

Oracle Inequalities for High-dimensional Prediction

JOHANNES LEDERER¹, LU YU², and IRINA GAYNANOVA³

¹*Department of Statistics
University of Washington
Box 354322
Seattle, WA 98195
E-mail: ledererj@uw.edu*

²*Department of Statistics
University of Washington
Box 354322
Seattle, WA 98195
E-mail: luyu92@uw.edu*

³*Department of Statistics
Texas A&M University
458D Blocker
College Station, TX 77843
E-mail: irinag@stat.tamu.edu*

Abstract The abundance of high-dimensional data in the modern sciences has generated tremendous interest in penalized estimators such as the lasso, scaled lasso, square-root lasso, elastic net, and many others. In this paper, we establish a general oracle inequality for prediction in high-dimensional linear regression with such methods. Since the proof relies only on convexity and continuity arguments, the result holds irrespective of the design matrix and applies to a wide range of penalized estimators. Overall, the bound demonstrates that generic estimators can provide consistent prediction with any design matrix. From a practical point of view, the bound can help to identify the potential of specific estimators, and they can help to get a sense of the prediction accuracy in a given application.

Keywords: Oracle inequalities, high-dimensional regression, prediction.

1. Introduction

Oracle inequalities are the standard theoretical framework for measuring the accuracy of high-dimensional estimators [9]. Two main benefits of oracle inequalities are that they hold for finite sample sizes and that they adapt to the underlying model parameters. Oracle inequalities are thus used, for example, to compare estimators and to obtain an idea of the sample size needed in a specific application.

For high-dimensional prediction, there are two types of oracle inequalities: so-called fast rate bounds and so-called slow rate bounds. Fast rate bounds hold for near orthogonal designs and bound the prediction error in terms of the sparsity of the regression vectors. Such bounds have been derived for a number of methods, including the lasso, the square-root lasso, and their extensions to grouped variables, see [3, 9–11, 19, 42] and others. Slow rate bounds, on the other hand, hold for any design and bound the prediction error in terms of the penalty value of the regression vectors. Some examples of such bounds have been developed [26, 29, 30], but a general theory has not been established. Importantly, unlike the unfortunate naming suggests, slow rate bounds are of great interest. In particular, slow rate bounds are not inferior to fast rate bounds, quite in contrast [16, 20]: (i) Slow rate bounds hold for any design, while fast rate bounds impose strong and in practice unverifiable assumptions on the correlations in the design. (ii) Even if the assumptions hold, fast rate bounds can contain unfavorable factors, while the factors in slow rate bounds are small, global constants. (iii) Also in terms of rates, slow rate bounds can outmatch even the most favorable fast rate bounds. See [16] and references therein for a detailed comparison of the two types of bounds. To avoid confusion in the following, we will use the terms *penalty bounds* instead of slow rate bounds and *sparsity bounds* instead of fast rate bounds.

In this paper, we develop a general penalty bound for prediction in high-dimensional linear regression. This oracle inequality holds for any sample size, design, and noise distribution, and it applies to a very general family of estimators. For established estimators such as the lasso and the square-root lasso, the result does not imply new guarantees, but it unifies existing bounds in a concise fashion. In general, the result is a convenient resource for penalty bounds, and it demonstrates that prediction guarantees hold broadly in high-dimensional regression.

The organization of the paper is as follows. Below, we introduce the setting and notation and establish relationships to existing work. In Section 2, we state the general result. In Section 3, we specialize this result to specific estimators, including lasso, square-root lasso, group square-root lasso, and others. In Section 4, we conclude with a brief discussion. The proofs are deferred to the Appendix.

The ordering of the following sections is geared towards readers that wish to dive into the technical aspects right away. For getting a first overview instead, one can have a quick glance at the model and the estimators in Displays (1) and (2), respectively, and then skip directly to the examples in Section 3.

Setting and Notation

Model We consider linear regression models of the form

$$Y = X\beta^* + \varepsilon \tag{1}$$

with outcome $Y \in \mathbb{R}^n$, design matrix $X \in \mathbb{R}^{n \times p}$, regression vector $\beta^* \in \mathbb{R}^p$, and noise vector $\varepsilon \in \mathbb{R}^n$. Our goal is prediction, that is, estimation of $X\beta^*$. We allow for general design matrices X , that is, we do not impose conditions on the correlations in X .

Moreover, we allow for general noise ε , that is, we do not restrict ourselves to specific distributions for ε .

Estimators We are particularly interested in high-dimensional settings, where the number of parameters p rivals or even exceeds the number of samples n . As needed in such settings, we assume that the regression vector β^* has some additional structure. This structure can be exploited by penalized estimators, which are the most standard methods for prediction in this context. We thus consider estimators of the form

$$\hat{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \{g(\|Y - X\beta\|_2^2) + \text{pen}(\lambda, \beta)\},$$

where g is some real-valued link function, the mapping $\text{pen}(\lambda, \beta) : \mathbb{R}^k \times \mathbb{R}^p \rightarrow \mathbb{R}$ accounts for the structural assumptions, and λ is a vector-valued tuning parameter. More specifically, we consider assumptions on β^* that are captured by (semi-)norms; the corresponding estimators then read

$$\hat{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\}. \quad (2)$$

To derive results that are specific enough to be useful in concrete examples, we impose some additional conditions in the following.

Link Function The link function $g : \mathbb{R} \rightarrow [0, \infty)$ satisfies $g(0) = 0$, is continuous and strictly increasing on $[0, \infty)$, and is continuously differentiable on $(0, \infty)$ with strictly positive and non-increasing derivative $g'(x) := \frac{d}{dy} g(y)|_{y=x}$. Moreover, the function $\mathbb{R}^n \rightarrow [0, \infty) : \alpha \mapsto g(\|\alpha\|_2^2)$ is assumed to be strictly convex. The most important examples of link functions are $g(x) = x$ and $g(x) = \sqrt{x}$.

Form of the Penalties We assume that the penalties are composite norms. First, we assume that the tuning parameter $\lambda := (\lambda_1, \dots, \lambda_k)^\top$ is in $(0, \infty)^k$. We then assume that the matrices $M_1, \dots, M_k \in \mathbb{R}^{p \times p}$ satisfy $\bigcap_{j=1}^k \text{Ker}(M_j) = \{\mathbf{0}_{p \times p}\}$, where Ker denotes the kernel of a matrix. This assumption is mild, simply stating that the row space of all matrices M_1, \dots, M_k combined span \mathbb{R}^p , that is, each parameter is covered by some penalization. In the simplest cases, the matrices equal the identity matrix. In general, however, these matrices allow for the incorporation of complex structural assumptions. For example, group structures can be modeled by taking the matrices M_1, \dots, M_k equal to (arbitrarily overlapping) block-diagonal matrices with some of the blocks in each M_j equal to zero. Finally, the single norms $\|\cdot\|_{q_j}$ with $q_j \geq 1$ are the regular ℓ_{q_j} -norms on \mathbb{R}^p . Their dual norms are denoted by $\|\cdot\|_{q_j}^*$, and it holds that $\|\cdot\|_{q_j}^* = \|\cdot\|_{p_j}$ for $p_j \in [1, \infty]$ such that $1/p_j + 1/q_j = 1$. Since each $\|\cdot\|_{q_j}$ is a norm, and since the rows of the matrices M_j span the entire \mathbb{R}^p , the penalty is indeed a norm.

Treatment of Overlap We introduce some further notation to make our result sharp also in cases where variables are subject to more than one penalty term, such as in the overlapping group (square-root) lasso. For this, we first denote by A^+ the Moore-Penrose pseudoinverse of a matrix A . We then note that by the rank assumption on the matrices M_1, \dots, M_k , there are projection matrices $P_1, \dots, P_k \in \mathbb{R}^{p \times p}$ such that

$$\sum_{j=1}^k P_j M_j^+ M_j = I_{p \times p} . \quad (3)$$

The projection matrices enter the “empirical process” terms associated with the tuning parameters and the prediction bounds. Our results hold for any P_1, \dots, P_k that satisfy the above equality; however, appropriate choices are needed to obtain sharp bounds. In generic examples, the choice of P_1, \dots, P_k is straightforward: if $k = 1$ (see, for example, the lasso, square-root lasso, and fused lasso) or if the row spaces of the matrices M_1, \dots, M_k are disjoint (see, for example, the group lasso with non-overlapping groups), one can select $P_1, \dots, P_k = I_{p \times p}$. More generally, if $k > 1$ and some variables are penalized twice (see, for example, the group lasso with overlapping groups), slightly more complicated choices lead to optimal bounds.

Technical Assumption on the Noise Distributions We consider general noise distributions; for example, we allow for heavy-tailed noise and for correlations within ε and between ε and X . However, we exclude non-generic noise distributions for technical ease. More specifically, we assume that $Y \neq \mathbf{0}_n$ and $\min_{j \in \{1, \dots, k\}} \|(X P_j M_j^+)^{\top} \varepsilon\|_{q_j}^* > 0$ with probability one. This implies in particular that $g'(\|Y - X \widehat{\beta}^\lambda\|_2^2) > 0$ with probability one, see Lemma A.3 in the Appendix. To illustrate that the assumptions hold in generic cases, note that the $\|\cdot\|_{q_j}^*$ ’s are norms, so that the second condition simply states that $(X P_j M_j^+)^{\top} \varepsilon \neq \mathbf{0}_p$, $j \in \{1, \dots, k\}$, with probability one. As an example, consider now the non-overlapping group penalty, which corresponds to standard group lasso/square-root lasso. One can then check readily that the assumption is equivalent to $X^{\top} \varepsilon \neq \mathbf{0}_p$ holding with probability one, which is satisfied for any generic continuous distributions of X and ε . It is also straightforward to relax the condition to hold only with probability $1 - \kappa$, $\kappa \equiv \kappa(n) \rightarrow 0$ as $n \rightarrow \infty$, to include discrete noise via standard concentration bounds; we omit the details.

Relations to Existing Literature

Statistical guarantees for high-dimensional prediction with penalized estimators are typically formulated in terms of oracle inequalities. A variety of oracle inequalities is known, we refer to [8, 9, 19] and references therein, and for some cases, also corresponding concentration inequalities are available [13, Theorem 1.1]. However, most of these bounds impose severe constraints on the model, such as eigenvalue-type conditions on the design matrix [11, 40] — see [16, 20] for in-depth comparisons of penalty bounds and sparsity bounds in the lasso case. In strong contrast, we are interested in oracle inequalities that

do not involve additional constraints. Such results are known both for the penalized formulation of the lasso [26, 29, 30] and for the constraint formulation [12]. We are interested in studying whether this type of guarantees can be established more generally in high-dimensional regression.

To achieve this generality, we introduce arguments mainly based on convexity and continuity. Convexity has been used previously to establish ℓ_∞ - and support recovery guarantees for the lasso [45], prediction bounds for the constraint version of the lasso [12] and for the lasso in transductive and semi-supervised learning [1], guarantees for low-rank matrix completion [26], and bounds for sparse recovery via entropy penalization [24]. The arguments in [45] are different from ours in that they have different objectives and lead to stringent assumptions on the design. Intermediate results in [12, Page 6] for the constraint lasso can be related to parts of our approach when our proof is specialized to the penalized lasso (cf. our Pages 27-31); yet, the strategy in [12] based on a projection argument is different from ours, and extensions of that argument to the penalized lasso and, more generally, to our framework with multiple tuning parameters/penalty terms and different link functions do not seem straightforward. Finally, the convexity arguments in [1, Proof of Lemma 1], [26, Equation (2.6)], and [24, Inequalities (3.1) and (3.2)] can be related to some of the techniques on Pages 27-31.

Our continuity arguments evolve around Brouwer's fix-point theorem. As intermediate steps, we show that suitable tuning parameters exist in the first place and that the unbounded set $(0, \infty)^k$ of tuning parameters can be replaced by a bounded set. These facts are known for the lasso and the square-root lasso, see [3] and others, but they are novel and non-trivial in the general case and might thus be of interest by themselves.

Our result specializes correctly and confirms existing expectations. As one example, our bounds for the penalized version of the lasso match the corresponding bounds in the literature [26, 30] and relate to similar expectation-type results [29]. As another example, the bounds for the (group) square-root lasso match the results of the (group) lasso, complementing previous findings that show the correspondence of the two methods in oracle inequalities under additional constraints [3, 10].

We also mention that our bounds are near rate-optimal in the absence of further assumptions. Indeed, it has been shown that the rates for lasso prediction [16, Proposition 4] (even when the noise is Gaussian) can not be improved in general beyond $1/\sqrt{n}$, which corresponds to our bound up to log-factors — see the Examples section. Under RE-type assumptions [41], one can find the rate $s \log p/n$ for the lasso prediction error, where s is the number of non-zero elements in β^* . In the case where s is small, this can be a substantial improvement over the $\sqrt{\log p/n} \|\beta^*\|_1$ -rate. Refined versions, allowing for a potentially large number of small entries in β^* , can be found in [42]. However, RE-type assumptions are very restrictive and often seem unlikely to hold in practice. We come back to this issue in the Discussion section.

We finally relate to oracle inequalities for objectives different from prediction. Besides prediction, standard goals include variable selection and ℓ_1 -, ℓ_2 -, and ℓ_∞ -estimation, we refer again to [9, 19] and references therein. These objectives necessarily involve strict assumptions on the design and are thus not of major relevance here. More closely related to our work is out-of-sample prediction, which — on a high level — can be thought of

as being somewhere between estimation and prediction. Classical results demonstrate that the lasso can achieve out-of-sample prediction without further assumptions on the design [18]. Another similarity to prediction is that consistency guarantees for constraint lasso in out-of-sample prediction can be formulated in terms of ℓ_1 -balls of predictors [18, Theorems 1 and 3]. (The ℓ_1 -balls in [18] are, however, more restrictive than the ℓ_1 -balls needed for prediction). On the other hand, the optimality of the ℓ_1 -related rates suggested in [18] and the applicability of the results to other estimators considered in our paper remain open questions.

2. General Result

We now state the general oracle inequality for the framework described on Pages 2–4. For this, we first have to discuss the existence of suitable tuning parameters. The valid tuning parameters in standard results for the lasso, for example, are of the form $c\|X^\top \varepsilon\|_\infty$, where the factor $c \geq 2$ depends on the specific type of oracle inequality ($c = 2$ for the standard penalty bounds; $c > 2$ for the standard sparsity prediction or estimation bounds, see [4, 15] and others). To show that there is a corresponding range for the tuning parameters in our general bounds, we derive the following result.

Lemma 2.1 (Existence). *Consider fixed constants $c_1, \dots, c_k \in (0, \infty)$. With probability one, there is a tuning parameter $\lambda \equiv \lambda(c_1, \dots, c_k) \in (0, \infty)^k$ such that*

$$\frac{\lambda}{2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2)} = (c_1\|(XP_1M_1^+)^\top \varepsilon\|_{q_1}^*, \dots, c_k\|(XP_kM_k^+)^\top \varepsilon\|_{q_k}^*)^\top.$$

This proof of existence ensures that suitable tuning parameters exist for arbitrary estimators of the form (2). If $g : x \mapsto x$, Lemma 2.1 can be verified easily. In particular, the above equation simplifies to $\lambda = 2c_1\|X^\top \varepsilon\|_\infty$ for the lasso. In general, however, the statement is more intricate, and there might be several tuning parameters that satisfy the equality. Our proof is, therefore, more involved, invoking continuity arguments and Brouwer’s fixed-point theorem, see Appendix.

One can also replace the implicit equalities in Lemma 2.1 by explicit bounds on the tuning parameters. Note first that $g'(\|Y - X\widehat{\beta}^\lambda\|_2^2)$ is not monotone in the tuning parameter in general. However, this problem can be circumvented by deriving concentration bounds. More specifically, for any link function g and even for heavy-tailed noise distributions, one can use empirical process theory (such as [28], for example) to derive concentration bounds for $|g'(\|Y - X\widehat{\beta}^\lambda\|_2^2) - g'(\|\varepsilon\|_2^2)|$ (very rough bounds are sufficient). This then implies bounds for the left-hand side in Lemma 2.1, and consequently, allows one to replace the implicit inequalities by explicit lower bounds on the tuning parameters.

We are now ready to state the main result.

Theorem 2.1 (Penalty bound). *For any choice of $c_1, \dots, c_k \in (0, \infty)$, with probability one, the estimator $\widehat{\beta}^\lambda$ defined by (2) with tuning parameter λ as in Lemma 2.1 above satisfies the prediction bound*

$$\begin{aligned} \frac{1}{n} \|X(\beta^* - \widehat{\beta}^\lambda)\|_2^2 &\leq \inf_{\substack{u \in (0,1) \\ \beta \in \mathbb{R}^p}} \left\{ \frac{1}{4u(1-u)n} \|X(\beta^* - \beta)\|_2^2 \right. \\ &\left. + \frac{1}{n} \sum_{j=1}^k \frac{1+c_j}{1-u} \|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* \|M_j \beta\|_{q_j} - \frac{1}{n} \sum_{j=1}^k \frac{c_j-1}{1-u} \|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* \|M_j \widehat{\beta}^\lambda\|_{q_j} \right\}. \end{aligned}$$

This oracle inequality provides bounds for the prediction errors of the estimators (2). Our proofs are based only on convexity and continuity arguments, making the result very general and sharp in its constants.

Let us provide some interpretations of the result. Note first that the bounds apply to any positive c_1, \dots, c_k , but the most interesting case is $c_1, \dots, c_k \geq 1$, since the terms on the right-hand side that depend on the estimator itself can then be dropped. We first consider a tuning parameter $\bar{\lambda} \in (0, \infty)^k$ that satisfies the equality in Lemma 2.1 with $c_1, \dots, c_k = 1$. We set $\bar{\beta} := \widehat{\beta}^{\bar{\lambda}}$ for ease of notation. Now, choosing $u = 0.5$ in the above functional implies

$$\frac{1}{n} \|X(\beta^* - \bar{\beta})\|_2^2 \leq \min_{\beta \in \mathbb{R}^p} \left\{ \frac{1}{n} \|X(\beta^* - \beta)\|_2^2 + \frac{4}{n} \sum_{j=1}^k \|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* \|M_j \beta\|_{q_j} \right\}. \quad (4)$$

To bring this on a more abstract level, we denote the prediction loss by $L(\beta) := \|X(\beta^* - \beta)\|_2^2/n$ and introduce model classes $\mathcal{F}_v := \{\beta \in \mathbb{R}^p : 4 \sum_{j=1}^k \|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* \|M_j \beta\|_{q_j} / n \leq v\}$ indexed by $v \in [0, \infty)$. Inequality (4) then reads

$$L(\bar{\beta}) \leq \min_{v \in [0, \infty)} \min_{\beta \in \mathcal{F}_v} \{L(\beta) + v\}.$$

Thus, the estimator $\bar{\beta}$ performs as well as the minimizer of the loss over the class \mathcal{F}_v — up to a complexity penalty of \mathcal{F}_v . For another abstract view on the theorem, we define a loss for any $\mathbf{a} \in (0, \infty)^k$ by

$$L_{\mathbf{a}}(\beta) := \frac{1}{n} \|X(\beta^* - \beta)\|_2^2 + \frac{1}{n} \sum_{j=1}^k a_j \|M_j \beta\|_{q_j}.$$

This loss balances prediction accuracy against model complexity. Now, we observe that if $2(c_j - 1)\|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* = a_j$, choosing again $u = 0.5$ in the initial functional yields

$$L_{\mathbf{a}}(\widehat{\beta}^\lambda) \leq \left(1 + \max_{j \in \{1, \dots, k\}} \frac{4\|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^*}{a_j}\right) \min_{\beta \in \mathbb{R}^p} L_{\mathbf{a}}(\beta).$$

This means that the estimator $\widehat{\beta}^\lambda$ performs as well — again up to constants — as the minimizer of the loss $L_{\mathbf{a}}$. These forms of our bound fit the classical notions of oracle inequalities (with sharp leading constant) in empirical risk minimization [25, Chapter 1.1] and non-parametric estimation [39, Chapter 1.8]. Finally, we consider again $\bar{\lambda}$, then set $\beta = \beta^*$ and take the limit $u \rightarrow 0$ in Theorem 2.1. We find

$$\frac{1}{n} \|X(\beta^* - \bar{\beta})\|_2^2 \leq \frac{2}{n} \sum_{j=1}^k \|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* \|M_j \beta^*\|_{q_j}. \quad (5)$$

This form of our bound commensurates with typical formulations of oracle inequalities in high dimensions [9, Chapters 2.4.2 and 6.2]; in particular, the bound implies known prediction bounds with correct constants — see the following section.

The tuning parameter $\bar{\lambda}$ minimizes $\sum_{j=1}^k (1 + c_j) \|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* \|M_j \beta\|_{q_j} / (n(1 - u))$, the first term on the right-hand side of the bound, as a function of c_1, \dots, c_k under the constraint that $c_1, \dots, c_k \geq 1$ (which implies that the terms on the right-hand side that depend on the estimator can be dropped). This choice of the tuning parameter also leads to rates that have been shown to be near-optimal in certain cases [16, Proposition 4]. However, this does not necessarily mean that $\bar{\lambda}$ minimizes the prediction loss $\|X(\beta^* - \widehat{\beta}^\lambda)\|_2^2$. Explicit formulations of tuning parameters that minimize the prediction loss are to date unknown. Nevertheless, some insights have been developed: for example, [16, 20] discuss lasso tuning parameters as a function of the correlations in X , and [2] and [35, Section 4] discuss lasso/scaled lasso tuning parameters that can lead to minimax rates in the case of sparsity and small correlations. Furthermore, our essay does not provide any guidance on how to select tuning parameters in practice; indeed, the tuning parameters in Lemma 2.1 depend on the noise ε , which is unknown in practice. For ideas on the practical selection of the lasso tuning parameter with finite sample guarantees, we refer to [14, 15]. For ideas on how to make the selection of tuning parameters independent of unknown model aspects, we refer to [17, 27] and the square-root/scaled lasso example in the following section.

We conclude this section highlighting five other properties of Theorem 2.1 (much of this becomes more lucid in the context of the specific examples discussed in the next section): First, the bound involves the values of the tuning parameters to the power one and holds for any design matrix X . The bound contains the penalty values of the regression vectors. Hence, the bounds are penalty bounds. Second, the oracle inequality holds for any distribution of the noise ε . Third, the bounds hold for any sample size n ; in particular, the bounds are non-asymptotic. Fourth, the bounds become smaller if the correlations in X become larger, cf. [20]. Fifth, the link function g appears in the existence result on tuning parameters but not in the prediction bound. This last, interesting point clarifies the role of the link function: its essential purpose is to “reshuffle” the tuning parameter path. One can relate this observation to the discussion of the lasso/square-root lasso below.

3. Examples

We now apply our general results to some specific estimators of the form (2).

Lasso The lasso [36] is defined as

$$\widehat{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \{\|Y - X\beta\|_2^2 + \lambda\|\beta\|_1\}.$$

We first show that we can recover the standard penalty bounds that have been derived in the lasso literature, see for example [26, Eq. (2.3) in Theorem 1] and [20, Equation (3)]. The proofs use that lasso objective function is minimal at $\widehat{\beta}^\lambda$, that is, for any $\beta \in \mathbb{R}^p$, it holds that

$$\|Y - X\widehat{\beta}^\lambda\|_2^2 + \lambda\|\widehat{\beta}^\lambda\|_1 \leq \|Y - X\beta\|_2^2 + \lambda\|\beta\|_1.$$

This is equivalent to

$$\|Y - X\beta^* + X\beta^* - X\widehat{\beta}^\lambda\|_2^2 + \lambda\|\widehat{\beta}^\lambda\|_1 \leq \|Y - X\beta^* + X\beta^* - X\beta\|_2^2 + \lambda\|\beta\|_1$$

and

$$\begin{aligned} & \|Y - X\beta^*\|_2^2 + 2\langle Y - X\beta^*, X\beta^* - X\widehat{\beta}^\lambda \rangle + \|X\beta^* - X\widehat{\beta}^\lambda\|_2^2 + \lambda\|\widehat{\beta}^\lambda\|_1 \\ & \leq \|Y - X\beta^*\|_2^2 + 2\langle Y - X\beta^*, X\beta^* - X\beta \rangle + \|X\beta^* - X\beta\|_2^2 + \lambda\|\beta\|_1. \end{aligned}$$

Invoking the model (1) and consolidating, this yields

$$\|X\beta^* - X\widehat{\beta}^\lambda\|_2^2 \leq \|X\beta^* - X\beta\|_2^2 + 2\langle \varepsilon, X\widehat{\beta}^\lambda - X\beta \rangle + \lambda\|\beta\|_1 - \lambda\|\widehat{\beta}^\lambda\|_1.$$

Hölder's inequality and the triangle inequality then lead to

$$\|X\beta^* - X\widehat{\beta}^\lambda\|_2^2 \leq \|X\beta^* - X\beta\|_2^2 + 2\|X^\top \varepsilon\|_\infty (\|\widehat{\beta}^\lambda\|_1 + \|\beta\|_1) + \lambda\|\beta\|_1 - \lambda\|\widehat{\beta}^\lambda\|_1.$$

Hence, for $\lambda = \bar{\lambda} = 2\|X^\top \varepsilon\|_\infty$, we eventually find

$$\frac{1}{n}\|X(\beta^* - \bar{\beta})\|_2^2 \leq \min_{\beta \in \mathbb{R}^p} \left\{ \frac{1}{n}\|X(\beta^* - \beta)\|_2^2 + \frac{4}{n}\|X^\top \varepsilon\|_\infty \|\beta\|_1 \right\}.$$

Observing that $k = 1$ and $M_1 = P_1 = I_{p \times p}$ in this example, one can check that (4) recovers this bound.

In the case $\beta = \beta^*$, Inequality (5) does slightly better. Indeed, our results imply

$$\frac{1}{n}\|X(\beta^* - \bar{\beta})\|_2^2 \leq \frac{2}{n}\|X^\top \varepsilon\|_\infty \|\beta^*\|_1,$$

which is smaller by a factor 2 than the above right-hand side at $\beta = \beta^*$. The same bound also follows from [34, Inequality (23)]. The key property that allows one to derive bounds with the improved factor is convexity; therefore, in the lasso case, one can also

find the above bound with the techniques in the papers mentioned in the corresponding discussion on Page 5, such as [1, 24, 26].

To provide a sense for the rates, we mention that if $\varepsilon_1, \dots, \varepsilon_n \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2)$ and $(X^\top X)_{jj} = n$ for all $j \in \{1, \dots, p\}$, it holds that $\bar{\lambda} \approx \sigma \sqrt{n \log(p)}$ and $\|X(\beta^* - \bar{\beta})\|_2^2/n \lesssim \sigma \sqrt{\log(p)/n} \|\beta^*\|_1$.¹ Importantly, the rate is lower bounded by $1/\sqrt{n}$ — unless further assumptions on the design matrix X are imposed [16, Proposition 4].

We finally note that one can also include “tailored” tuning parameters, which corresponds to considering the lasso as a specification of our general framework with $k = p$ and not necessarily equal tuning parameters $\lambda_1, \dots, \lambda_p$. One can check easily that in particular if $(X^\top \varepsilon)_1, \dots, (X^\top \varepsilon)_p$ i.i.d., one obtains the same bounds as above.

Square-root/scaled lasso The square-root lasso [3] reads

$$\hat{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \{ \|Y - X\beta\|_2 + \lambda \|\beta\|_1 \}.$$

In our framework, $k = 1$ and $M_1 = P_1 = I_{p \times p}$, so that $\bar{\lambda} = \|X^\top \varepsilon\|_\infty / \|Y - X\hat{\beta}^\lambda\|_2$ and the prediction bound (5) reads

$$\frac{1}{n} \|X(\beta^* - \bar{\beta})\|_2^2 \leq \frac{2}{n} \|X^\top \varepsilon\|_\infty \|\beta^*\|_1.$$

A similar bound is implied by [34, Theorem 1], and Inequality (4) and the general bound in our main theorem translate accordingly (for ease of comparison, we focus on (5) in the following). The bounds match the corresponding ones for the lasso, but the tuning parameters differ. The crux of the square-root lasso, and similarly, the scaled lasso [34], is that their tuning parameters can be essentially independent of the noise variance. For example, if $\varepsilon_1, \dots, \varepsilon_n \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2)$ and $(X^\top X)_{jj} = n$ for all $j \in \{1, \dots, p\}$, it holds² that $\bar{\lambda} \approx \|X^\top \varepsilon\|_\infty / \|\varepsilon\|_2 \approx \sqrt{\log(p)}$, which is independent of σ . Since σ is typically unknown in practice, the square-root/scaled lasso can thus facilitate the tuning of λ .

Slope estimator The slope estimator [6] can be written as

$$\hat{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ \|Y - X\beta\|_2^2 + \lambda \sum_{j=1}^p \omega_j |\beta|_{(j)} \right\},$$

where $|\beta|_{(j)}$ denotes the j th largest entry of β in absolute value, $\omega_1 \geq \dots \geq \omega_p > 0$ is a non-increasing sequence of weights, and $\lambda > 0$ is a tuning parameter. A promising case for the slope estimator is the one where $\varepsilon_1, \dots, \varepsilon_n \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2)$ and $(X^\top X)_{jj} = n$ for all $j \in \{1, \dots, p\}$. The weights can then be chosen as $\omega_j := 2\sigma \sqrt{n \log(2p/j)}$ in the spirit of the Benjamini-Hochberg procedure [6, 33], and a theoretically justified

¹The wiggles indicate that we are interested only in the rough shapes and neglect constants, for example.

²see Page 16 for some hints

choice of the tuning parameter is $\lambda > 4 + \sqrt{2}$ [2, Equation (2.5)]. In particular, this choice works in the sense of the first part of Theorem 2.1, and one finds the bound $\|X(\beta^* - \bar{\beta})\|_2^2/n \lesssim \lambda\sigma\sqrt{\log(p)}/n \|\beta^*\|_1$, for example, which coincides with the bounds above. Similar considerations apply to the oscar penalty [7].

Elastic net The elastic net [47] reads

$$\hat{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ \|Y - X\beta\|_2^2 + \lambda_1 \|\beta\|_1 + \lambda_2 \|\beta\|_2^2 \right\}.$$

This is not directly in the form (2). However, one can use the usual trick writing the estimator as a lasso with augmented data, cf. [21, Lemma 1]. Using $M_1 = P_1 = I_{p \times p}$, our results then hold for any tuning parameters that satisfy $\lambda_1 = 2\|X^\top \varepsilon - \lambda_2 \beta^*\|_\infty$. For example, we can set $\bar{\lambda}_2 = \arg \min_{\lambda_2} \|X^\top \varepsilon - \lambda_2 \beta^*\|_\infty$ and $\bar{\lambda}_1 = 2\|X^\top \varepsilon - \bar{\lambda}_2 \beta^*\|_\infty$. The bound in (5), for example, then reads

$$\frac{1}{n} \|X(\beta^* - \bar{\beta})\|_2^2 \leq \frac{2}{n} \|X^\top \varepsilon - \bar{\lambda}_2 \beta^*\|_\infty \|\beta^*\|_1 \leq \frac{2}{n} \|X^\top \varepsilon\|_\infty \|\beta^*\|_1.$$

Similar results hold if, for example, the settings with the normal noise vectors described in the two examples above apply and $\bar{\lambda}_2 = \mathcal{O}(\sqrt{n})$. The main intent of the elastic net is to improve variable selection. However, our results show that the elastic net with well-chosen tuning parameters also has similar penalty guarantees for prediction as the lasso.

Lasso and square-root lasso with group structures The estimators considered so far are based on a simple notion of sparsity. In practice, however, it can be reasonable to assume more complex sparsity structures in the regression vector β^* . Estimators that take such structures into account include the group lasso [46], group square-root lasso [10], hierarchical group lasso [5], lasso with overlapping groups [22], and sparse group lasso [32]. They all fit our framework.

In contrast to the examples above, the matrices $M_1, \dots, M_k, P_1, \dots, P_k$ play a non-trivial role in these examples. On a high level, the matrix M_j specifies which variables are incorporated in the j th group. If groups overlap, the matrix P_j specifies which parts of $X^\top \varepsilon$ are attributed to the tuning parameter $\bar{\lambda}_j$. For example, if the m th variable is in the j th and l th group, the matrices P_j and P_l can be chosen such that the corresponding element $(X^\top \varepsilon)_m$ is part of either $\bar{\lambda}_j$ or $\bar{\lambda}_l$ and not in both of them.

As an illustration, let us consider the group lasso with non-overlapping groups:

$$\hat{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ \|Y - X\beta\|_2^2 + \lambda \sum_{j=1}^k \|\beta_{G_j}\|_2 \right\}.$$

Here, G_1, \dots, G_k is a partition of $\{1, \dots, p\}$ and $(\beta_{G_j})_i := \beta_i \mathbb{1}\{i \in G_j\}$. To put the estimator in framework (2), one can either generalize the estimator to incorporate possibly different tuning parameters for each group, set M_j to $(M_j)_{st} = \mathbb{1}\{s = t, s \in G_j\}$, and then choose a dominating tuning parameter, or one can directly extend our

results (as mentioned earlier) to arbitrary norm penalties. In any case, the projection matrices/matrix equal the identity matrix, our standard tuning parameter is $\bar{\lambda} = 2 \max_{l \in \{1, \dots, k\}} \|(X^\top \varepsilon)_{G_l}\|_2$, and the bound (5) gives

$$\frac{1}{n} \|X(\beta^* - \bar{\beta})\|_2^2 \leq \frac{2}{n} \max_{l \in \{1, \dots, k\}} \|(X^\top \varepsilon)_{G_l}\|_2 \sum_{j=1}^k \|\beta_{G_j}^*\|_2.$$

Trend filtering and total variation/fused penalty Trend filtering [23, 37] reads

$$\hat{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \{ \|Y - \beta\|_2^2 + \lambda \|M\beta\|_1 \},$$

where for given $l \in \{1, 2, \dots\}$, the matrix $M \in \mathbb{R}^{p \times p}$ is defined as $M := \underbrace{D \times \dots \times D}_{l \text{ times}}$

with $D \in \mathbb{R}^{p \times p}$ given by

$$D_{ij} := \begin{cases} -1 & \text{if } i < p \text{ and } i = j \\ 1 & \text{if } i < p \text{ and } i = j - 1 \\ 0 & \text{otherwise} \end{cases}.$$

We find $\bar{\lambda} = 2 \|M^{+\top} \varepsilon\|_\infty$, and (5) implies³

$$\frac{1}{n} \|\beta^* - \bar{\beta}\|_2^2 \leq \frac{2}{n} \|M^{+\top} \varepsilon\|_\infty \|M\beta^*\|_1.$$

In the case $l = 1$, the estimator becomes

$$\hat{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \{ \|Y - \beta\|_2^2 + \lambda \sum_{j=2}^p |\beta_j - \beta_{j-1}| \},$$

which corresponds to the total variation [31] and fused lasso penalizations [38]. Moreover, one can check that the Moore-Penrose inverse of $M = D$ is then given by D^+ with entries

$$D_{ij}^+ := \begin{cases} (j-p)/p & \text{if } i \leq j < p \\ j/p & \text{if } i > j, j < p \\ 0 & \text{if } j = p \end{cases}.$$

We find $\bar{\lambda} = 2 \|D^{+\top} \varepsilon\|_\infty$ and the corresponding bound

$$\frac{1}{n} \|\beta^* - \bar{\beta}\|_2^2 \leq \frac{2}{n} \|D^{+\top} \varepsilon\|_\infty \|D\beta^*\|_1.$$

³Unlike assumed earlier, $\text{Ker}(M) \neq \{\mathbf{0}_{p \times p}\}$ in this example, but one can replace matrix M by the invertible matrix $M + \epsilon \mathbf{I}_{p \times p}$ and then take the limit $\epsilon \rightarrow 0$.

4. Discussion

Sparsity bounds have been derived for many high-dimensional estimators. In this paper, we complement these bounds with corresponding penalty bounds. Which type of bound is sharper depends on the underlying model. As a general rule, penalty bounds improve with increasing correlations in the design matrix, while sparsity bounds deteriorate with increasing correlations and are eventually infinite once the design matrix is too far from an orthogonal matrix [16].

Without making assumptions on the design, and for a wide range of penalized estimators, our results imply non-trivial rates of convergence for prediction. This is of direct practical relevance, since the assumptions inflicted by sparsity bounds are often unrealistic in applications and, in any case, depend on inaccessible model parameters and thus cannot be verified in practice. For example, sparsity bounds for the lasso have been derived under a variety of assumptions on X , including RIP, restricted eigenvalue condition, and compatibility condition, see [41] for an overview of these concepts. Results from random matrix theory show that these assumptions are fulfilled with high probability if the data generating process is “nice” (sub-Gaussian, isotropic, ...) and the sample size n is large enough, see [43] for a recent result. Unfortunately, in practice, the data generating processes are not necessarily nice, and sample sizes can be small — not only in comparison with the number of parameters, but also in absolute terms. Moreover, even if the conditions of fast rates bounds are satisfied, these bounds can contain very large factors and are then only interesting from an asymptotic point of view.

Acknowledgements

We thank Jacob Bien and Mohamed Hebiri for their insightful comments and the Editors and Reviewers for their helpful suggestions.

References

- [1] BELLEC, P., DALALYAN, A., GRAPPIN, E. and PARIS, Q. (2016). On the prediction loss of the lasso in the partially labeled setting. *arXiv:1606.06179*.
- [2] BELLEC, P., LECUÉ, G. and TSYBAKOV, A. (2016). Slope meets Lasso: improved oracle bounds and optimality. *arXiv:1605.08651*.
- [3] BELLONI, A., CHERNOZHUKOV, V. and WANG, L. (2011). Square-root lasso: pivotal recovery of sparse signals via conic programming. *Biometrika* **98** 791–806.
- [4] BICKEL, P., RITOV, Y. and TSYBAKOV, A. (2009). Simultaneous analysis of lasso and Dantzig selector. *Ann. Statist.* **37** 1705–1732.
- [5] BIEN, J., TAYLOR, J. and TIBSHIRANI, R. (2013). A lasso for hierarchical interactions. *Ann. Statist.* **41** 1111.
- [6] BOGDAN, M., VAN DEN BERG, E., SABATTI, C., SU, W. and CANDÈS, E. (2015). SLOPE - adaptive variable selection via convex optimization. *Ann. Appl. Stat.* **9** 1103.

- [7] BONDELL, H. and REICH, B. (2008). Simultaneous regression shrinkage, variable selection, and supervised clustering of predictors with OSCAR. *Biometrics* **64** 115–123.
- [8] BÜHLMANN, P. (2013). Statistical significance in high-dimensional linear models. *Bernoulli* **19** 1212–1242.
- [9] BÜHLMANN, P. and VAN DE GEER, S. (2011). *Statistics for high-dimensional data: Methods, theory and applications*. Springer.
- [10] BUNEA, F., LEDERER, J. and SHE, Y. (2014). The Group Square-Root Lasso: Theoretical Properties and Fast Algorithms. *IEEE Trans. Inform. Theory* **60** 1313–1325.
- [11] BUNEA, F., TSYBAKOV, A. and WEGKAMP, M. (2007). Sparsity oracle inequalities for the Lasso. *Electron. J. Stat.* **1** 169–194.
- [12] CHATTERJEE, S. (2013). Assumptionless consistency of the lasso. *arXiv:1303.5817*.
- [13] CHATTERJEE, S. (2014). A new perspective on least squares under convex constraint. *Ann. Statist.* **42** 2340–2381.
- [14] CHÉTELAT, D., LEDERER, J. and SALMON, J. (2017). Optimal two-step prediction in regression. *Electron. J. Statist.* **11** 2519–2546.
- [15] CHICHIGNOUD, M., LEDERER, J. and WAINWRIGHT, M. (2016). A Practical Scheme and Fast Algorithm to Tune the Lasso With Optimality Guarantees. *J. Mach. Learn. Res.* **17** 1–20.
- [16] DALALYAN, A., HEBIRI, M. and LEDERER, J. (2014). On the Prediction Performance of the Lasso. *To appear in Bernoulli*.
- [17] GIRAUD, C., HUET, S. and VERZELEN, N. (2012). High-dimensional regression with unknown variance. *Statist. Sci.* **27** 500–518.
- [18] GREENSHTEIN, E. and RITOV, Y. (2004). Persistence in high-dimensional linear predictor selection and the virtue of overparametrization. *Bernoulli* **10** 971–988.
- [19] HASTIE, T., TIBSHIRANI, R. and WAINWRIGHT, M. (2015). *Statistical Learning with Sparsity: The Lasso and Generalizations*. Chapman and Hall.
- [20] HEBIRI, M. and LEDERER, J. (2013). How Correlations Influence Lasso Prediction. *IEEE Trans. Inform. Theory* **59** 1846–1854.
- [21] HEBIRI, M. and VAN DE GEER, S. (2011). The Smooth-Lasso and other $\ell_1 + \ell_2$ -penalized methods. *Electron. J. Stat.* **5** 1184–1226.
- [22] JACOB, L., OBOZINSKI, G. and VERT, J. P. (2009). Group lasso with overlap and graph lasso. In *ICML 2009* 433–440.
- [23] KIM, S. J., KOH, K., BOYD, S. and GORINEVSKY, D. (2009). ℓ_1 Trend Filtering. *SIAM Rev.* **51** 339–360.
- [24] KOLTCHINSKII, V. (2009). Sparse recovery in convex hulls via entropy penalization. *The Annals of Statistics* **37** 1332–1359.
- [25] KOLTCHINSKII, V. (2011). *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems*. Springer.
- [26] KOLTCHINSKII, V., LOUNICI, K. and TSYBAKOV, A. (2011). Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Ann. Statist.* **39** 2302–2329.
- [27] LEDERER, J. and MÜLLER, C. (2015). Don’t fall for tuning parameters: Tuning-free

- variable selection in high dimensions with the TREX. *In AAAI-15* 2729–2735.
- [28] LEDERER, J. and VAN DE GEER, S. (2014). New concentration inequalities for empirical processes. *Bernoulli* **20** 2020–2038.
- [29] MASSART, P. and MEYNET, C. (2011). The Lasso as an ℓ_1 -ball model selection procedure. *Electron. J. Stat.* **5** 669–687.
- [30] RIGOLLET, P. and TSYBAKOV, A. (2011). Exponential Screening and optimal rates of sparse estimation. *Ann. Statist.* **39** 731–771.
- [31] RUDIN, L., OSHER, S. and FATEMI, E. (1992). Nonlinear total variation based noise removal algorithms. *Phys. D* **60** 259–268.
- [32] SIMON, N., FRIEDMAN, J., HASTIE, T. and TIBSHIRANI, R. (2013). A sparse-group lasso. *J. Comput. Graph. Statist.* **22** 231–245.
- [33] SU, W. and CANDÈS, E. (2016). SLOPE is adaptive to unknown sparsity and asymptotically minimax. *Ann. Statist.* **44** 1038–1068.
- [34] SUN, T. and ZHANG, C. (2012). Scaled sparse linear regression. *Biometrika* **99** 879–898.
- [35] SUN, T. and ZHANG, C. H. (2013). Sparse matrix inversion with scaled lasso. *J. Mach. Learn. Res.* **14** 3385–3418.
- [36] TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **58** 267–288.
- [37] TIBSHIRANI, R. (2014). Adaptive piecewise polynomial estimation via trend filtering. *Ann. Statist.* **42** 285–323.
- [38] TIBSHIRANI, R., SAUNDERS, M., ROSSET, S., ZHU, J. and KNIGHT, K. (2005). Sparsity and smoothness via the fused lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **67** 91–108.
- [39] TSYBAKOV, A. (2009). *Introduction to nonparametric estimation*. Springer.
- [40] VAN DE GEER, S. (2007). The Deterministic Lasso. *2007 Proc. Amer. Math. Soc. [CD-ROM]*, see also www.stat.math.ethz.ch/~geer/lasso.pdf.
- [41] VAN DE GEER, S. and BÜHLMANN, P. (2009). On the conditions used to prove oracle results for the Lasso. *Electron. J. Stat.* **3** 1360–1392.
- [42] VAN DE GEER, S. and LEDERER, J. (2013). The Lasso, correlated design, and improved oracle inequalities. *IMS Collections* **9** 303–316.
- [43] VAN DE GEER, S. and MURO, A. (2014). On higher order isotropy conditions and lower bounds for sparse quadratic forms. *Electron. J. Stat.* **8** 3031–3061.
- [44] VAN DER VAART, A. and WELLNER, J. (1996). *Weak Convergence and Empirical Processes*. Springer.
- [45] WAINWRIGHT, M. (2009). Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (Lasso). *IEEE Trans. Inform. Theory* **55** 2183–2202.
- [46] YUAN, M. and LIN, Y. (2006). Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **68** 49–67.
- [47] ZOU, H. and HASTIE, T. (2005). Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **67** 301–320.

Appendix A: Proofs

We start with four auxiliary results, Lemmas A.1-A.3. We then prove Lemma 2.1 and Theorem 2.1. Figure 1 depicts the dependence structure of the results.




Figure 1. Dependencies between the results. For example, the arrow between Lemma A.1 and Lemma 2.1 depicts that the proof of Lemma 2.1 makes use of Lemma A.1.

Hints for Page 10: Note that (see [44, Section 2.2] for maximal inequalities that can be used for the last line)

$$\begin{aligned}
 \|Y - X\bar{\beta}\|_2 &= \|Y - X\beta^* + X\beta^* - X\bar{\beta}\|_2 \\
 &\leq \|Y - X\beta^*\|_2 + \|X\beta^* - X\bar{\beta}\|_2 \\
 &\leq \|\varepsilon\|_2 + \sqrt{\frac{2\|X^\top \varepsilon\|_\infty \|\beta^*\|_1}{n}} \sqrt{n} \\
 &\lesssim \sqrt{n} + \sqrt{\frac{\log(p) \|\beta^*\|_1}{\sqrt{n}}} \sqrt{n}
 \end{aligned}$$

and similarly

$$\|Y - X\bar{\beta}\|_2 \gtrsim \sqrt{n} - \sqrt{\frac{\log(p) \|\beta^*\|_1}{\sqrt{n}}} \sqrt{n}.$$

Thus, as long as $\|\beta^*\|_1 = o(\sqrt{n}/\log(p))$, it holds that $\bar{\lambda} \approx \|X^\top \varepsilon\|_\infty / \|\varepsilon\|_2 \approx \sqrt{\log(p)}$.

A.1. Auxiliary Lemmas

Lemma A.1. For any $\hat{\beta}^\lambda, \tilde{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \{g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j}\}$ and $\alpha \in [0, 1]$, it holds that $X\hat{\beta}^\lambda = X\tilde{\beta}^\lambda$ and

$$\alpha \hat{\beta}^\lambda + (1 - \alpha) \tilde{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\}.$$

Lemma A.2. *Let \mathbb{R}^k be equipped with the Euclidean norm, and \mathbb{R} be equipped with the absolute value norm. Then, the function*

$$(0, \infty)^k \rightarrow \mathbb{R} \\ \lambda \mapsto g(\|Y - X\widehat{\beta}^\lambda\|_2^2)$$

and the function

$$(0, \infty)^k \rightarrow \mathbb{R} \\ \lambda \mapsto \|Y - X\widehat{\beta}^\lambda\|_2^2$$

are both continuous.

Lemma A.3. *With probability one, it holds that $\|Y - X\widehat{\beta}^\lambda\|_2^2 > 0$ and $g'(\|Y - X\widehat{\beta}^\lambda\|_2^2) > 0$ for any tuning parameter $\lambda \in (0, \infty)^k$.*

A.2. Proofs of the Auxilliary Lemmas

Proof of Lemma A.1. The case $\alpha \in \{0, 1\}$ is straightforward, so that we consider a given $\alpha \in (0, 1)$. We first show that $X\widehat{\beta}^\alpha = X\widetilde{\beta}^\alpha$. Since the function $\alpha \mapsto g(\|\alpha\|_2^2)$ is strictly convex by assumption, it follows for any vectors $a, \tilde{a} \in \mathbb{R}^n$ that

$$g(\|\alpha a + (1 - \alpha)\tilde{a}\|_2^2) \leq \alpha g(\|a\|_2^2) + (1 - \alpha)g(\|\tilde{a}\|_2^2)$$

with strict inequality if $a \neq \tilde{a}$. Using this with $a = Y - X\widehat{\beta}^\alpha$ and $\tilde{a} = Y - X\widetilde{\beta}^\alpha$, and invoking the convexity of the norms $\|\cdot\|_{q_j}$, we find

$$\begin{aligned} & g(\|Y - X(\alpha\widehat{\beta}^\alpha + (1 - \alpha)\widetilde{\beta}^\alpha)\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j(\alpha\widehat{\beta}^\alpha + (1 - \alpha)\widetilde{\beta}^\alpha)\|_{q_j} \\ &= g(\|\alpha(Y - X\widehat{\beta}^\alpha) + (1 - \alpha)(Y - X\widetilde{\beta}^\alpha)\|_2^2) + \sum_{j=1}^k \lambda_j \|\alpha M_j\widehat{\beta}^\alpha + (1 - \alpha)M_j\widetilde{\beta}^\alpha\|_{q_j} \\ &\leq \alpha g(\|Y - X\widehat{\beta}^\alpha\|_2^2) + (1 - \alpha)g(\|Y - X\widetilde{\beta}^\alpha\|_2^2) + \alpha \sum_{j=1}^k \lambda_j \|M_j\widehat{\beta}^\alpha\|_{q_j} + (1 - \alpha) \sum_{j=1}^k \lambda_j \|M_j\widetilde{\beta}^\alpha\|_{q_j} \\ &= \alpha \left(g(\|Y - X\widehat{\beta}^\alpha\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j\widehat{\beta}^\alpha\|_{q_j} \right) + (1 - \alpha) \left(g(\|Y - X\widetilde{\beta}^\alpha\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j\widetilde{\beta}^\alpha\|_{q_j} \right) \end{aligned}$$

with strict inequality if $X\widehat{\beta}^\alpha \neq X\widetilde{\beta}^\alpha$. Moreover, we note that

$$\widehat{\beta}^\alpha, \widetilde{\beta}^\alpha \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j\beta\|_{q_j} \right\}$$

implies

$$g(\|Y - X\tilde{\beta}^\lambda\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \tilde{\beta}^\lambda\|_{q_j} = g(\|Y - X\hat{\beta}^\lambda\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \hat{\beta}^\lambda\|_{q_j}.$$

Combining the results yields

$$\begin{aligned} & g(\|Y - X(\alpha\hat{\beta}^\lambda + (1-\alpha)\tilde{\beta}^\lambda)\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j(\alpha\hat{\beta}^\lambda + (1-\alpha)\tilde{\beta}^\lambda)\|_{q_j} \\ & \leq g(\|Y - X\hat{\beta}^\lambda\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \hat{\beta}^\lambda\|_{q_j} \end{aligned}$$

with strict inequality if $X\hat{\beta}^\lambda \neq X\tilde{\beta}^\lambda$. Using again that

$$\hat{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\},$$

we find that the above inequality is actually an equality, so that

$$X\hat{\beta}^\lambda = X\tilde{\beta}^\lambda.$$

and

$$\alpha\hat{\beta}^\lambda + (1-\alpha)\tilde{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\}$$

as desired. □

Proof of Lemma A.2. To show that the function $\lambda \mapsto g(\|Y - X\hat{\beta}^\lambda\|_2^2)$ is continuous, we first show that the function

$$\begin{aligned} & \mathbb{R}^k \rightarrow \mathbb{R} \\ & \lambda \mapsto g(\|Y - X\hat{\beta}^\lambda\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \hat{\beta}^\lambda\|_{q_j} \end{aligned}$$

is continuous.

For this, we show the continuity at any fixed $\lambda \in (0, \infty)^k$. Given an $\epsilon > 0$, define $m := 2 \max_{1 \leq j \leq k} \left\{ \frac{g(\|Y\|_2^2)}{\lambda_j} \right\}$ and $\delta := \min_{1 \leq j \leq k} \frac{\lambda_j}{2} \wedge \frac{\epsilon}{\sqrt{km}}$, where \wedge denotes the minimum. Consider now an arbitrary $\lambda' \in (0, \infty)^k$ with $\|\lambda - \lambda'\|_2 < \delta$.

As a next step, note that Criterion (2) implies

$$g(\|Y - X\hat{\beta}^\lambda\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \hat{\beta}^\lambda\|_{q_j} \leq g(\|Y\|_2^2).$$

In particular, since the function g is non-negative on $[0, \infty)$ by assumption, it holds that

$$\sum_{j=1}^k \lambda_j \|M_j \widehat{\beta}^\lambda\|_{q_j} \leq g(\|Y\|_2^2).$$

Thus, we have

$$\lambda_j \|M_j \widehat{\beta}^\lambda\|_{q_j} \leq g(\|Y\|_2^2) \quad \text{for } j \in \{1, 2, \dots, k\}.$$

Hence, if $\lambda_j > 0$, it holds that

$$\|M_j \widehat{\beta}^\lambda\|_{q_j} \leq \frac{g(\|Y\|_2^2)}{\lambda_j} \quad \text{for } j \in \{1, 2, \dots, k\}. \quad (6)$$

Also note that by Criterion (2), it holds that

$$g(\|Y - X \widehat{\beta}^\lambda\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \widehat{\beta}^\lambda\|_{q_j} \leq g(\|Y - X \widehat{\beta}^{\lambda'}\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \widehat{\beta}^{\lambda'}\|_{q_j}$$

and

$$g(\|Y - X \widehat{\beta}^{\lambda'}\|_2^2) + \sum_{j=1}^k \lambda'_j \|M_j \widehat{\beta}^{\lambda'}\|_{q_j} \leq g(\|Y - X \widehat{\beta}^\lambda\|_2^2) + \sum_{j=1}^k \lambda'_j \|M_j \widehat{\beta}^\lambda\|_{q_j}.$$

Rearranging these two inequalities, we obtain

$$g(\|Y - X \widehat{\beta}^\lambda\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \widehat{\beta}^\lambda\|_{q_j} - g(\|Y - X \widehat{\beta}^{\lambda'}\|_2^2) - \sum_{j=1}^k \lambda'_j \|M_j \widehat{\beta}^{\lambda'}\|_{q_j} \leq \sum_{j=1}^k (\lambda_j - \lambda'_j) \|M_j \widehat{\beta}^{\lambda'}\|_{q_j}$$

and

$$g(\|Y - X \widehat{\beta}^\lambda\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \widehat{\beta}^\lambda\|_{q_j} - g(\|Y - X \widehat{\beta}^{\lambda'}\|_2^2) - \sum_{j=1}^k \lambda'_j \|M_j \widehat{\beta}^{\lambda'}\|_{q_j} \geq \sum_{j=1}^k (\lambda_j - \lambda'_j) \|M_j \widehat{\beta}^\lambda\|_{q_j}.$$

By Hölder's inequality and Inequality (6), it holds that

$$\begin{aligned} \sum_{j=1}^k (\lambda_j - \lambda'_j) \|M_j \widehat{\beta}^\lambda\|_{q_j} &\geq -\|\lambda - \lambda'\|_1 \max_{1 \leq j \leq k} \{ \|M_j \widehat{\beta}^\lambda\|_{q_j} \} \\ &\geq -\|\lambda - \lambda'\|_1 \max_{1 \leq j \leq k} \left\{ \frac{g(\|Y\|_2^2)}{\lambda_j} \right\} \end{aligned}$$

and similarly

$$\sum_{j=1}^k (\lambda_j - \lambda'_j) \|M_j \widehat{\beta}^{\lambda'}\|_{q_j} \leq \|\lambda - \lambda'\|_1 \max_{1 \leq j \leq k} \left\{ \frac{g(\|Y\|_2^2)}{\lambda'_j} \right\}.$$

Now trivially, $\max_{1 \leq j \leq k} \left\{ \frac{g(\|Y\|_2^2)}{\lambda_j} \right\} \leq m$. Moreover,

$$\max_{1 \leq j \leq k} \left\{ \frac{g(\|Y\|_2^2)}{\lambda'_j} \right\} \leq \max_{1 \leq j \leq k} \left\{ \frac{g(\|Y\|_2^2)}{\lambda_j - \|\lambda - \lambda'\|_2} \right\} \leq m$$

by definition of m and δ . Using this and again the definition of δ , we then find

$$\left| g(\|Y - X\widehat{\beta}^\lambda\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \widehat{\beta}^\lambda\|_{q_j} - g(\|Y - X\widehat{\beta}^{\lambda'}\|_2^2) - \sum_{j=1}^k \lambda'_j \|M_j \widehat{\beta}^{\lambda'}\|_{q_j} \right| < \epsilon.$$

This implies the desired continuity.

Now we show that function $\lambda \mapsto g(\|Y - X\widehat{\beta}^\lambda\|_2^2)$ is continuous. We proceed with contradiction. Thus, we assume there exist $\lambda' \in (0, \infty)^k$ and $\epsilon_0 > 0$, so that for any $\delta > 0$, there exists $\lambda \in (0, \infty)^k$ satisfying

$$\|\lambda - \lambda'\|_2 < \delta \quad \text{and} \quad \left| g(\|Y - X\widehat{\beta}^\lambda\|_2^2) - g(\|Y - X\widehat{\beta}^{\lambda'}\|_2^2) \right| \geq \epsilon_0.$$

We note that by Lemma A.1, the value of $g(\|Y - X\widehat{\beta}^\lambda\|_2^2)$ does not depend on the specific choice of the estimator $\widehat{\beta}^\lambda$. Since $x \mapsto g(x)$ is strictly increasing, there exists $\epsilon'_0 \equiv \epsilon'_0(\epsilon_0) > 0$ such that

$$\left| \|X\widehat{\beta}^\lambda\|_2^2 - \|X\widehat{\beta}^{\lambda'}\|_2^2 \right| > \epsilon'_0. \quad (7)$$

Define the set $B := \left\{ \beta \in \mathbb{R}^p : \left| \|X\widehat{\beta}^{\lambda'}\|_2^2 - \|X\beta\|_2^2 \right| > \epsilon'_0 \right\}$. It follows directly that $\widehat{\beta}^{\lambda'} \notin B$, and due to Inequality (7) above, it follows that $\widehat{\beta}^\lambda \in B$. Let $\eta > 0$ be arbitrary. Invoking Criterion (2) and the continuity of the objective function (note that taking minima does not affect the continuity), we obtain

$$\begin{aligned} & g(\|Y - X\widehat{\beta}^\lambda\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \widehat{\beta}^\lambda\|_{q_j} \\ &= \min_{\beta \in B} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\} \\ &\geq \min_{\beta \in B} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda'_j \|M_j \beta\|_{q_j} \right\} - \eta \end{aligned}$$

if δ is sufficiently small. Moreover, using $\widehat{\beta}^{\lambda'} \notin B$ and again the continuity, it holds for δ

sufficiently small that

$$\begin{aligned}
& \min_{\beta \in B} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda'_j \|M_j \beta\|_{q_j} \right\} \\
& > \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda'_j \|M_j \beta\|_{q_j} \right\} + \xi \\
& \geq \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\} + \xi/2
\end{aligned}$$

for a $\xi \equiv \xi(\lambda', \epsilon'_0) > 0$. Choosing $\eta = \xi/4$, we find

$$\begin{aligned}
& \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\} \\
& > \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\} + \xi/4,
\end{aligned}$$

which is a contradiction and thus concludes the proof of the continuity of the function $\lambda \mapsto g(\|Y - X\hat{\beta}^\lambda\|_2^2)$. The continuity of the function $\lambda \mapsto \|Y - X\hat{\beta}^\lambda\|_2^2$ then follows from the assumption that the link function g is continuous and increasing. This concludes the proof of the lemma. \square

Proof of Lemma A.3. Since $Y \neq \mathbf{0}_n$ with probability one, we assume $Y \neq \mathbf{0}_n$ in the following.

We then show that

$$\|Y - X\hat{\beta}^\lambda\|_2 > 0.$$

We do this by contradiction, that is, we assume

$$\|Y - X\hat{\beta}^\lambda\|_2 = 0.$$

This implies

$$Y - X\hat{\beta}^\lambda = \mathbf{0}_n. \quad (8)$$

Since $\beta \mapsto g(\|Y - X\beta\|_2^2)$ is convex, the subdifferential $\partial_\beta g(\|Y - X\beta\|_2^2) \Big|_{\beta=\hat{\beta}^\lambda}$ exists. Thus, the KKT conditions imply

$$\mathbf{0}_p \in \partial_\beta g(\|Y - X\beta\|_2^2) \Big|_{\beta=\hat{\beta}^\lambda} + \sum_{j=1}^k \lambda_j \partial_\beta \|M_j \beta\|_{q_j} \Big|_{\beta=\hat{\beta}^\lambda},$$

which implies by the chain rule

$$\mathbf{0}_p \in \partial_x g(x) \Big|_{x=\|Y - X\hat{\beta}^\lambda\|_2^2} (-2X^\top (Y - X\hat{\beta}^\lambda)) + \sum_{j=1}^k \lambda_j \partial_\beta \|M_j \beta\|_{q_j} \Big|_{\beta=\hat{\beta}^\lambda}.$$

Plugging Equality (8) into this display yields

$$\mathbf{0}_p \in \sum_{j=1}^k \lambda_j \partial_{\beta} \|M_j \beta\|_{q_j} \Big|_{\beta=\hat{\beta}^\lambda}.$$

This means that the vector $\hat{\beta}^\lambda$ minimizes the function $\beta \mapsto \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j}$. However, since $\lambda_1, \dots, \lambda_k > 0$, and by assumption on the matrices M_1, \dots, M_k , this mapping is a norm and is thus minimized only at $\mathbf{0}_p$. Consequently, $\hat{\beta}^\lambda = \mathbf{0}_p$. However, Equality (8) then gives

$$Y = X \hat{\beta}^\lambda = X \mathbf{0}_p = \mathbf{0}_n,$$

which contradicts $Y \neq \mathbf{0}_n$. Thus, $Y - X \hat{\beta}^\lambda \neq \mathbf{0}_n$, and it follows that $\|Y - X \hat{\beta}^\lambda\|_2^2 \neq 0$.

Since the function $x \mapsto g(x)$ is continuously differentiable on $(0, \infty)$ with strictly positive derivative, we finally obtain

$$g'(\|Y - X \hat{\beta}^\lambda\|_2^2) > 0$$

as desired. \square

A.3. Proof of Lemma 2.1

Proof of Lemma 2.1. The proof consists of three steps. First, we show that the solution equals zero if the tuning parameters are large enough. Second, we show that if one element of the tuning parameter is sufficiently large, increasing that element does not affect the estimator. Finally, we use these results to show the existence of suitable tuning parameters.

Let us start with some notation. For each $j \in \{1, \dots, k\}$, we define the set $A_j \subset \{1, 2, \dots, p\}$ such that for any $u \in A_j$, the u th row of M_j is not zero, that is, $A_j := \{u \in \{1, \dots, p\} : \max_{v \in \{1, \dots, p\}} |(M_j)_{uv}| > 0\}$. By assumption on the sequence M_1, \dots, M_k , it holds that $\bigcup_{j=1}^k A_j = \{1, \dots, p\}$. For $j \in \{1, \dots, k\}$, define the vector $(X^\top Y)^j \in \mathbb{R}^p$ via $(X^\top Y)_i^j := (X^\top Y)_i \mathbb{1}\{i \in A_j\}$, and set $m := \max_{j \in \{1, \dots, k\}} 2g'(\|Y\|_2^2) \|(X^\top Y)^j\|_{q_j}^* \vee 1$, where \vee denotes the maximum.

The three mentioned steps now read in detail:

1. Show that for any $\lambda \in [m, \infty)^k$, it holds that

$$\{\mathbf{0}_p\} = \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\}.$$

2. Show that $X \hat{\beta}^{\tilde{\lambda}} = X \hat{\beta}^\lambda$ if $\tilde{\lambda}, \lambda \in (0, \infty)^k$ satisfy

$$\begin{aligned} \tilde{\lambda}_j &> \lambda_j = m && \text{if } j \in B \\ \tilde{\lambda}_j &= \lambda_j && \text{if } j \notin B \end{aligned}$$

for a non-empty subset $B \subset \{1, 2, \dots, k\}$.

3. Show that with probability one, there exists a vector $\lambda \in (0, \infty)^k$ that satisfies

$$\frac{\lambda_j}{2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2)} = c_j \|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* \quad (9)$$

for all $j \in \{1, \dots, k\}$.

Step 1: We first show that

$$\mathbf{0}_p \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\}$$

implies

$$\{\mathbf{0}_p\} = \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\}.$$

Assume for a $\lambda \in [m, \infty)^k$, it holds that

$$\mathbf{0}_p \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\}.$$

then by Lemma A.1, for any $\widehat{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\}$, it holds that $X\widehat{\beta}^\lambda = \mathbf{0}_n$. Hence, by assumption on the matrices M_1, \dots, M_k , it holds for any $\widehat{\beta}^\lambda \neq \mathbf{0}_p$,

$$\begin{aligned} & g(\|Y - X\widehat{\beta}^\lambda\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \widehat{\beta}^\lambda\|_{q_j} \\ &= g(\|Y\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \widehat{\beta}^\lambda\|_{q_j} \\ &> g(\|Y\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \mathbf{0}_p\|_{q_j} \\ &= g(\|Y - X\mathbf{0}_p\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \mathbf{0}_p\|_{q_j}. \end{aligned}$$

This contradicts $\widehat{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\}$, and thus, $\widehat{\beta}^\lambda = \mathbf{0}_p$. Thus,

$$\{\mathbf{0}_p\} = \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\}$$

as desired. It is left to show that for any vector $\lambda \in [m, \infty)^k$, it holds that

$$\mathbf{0}_p \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\}.$$

By the KKT conditions, we have to show that there are vectors $\kappa(j) \in \partial_{\beta} \|M_j \beta\|_{q_j} \Big|_{\beta=\mathbf{0}_p}$ such that

$$-2g'(\|Y\|_2^2)X^{\top}Y + \sum_{j=1}^k \lambda_j \kappa(j) = \mathbf{0}_p. \quad (10)$$

Define $\tilde{A}_1 := A_1$, $\tilde{A}_j := A_j \setminus \{A_1, \dots, A_{j-1}\}$ for $j = 2, \dots, k$. In particular, $\tilde{A}_j \subset A_j$, the \tilde{A}_j 's are disjoint, and $\bigcup_{j=1}^k \tilde{A}_j = \{1, \dots, p\}$. With this notation, we need to show that for any $v \in \tilde{A}_j$ and $j \in \{1, \dots, k\}$, it holds that

$$-2g'(\|Y\|_2^2)(X^{\top}Y)_v + \sum_{j=1}^k \lambda_j \kappa(j)_v = 0.$$

Define the vector $\kappa(j)$ for $j \in \{1, \dots, k\}$ via

$$\kappa(j)_v := \begin{cases} \frac{2g'(\|Y\|_2^2)(X^{\top}Y)_v^j}{\lambda_j} & \text{if } v \in \tilde{A}_j \\ 0 & \text{if } v \notin \tilde{A}_j. \end{cases}$$

Since $\tilde{A}_j \subset A_j$ and \tilde{A}_j 's are disjoint, we find for all $j \in \{1, \dots, k\}$ and $v \in \tilde{A}_j$, that

$$-2g'(\|Y\|_2^2)(X^{\top}Y)_v + \sum_{j=1}^k \lambda_j \kappa(j)_v = -2g'(\|Y\|_2^2)(X^{\top}Y)_v^j + \lambda_j \kappa(j)_v.$$

By definition of the vectors $\kappa(j)$ for $j \in \{1, \dots, k\}$, we thus have

$$-2g'(\|Y\|_2^2)(X^{\top}Y)_v^j + \lambda_j \kappa(j)_v = 0,$$

which implies

$$-2g'(\|Y\|_2^2)(X^{\top}Y)_v + \sum_{j=1}^k \lambda_j \kappa(j)_v = 0.$$

Since $\lambda_j \geq m \geq 2g'(\|Y\|_2^2)\|(X^{\top}Y)^j\|_{q_j}^*$, it also follows by taking the dual norm on both sides that

$$\|\kappa(j)\|_{q_j}^* \leq \frac{2g'(\|Y\|_2^2)\|(X^{\top}Y)^j\|_{q_j}^*}{\lambda_j} \leq 1, \quad (11)$$

and hence, $\kappa(j) \in \partial_{\beta} \|M_j \beta\|_{q_j} \Big|_{\beta=\mathbf{0}_p}$ for all $j \in \{1, \dots, k\}$. So we get

$$\mathbf{0}_p \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\}.$$

as desired. We conclude that

$$\{\mathbf{0}_p\} = \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\}.$$

Step 2: Consider a pair of vectors $\lambda, \tilde{\lambda} \in (0, \infty)^k$, that satisfy $\tilde{\lambda}_j > \lambda_j = m$ for $j \in B$ and $\tilde{\lambda}_j = \lambda_j$ for $j \in \{1, \dots, k\} \setminus B$. For λ , fix a solution

$$\hat{\beta}^{\lambda} \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j} \right\}$$

with corresponding subdifferentials $\kappa(j) \in \partial_{\beta} \|M_j \beta\|_{q_j} \Big|_{\beta=\hat{\beta}^{\lambda}}$ that satisfy the KKT conditions

$$-2g'(\|Y - X\hat{\beta}^{\lambda}\|_2^2)X^{\top}(Y - X\hat{\beta}^{\lambda}) + \sum_{j=1}^k \lambda_j \kappa(j) = \mathbf{0}_p.$$

We first need to show that

$$\hat{\beta}^{\lambda} \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2) + \sum_{j=1}^k \tilde{\lambda}_j \|M_j \beta\|_{q_j} \right\}.$$

By the KKT conditions, we have to show that there are the vectors $\tilde{\kappa}(j) \in \partial_{\beta} \|M_j \beta\|_{q_j} \Big|_{\beta=\hat{\beta}^{\lambda}}$ such that

$$-2g'(\|Y - X\hat{\beta}^{\lambda}\|_2^2)X^{\top}(Y - X\hat{\beta}^{\lambda}) + \sum_{j=1}^k \tilde{\lambda}_j \tilde{\kappa}(j) = \mathbf{0}_p.$$

Define $\tilde{\kappa}(j)$ for $j \in \{1, \dots, k\}$ via

$$\tilde{\kappa}(j) := \frac{\lambda_j}{\tilde{\lambda}_j} \kappa(j).$$

Plugging this into the previous display yields

$$-2g'(\|Y - X\hat{\beta}^{\lambda}\|_2^2)X^{\top}(Y - X\hat{\beta}^{\lambda}) + \sum_{j=1}^k \tilde{\lambda}_j \tilde{\kappa}(j) = -2g'(\|Y - X\hat{\beta}^{\lambda}\|_2^2)X^{\top}(Y - X\hat{\beta}^{\lambda}) + \sum_{j=1}^k \lambda_j \kappa_j.$$

Therefore, it holds that

$$-2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2)X^\top(Y - X\widehat{\beta}^\lambda) + \sum_{j=1}^k \tilde{\lambda}_j \tilde{\kappa}(j) = \mathbf{0}_p.$$

Moreover, by definition of $\tilde{\kappa}(j)$ and Inequality (11), we have that

$$\|\tilde{\kappa}(j)\|_{q_j}^* = \frac{\lambda_j}{\tilde{\lambda}_j} \|\kappa(j)\|_{q_j}^* \leq \|\kappa(j)\|_{q_j}^* \leq 1,$$

for all $j \in B$. So, $\tilde{\kappa}(j) \in \partial_\beta \|M_j \beta\|_{q_j} \Big|_{\beta = \widehat{\beta}^\lambda}$ for any $j \in \{1, \dots, k\}$. Hence, it holds that

$$\widehat{\beta}^\lambda \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ g(\|Y - X\beta\|_2) + \sum_{j=1}^k \tilde{\lambda}_j \|(M_j \beta)^j\|_{q_j} \right\}.$$

This gives $X\widehat{\beta}^{\tilde{\lambda}} = X\widehat{\beta}^\lambda$ by Lemma A.1.

Step 3: Finally, we show the existence of a $\lambda \in (0, \infty)^k$ that satisfies Equality (9). For this, we define

$$a := m \wedge \left(2g'(\|Y\|_2^2) \min_{j \in \{1, \dots, k\}} c_j \|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* \right).$$

By assumption on the noise ε , it holds that $a > 0$ with probability one. Next, we consider the function

$$f : [a, m]^k \rightarrow \mathbb{R}^k$$

$$\lambda \mapsto f(\lambda) := 2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2) (c_1 \|(XP_1 M_1^+)^\top \varepsilon\|_{q_1}^*, \dots, c_k \|(XP_k M_k^+)^\top \varepsilon\|_{q_k}^*)^\top.$$

Note first that $g(\|Y - X\widehat{\beta}^\lambda\|_2^2) \leq g(\|Y\|_2^2)$ by definition of the estimator $\widehat{\beta}^\lambda$. Hence, since g' is non-increasing, we find

$$\min_{\lambda \in [a, m]^k} \min_{j \in \{1, \dots, k\}} f(\lambda)_j \geq a.$$

Note also that $[a, m]^k$ is compact and that f is continuous by Lemmas A.2 and A.3 and the assumption that $g'(x)$ is continuous on $(0, \infty)$. It therefore holds that

$$\sup_{\lambda \in [a, m]^k} \|f(\lambda)\|_\infty \leq b$$

for some $b \in (0, \infty)$. Using this and Step 2, we find that the function f and the function

$$h : [a, b \vee m]^k \rightarrow [a, b \vee m]^k$$

$$\lambda \mapsto h(\lambda) := 2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2) (c_1 \|(XP_1 M_1^+)^\top \varepsilon\|_{q_1}^*, \dots, c_k \|(XP_k M_k^+)^\top \varepsilon\|_{q_k}^*)^\top$$

have the same images, that is,

$$\{y : y = f(\lambda), \lambda \in [a, m]^k\} = \{y : y = h(\lambda), \lambda \in [a, b \vee m]^k\}.$$

The function h is continuous. Moreover, $[a, b \vee m]^k$ is a compact and convex subset of \mathbb{R}^k . We can thus apply Brouwer's fixed-point theorem to deduce that

$$\lambda = 2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2) (c_1\|(XP_1M_1^+)^\top \varepsilon\|_{q_1}^*, \dots, c_k\|(XP_kM_k^+)^\top \varepsilon\|_{q_k}^*)^\top$$

for a vector $\lambda \in [a, b \vee m]^k$. According to Lemma A.3, it holds that $g'(\|Y - X\widehat{\beta}^\lambda\|_2^2) > 0$ with probability one, so that

$$\frac{\lambda}{2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2)} = (c_1\|(XP_1M_1^+)^\top \varepsilon\|_{q_1}^*, \dots, c_k\|(XP_kM_k^+)^\top \varepsilon\|_{q_k}^*)^\top$$

as desired. \square

A.4. Proof of Theorem 2.1

Proof of Theorem 2.1. Consider the function

$$\begin{aligned} \mathbb{R}^p &\rightarrow \mathbb{R} \\ \beta &\mapsto f(\beta) := g(\|Y - X\beta\|_2^2) + \sum_{j=1}^k \lambda_j \|M_j \beta\|_{q_j}. \end{aligned}$$

According to our assumptions on the link function and the penalties, f is convex, and $\widehat{\beta}^\lambda$ minimizes f . Therefore, $\mathbf{0}_p \in \partial f(\beta)|_{\beta=\widehat{\beta}^\lambda}$. Subdifferentials are additive, so that we can write $\partial f(\beta)|_{\beta=\widehat{\beta}^\lambda}$ as a sum of subdifferentials of the individual parts of f . In particular, we can decompose $\mathbf{0}_p \in \partial f(\beta)|_{\beta=\widehat{\beta}^\lambda}$ as

$$\mathbf{0}_p = \frac{\partial}{\partial \beta} g(\|Y - X\beta\|_2^2) \Big|_{\beta=\widehat{\beta}^\lambda} + \sum_{j=1}^k \lambda_j \kappa_j,$$

where $\kappa_j \in \partial\{\alpha \mapsto \|M_j \alpha\|_{q_j}\}|_{\alpha=\widehat{\beta}^\lambda}$ and, using the assumption that the function g is differentiable on $(0, \infty)$,

$$\frac{\partial}{\partial \beta} g(\|Y - X\beta\|_2^2) \Big|_{\beta=\widehat{\beta}^\lambda} = -2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2) (X^\top (Y - X\widehat{\beta}^\lambda)).$$

Adding the pieces together implies for all $\beta \in \mathbb{R}^p$

$$0 = \mathbf{0}_p^\top (\beta - \widehat{\beta}^\lambda) = (-2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2) (X^\top (Y - X\widehat{\beta}^\lambda)) + \sum_{j=1}^k \lambda_j \kappa_j)^\top (\beta - \widehat{\beta}^\lambda).$$

Now, $\kappa_j \in \partial\{\alpha \mapsto \|M_j \alpha\|_{q_j}\}_{\alpha=\hat{\beta}^\lambda}$, which means by the definition of subdifferentials for convex functions that for all $\beta \in \mathbb{R}^p$

$$\|M_j \beta\|_{q_j} \geq \|M_j \hat{\beta}^\lambda\|_{q_j} + \langle \kappa_j, \beta - \hat{\beta}^\lambda \rangle,$$

which is equivalent to

$$\|M_j \beta\|_{q_j} - \|M_j \hat{\beta}^\lambda\|_{q_j} \geq \kappa_j^\top (\beta - \hat{\beta}^\lambda).$$

Combining with the above yields

$$-2g'(\|Y - X\hat{\beta}^\lambda\|_2^2)(X^\top(Y - X\hat{\beta}^\lambda))^\top(\beta - \hat{\beta}^\lambda) + \sum_{j=1}^k \lambda_j (\|M_j \beta\|_{q_j} - \|M_j \hat{\beta}^\lambda\|_{q_j}) \geq 0.$$

According to the model (1), we can replace Y with $X\beta^* + \varepsilon$ to obtain

$$\begin{aligned} & (X^\top(Y - X\hat{\beta}^\lambda))^\top(\beta - \hat{\beta}^\lambda) \\ &= (X^\top(X\beta^* + \varepsilon - X\hat{\beta}^\lambda))^\top(\beta - \hat{\beta}^\lambda) \\ &= (X(\beta^* - \hat{\beta}^\lambda))^\top X(\beta - \hat{\beta}^\lambda) + \varepsilon^\top X(\beta - \hat{\beta}^\lambda) \\ &= (X(\beta^* - \hat{\beta}^\lambda))^\top X(\beta^* - \hat{\beta}^\lambda + \beta - \beta^*) + \varepsilon^\top X(\beta - \hat{\beta}^\lambda) \\ &= \|X(\beta^* - \hat{\beta}^\lambda)\|_2^2 + (X(\beta^* - \hat{\beta}^\lambda))^\top X(\beta - \beta^*) + \varepsilon^\top X(\beta - \hat{\beta}^\lambda). \end{aligned}$$

Now, we note that for any $u > 0$,

$$(X(\beta^* - \hat{\beta}^\lambda))^\top X(\beta - \beta^*) \geq -u\|X(\beta^* - \hat{\beta}^\lambda)\|_2^2 - \frac{\|X(\beta - \beta^*)\|_2^2}{4u}.$$

Using this in the foregoing display and consolidating gives us

$$(X^\top(Y - X\hat{\beta}^\lambda))^\top(\beta - \hat{\beta}^\lambda) \geq (1-u)\|X(\beta^* - \hat{\beta}^\lambda)\|_2^2 - \frac{\|X(\beta - \beta^*)\|_2^2}{4u} + \varepsilon^\top X(\beta - \hat{\beta}^\lambda).$$

Plugging this back into the earlier display yields, noting that g' is a positive function by assumption,

$$\begin{aligned} & -2g'(\|Y - X\hat{\beta}^\lambda\|_2^2)((1-u)\|X(\beta^* - \hat{\beta}^\lambda)\|_2^2 - \frac{\|X(\beta - \beta^*)\|_2^2}{4u} + \langle \varepsilon, X(\beta - \hat{\beta}^\lambda) \rangle) \\ & \quad + \sum_{j=1}^k \lambda_j (\|M_j \beta\|_{q_j} - \|M_j \hat{\beta}^\lambda\|_{q_j}) \geq 0. \end{aligned}$$

Rearranging this inequality, we obtain

$$\begin{aligned} & 2g'(\|Y - X\hat{\beta}^\lambda\|_2^2)((1-u)\|X(\beta^* - \hat{\beta}^\lambda)\|_2^2 + \langle \varepsilon, X(\beta - \hat{\beta}^\lambda) \rangle) \\ & \leq \sum_{j=1}^k \lambda_j (\|M_j \beta\|_{q_j} - \|M_j \hat{\beta}^\lambda\|_{q_j}) + 2g'(\|Y - X\hat{\beta}^\lambda\|_2^2) \frac{\|X(\beta - \beta^*)\|_2^2}{4u}. \end{aligned}$$

According to Lemma A.3, we can divide both sides by $2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2)$, so that

$$\begin{aligned} (1-u)\|X(\beta^* - \widehat{\beta}^\lambda)\|_2^2 + \langle \varepsilon, X(\beta - \widehat{\beta}^\lambda) \rangle \\ \leq \sum_{j=1}^k \frac{\lambda_j}{2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2)} (\|M_j\beta\|_{q_j} - \|M_j\widehat{\beta}^\lambda\|_{q_j}) + \frac{\|X(\beta - \beta^*)\|_2^2}{4u} \end{aligned}$$

and

$$\begin{aligned} (1-u)\|X(\beta^* - \widehat{\beta}^\lambda)\|_2^2 \\ \leq \langle \varepsilon, X(\widehat{\beta}^\lambda - \beta) \rangle + \sum_{j=1}^k \frac{\lambda_j}{2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2)} (\|M_j\beta\|_{q_j} - \|M_j\widehat{\beta}^\lambda\|_{q_j}) + \frac{\|X(\beta - \beta^*)\|_2^2}{4u} \end{aligned}$$

with probability one. Recall that by Equation (3), the vector $\widehat{\beta}^\lambda - \beta$ can be rewritten as

$$\widehat{\beta}^\lambda - \beta = \sum_{j=1}^k P_j M_j^+ M_j (\widehat{\beta}^\lambda - \beta).$$

So we can reorganize the inner product $\langle \varepsilon, X(\widehat{\beta}^\lambda - \beta) \rangle$ according to

$$\begin{aligned} \langle \varepsilon, X(\widehat{\beta}^\lambda - \beta) \rangle &= \langle \varepsilon, X \sum_{j=1}^k P_j M_j^+ M_j (\widehat{\beta}^\lambda - \beta) \rangle \\ &= \sum_{j=1}^k \langle \varepsilon, X P_j M_j^+ M_j (\widehat{\beta}^\lambda - \beta) \rangle \\ &= \sum_{j=1}^k \langle (X P_j M_j^+)^T \varepsilon, M_j (\widehat{\beta}^\lambda - \beta) \rangle. \end{aligned}$$

Plugging this into the previous inequality yields

$$\begin{aligned} (1-u)\|X(\beta^* - \widehat{\beta}^\lambda)\|_2^2 \\ \leq \sum_{j=1}^k \langle (X P_j M_j^+)^T \varepsilon, M_j (\widehat{\beta}^\lambda - \beta) \rangle + \frac{\lambda_j}{2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2)} (\|M_j\beta\|_{q_j} - \|M_j\widehat{\beta}^\lambda\|_{q_j}) + \frac{\|X(\beta - \beta^*)\|_2^2}{4u}. \end{aligned}$$

Using Hölder's Inequality, we can rewrite the first term in the second line according to

$$\langle (X P_j M_j^+)^T \varepsilon, M_j (\widehat{\beta}^\lambda - \beta) \rangle \leq \|(X P_j M_j^+)^T \varepsilon\|_{q_j^*} \|M_j \widehat{\beta}^\lambda - M_j \beta\|_{q_j}.$$

Plugging this back into the previous display gives

$$\begin{aligned} (1-u)\|X(\beta^* - \widehat{\beta}^\lambda)\|_2^2 &\leq \sum_{j=1}^k (\|(X P_j M_j^+)^T \varepsilon\|_{q_j^*} \|M_j \widehat{\beta}^\lambda - M_j \beta\|_{q_j}) \\ &\quad + \frac{\lambda_j}{2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2)} (\|M_j\beta\|_{q_j} - \|M_j\widehat{\beta}^\lambda\|_{q_j}) + \frac{\|X(\beta - \beta^*)\|_2^2}{4u}. \end{aligned}$$

We can modify this further by applying the triangle inequality and by reorganizing the terms. We find

$$\begin{aligned}
& (1-u)\|X(\beta^* - \widehat{\beta}^\lambda)\|_2^2 \\
& \leq \sum_{j=1}^k \left(\|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* (\|M_j \widehat{\beta}^\lambda\|_{q_j} + \|M_j \beta\|_{q_j}) + \frac{\lambda_j}{2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2)} (\|M_j \beta\|_{q_j} - \|M_j \widehat{\beta}^\lambda\|_{q_j}) \right) \\
& \quad + \frac{\|X(\beta - \beta^*)\|_2^2}{4u} \\
& = \sum_{j=1}^k \left(\|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* + \frac{\lambda_j}{2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2)} \right) \|M_j \beta\|_{q_j} \\
& \quad + \sum_{j=1}^k \left(\|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* - \frac{\lambda_j}{2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2)} \right) \|M_j \widehat{\beta}^\lambda\|_{q_j} + \frac{\|X(\beta - \beta^*)\|_2^2}{4u}.
\end{aligned}$$

We now set λ according to Lemma 2.1. It then holds that

$$\frac{\lambda_j}{2g'(\|Y - X\widehat{\beta}^\lambda\|_2^2)} = c_j \|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^*$$

for all $j \in \{1, \dots, k\}$. Then it follows that

$$\begin{aligned}
& (1-u)\|X(\beta^* - \widehat{\beta}^\lambda)\|_2^2 \\
& \leq \sum_{j=1}^k (1+c_j) \|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* \|M_j \beta\|_{q_j} + \sum_{j=1}^k (1-c_j) \|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* \|M_j \widehat{\beta}^\lambda\|_{q_j} + \frac{\|X(\beta^* - \beta)\|_2^2}{4u}.
\end{aligned}$$

To bring this into the standard form, assuming $u < 1$, we finally divide both sides by $(1-u)n$ and find

$$\begin{aligned}
\frac{1}{n}\|X(\beta^* - \widehat{\beta}^\lambda)\|_2^2 & \leq \frac{1}{n} \sum_{j=1}^k \frac{1+c_j}{1-u} \|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* \|M_j \beta\|_{q_j} \\
& \quad + \frac{1}{n} \sum_{j=1}^k \frac{1-c_j}{1-u} \|(XP_j M_j^+)^\top \varepsilon\|_{q_j}^* \|M_j \widehat{\beta}^\lambda\|_{q_j} + \frac{1}{4u(1-u)n} \|X(\beta^* - \beta)\|_2^2.
\end{aligned}$$

Since $\beta \in \mathbb{R}^p$ and $u \in (0, 1)$ were arbitrary, this inequality provides us with the desired statement. \square