A. Staicu. Variable selection in generalized functional linear models. *Stat*, 2(1):86–101, 2013.
- T. Hastie, R. Tibshirani, and J. Friedman. *Elements of Statistical Learning*. Springer, 2001.
- D. Hoover, J. Rice, C. Wu, and L.-P. Yang. Nonparametric smoothing estimates of time-varying coefficient models with longitudinal data. *Biometrika*, 85(4):809–822, 1998.
- L. Horváth and P. Kokoszka. *Inference for Functional Data with Applications*. Springer, New York, 2012.
- D. Hsu, S. Kakade, and T. Zhang. A tail inequality for quadratic forms of subgaussian random vectors. Technical report, Technical report, Preprint, 2011.
- G. James, T. Hastie, D. Witten, and R. Tibshirani. *An Introduction to Statistical Learning: With Applications in R*. Springer Texts in Statistics. Springer London, Limited, 2013. ISBN 9781461471370. URL <http://books.google.com/books?id=at1bmAECAAJ>.
- G. M. James, J. Wang, and J. Zhu. Functional linear regression that’s interpretable. *The Annals of Statistics*, 37:2083–2595, 2009.
- S. Kathiresan, A. K. Manning, S. Demissie, R. B. D’agostino, A. Surti, C. Guiducci, L. Gianniny, N. P. Burt, O. Melander, and M. Orho-Melander. A genome-wide association study for blood lipid phenotypes in the framingham heart study. *BMC medical genetics*, 8(1):1, 2007.
- P. Kokoszka, I. Maslova, J. Sojka, and L. Zhu. Testing for lack of dependence in the functional linear model. *Canadian Journal of Statistics*, 36:207–222, 2008.
- D. Kong, K. Xue, F. Yao, and H. H. Zhang. Partially functional linear regression in high dimensions. *Biometrika*, page asv062, 2016.
- M. G. Larson, L. D. Atwood, E. J. Benjamin, L. A. Cupples, R. B. D’Agostino, C. S. Fox, D. R. Govindaraju, C.-Y. Guo, N. L. Heard-Costa, and S.-J. Hwang. Framingham heart study 100k project: genome-wide associations for cardiovascular disease outcomes. *BMC medical genetics*, 8(1):1, 2007.
- B. Laurent and P. Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, 28:1302–1338, 2000.
- Y. Li and T. Hsing. Uniform convergence rates for nonparametric regression and principal component analysis in functional/longitudinal data. *The Annals of Statistics*, pages 3321–3351, 2010.
- H. Lian. Shrinkage estimation and selection for multiple functional regression. *Statistica Sinica*, 23:51–74, 2013.

- K. Lounici, M. Pontil, S. Van de Geer, and A. Tsybakov. Oracle inequalities and optimal inference under group sparsity. *The Annals of Statistics*, 39(4):2164–2204, 2011.
- S. S. Mahmood, D. Levy, R. S. Vasan, and T. J. Wang. The framingham heart study and the epidemiology of cardiovascular disease: a historical perspective. *The Lancet*, 383(9921):999–1008, 2014.
- H. Matsui and S. Konishi. Variable selection for functional regression models via the L1 regularization. *Computational Statistics & Data Analysis*, 55(12):3304–3310, 2011.
- H. Müller. Functional modeling of longitudinal data. In *Longitudinal Data Analysis*, pages 223–252. CRC Press, 2008.
- C. J. O’Donnell and R. Elosua. Cardiovascular risk factors. insights from framingham heart study. *Revista Espanola de Cardiologia (English Edition)*, 61(3):299–310, 2008.
- R. Oliveira. The lower tail of random quadratic forms, with applications to ordinary least squares and restricted eigenvalue properties. *arXiv preprint arXiv:1312.2903*, 2013.
- J. Ramsay and B. Silverman. *Functional Data Analysis*. Springer, 2005.
- M. Reimherr and D. Nicolae. A functional data analysis approach for genetic association studies. *The Annals of Applied Statistics*, 8:406–429, 2014.
- P. Reiss, M. Mennes, and L. Huang. Fast function-on-scalar regression with penalized basis expansions. *International Journal of Biostatistics*, 6:Article 28, 2010.
- N. J. Samani, J. Erdmann, A. S. Hall, C. Hengstenberg, M. Mangino, B. Mayer, R. J. Dixon, T. Meitinger, P. Braund, and H.-E. Wichmann. Genomewide association analysis of coronary artery disease. *New England Journal of Medicine*, 357(5):443–453, 2007.
- L. J. Scott, K. L. Mohlke, L. L. Bonnycastle, C. J. Willer, Y. Li, W. L. Duren, M. R. Erdos, H. M. Stringham, P. S. Chines, and A. U. Jackson. A genome-wide association study of type 2 diabetes in finns detects multiple susceptibility variants. *science*, 316(5829):1341–1345, 2007.
- P. Shi. *Weak signal identification and inference in penalized model selection*. PhD thesis, University of Illinois at Urbana-Champaign, 2015.
- R. Tibshirani. The lasso problem and uniqueness. *Electronic Journal of Statistics*, 7:1456–1490, 2013.
- J. Wang, B. Lin, P. Gong, P. Wonka, and J. Ye. Lasso screening rules via dual polytope projection. In *Advances in Neural Information Processing Systems*, pages 1070–1078, 2013.
- L. Wang, H. Li, and J. Huang. Variable selection in nonparametric varying-coefficient models for analysis of repeated measurements. *Journal of the American Statistical Association*, 103(484):1556–1569, 2008.

- F. Wei, J. Huang, and H. Li. Variable selection and estimation in high-dimensional varying-coefficient models. *Statistica Sinica*, 21(4):1515–1540, 2011.
- F. Yao, H.-G. Müller, and J.-L. Wang. Functional linear regression analysis for longitudinal data. *The Annals of Statistics*, 33:2873–2903, 2005.
- X. Zhang and J. Wang. From sparse to dense functional data and beyond. Technical report, Preprint, 2015+.
- P. Zhao and L. Xue. Variable selection for semiparametric varying coefficient partially linear errors-in-variables models. *Journal of Multivariate Analysis*, 101(8):1872–1883, 2010.
- H. Zhu, R. Li, and L. Kong. Multivariate varying coefficient model for functional responses. *Annals of Statistics*, 40(5):2634–2666, 2012.

A Binomial Simulations

Here we carry out a nearly identical simulation study as in Section 5 except the normally distributed predictors are replaced with $B_{in} \sim \text{Binom}(2, p_i)$ random variables to more closely emulate the structure of the application. To generate B_{in} which are correlated, we use a probit model based on the the same simulation structure as in Section 5; we take those normal random variables as the input in probit model where $p_i = \text{logit}^{-1}(X_{in})$. The MSE curves are summarized in Figure 7. If we compare these results with those from the bottom right panel from Figure 2, here we also use CV, we see that the pattern is very similar. The error does not change dramatically between the two correlations, but the sparse method has a noticeably lower MSE.

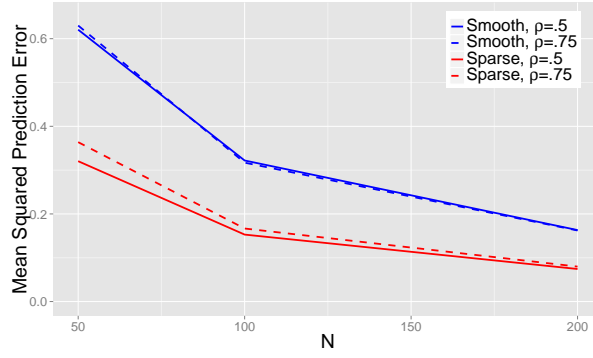


Figure 7: MSE curves comparing the smooth and sparse methods with $\text{Binom}(2, 1/2)$ predictors. Tuning parameter is selected with cross-validation.

B Proofs for Section 2

B.1 Proof of Theorem 1

Our analysis for this Theorem closely follows standard techniques for sparse and group sparse regression under restricted eigenvalue conditions, such as the work by Lounici et al. (2011) and Bickel et al. (2009).

Recall that $\hat{\mathbf{B}}$ is a minimizer of

$$L(\mathbf{B}) = \frac{1}{2} \|Y - \mathbf{A} \text{vec}(\mathbf{B}^\top)\|_2^2 + \lambda \|\mathbf{B}\|_{\ell_1/\ell_2}.$$

Then in particular, we must have $L(\hat{\mathbf{B}}) \leq L(\mathbf{B}^*)$, that is,

$$\frac{1}{2} \|Y - \mathbf{A} \text{vec}(\hat{\mathbf{B}}^\top)\|_2^2 + \lambda \|\hat{\mathbf{B}}\|_{\ell_1/\ell_2} \leq \frac{1}{2} \|Y - \mathbf{A} \text{vec}(\mathbf{B}^{*\top})\|_2^2 + \lambda \|\mathbf{B}^*\|_{\ell_1/\ell_2}.$$

After rearranging some terms,

$$\frac{1}{2} \|\mathbf{A} \text{vec}(\hat{\mathbf{B}}^\top - \mathbf{B}^{*\top})\|_2^2 \leq \left\langle Y - \mathbf{A} \text{vec}(\mathbf{B}^{*\top}), \mathbf{A} \text{vec}(\hat{\mathbf{B}}^\top - \mathbf{B}^*) \right\rangle + \lambda \left(\|\mathbf{B}^*\|_{\ell_1/\ell_2} - \|\hat{\mathbf{B}}\|_{\ell_1/\ell_2} \right).$$

Let $S_0 \subset [I]$ indicate the row support of \mathbf{B}^* , with $|S_0| \leq I_0$. Write $\Delta = \widehat{\mathbf{B}} - \mathbf{B}^*$. Then standard arguments show that

$$\|\mathbf{B}^*\|_{\ell_1/\ell_2} - \|\widehat{\mathbf{B}}\|_{\ell_1/\ell_2} \leq \|\Delta_{S_0}\|_{\ell_1/\ell_2} - \|\Delta_{S_0^c}\|_{\ell_1/\ell_2}.$$

Furthermore,

$$\begin{aligned} \left\langle Y - \mathbf{A} \text{vec}(\mathbf{B}^{\star\top}), \mathbf{A} \text{vec}(\widehat{\mathbf{B}}^\top - \mathbf{B}^{\star\top}) \right\rangle &= \left\langle \mathbf{A}^\top (Y - \mathbf{A} \text{vec}(\mathbf{B}^{\star\top})), \text{vec}(\Delta^\top) \right\rangle \\ &= \sum_i \left\langle \mathbf{A}_i^\top (Y - \mathbf{A} \text{vec}(\mathbf{B}^{\star\top})), \Delta_{i,*} \right\rangle \\ &\leq \sum_i (\|\mathbf{A}_i^\top (Y - \mathbf{A} \text{vec}(\mathbf{B}^{\star\top}))\|_2 \cdot \|\Delta_{i,*}\|_2) \\ &\leq \max_i \|\mathbf{A}_i^\top (Y - \mathbf{A} \text{vec}(\mathbf{B}^{\star\top}))\|_2 \cdot \|\Delta\|_{\ell_1/\ell_2}, \end{aligned}$$

where \mathbf{A}_i is the matrix with entries $(\mathbf{A}_i)_{nm,j} = \mathbf{A}_{nm,ij} = X_{ni} \cdot e_j(t_{nm})$.

Consider the event $\lambda \geq 2 \max_i \|\mathbf{A}_i^\top (Y - \mathbf{A} \text{vec}(\mathbf{B}^{\star\top}))\|_2$. On this event, we would then have

$$\begin{aligned} \frac{1}{2} \|\mathbf{A} \text{vec}(\Delta^\top)\|^2 &\leq \|\Delta\|_{\ell_1/\ell_2} \cdot \max_i \|\mathbf{A}_i^\top (Y - \mathbf{A} \text{vec}(\mathbf{B}^{\star\top}))\|_2 + \lambda (\|\Delta_{S_0}\|_{\ell_1/\ell_2} - \|\Delta_{S_0^c}\|_{\ell_1/\ell_2}) \\ &\leq \|\Delta_{S_0}\|_{\ell_1/\ell_2} \cdot 1.5\lambda - \|\Delta_{S_0^c}\|_{\ell_1/\ell_2} \cdot 0.5\lambda. \end{aligned}$$

Since the left-hand side is nonnegative, therefore, $\|\Delta_{S_0^c}\|_{\ell_1/\ell_2} \leq 3\|\Delta_{S_0}\|_{\ell_1/\ell_2}$. Applying the $\text{RE}_{\text{group}}(I_0, \alpha)$ condition, then,

$$\|\mathbf{A} \text{vec}(\Delta^\top)\|^2 \geq N\alpha \|\Delta\|_{\text{F}}^2.$$

Combining everything,

$$\frac{1}{2} N\alpha \|\Delta\|_{\text{F}}^2 \leq \frac{1}{2} \|\mathbf{A} \text{vec}(\Delta^\top)\|^2 \leq \|\Delta_{S_0}\|_{\ell_1/\ell_2} \cdot 1.5\lambda - \|\Delta_{S_0^c}\|_{\ell_1/\ell_2} \cdot 0.5\lambda \leq \|\Delta_{S_0}\|_{\ell_1/\ell_2} \cdot 1.5\lambda.$$

For the Frobenius norm result, we further write

$$\|\Delta_{S_0}\|_{\ell_1/\ell_2} \cdot 1.5\lambda \leq \sqrt{|S_0|} \cdot \|\Delta_{S_0}\|_{\text{F}} \cdot 1.5\lambda \leq \sqrt{|S_0|} \cdot \|\Delta\|_{\text{F}} \cdot 1.5\lambda$$

and so

$$\|\Delta\|_{\text{F}} \leq \sqrt{|S_0|} \cdot 1.5\lambda \cdot \frac{2}{N\alpha}$$

while for the ℓ_1/ℓ_2 -norm result we have

$$\|\Delta\|_{\ell_1/\ell_2} = \|\Delta_{S_0}\|_{\ell_1/\ell_2} + \|\Delta_{S_0^c}\|_{\ell_1/\ell_2} \leq 4\|\Delta_{S_0}\|_{\ell_1/\ell_2} \leq 4\sqrt{|S_0|} \|\Delta_{S_0}\|_{\text{F}} \leq 4\sqrt{|S_0|} \|\Delta\|_{\text{F}}.$$

This proves the theorem, as long as we can show that with high probability,

$$\lambda \geq 2 \max_i \|\mathbf{A}_i^\top (Y - \mathbf{A} \text{vec}(\mathbf{B}^{\star\top}))\|_2.$$

To prove this, take any $i \in [I]$. Then

$$\|\mathbf{A}_i^\top (Y - \mathbf{A} \text{vec}(\mathbf{B}^{\star\top}))\|_2 = \|\mathbf{A}_i^\top (T + \epsilon)\|_2 \leq \|\mathbf{A}_i^\top T\|_2 + \|\mathbf{A}_i^\top \epsilon\|_2 ,$$

where T is the truncation error defined earlier while ϵ has entries $\epsilon_{nm} = \epsilon_n(t_{nm})$. Now we bound the two pieces separately. First, we bound the operator norm of \mathbf{A}_i . By definition, $(\mathbf{A}_i)_{nm,j} = X_{ni} \cdot e_j(t_{nm})$ and so we have

$$\mathbf{A}_i^\top \mathbf{A}_i = \sum_n X_{ni}^2 \mathbf{E}_n^\top \mathbf{E}_n \preceq \|X\|_\infty^2 \sum_n \mathbf{E}_n^\top \mathbf{E}_n = \|X\|_\infty^2 \cdot N \cdot \mathbf{F} ,$$

proving that $\|\mathbf{A}_i\|_{\text{op}} \leq \|X\|_\infty \cdot \sqrt{N\|\mathbf{F}\|_{\text{op}}}$. Then

$$\|\mathbf{A}_i^\top T\|_2 \leq \|\mathbf{A}_i\|_{\text{op}} \cdot \|T\|_2 \leq \|X\|_\infty \cdot \sqrt{N\|\mathbf{F}\|_{\text{op}}} \cdot \|T\|_2 .$$

Next, to bound $\|\mathbf{A}_i^\top \epsilon\|_2$, observe that $\mathbf{A}_i^\top \epsilon \sim N(0, \mathbf{A}_i^\top \Sigma \mathbf{A}_i)$, where Σ is a block-diagonal matrix with blocks given by $\Sigma_1, \dots, \Sigma_N$. Then, applying (Hsu et al., 2011, Proposition 1.1),

$$\mathbb{P} \left\{ \|\mathbf{A}_i^\top \epsilon\|_2^2 > \text{trace}(\mathbf{A}_i^\top \Sigma \mathbf{A}_i) + 2\sqrt{\text{trace}((\mathbf{A}_i^\top \Sigma \mathbf{A}_i)^2) \cdot t} + 2\|\mathbf{A}_i^\top \Sigma \mathbf{A}_i\|_{\text{op}} \cdot t \right\} \leq e^{-t}$$

for any $t \geq 0$. Since $\text{trace}(M) \leq J\|M\|_{\text{op}}$ for any $J \times J$ matrix, and using the simple identity $2\sqrt{Jt} \leq J + t$, we simplify this to the weaker statement

$$\mathbb{P} \left\{ \|\mathbf{A}_i^\top \epsilon\|_2^2 > \|\mathbf{A}_i^\top \Sigma \mathbf{A}_i\|_{\text{op}} \cdot (2J + 3t) \right\} \leq e^{-t}$$

We calculate

$$\|\mathbf{A}_i^\top \Sigma \mathbf{A}_i\|_{\text{op}} \leq \|\mathbf{A}_i\|_{\text{op}}^2 \cdot \|\Sigma\|_{\text{op}} \leq \|X\|_\infty^2 \cdot N\|\mathbf{F}\|_{\text{op}} \cdot \max_n \|\Sigma_n\|_{\text{op}} .$$

Combining everything, and setting $t = \log(I/\delta)$, for each I we have

$$\mathbb{P} \left\{ \|\mathbf{A}_i^\top \epsilon\|_2^2 > \|X\|_\infty^2 \cdot N\|\mathbf{F}\|_{\text{op}} \cdot \max_n \|\Sigma_n\|_{\text{op}} \cdot (2J + 3\log(I/\delta)) \right\} \leq \frac{\delta}{I} .$$

Therefore, with probability at least $1 - \delta$, for all $i = 1, \dots, I$,

$$\|\mathbf{A}_i^\top \epsilon\|_2^2 \leq \|X\|_\infty^2 \cdot N\|\mathbf{F}\|_{\text{op}} \cdot \max_n \|\Sigma_n\|_{\text{op}} \cdot (2J + 3\log(I/\delta))$$

and so for all i ,

$$\begin{aligned} \|\mathbf{A}_i^\top (Y - \mathbf{A} \text{vec}(\mathbf{B}^{\star\top}))\|_2 &\leq \|\mathbf{A}_i^\top T\|_2 + \|\mathbf{A}_i^\top \epsilon\|_2 \\ &\leq \|X\|_\infty \sqrt{N\|\mathbf{F}\|_{\text{op}}} \cdot \left(\|T\|_2 + \sqrt{\max_n \|\Sigma_n\|_{\text{op}} \cdot (2J + 3\log(I/\delta))} \right) \\ &\leq \lambda/2 , \end{aligned}$$

as desired.

B.2 Proof of Theorem 2

We use the Transfer Principle of Oliveira (2013, Lemma 5.1), which connects sparse eigenvalues of \mathbf{A} to a restricted eigenvalue property for \mathbf{A} . This result is stated for sparsity but can be straightforwardly extended to group-sparsity. We restate the group-sparse form here as a Lemma without proof.

Lemma 1 (Adapted from Oliveira (2013, Lemma 5.1)). *Suppose \mathbf{A} satisfies that, for any k_1 -row-sparse $\mathbf{W} \in \mathbb{R}^{I \times J}$,*

$$(7) \quad \|\mathbf{A} \text{vec}(\mathbf{W}^\top)\|_2^2 \geq \lambda_0 \|\mathbf{W}\|_{\text{F}}^2 .$$

Then for all $\mathbf{W} \in \mathbb{R}^{I \times J}$,

$$\|\mathbf{A} \text{vec}(\mathbf{W}^\top)\|_2^2 \geq \lambda_0 \|\mathbf{W}\|_{\text{F}}^2 - \frac{\max_i \|\mathbf{A}_i\|_{\text{op}}^2 \cdot \|\mathbf{W}\|_{\ell_1/\ell_2}^2}{k_1 - 1} .$$

To apply this result to our work, we need to first find a λ_0 such that (7) is satisfied (with high probability), and then we need to compute a bound (holding with high probability) on each $\|\mathbf{A}_i\|_{\text{op}}$.

Step 1: finding λ_0 for condition (7) We first give another lemma, proved below:

Lemma 2. *Suppose that the assumptions of Theorem 2 hold, and take any fixed sequence $\{\mathbf{E}_n\}$ satisfying (4). Choose any $\delta > 0$ and any $k_1 \geq 1$. Then there are constants $a_1, a_2 > 0$ depending only on $\nu, \lambda_{\min}(\Sigma), \gamma_0, \gamma_1$, such that if*

$$(8) \quad N \geq a_1(k_1 J + k_1 \log(I) + \log(1/\delta)),$$

then with probability at least $1 - \delta/2$, for all k_1 -row-sparse \mathbf{W} ,

$$\frac{\sum_n \|\mathbf{E}_n^\top \mathbf{W}^\top X_n\|_2}{N \sqrt{\text{trace}(\mathbf{W}^\top \Sigma \mathbf{W})}} \geq a_2 .$$

To apply these results to our work, first note that by Lemma 2, if N satisfies (8), then with probability at least $1 - \delta/2$, for all k_1 -sparse \mathbf{W} (where we specify k_1 later),

$$\begin{aligned} \|\mathbf{A} \cdot \text{vec}(\mathbf{W})\|_2^2 &= \sum_n \|\mathbf{E}_n^\top \mathbf{W}^\top X_n\|_2^2 \geq \frac{1}{N} \left(\sum_n \|\mathbf{E}_n^\top \mathbf{W}^\top X_n\|_2 \right)^2 \\ &\geq N \text{trace}(\mathbf{W}^\top \Sigma \mathbf{W}) \cdot (a_2)^2 . \end{aligned}$$

In particular, the assumption (7) in Lemma 1 holds with $\lambda_0 = N(a_2)^2 \lambda_{\min}(\Sigma)$.

Step 2: bounding $\|\mathbf{A}_i\|_{\text{op}}$ Next we bound $\max_i \|\mathbf{A}_i\|_{\text{op}}^2$. From the proof of Theorem 1, we know that $\|\mathbf{A}_i\|_{\text{op}}^2 \leq \|X\|_\infty^2 \cdot N \|\mathbf{F}\|_{\text{op}}$, where $\mathbf{F} = \frac{1}{N} \sum_n \mathbf{E}_n^\top \mathbf{E}_n$. By (4), we see that $\|\mathbf{F}\|_{\text{op}} \leq \gamma_1^2$. Furthermore, X_{ij} is mean zero and ν -subgaussian, and so, by standard subgaussian tail bounds, with probability at least $1 - \delta/2$,

$$\|X\|_\infty = \max_{ni} |X_{ni}| \leq \nu \sqrt{2 \log(4IJ/\delta)} .$$

Step 3: applying Lemma 1 Combining the results from Step 1 and Step 2, and applying Lemma 1, we see that with probability at least $1 - \delta$, for all $\mathbf{W} \in \mathbb{R}^{I \times J}$,

$$\|\mathbf{A} \text{vec}(\mathbf{W}^\top)\|_2^2 \geq N \left[(a_1)^2 \lambda_{\min}(\Sigma) \cdot \|\mathbf{W}\|_{\mathbf{F}}^2 - \frac{2\nu^2 \log(4IJ/\delta) \cdot \gamma_1^2 \cdot \|\mathbf{W}\|_{\ell_1/\ell_2}^2}{k_1 - 1} \right].$$

For any \mathbf{W} with $\|\mathbf{W}_{S^c}\|_{\ell_1/\ell_2} \leq 3 \|\mathbf{W}_S\|_{\ell_1/\ell_2}$ (where $|S| \leq k$), we then have

$$\|\mathbf{W}\|_{\ell_1/\ell_2} \leq 4 \|\mathbf{W}_S\|_{\ell_1/\ell_2} \leq 4\sqrt{k} \|\mathbf{W}_S\|_{\mathbf{F}} \leq 4\sqrt{k} \|\mathbf{W}\|_{\mathbf{F}},$$

and so the result above gives

$$\|\mathbf{A} \text{vec}(\mathbf{W}^\top)\|_2^2 \geq N \|\mathbf{W}\|_{\mathbf{F}}^2 \left[(a_2)^2 \lambda_{\min}(\Sigma) - \frac{2\nu^2 \log(4IJ/\delta) \cdot \gamma_1^2 \cdot 16k}{k_1 - 1} \right].$$

Taking

$$k_1 \geq 1 + \frac{2\nu^2 \log(4IJ/\delta) \cdot \gamma_1^2 \cdot 16k}{0.5(a_2)^2 \lambda_{\min}(\Sigma)},$$

we obtain

$$\|\mathbf{A} \text{vec}(\mathbf{W}^\top)\|_2^2 \geq N \|\mathbf{W}\|_{\mathbf{F}}^2 \cdot 0.5(a_2)^2 \lambda_{\min}(\Sigma)$$

for all $\mathbf{W} \in \mathbb{R}^{I \times J}$ with $\|\mathbf{W}_{S^c}\|_{\ell_1/\ell_2} \leq 3 \|\mathbf{W}_S\|_{\ell_1/\ell_2}$ for any $|S| \leq k$, as desired. The sample size requirement (5) in Theorem 2 ensures that the sample size assumption in Lemma 2 will hold for the specified choice of k_1 .

B.2.1 Proof of Lemma 2

For the proof of this lemma, we'll need several supporting lemmas, proved below. The first lemma shows that $\sum_n \|\mathbf{E}_n^\top \mathbf{W}^\top X_n\|_2$ satisfies upper and lower bounds with high probability for any fixed *single* matrix \mathbf{W} .

Lemma 3. *For any fixed \mathbf{W} ,*

$$\mathbb{P} \left\{ c_1 \leq \frac{\sum_n \|\mathbf{E}_n^\top \mathbf{W}^\top X_n\|_2}{N \sqrt{\text{trace}(\mathbf{W}^\top \Sigma \mathbf{W})}} \leq c_2 \right\} \geq 1 - 2e^{-c_3 N},$$

where $c_1, c_2, c_3 > 0$ depend only on $\nu, \lambda_{\min}(\Sigma), \gamma_0, \gamma_1$.

The next lemma shows that proving a lower bound on $\sum_n \|\mathbf{E}_n^\top \mathbf{W}^\top X_n\|_2$ over all row-sparse \mathbf{W} can be reduced to proving a lower bound on a finite covering set. This result is a simple extension of Baraniuk et al. (2008, Lemma 5.1), from the sparse setting to the group-sparse setting, and so we do not give the proof here.

Lemma 4 (Adapted from Baraniuk et al. (2008, Lemma 5.1)). *Let*

$$\mathcal{W} = \{ \mathbf{W} \in \mathbb{R}^{I \times J} : \mathbf{W} \text{ is } k_1\text{-row-sparse, } \text{trace}(\mathbf{W}^\top \Sigma \mathbf{W}) = 1 \}.$$

Let $f(\mathbf{W})$ be any seminorm (that is, $f(c \cdot \mathbf{W}) = |c| \cdot f(\mathbf{W})$ and $f(\mathbf{W}_1 + \mathbf{W}_2) \leq f(\mathbf{W}_1) + f(\mathbf{W}_2)$). Choose any $\epsilon \in (0, 1/2)$. Then there exists $\mathcal{W}' \subset \mathcal{W}$ with $|\mathcal{W}'| \leq I^{k_1}(3/\epsilon)^{k_1 J}$ such that

$$\inf_{\mathbf{W} \in \mathcal{W}} f(\mathbf{W}) \geq \inf_{\mathbf{W} \in \mathcal{W}'} f(\mathbf{W}) - \frac{\epsilon}{1 - \epsilon} \sup_{\mathbf{W} \in \mathcal{W}'} f(\mathbf{W}).$$

Proof of Lemma 2. By Lemma 3, for any fixed \mathbf{W} , with probability at least $1 - 2e^{-c_3 N}$,

$$c_1 \leq \frac{\sum_n \|\mathbf{E}_n^\top \mathbf{W}^\top X_n\|_2}{N \sqrt{\text{trace}(\mathbf{W}^\top \Sigma \mathbf{W})}} \leq c_2$$

where $c_1, c_2, c_3 > 0$ depend only on $\nu, \lambda_{\min}(\Sigma), \gamma_0, \gamma_1$. Now choose $\epsilon = \frac{c_1}{4c_2}$. In the notation of Lemma 4, let $f(\mathbf{W}) = \sum_n \|\mathbf{E}_n^\top \mathbf{W}^\top X_n\|_2$, and take \mathcal{W}' as in the lemma. Since $|\mathcal{W}'| \leq I^{k_1}(3/\epsilon)^{k_1 J}$, the above bound is true for all $\mathbf{W} \in \mathcal{W}'$ with probability at least

$$1 - 2e^{-c_3 N + \log(I^{k_1}(3/\epsilon)^{k_1 J})} \geq 1 - \delta/2,$$

by our lower bound $N \geq a_1(k_1 J + k_1 \log(I) + \log(1/\delta))$ as long as the constant a_1 is chosen appropriately. Next, applying Lemma 4, for all $\mathbf{W} \in \mathcal{W}$,

$$\frac{\sum_n \|\mathbf{E}_n^\top \mathbf{W}^\top X_n\|_2}{N \sqrt{\text{trace}(\mathbf{W}^\top \Sigma \mathbf{W})}} \geq c_1 - \frac{\epsilon}{1 - \epsilon} \cdot c_2 \geq \frac{c_1}{2}.$$

After defining $a_2 = c_1/2$, this proves the lemma. \square

B.2.2 Proof of Lemma 3

We begin by stating a supporting lemma.

Lemma 5. Let $B_1 \in \mathbb{R}^{M_1 \times d}, \dots, B_N \in \mathbb{R}^{M_N \times d}$ be fixed matrices with

$$\sum_{n=1}^N \|B_n\|_F \geq c_1 N$$

and

$$\sqrt{\sum_{n=1}^N \|B_n\|_F^2} \leq c_2 \sqrt{N}$$

for some $c_1, c_2 > 0$. Let $X_1, \dots, X_N \in \mathbb{R}^d$ be iid random vectors with $\mathbb{E}[X_n] = 0$ and $\mathbb{E}[X_n X_n^\top] = \mathbb{I}_d$, which are σ -subgaussian, meaning that $\mathbb{E}[e^{v^\top X_n}] \leq e^{\sigma^2 \|v\|_2^2 / 2}$ for all fixed $v \in \mathbb{R}^d$. Then

$$\mathbb{P} \left\{ c_3 N \leq \sum_{n=1}^N \|B_n X_n\|_2 \leq c_4 N \right\} \geq 1 - 2e^{-c_5 N},$$

where $c_3, c_4, c_5 > 0$ depend only on c_1, c_2, σ .

Proof of Lemma 3. First we calculate

$$\begin{aligned}
\sum_n \|\mathbf{E}_n^\top \mathbf{W}^\top \Sigma^{1/2}\|_F^2 &= \sum_n \text{trace}(\mathbf{E}_n^\top \mathbf{W}^\top \Sigma \mathbf{W} \mathbf{E}_n) \\
&= \text{trace}(\Sigma^{1/2} \mathbf{W} (\sum_n \mathbf{E}_n \mathbf{E}_n^\top) \mathbf{W}^\top \Sigma^{1/2}) \leq \text{trace}(\Sigma^{1/2} \mathbf{W} \mathbf{W}^\top \Sigma^{1/2}) \cdot N \left\| \frac{1}{N} \sum_n \mathbf{E}_n \mathbf{E}_n^\top \right\| \\
&= \text{trace}(\mathbf{W}^\top \Sigma \mathbf{W}) \cdot N \left\| \frac{1}{N} \sum_n \mathbf{E}_n \mathbf{E}_n^\top \right\| \leq \gamma_1 N \text{trace}(\mathbf{W}^\top \Sigma \mathbf{W}).
\end{aligned}$$

And,

$$\begin{aligned}
\sum_n \|\mathbf{E}_n^\top \mathbf{W}^\top \Sigma^{1/2}\|_F &= \sum_n \sqrt{\sum_i \|(\Sigma^{1/2} \mathbf{W})_{i,*} \mathbf{E}_n\|_2^2} \geq \sum_n \sum_i \frac{\|(\Sigma^{1/2} \mathbf{W})_{i,*}\|_2}{\sqrt{\text{trace}(\mathbf{W}^\top \Sigma \mathbf{W})}} \|(\Sigma^{1/2} \mathbf{W})_{i,*} \mathbf{E}_n\|_2 \\
&\geq \gamma_0 N \sum_i \frac{\|(\Sigma^{1/2} \mathbf{W})_{i,*}\|_2^2}{\sqrt{\text{trace}(\mathbf{W}^\top \Sigma \mathbf{W})}} = \gamma_0 N \sqrt{\text{trace}(\mathbf{W}^\top \Sigma \mathbf{W})}.
\end{aligned}$$

Now, writing $\tilde{X}_n = \Sigma^{-1/2} X_n$ and defining

$$B_n = \frac{\mathbf{E}_n^\top \mathbf{W}^\top}{\sqrt{\text{trace}(\mathbf{W}^\top \Sigma \mathbf{W})}},$$

we see that the \tilde{X}_n 's are iid with mean zero, identity covariance, and are $\nu/\sqrt{\lambda_{\min}(\Sigma)}$ -subgaussian. And, $\|\mathbf{E}_n^\top \mathbf{W}^\top X_n\|_2 = \sqrt{\text{trace}(\mathbf{W}^\top \Sigma \mathbf{W})} \cdot \|B_n \Sigma^{1/2} \tilde{X}_n\|_2$. Applying Lemma 5, then, for some constants c_1, c_2, c_3 depending only on $\nu, \lambda_{\min}(\Sigma), \gamma_0, \gamma_1$,

$$\mathbb{P} \left\{ c_1 N \sqrt{\text{trace}(\mathbf{W}^\top \Sigma \mathbf{W})} \leq \sum_n \|\mathbf{E}_n^\top \mathbf{W}^\top X_n\|_2 \leq c_2 N \sqrt{\text{trace}(\mathbf{W}^\top \Sigma \mathbf{W})} \right\} \geq 1 - 2e^{-c_3 N}.$$

□

B.2.3 Proof of Lemma 5

Before we prove Lemma 5, we first state two additional supporting lemmas, proved below:

Lemma 6. *Let $x \in \mathbb{R}^N$ be a vector with $\|x\|_1 \geq c_1 N$, $\|x\|_2 \leq c_2 \sqrt{N}$. Then*

$$\left| \left\{ n : |x_n| \geq \frac{c_1}{2} \right\} \right| \geq \frac{(c_1)^2}{4(c_2)^2} N.$$

Lemma 7. *Suppose Z_1, \dots, Z_N are independent non-negative random variables with $\mathbb{E}[(Z_n)^2] \geq a_1$ and $\mathbb{E}[e^{(Z_n)^2}] \leq a_2$ for all $n = 1, \dots, N$, where $a_1, a_2 > 0$. Then*

$$\mathbb{P} \left\{ \sum_n Z_n < a_3 N \right\} \leq e^{-a_4 N} \quad \text{and} \quad \mathbb{P} \left\{ \sum_n (Z_n)^2 > a_5 N \right\} \leq e^{-a_6 N}$$

where $a_3, a_4, a_5, a_6 > 0$ are constants depending only on a_1, a_2 .

Now we prove Lemma 5.

Proof of Lemma 5. For each n , define $Z_n = \frac{\|B_n X_n\|_2}{2\sigma\|B_n\|_F}$. We have

$$\mathbb{E} [\|B_n X_n\|_2^2] = \mathbb{E} [X_n^\top B_n^\top B_n X_n] = \mathbb{E} [\text{trace}(B_n^\top B_n X_n X_n^\top)] = \text{trace}(B_n^\top B_n) = \|B_n\|_F^2,$$

and so $\mathbb{E} [(Z_n)^2] = \frac{1}{4\sigma^2} =: a_1$. Next, it is known from Hsu et al. (2011, Remark 2.3) that

$$\mathbb{E} \left[e^{\eta\|B_n X_n\|_2^2} \right] \leq \sigma^2 \text{trace}(B_n^\top B_n) \eta + \frac{\sigma^4 \text{trace}((B_n^\top B_n)^2) \eta^2}{1 - 2\sigma^2 \|B_n^\top B_n\| \eta}$$

for all $0 \leq \eta < \frac{1}{2\sigma^2 \|B_n^\top B_n\|}$. Using basic facts about matrix norms we have $\|B_n^\top B_n\| = \|B_n\|^2 \leq \|B_n\|_F^2$ and $\text{trace}(B_n^\top B_n) = \|B_n\|_F^2$ and $\text{trace}((B_n^\top B_n)^2) = \|(B_n^\top B_n)^2\|_F \leq \|B_n^\top B_n\|_F^2$. Setting $\eta = \frac{1}{4\sigma^2 \|B_n\|_F^2}$ we obtain

$$\mathbb{E} \left[e^{(Z_n)^2} \right] = \mathbb{E} \left[e^{\|B_n X_n\|_2^2 / (4\sigma^2 \|B_n\|_F^2)} \right] \leq e^{3/8} =: a_2.$$

Applying Lemma 7 to $\{Z_n : n = 1, \dots, n\}$, we see that for constants $a_5, a_6 > 0$ depending only on a_1, a_2 ,

$$\mathbb{P} \left\{ \sum_n \frac{\|B_n X_n\|_2^2}{4\sigma^2 \|B_n\|_F^2} \leq a_5 N \right\} \geq 1 - e^{-a_6 N}.$$

Furthermore, setting $\mathcal{I} = \{n : \|B_n\|_F \geq c_1/2\}$, by applying Lemma 6 to the vector $(\|B_n\|_F)_{n=1, \dots, N}$, we see that $|\mathcal{I}| \geq \frac{(c_1)^2}{4(c_2)^2} N$. Applying Lemma 7 to $\{Z_n : n \in \mathcal{I}\}$ we see that for constants $a_3, a_4 > 0$ depending only on a_1, a_2 ,

$$\mathbb{P} \left\{ \sum_{n \in \mathcal{I}} \frac{\|B_n X_n\|_2}{2\sigma \|B_n\|_F} \geq a_3 |\mathcal{I}| \right\} \geq 1 - e^{-a_4 |\mathcal{I}|} \geq 1 - e^{-a_4 \frac{(c_1)^2}{4(c_2)^2} N}.$$

Next, assume these events hold. We have

$$\sum_n \|B_n X_n\|_2 = \sum_n \frac{\|B_n X_n\|_2}{\|B_n\|_F} \cdot \|B_n\|_F \leq \sqrt{\sum_n \frac{\|B_n X_n\|_2^2}{\|B_n\|_F^2}} \cdot \sqrt{\sum_n \|B_n\|_F^2} \leq 2\sigma \sqrt{a_5 N} \cdot c_2 \sqrt{N}.$$

Furthermore,

$$\begin{aligned} \sum_n \|B_n X_n\|_2 &\geq \sum_{n \in \mathcal{I}} \|B_n X_n\|_2 = \sum_{n \in \mathcal{I}} 2\sigma \|B_n\|_F \cdot \frac{\|B_n X_n\|_2}{2\sigma \|B_n\|_F} \\ &\geq c_1 \cdot \sum_{n \in \mathcal{I}} \frac{\|B_n X_n\|_2}{2\sigma \|B_n\|_F} \geq c_1 \cdot a_3 |\mathcal{I}| \geq c_1 \cdot a_3 \cdot \frac{(c_1)^2}{4(c_2)^2} N. \end{aligned}$$

Setting

$$c_3 = c_1 \cdot a_3 \cdot \frac{(c_1)^2}{4(c_2)^2}, \quad c_4 = 2\sigma c_2 \sqrt{a_5}, \quad c_5 = \min \left\{ a_6, a_4 \frac{(c_1)^2}{4(c_2)^2} \right\},$$

we have proved the lemma. \square

Finally, we prove the two supporting results, Lemmas 6 and 7.

Proof of Lemma 6. We have

$$\begin{aligned}
c_1 N &\leq \sum_{n=1}^N |x_n| = \sum_{n=1}^N |x_n| \cdot \mathbf{1}\{|x_n| \geq c_1/2\} + \sum_{n=1}^N |x_n| \cdot \mathbf{1}\{|x_n| < c_1/2\} \\
&\leq \sum_{n=1}^N |x_n| \cdot \mathbf{1}\{|x_n| \geq c_1/2\} + N \cdot c_1/2 \\
&\leq \sqrt{\sum_{n=1}^N |x_n|^2} \cdot \sqrt{\sum_{n=1}^N \mathbf{1}\{|x_n| \geq c_1/2\}^2} + N \cdot c_1/2 \\
&\leq c_2 \sqrt{N} \cdot \sqrt{|\{n : |x_n| \geq c_1/2\}|} + N \cdot c_1/2,
\end{aligned}$$

and therefore, $|\{n : |x_n| \geq c_1/2\}| \geq \frac{(c_1)^2}{4(c_2)^2} N$. □

Proof of Lemma 7. For the upper bound, we have

$$\mathbb{P} \left\{ \sum_n (Z_n)^2 > 2a_2 N \right\} \leq \mathbb{E} \left[e^{\sum_n (Z_n)^2 - 2a_2 N} \right] = \left(\prod_n \mathbb{E} \left[e^{(Z_n)^2} \right] \right) \cdot e^{-2a_2 N} = e^{-a_2 N}.$$

For the lower bound, we have, for each n ,

$$\begin{aligned}
a_1 &\leq \mathbb{E} [(Z_n)^2] = \mathbb{E} [(Z_n)^{0.5} (Z_n)^{1.5}] \leq \sqrt{\mathbb{E} [Z_n]} \cdot \sqrt{\mathbb{E} [(Z_n)^3]} \quad \text{since } Z_n \geq 0 \\
&\leq \sqrt{\mathbb{E} [Z_n]} \cdot \sqrt[4]{\mathbb{E} [(Z_n)^6]} \leq \sqrt{\mathbb{E} [Z_n]} \cdot \sqrt[4]{\mathbb{E} [6e^{(Z_n)^2}]} \leq \sqrt{\mathbb{E} [Z_n]} \cdot \sqrt[4]{6a_2}
\end{aligned}$$

and so $\mathbb{E} [Z_n] \geq \frac{a_1^2}{\sqrt{6a_2}}$. Next, taking any $t \in [0, 1]$,

$$\begin{aligned}
\mathbb{E} [e^{-tZ_n}] &= 1 - t\mathbb{E} [Z_n] + \sum_{k \geq 2} \frac{(-1)^k t^k \mathbb{E} [(Z_n)^k]}{k!} \\
&\leq -t\mathbb{E} [Z_n] + \left(1 + \sum_{k \geq 1} \frac{t^{2k} \mathbb{E} [(Z_n)^{2k}]}{(2k)!} \right) \quad \text{by removing negative terms} \\
&= -t\mathbb{E} [Z_n] + \mathbb{E} \left[e^{t^2 (Z_n)^2} \right] \quad \text{since } k! \leq (2k)! \\
&\leq -t\mathbb{E} [Z_n] + \mathbb{E} \left[t^2 e^{(Z_n)^2} + (1 - t^2) \cdot 1 \right] \quad \text{by convexity of } z \mapsto e^z \\
&\leq 1 - t \frac{a_1^2}{\sqrt{6a_2}} + t^2 (a_2 - 1).
\end{aligned}$$

Setting $t = \min \left\{ 1, \frac{a_1^2}{2(a_2 - 1)\sqrt{6a_2}} \right\}$ we get

$$\mathbb{E} [e^{-tZ_n}] \leq 1 - t \frac{(a_1)^2}{2\sqrt{6a_2}} \leq e^{-t \frac{(a_1)^2}{2\sqrt{6a_2}}}.$$

Then

$$\begin{aligned} \mathbb{P} \left\{ \sum_n Z_n < \frac{(a_1)^2}{4\sqrt{6a_2}} N \right\} &\leq \mathbb{E} \left[e^{tN \frac{(a_1)^2}{4\sqrt{6a_2}} - t \sum_n Z_n} \right] \\ &= \mathbb{E} \left[e^{tN \frac{(a_1)^2}{4\sqrt{6a_2}}} \right] \left(\prod_n \mathbb{E} [e^{-tZ_n}] \right) = e^{tN \frac{(a_1)^2}{4\sqrt{6a_2}} - tN \frac{(a_1)^2}{2\sqrt{6a_2}}} = e^{-tN \frac{(a_1)^2}{4\sqrt{6a_2}}}. \end{aligned}$$

Setting

$$a_3 = \frac{(a_1)^2}{4\sqrt{6a_2}}, \quad a_4 = \min \left\{ 1, \frac{a_1^2}{2(a_2 - 1)\sqrt{6a_2}} \right\} \cdot \frac{(a_1)^2}{4\sqrt{6a_2}}, \quad a_5 = 2a_2, \quad a_6 = a_2,$$

we have proved the lemma. \square

C Proofs for Section 3

We begin this section with the proof of Theorem 3. After, we will divide the proof of Theorem 4 into a sequence of lemmas.

C.1 Proof of Theorem 3

Let e_1, e_2, \dots be an orthonormal basis of \mathcal{H} . Then each coordinate x_i of $x \in \mathcal{H}^I$, can be expressed as

$$x_i = \sum_k x_{i,k} e_k.$$

We denote $x_{i,k} \in \mathbb{R}$ as the coordinates with respect to the e_1, e_2, \dots basis, and we will also let $x^{(k)}$ represent the vector $\{x_{1,k}, \dots, x_{I,k}\}$. We can then write

$$\begin{aligned} \|\mathbf{X}x\|_{\mathcal{H}}^2 &= \langle \mathbf{X}x, \mathbf{X}x \rangle_{\mathcal{H}} = \sum_{n=1}^N \langle X_n x, X_n x \rangle \\ &= \sum_{n=1}^N \sum_{i=1}^I \sum_{j=1}^I X_{ni} X_{nj} \langle x_i, x_j \rangle_{\mathcal{H}} \end{aligned}$$

Expressing the x_i with respect to the orthonormal basis e_1, e_2, \dots ,

$$\begin{aligned} &= \sum_n \sum_i \sum_j \sum_k \sum_l X_{ni} X_{nj} x_{i,k} x_{j,l} \langle e_k, e_l \rangle \\ &= \sum_n \sum_i \sum_j \sum_k X_{ni} X_{nj} x_{i,k} x_{j,k} \\ &= \sum_k x^{(k)\top} \mathbf{X}^\top \mathbf{X} x^{(k)} \\ &= \sum_k \|\mathbf{X}x^{(k)}\|_2^2. \end{aligned}$$

Note that the last norm above is simply the Euclidean norm on \mathbb{R}^N . By assumption in the Theorem, we have that the above is bounded from below as follows:

$$\begin{aligned}
\|\mathbf{X}x\|_{\mathcal{H}}^2 &= \sum_k \|\mathbf{X}x^{(k)}\|_2^2 \\
&\geq \sum_k \left(c_1 \sqrt{N} \|x^{(k)}\|_2 - c_2 \sqrt{\log(I)} \|x^{(k)}\|_1 \right)^2 \\
&= (c_1)^2 N \sum_k \|x^{(k)}\|_2^2 - 2c_1 c_2 \sqrt{N \log(I)} \sum_k \|x^{(k)}\|_2 \|x^{(k)}\|_1 + (c_2)^2 \log(I) \sum_k \|x^{(k)}\|_1^2.
\end{aligned}$$

Examining the first sum in this last line, we have by Parseval's identity

$$\|x\|_{\mathcal{H}}^2 = \sum_i \|x_i\|_{\mathcal{H}}^2 = \sum_i \sum_k |x_{i,k}|^2 = \sum_k \|x^{(k)}\|_2^2.$$

For the second sum, the Cauchy-Schwartz inequality gives

$$\sum_k \|x^{(k)}\|_2 \|x^{(k)}\|_1 \leq \|x\|_{\mathcal{H}} \sqrt{\sum_k \|x^{(k)}\|_1^2}.$$

So we have that

$$\|\mathbf{X}x\|_{\mathcal{H}}^2 \geq \left(c_1 \sqrt{N} \|x\|_{\mathcal{H}} - c_2 \sqrt{\log(I)} \sqrt{\sum_k \|x^{(k)}\|_1^2} \right)^2$$

or equivalently

$$\|\mathbf{X}x\|_{\mathcal{H}} \geq c_1 \sqrt{N} \|x\|_{\mathcal{H}} - c_2 \sqrt{\log(I)} \sqrt{\sum_k \|x^{(k)}\|_1^2}.$$

Now for the last piece observe that (by Cauchy-Schwarz)

$$\begin{aligned}
\sum_k \|x^{(k)}\|_1^2 &= \sum_k \sum_i \sum_j |x_{i,k}| |x_{j,k}| \\
&\leq \sum_i \sum_j \left(\sum_k |x_{i,k}|^2 \right)^{1/2} \left(\sum_k |x_{j,k}|^2 \right)^{1/2} \\
&= \left(\sum_i \left(\sum_k |x_{i,k}|^2 \right)^{1/2} \right)^2 \\
&= \left(\sum_i \|x_i\|_{\mathcal{H}} \right)^2 = \|x\|_{\ell_1/\mathcal{H}}^2.
\end{aligned}$$

We can therefore conclude that

$$\|\mathbf{X}x\|_{\mathcal{H}} \geq c_1 \sqrt{N} \|x\|_{\mathcal{H}} - c_2 \sqrt{\log(I)} \|x\|_{\ell_1/\mathcal{H}}$$

as claimed.

Next suppose that $c_1 > 4c_2\sqrt{\frac{I_0 \log(I)}{N}}$. We will show that \mathbf{X} satisfies the $\text{RE}_F(I_0, \alpha)$ property with $\alpha = \left(c_1 - 4c_2\sqrt{\frac{I_0 \log(I)}{N}}\right)^2$. To see this, take any set S with $|S| \leq I_0$, and any x with $\|x_{S^c}\|_{\ell_1/\mathcal{H}} \leq 3\|x_S\|_{\ell_1/\mathcal{H}}$. By the Cauchy-Schwarz inequality, which relates the ℓ_1 and ℓ_2 norms of any vector (including a vector of functions x_S), we have

$$\|x_S\|_{\ell_1/\mathcal{H}} \leq \|x_S\|_{\mathcal{H}} \cdot \sqrt{|\text{Support}(x_S)|} \leq \|x_S\|_{\mathcal{H}} \cdot \sqrt{I_0}.$$

By the triangle inequality,

$$\|x\|_{\ell_1/\mathcal{H}} \leq \|x_S\|_{\ell_1/\mathcal{H}} + \|x_{S^c}\|_{\ell_1/\mathcal{H}} \leq 4\|x_S\|_{\ell_1/\mathcal{H}} \leq 4\|x_S\|_{\mathcal{H}} \cdot \sqrt{I_0} \leq 4\|x\|_{\mathcal{H}} \cdot \sqrt{I_0}.$$

Then returning to the above,

$$\begin{aligned} \|\mathbf{X}x\|_{\mathcal{H}} &\geq c_1\sqrt{N}\|x\|_{\mathcal{H}} - c_2\sqrt{\log(I)}\|x\|_{\ell_1/\mathcal{H}} \geq c_1\sqrt{N}\|x\|_{\mathcal{H}} - c_2\sqrt{\log(I)} \left(4\|x\|_{\mathcal{H}} \cdot \sqrt{I_0}\right) \\ &= \|x\|_{\mathcal{H}} \cdot \sqrt{N} \cdot \left(c_1 - 4c_2\sqrt{\frac{I_0 \log(I)}{N}}\right) = \|x\|_{\mathcal{H}} \cdot \sqrt{\alpha N}, \end{aligned}$$

proving that \mathbf{X} satisfies the $\text{RE}_F(I_0, \alpha)$ property.

C.2 Proof of Theorem 4

We follow a very similar structure to the one found in Chapter 6 of Bühlmann and Van de Geer (2011), adapting the arguments for function spaces as needed. We break the proof into a sequence of Lemmas which can be thought of as a basic inequality, a functional concentration inequality, applying that inequality, a last lemma to set up the final proof. We then combine all the lemmas to prove Theorem 4.

Lemma 8 (Basic Inequality). *With probability one, we always have the inequality*

$$\frac{1}{2}\|\mathbf{X}(\hat{\beta} - \beta^*)\|_{\mathcal{H}}^2 + \lambda\|\hat{\beta}\|_{\ell_1/\mathcal{H}} \leq \left(\max_{1 \leq i \leq I} \|\varepsilon^\top \mathbf{X}^{(i)}\|\right) \|\hat{\beta} - \beta^*\|_{\ell_1/\mathcal{H}} + \lambda\|\beta^*\|_{\ell_1/\mathcal{H}}.$$

Proof. We can set up a *basic inequality* using

$$\frac{1}{2}\|Y - \mathbf{X}\hat{\beta}\|^2 + \lambda\|\hat{\beta}\|_{\ell_1/\mathcal{H}} \leq \frac{1}{2}\|Y - \mathbf{X}\beta^*\|^2 + \lambda\|\beta^*\|_{\ell_1/\mathcal{H}}.$$

We can rewrite

$$\begin{aligned} \frac{1}{2}\|Y - \mathbf{X}\hat{\beta}\|^2 &= \frac{1}{2}\|Y - \mathbf{X}\beta^* - \mathbf{X}(\hat{\beta} - \beta^*)\|^2 \\ &= \frac{1}{2}\|Y - \mathbf{X}\beta^*\|^2 + \frac{1}{2}\|\mathbf{X}(\hat{\beta} - \beta^*)\|^2 - \frac{2}{2}\langle Y - \mathbf{X}\beta^*, \mathbf{X}(\hat{\beta} - \beta^*) \rangle. \end{aligned}$$

This gives

$$\begin{aligned}
\frac{1}{2}\|\mathbf{X}(\widehat{\beta} - \beta^*)\|^2 + \lambda\|\widehat{\beta}\|_{\ell_1/\mathcal{H}} &\leq \langle Y - \mathbf{X}\beta^*, \mathbf{X}(\widehat{\beta} - \beta^*) \rangle + \lambda\|\beta^*\|_{\ell_1/\mathcal{H}} \\
&= \sum_{i=1}^I \langle \varepsilon^\top \mathbf{X}^{(i)}, \widehat{\beta}_i - \beta_i^* \rangle + \lambda\|\beta^*\|_{\ell_1/\mathcal{H}} \\
&\leq \sum_{i=1}^I \|\varepsilon^\top \mathbf{X}^{(i)}\| \|\widehat{\beta}_i - \beta_i^*\| + \lambda\|\beta^*\|_{\ell_1/\mathcal{H}} \\
&\leq \max_{1 \leq i \leq I} \|\varepsilon^\top \mathbf{X}^{(i)}\| \|\widehat{\beta} - \beta^*\|_{\ell_1/\mathcal{H}} + \lambda\|\beta^*\|_{\ell_1/\mathcal{H}},
\end{aligned}$$

which is the desired result. \square

Lemma 9. *Let X be an \mathcal{H} valued Gaussian process with mean zero and covariance operator C . Let $\Lambda^\top = (\lambda_1, \lambda_2, \dots)$ be a vector of the eigenvalues of C (in decreasing order). Then we have the bound*

$$P\left\{\|X\|^2 \geq \|\Lambda\|_1 + 2\|\Lambda\|_2\sqrt{t} + 2\|\Lambda\|_\infty t\right\} \leq \exp(-t).$$

Proof. Using the Karhunen-Loeve expansion, we can express

$$\|X\|^2 \stackrel{\mathcal{D}}{=} \sum_{j=1}^{\infty} \lambda_j Z_j^2,$$

where $\{Z_j\}$ are iid standard normal. Note that since X is square integrable, we have that $\sum_i \lambda_i = \|\Lambda\|_1 < \infty$, and therefore $\|\Lambda\|_2 < \infty$ as well. Define the events, for $J = 1, 2, \dots$

$$\mathcal{A}_J = \left\{ \sum_{j=1}^J \lambda_j Z_j^2 \geq \|\Lambda\|_1 + 2\|\Lambda\|_2\sqrt{t} + 2\lambda_1 t \right\}.$$

Since $\|\Lambda\|_1 \geq \sum_{j=1}^J \lambda_j$ and $\|\Lambda\|_2^2 \geq \sum_{j=1}^J \lambda_j^2$, it follows from Lemma 1 in Laurent and Massart (2000) that

$$P(\mathcal{A}_J) \leq \exp(-t),$$

for all J . Since $\sum_{j=1}^J \lambda_j Z_j^2$ is a strictly increasing sequence of J we have that

$$\mathcal{A}_1 \subseteq \mathcal{A}_2 \subseteq \mathcal{A}_3 \dots$$

Using the continuity from below of probability measures, we can pass the limit $J \rightarrow \infty$ to obtain the desired result. \square

Lemma 10. *Assume that $\varepsilon = (\varepsilon_1, \dots, \varepsilon_N)$ are iid Gaussian elements of \mathcal{H} and \mathbf{X} is a design matrix with standardized columns. Define the event*

$$\mathcal{F} = \left\{ \max_{1 \leq i \leq I} \|\varepsilon^\top \mathbf{X}^{(i)}\| \leq \lambda_0 \right\}$$

where, for $\delta > 0$,

$$\lambda_0 = \sqrt{N} \sqrt{\|\Lambda\|_1 + 2\|\Lambda\|_2 \sqrt{\log(I/\delta)} + 2\lambda_1 \log(I/\delta)}.$$

Then we have that

$$P(\mathcal{F}) \geq 1 - \delta.$$

Proof. Notice that if $\{\varepsilon_n\}$ are iid Gaussian with mean zero and covariance operator C , then $\varepsilon^\top \mathbf{X}^{(i)}/\sqrt{N}$ is also mean zero Gaussian, and the covariance operator is given by

$$\frac{1}{N} \text{Cov}(\varepsilon^\top \mathbf{X}^{(i)}) = \frac{1}{N} \sum_{n=1}^N \text{Cov}(\varepsilon_n X_{ni}) = C,$$

since the columns of \mathbf{X} are standardized. Therefore, $V_i = \varepsilon^\top \mathbf{X}^{(i)}/\sqrt{N}$ is also Gaussian with mean zero and covariance operator C . Now apply Lemma 9 with $t = -\log(\delta)$ to obtain

$$\begin{aligned} P \left\{ \max_{1 \leq i \leq I} \|V_i\|^2 \geq \|\Lambda\|_1 + 2\|\Lambda\|_2 \sqrt{t + \log(I)} + 2\|\Lambda\|_\infty(t + \log(I)) \right\} \\ \leq I \exp(-(t + \log(I))) = \exp(-t). \end{aligned}$$

Since $\exp(-t) = \delta$, the claim holds. \square

Lemma 11. On \mathcal{F} with $\lambda \geq 2\lambda_0$ we have that

$$\|\mathbf{X}(\beta^* - \hat{\beta})\|^2 + \lambda \|\hat{\beta}_{S_0^c}\|_{\ell_1/\mathcal{H}} \leq 3\lambda \|(\hat{\beta} - \beta^*)_{S_0}\|_{\ell_1/\mathcal{H}}.$$

Proof. Applying Lemma 8 and multiplying both sides by 2, we have that, on \mathcal{F} ,

$$\|\mathbf{X}(\beta^* - \hat{\beta})\|^2 + 2\lambda \|\hat{\beta}\|_{\ell_1/\mathcal{H}} \leq \lambda \|\hat{\beta} - \beta^*\|_{\ell_1/\mathcal{H}} + 2\lambda \|\beta^*\|_{\ell_1/\mathcal{H}}.$$

Analyzing the penalty term on the LHS, we can apply the reverse triangle inequality to obtain

$$\|\hat{\beta}\|_{\ell_1/\mathcal{H}} = \|\hat{\beta}_{S_0}\|_{\ell_1/\mathcal{H}} + \|\hat{\beta}_{S_0^c}\|_{\ell_1/\mathcal{H}} \geq \|\beta_{S_0}^*\|_{\ell_1/\mathcal{H}} - \|\hat{\beta}_{S_0} - \beta_{S_0}^*\|_{\ell_1/\mathcal{H}} + \|\hat{\beta}_{S_0^c}\|_{\ell_1/\mathcal{H}}$$

and on the RHS we have

$$\|\hat{\beta} - \beta^*\|_{\ell_1/\mathcal{H}} = \|\hat{\beta}_{S_0} - \beta_{S_0}^*\|_{\ell_1/\mathcal{H}} + \|\hat{\beta}_{S_0^c}\|_{\ell_1/\mathcal{H}}.$$

Combining these two calculations we arrive at the bound

$$\begin{aligned} \|\mathbf{X}(\beta^* - \hat{\beta})\|^2 &\leq \lambda \|\hat{\beta} - \beta^*\|_{\ell_1/\mathcal{H}} + 2\lambda \|\beta^*\|_{\ell_1/\mathcal{H}} - 2\lambda \|\hat{\beta}\|_{\ell_1/\mathcal{H}} \\ &\leq \lambda (\|\hat{\beta}_{S_0} - \beta_{S_0}^*\|_{\ell_1/\mathcal{H}} + \|\hat{\beta}_{S_0^c}\|_{\ell_1/\mathcal{H}}) + 2\lambda \|\beta_{S_0}^*\|_{\ell_1/\mathcal{H}} \\ &\quad - 2\lambda (\|\beta_{S_0}^*\|_{\ell_1/\mathcal{H}} - \|\hat{\beta}_{S_0} - \beta_{S_0}^*\|_{\ell_1/\mathcal{H}} + \|\hat{\beta}_{S_0^c}\|_{\ell_1/\mathcal{H}}) \\ &= 3\lambda \|\hat{\beta}_{S_0} - \beta_{S_0}^*\|_{\ell_1/\mathcal{H}} - \lambda \|\hat{\beta}_{S_0^c}\|_{\ell_1/\mathcal{H}}, \end{aligned}$$

which is the desired result. \square

Proof of Theorem 4. We examine both terms together to obtain

$$\|\mathbf{X}(\widehat{\beta} - \beta^*)\|^2 + \lambda\|\widehat{\beta} - \beta^*\|_{\ell_1/\mathcal{H}} = \|\mathbf{X}(\widehat{\beta} - \beta^*)\|^2 + \lambda\|(\widehat{\beta} - \beta^*)_{S_0}\|_{\ell_1/\mathcal{H}} + \lambda\|\widehat{\beta}_{S_0^c}\|_{\ell_1/\mathcal{H}}.$$

Applying Lemmas 10 and 11 we have that, with probability $1 - \delta$,

$$\|\mathbf{X}(\widehat{\beta} - \beta^*)\|^2 + \lambda\|(\widehat{\beta} - \beta^*)_{S_0}\|_{\ell_1/\mathcal{H}} + \lambda\|\widehat{\beta}_{S_0^c}\| \leq 4\lambda\|(\widehat{\beta} - \beta^*)_{S_0}\|_{\ell_1/\mathcal{H}} \leq 4\lambda\sqrt{I_0}\|(\widehat{\beta} - \beta^*)_{S_0}\|.$$

Applying Definition 3.1, we have that the RHS above is bounded by

$$4\lambda\sqrt{I_0}\|(\widehat{\beta} - \beta^*)_{S_0}\| \leq \frac{4\lambda\sqrt{I_0}}{\sqrt{\alpha N}}\|\mathbf{X}(\beta^* - \widehat{\beta})\| \leq \frac{1}{2}\|\mathbf{X}(\beta^* - \widehat{\beta})\|^2 + \frac{8\lambda^2 I_0}{\alpha N}.$$

Note the last inequality follows from the simple bound $4uv \leq \frac{1}{2}u^2 + 8v^2$ for any real u and v (the LHS just completes the square). We therefore have that

$$\|\mathbf{X}(\beta^* - \widehat{\beta})\|^2 + \lambda\|(\widehat{\beta} - \beta^*)_{S_0}\|_{\ell_1/\mathcal{H}} + \lambda\|\widehat{\beta}_{S_0^c}\|_{\ell_1/\mathcal{H}} \leq \frac{1}{2}\|\mathbf{X}(\beta^* - \widehat{\beta})\|^2 + \frac{8\lambda^2 I_0}{\alpha N}.$$

This implies

$$\frac{1}{2}\|\mathbf{X}(\beta^* - \widehat{\beta})\|^2 + \lambda\|\widehat{\beta} - \beta^*\|_{\ell_1/\mathcal{H}} \leq \frac{8\lambda^2 I_0}{\alpha N},$$

which proves the claim. \square