

A unified framework for high-dimensional analysis of M -estimators with decomposable regularizers

Sahand N. Negahban¹, Pradeep Ravikumar²,
Martin J. Wainwright^{3,4} and Bin Yu^{3,4}

MIT, Department of EECS¹

UT Austin, Department of CS²

UC Berkeley, Department of EECS^{3,4} and Statistics^{3,4}

Abstract. High-dimensional statistical inference deals with models in which the the number of parameters p is comparable to or larger than the sample size n . Since it is usually impossible to obtain consistent procedures unless $p/n \rightarrow 0$, a line of recent work has studied models with various types of low-dimensional structure, including sparse vectors, sparse and structured matrices, low-rank matrices, and combinations thereof. In such settings, a general approach to estimation is to solve a regularized optimization problem, which combines a loss function measuring how well the model fits the data with some regularization function that encourages the assumed structure. This paper provides a unified framework for establishing consistency and convergence rates for such regularized M -estimators under high-dimensional scaling. We state one main theorem and show how it can be used to re-derive some existing results, and also to obtain a number of new results on consistency and convergence rates, in both ℓ_2 -error and related norms. Our analysis also identifies two key properties of loss and regularization functions, referred to as restricted strong convexity and decomposability, that ensure corresponding regularized M -estimators have fast convergence rates, and which are optimal in many well-studied cases.

Sahand Negahban, Department of EECS, Massachusetts Institute of Technology, Cambridge MA 02139 (e-mail: sahandn@mit.edu). Pradeep Ravikumar, Department of CS, University of Texas, Austin, Austin, TX 78701 (e-mail: pradeepr@cs.utexas.edu). Martin J. Wainwright, Department of EECS and Statistics, University of California Berkeley, Berkeley CA 94720 (e-mail: wainwrig@eecs.berkeley.edu). Bin Yu, Department of Statistics, University of California Berkeley, Berkeley CA 94720 (e-mail: binyu@stat.berkeley.edu).

1. INTRODUCTION

High-dimensional statistics is concerned with models in which the ambient dimension of the problem p may be of the same order as—or substantially larger than—the sample size n . On one hand, its roots are quite old, dating back to work on random matrix theory and high-dimensional testing problems (e.g. [24, 44, 56, 77]). On the other hand, the past decade has witnessed a tremendous surge of research activity. Rapid development of data collection technology is a major driving force: it allows for more observations to be collected (larger n), and also for more variables to be measured (larger p). Examples are ubiquitous throughout science: astronomical projects such as the Large Synoptic Survey Telescope [1] produce terabytes of data in a single evening; each sample is a high-resolution image, with several hundred megapixels, so that $p \gg 10^8$. Financial data is also of a high-dimensional nature, with hundreds or thousands of financial instruments being measured and tracked over time, often at very fine time intervals for use in high frequency trading. Advances in biotechnology now allow for measurements of thousands of genes or proteins, and lead to numerous statistical challenges (e.g., see the paper [7] and references therein). Various types of imaging technology, among them magnetic resonance imaging in medicine [42] and hyper-spectral imaging in ecology [38], also lead to high-dimensional data sets.

In the regime $p \gg n$, it is well known that consistent estimators cannot be obtained unless additional constraints are imposed on the model. Accordingly, there are now several lines of work within high-dimensional statistics, all of which are based on imposing some type of low-dimensional constraint on the model space, and then studying the behavior of different estimators. Examples include linear regression with sparsity constraints, estimation of structured covariance or inverse covariance matrices, graphical model selection, sparse principal component analysis, low-rank matrix estimation, matrix decomposition problems, and estimation of sparse additive non-parametric models. The classical technique of regularization has proven fruitful in all of these contexts. Many well-known estimators are based on solving a convex optimization problem formed by the sum of a loss function with a weighted regularizer; we refer to any such method as a *regularized M -estimator*. For instance, in application to linear models, the Lasso or basis pursuit approach [69, 20] is based on a combination of the least-squares loss with ℓ_1 -regularization, and so involves solving a quadratic program. Similar approaches have been applied to generalized linear models, resulting in more general (non-quadratic) convex programs with ℓ_1 -constraints. Several types of regularization have been used for estimating matrices, including standard ℓ_1 -regularization, a wide range of sparse group-structured regularizers, as well as regularization based on the nuclear norm (sum of singular values).

Past work: Within the framework of high-dimensional statistics, the goal is to obtain bounds on a given performance metric that hold with high probability for a finite sample size, and provide explicit control on the ambient dimension p , as well as other structural parameters such as the sparsity of a vector, degree of a graph, or rank of matrix. Typically, such bounds show that the ambient dimension and structural parameters can grow as some function of the sample size n , while still having the statistical error decrease to zero. The choice of performance metric is application-dependent; some examples include prediction error, parameter estimation error, and model selection error.

By now, there are a large number of theoretical results in place for various types of regularized M -estimators.¹ Sparse linear regression has perhaps been the most active area, and multiple bodies of work can be differentiated by the error metric under consideration. They include work on exact recovery for noiseless observations (e.g., [22, 21, 15]), prediction error consistency (e.g., [25, 13, 73, 80]), consistency of the parameter estimates in ℓ_2 or some other norm (e.g., [13, 12, 73, 80, 48, 9, 16]), as well as variable selection consistency (e.g., [47, 76, 82]). The information-theoretic limits of sparse linear regression are also well-understood, and ℓ_1 -based methods are known to be optimal for ℓ_q -ball sparsity [58], and near-optimal for model selection [75]. For generalized linear models (GLMs), estimators based on ℓ_1 -regularized maximum likelihood have also been studied, including results on risk consistency [74], consistency in ℓ_2 or ℓ_1 -norm [5, 31, 46], and model selection consistency [61, 10]. Sparsity has also proven useful in application to different types of matrix estimation problems, among them banded and sparse covariance matrices (e.g., [8, 14, 32]). Another line of work has studied the problem of estimating Gaussian Markov random fields, or equivalently inverse covariance matrices with sparsity constraints. Here there are a range of results, including convergence rates in Frobenius, operator and other matrix norms [66, 62, 37, 83], as well as results on model selection consistency [62, 37, 47]. Motivated by applications in which sparsity arises in a structured manner, other researchers have proposed different types of block-structured regularizers (e.g., [71, 34, 72, 81, 79, 3, 6, 29]), among them the group Lasso based on ℓ_1/ℓ_2 regularization. High-dimensional consistency results have been obtained for exact recovery based on noiseless observations [68, 6], convergence rates in ℓ_2 -norm (e.g., [49, 28, 41, 6]) as well as model selection consistency (e.g., [55, 53, 49]). Problems of low-rank matrix estimation also arise in numerous applications. Techniques based on nuclear norm regularization have been studied for different statistical models, including compressed sensing [64, 39], matrix completion [17, 33, 63, 54], multitask regression [78, 52, 65, 11, 4], and system identification [23, 52, 40]. Finally, although the primary emphasis of this paper is on high-dimensional parametric models, regularization methods have also proven effective for a class of high-dimensional non-parametric models that have a sparse additive decomposition (e.g., [60, 45, 35, 36]), and shown to achieve minimax-optimal rates [59].

Our contributions: As we have noted previously, almost all of these estimators can be seen as particular types of regularized M -estimators, with the choice of loss function, regularizer and statistical assumptions changing according to the model. This methodological similarity suggests an intriguing possibility: is there a *common set of theoretical principles* that underlies analysis of all these estimators? If so, it could be possible to gain a unified understanding of a large collection of techniques for high-dimensional estimation, and afford some insight into the literature.

The main contribution of this paper is to provide an affirmative answer to this question. In particular, we isolate and highlight two key properties of a regularized M -estimator—namely, a *decomposability property* for the regularizer, and a notion of *restricted strong convexity* that depends on the interaction between

¹Given the extraordinary number of papers that have appeared in recent years, it must be emphasized that our referencing is necessarily incomplete.

the regularizer and the loss function. For loss functions and regularizers satisfying these two conditions, we prove a general result (Theorem 1) about consistency and convergence rates for the associated estimators. This result provides a family of bounds indexed by subspaces, and each bound consists of the sum of approximation error and estimation error. This general result, when specialized to different statistical models, yields in a direct manner a large number of corollaries, some of them known and others novel. In concurrent work, a subset of the current authors have also used this framework to prove several results on low-rank matrix estimation using the nuclear norm [52], as well as minimax-optimal rates for noisy matrix completion [54] and noisy matrix decomposition [2]. Finally, en route to establishing these corollaries, we also prove some new technical results that are of independent interest, including guarantees of restricted strong convexity for group-structured regularization (Proposition 1).

The remainder of this paper is organized as follows. We begin in Section 2 by formulating the class of regularized M -estimators that we consider, and then defining the notions of decomposability and restricted strong convexity. Section 3 is devoted to the statement of our main result (Theorem 1), and discussion of its consequences. Subsequent sections are devoted to corollaries of this main result for different statistical models, including sparse linear regression (Section 4) and estimators based on group-structured regularizers (Section 5). A number of technical results are presented within the appendices in the supplementary file [51].

2. PROBLEM FORMULATION AND SOME KEY PROPERTIES

In this section, we begin with a precise formulation of the problem, and then develop some key properties of the regularizer and loss function.

2.1 A family of M -estimators

Let $Z_1^n := \{Z_1, \dots, Z_n\}$ denote n identically distributed observations with marginal distribution \mathbb{P} , and suppose that we are interested in estimating some parameter θ of the distribution \mathbb{P} . Let $\mathcal{L} : \mathbb{R}^p \times \mathcal{Z}^n \rightarrow \mathbb{R}$ be a convex and differentiable loss function that, for a given set of observations Z_1^n , assigns a cost $\mathcal{L}(\theta; Z_1^n)$ to any parameter $\theta \in \mathbb{R}^p$. Let $\theta^* \in \arg \min_{\theta \in \mathbb{R}^p} \bar{\mathcal{L}}(\theta)$ be any minimizer of the population risk $\bar{\mathcal{L}}(\theta) := \mathbb{E}_{Z_1^n}[\mathcal{L}(\theta; Z_1^n)]$. In order to estimate this quantity based on the data Z_1^n , we solve the convex optimization problem

$$(1) \quad \hat{\theta}_{\lambda_n} \in \arg \min_{\theta \in \mathbb{R}^p} \{\mathcal{L}(\theta; Z_1^n) + \lambda_n \mathcal{R}(\theta)\},$$

where $\lambda_n > 0$ is a user-defined regularization penalty, and $\mathcal{R} : \mathbb{R}^p \rightarrow \mathbb{R}_+$ is a norm. Note that this set-up allows for the possibility of mis-specified models as well.

Our goal is to provide general techniques for deriving bounds on the difference between any solution $\hat{\theta}_{\lambda_n}$ to the convex program (1) and the unknown vector θ^* . In this paper, we derive bounds on the quantity $\|\hat{\theta}_{\lambda_n} - \theta^*\|$, where the error norm $\|\cdot\|$ is induced by some inner product $\langle \cdot, \cdot \rangle$ on \mathbb{R}^p . Most often, this error norm will either be the Euclidean ℓ_2 -norm on vectors, or the analogous Frobenius norm for matrices, but our theory also applies to certain types of weighted norms. In addition, we provide bounds on the quantity $\mathcal{R}(\hat{\theta}_{\lambda_n} - \theta^*)$, which measures the error in the regularizer norm. In the classical setting, the ambient dimension p

stays fixed while the number of observations n tends to infinity. Under these conditions, there are standard techniques for proving consistency and asymptotic normality for the error $\widehat{\theta}_{\lambda_n} - \theta^*$. In contrast, the analysis of this paper is all within a high-dimensional framework, in which the tuple (n, p) , as well as other problem parameters, such as vector sparsity or matrix rank etc., are all allowed to tend to infinity. In contrast to asymptotic statements, our goal is to obtain explicit finite sample error bounds that hold with high probability.

2.2 Decomposability of \mathcal{R}

The first ingredient in our analysis is a property of the regularizer known as decomposability, defined in terms of a pair of subspaces $\mathcal{M} \subseteq \overline{\mathcal{M}}$ of \mathbb{R}^p . The role of the *model subspace* \mathcal{M} is to capture the constraints specified by the model; for instance, it might be the subspace of vectors with a particular support (see Example 1), or a subspace of low-rank matrices (see Example 3). The orthogonal complement of the space $\overline{\mathcal{M}}$, namely the set

$$(2) \quad \overline{\mathcal{M}}^\perp := \{v \in \mathbb{R}^p \mid \langle u, v \rangle = 0 \text{ for all } u \in \overline{\mathcal{M}}\}$$

is referred to as the *perturbation subspace*, representing deviations away from the model subspace \mathcal{M} . In the ideal case, we have $\overline{\mathcal{M}}^\perp = \mathcal{M}^\perp$, but our definition allows for the possibility that $\overline{\mathcal{M}}$ is strictly larger than \mathcal{M} , so that $\overline{\mathcal{M}}^\perp$ is strictly smaller than \mathcal{M}^\perp . This generality is needed for treating the case of low-rank matrices and nuclear norm, as discussed in Example 3 to follow.

DEFINITION 1. Given a pair of subspaces $\mathcal{M} \subseteq \overline{\mathcal{M}}$, a norm-based regularizer \mathcal{R} is *decomposable* with respect to $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ if

$$(3) \quad \mathcal{R}(\theta + \gamma) = \mathcal{R}(\theta) + \mathcal{R}(\gamma) \quad \text{for all } \theta \in \mathcal{M} \text{ and } \gamma \in \overline{\mathcal{M}}^\perp.$$

In order to build some intuition, let us consider the ideal case $\mathcal{M} = \overline{\mathcal{M}}$ for the time being, so that the decomposition (3) holds for all pairs $(\theta, \gamma) \in \mathcal{M} \times \mathcal{M}^\perp$. For any given pair (θ, γ) of this form, the vector $\theta + \gamma$ can be interpreted as perturbation of the model vector θ away from the subspace \mathcal{M} , and it is desirable that the regularizer penalize such deviations as much as possible. By the triangle inequality for a norm, we always have $\mathcal{R}(\theta + \gamma) \leq \mathcal{R}(\theta) + \mathcal{R}(\gamma)$, so that the decomposability condition (3) holds if and only if the triangle inequality is tight for all pairs $(\theta, \gamma) \in (\mathcal{M}, \mathcal{M}^\perp)$. It is exactly in this setting that the regularizer penalizes deviations away from the model subspace \mathcal{M} as much as possible.

In general, it is not difficult to find subspace pairs that satisfy the decomposability property. As a trivial example, any regularizer is decomposable with respect to $\mathcal{M} = \mathbb{R}^p$ and its orthogonal complement $\mathcal{M}^\perp = \{0\}$. As will be clear in our main theorem, it is of more interest to find subspace pairs in which the model subspace \mathcal{M} is “small”, so that the orthogonal complement \mathcal{M}^\perp is “large”. To formalize this intuition, let us define the projection operator

$$(4) \quad \Pi_{\mathcal{M}}(u) := \arg \min_{v \in \mathcal{M}} \|u - v\|,$$

with the projection $\Pi_{\mathcal{M}^\perp}$ defined in an analogous manner. To simplify notation, we frequently use the shorthand $u_{\mathcal{M}} = \Pi_{\mathcal{M}}(u)$ and $u_{\mathcal{M}^\perp} = \Pi_{\mathcal{M}^\perp}(u)$.

Of interest to us are the action of these projection operators on the unknown parameter $\theta^* \in \mathbb{R}^p$. In the most desirable setting, the model subspace \mathcal{M} can be chosen such that $\theta_{\mathcal{M}}^* \approx \theta^*$, or equivalently, such that $\theta_{\mathcal{M}^\perp}^* \approx 0$. If this can be achieved with the model subspace \mathcal{M} remaining relatively small, then our main theorem guarantees that it is possible to estimate θ^* at a relatively fast rate. The following examples illustrate suitable choices of the spaces \mathcal{M} and $\overline{\mathcal{M}}$ in three concrete settings, beginning with the case of sparse vectors.

EXAMPLE 1. *Sparse vectors and ℓ_1 -norm regularization.* Suppose the error norm $\|\cdot\|$ is the usual ℓ_2 -norm, and that the model class of interest is the set of s -sparse vectors in p dimensions. For any particular subset $S \subseteq \{1, 2, \dots, p\}$ with cardinality s , we define the model subspace

$$(5) \quad \mathcal{M}(S) := \{\theta \in \mathbb{R}^p \mid \theta_j = 0 \text{ for all } j \notin S\}.$$

Here our notation reflects the fact that \mathcal{M} depends explicitly on the chosen subset S . By construction, we have $\Pi_{\mathcal{M}(S)}(\theta^*) = \theta^*$ for any vector θ^* that is supported on S .

In this case, we may define $\overline{\mathcal{M}}(S) = \mathcal{M}(S)$, and note that the orthogonal complement with respect to the Euclidean inner product is given by

$$(6) \quad \overline{\mathcal{M}}^\perp(S) = \mathcal{M}^\perp(S) = \{\gamma \in \mathbb{R}^p \mid \gamma_j = 0 \text{ for all } j \in S\}.$$

This set corresponds to the perturbation subspace, capturing deviations away from the set of vectors with support S . We claim that for any subset S , the ℓ_1 -norm $\mathcal{R}(\theta) = \|\theta\|_1$ is decomposable with respect to the pair $(\mathcal{M}(S), \mathcal{M}^\perp(S))$. Indeed, by construction of the subspaces, any $\theta \in \mathcal{M}(S)$ can be written in the partitioned form $\theta = (\theta_S, 0_{S^c})$, where $\theta_S \in \mathbb{R}^s$ and $0_{S^c} \in \mathbb{R}^{p-s}$ is a vector of zeros. Similarly, any vector $\gamma \in \mathcal{M}^\perp(S)$ has the partitioned representation $(0_S, \gamma_{S^c})$. Putting together the pieces, we obtain

$$\|\theta + \gamma\|_1 = \|(\theta_S, 0) + (0, \gamma_{S^c})\|_1 = \|\theta\|_1 + \|\gamma\|_1,$$

showing that the ℓ_1 -norm is decomposable as claimed. \diamond

As a follow-up to the previous example, it is also worth noting that the same argument shows that for a strictly positive weight vector ω , the *weighted ℓ_1 -norm* $\|\theta\|_\omega := \sum_{j=1}^p \omega_j |\theta_j|$ is also decomposable with respect to the pair $(\mathcal{M}(S), \overline{\mathcal{M}}(S))$. For another natural extension, we now turn to the case of sparsity models with more structure.

EXAMPLE 2. *Group-structured norms.* In many applications, sparsity arises in a more structured fashion, with groups of coefficients likely to be zero (or non-zero) simultaneously. In order to model this behavior, suppose that the index set $\{1, 2, \dots, p\}$ can be partitioned into a set of $N_{\mathcal{G}}$ disjoint groups, say $\mathcal{G} = \{G_1, G_2, \dots, G_{N_{\mathcal{G}}}\}$. With this set-up, for a given vector $\vec{\alpha} = (\alpha_1, \dots, \alpha_{N_{\mathcal{G}}}) \in [1, \infty]^{N_{\mathcal{G}}}$, the associated $(1, \vec{\alpha})$ -group norm takes the form

$$(7) \quad \|\theta\|_{\mathcal{G}, \vec{\alpha}} := \sum_{t=1}^{N_{\mathcal{G}}} \|\theta_{G_t}\|_{\alpha_t}.$$

For instance, with the choice $\vec{\alpha} = (2, 2, \dots, 2)$, we obtain the group ℓ_1/ℓ_2 -norm, corresponding to the regularizer that underlies the group Lasso [79]. On the other hand, the choice $\vec{\alpha} = (\infty, \dots, \infty)$, corresponding to a form of block ℓ_1/ℓ_∞ regularization, has also been studied in past work [72, 53, 81]. Note that for $\vec{\alpha} = (1, 1, \dots, 1)$, we obtain the standard ℓ_1 penalty. Interestingly, our analysis shows that setting $\vec{\alpha} \in [2, \infty]^{N_G}$ can often lead to superior statistical performance.

We now show that the norm $\|\cdot\|_{\mathcal{G}, \vec{\alpha}}$ is again decomposable with respect to appropriately defined subspaces. Indeed, given any subset $S_G \subseteq \{1, \dots, N_G\}$ of group indices, say with cardinality $s_G = |S_G|$, we can define the subspace

$$(8) \quad \mathcal{M}(S_G) := \{\theta \in \mathbb{R}^p \mid \theta_{G_t} = 0 \text{ for all } t \notin S_G\},$$

as well as its orthogonal complement with respect to the usual Euclidean inner product

$$(9) \quad \mathcal{M}^\perp(S_G) = \overline{\mathcal{M}}^\perp(S_G) := \{\theta \in \mathbb{R}^p \mid \theta_{G_t} = 0 \text{ for all } t \in S_G\}.$$

With these definitions, for any pair of vectors $\theta \in \mathcal{M}(S_G)$ and $\gamma \in \overline{\mathcal{M}}^\perp(S_G)$, we have

$$(10) \quad \|\theta + \gamma\|_{\mathcal{G}, \vec{\alpha}} = \sum_{t \in S_G} \|\theta_{G_t} + 0_{G_t}\|_{\alpha_t} + \sum_{t \notin S_G} \|0_{G_t} + \gamma_{G_t}\|_{\alpha_t} = \|\theta\|_{\mathcal{G}, \vec{\alpha}} + \|\gamma\|_{\mathcal{G}, \vec{\alpha}},$$

thus verifying the decomposability condition. \diamond

In the preceding example, we exploited the fact that the groups were non-overlapping in order to establish the decomposability property. Therefore, some modifications would be required in order to choose the subspaces appropriately for overlapping group regularizers proposed in past work [29, 30].

EXAMPLE 3. *Low-rank matrices and nuclear norm.* Now suppose that each parameter $\Theta \in \mathbb{R}^{p_1 \times p_2}$ is a matrix; this corresponds to an instance of our general set-up with $p = p_1 p_2$, as long as we identify the space $\mathbb{R}^{p_1 \times p_2}$ with $\mathbb{R}^{p_1 p_2}$ in the usual way. We equip this space with the inner product $\langle\langle \Theta, \Gamma \rangle\rangle := \text{trace}(\Theta \Gamma^T)$, a choice which yields (as the induced norm) the *Frobenius norm*

$$(11) \quad \|\Theta\|_F := \sqrt{\langle\langle \Theta, \Theta \rangle\rangle} = \sqrt{\sum_{j=1}^{p_1} \sum_{k=1}^{p_2} \Theta_{jk}^2}.$$

In many settings, it is natural to consider estimating matrices that are low-rank; examples include principal component analysis, spectral clustering, collaborative filtering, and matrix completion. With certain exceptions, it is computationally expensive to enforce a rank-constraint in a direct manner, so that a variety of researchers have studied the *nuclear norm*, also known as the trace norm, as a surrogate for a rank constraint. More precisely, the nuclear norm is given by

$$(12) \quad \|\Theta\|_{\text{nuc}} := \sum_{j=1}^{\min\{p_1, p_2\}} \sigma_j(\Theta),$$

where $\{\sigma_j(\Theta)\}$ are the singular values of the matrix Θ .

The nuclear norm is decomposable with respect to appropriately chosen subspaces. Let us consider the class of matrices $\Theta \in \mathbb{R}^{p_1 \times p_2}$ that have rank $r \leq \min\{p_1, p_2\}$. For any given matrix Θ , we let $\text{row}(\Theta) \subseteq \mathbb{R}^{p_2}$ and $\text{col}(\Theta) \subseteq \mathbb{R}^{p_1}$ denote its row space and column space respectively. Let U and V be a given pair of r -dimensional subspaces $U \subseteq \mathbb{R}^{p_1}$ and $V \subseteq \mathbb{R}^{p_2}$; these subspaces will represent left and right singular vectors of the target matrix Θ^* to be estimated. For a given pair (U, V) , we can define the subspaces $\mathcal{M}(U, V)$ and $\overline{\mathcal{M}}^\perp(U, V)$ of $\mathbb{R}^{p_1 \times p_2}$ given by

$$(13a) \quad \mathcal{M}(U, V) := \{\Theta \in \mathbb{R}^{p_1 \times p_2} \mid \text{row}(\Theta) \subseteq V, \text{col}(\Theta) \subseteq U\}, \quad \text{and}$$

$$(13b) \quad \overline{\mathcal{M}}^\perp(U, V) := \{\Theta \in \mathbb{R}^{p_1 \times p_2} \mid \text{row}(\Theta) \subseteq V^\perp, \text{col}(\Theta) \subseteq U^\perp\}.$$

So as to simplify notation, we omit the indices (U, V) when they are clear from context. Unlike the preceding examples, in this case, the set \mathcal{M} is not² equal to $\overline{\mathcal{M}}$.

Finally, we claim that the nuclear norm is decomposable with respect to the pair $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$. By construction, any pair of matrices $\Theta \in \mathcal{M}$ and $\Gamma \in \overline{\mathcal{M}}^\perp$ have orthogonal row and column spaces, which implies the required decomposability condition—namely $\|\Theta + \Gamma\|_1 = \|\Theta\|_1 + \|\Gamma\|_1$. \diamond

A line of recent work (e.g., [19, 27, 18, 2, 26, 43]) has studied matrix problems involving the sum of a low-rank matrix with a sparse matrix, along with the regularizer formed by a weighted sum of the nuclear norm and the element-wise ℓ_1 -norm. By a combination of Examples 1 and Example 3, this regularizer also satisfies the decomposability property with respect to appropriately defined subspaces.

2.3 A key consequence of decomposability

Thus far, we have specified a class (1) of M -estimators based on regularization, defined the notion of decomposability for the regularizer and worked through several illustrative examples. We now turn to the statistical consequences of decomposability—more specifically, its implications for the error vector $\widehat{\Delta}_{\lambda_n} = \widehat{\theta}_{\lambda_n} - \theta^*$, where $\widehat{\theta} \in \mathbb{R}^p$ is any solution of the regularized M -estimation procedure (1). For a given inner product $\langle \cdot, \cdot \rangle$, the dual norm of \mathcal{R} is given by

$$(14) \quad \mathcal{R}^*(v) := \sup_{u \in \mathbb{R}^p \setminus \{0\}} \frac{\langle u, v \rangle}{\mathcal{R}(u)} = \sup_{\mathcal{R}(u) \leq 1} \langle u, v \rangle.$$

This notion is best understood by working through some examples.

Dual of ℓ_1 -norm: For the ℓ_1 -norm $\mathcal{R}(u) = \|u\|_1$ previously discussed in Example 1, let us compute its dual norm with respect to the Euclidean inner product on \mathbb{R}^p . For any vector $v \in \mathbb{R}^p$, we have

$$\sup_{\|u\|_1 \leq 1} \langle u, v \rangle \leq \sup_{\|u\|_1 \leq 1} \sum_{k=1}^p |u_k| |v_k| \leq \sup_{\|u\|_1 \leq 1} \left(\sum_{k=1}^p |u_k| \right) \max_{k=1, \dots, p} |v_k| = \|v\|_\infty.$$

²However, as is required by our theory, we do have the inclusion $\mathcal{M} \subseteq \overline{\mathcal{M}}$. Indeed, given any $\Theta \in \mathcal{M}$ and $\Gamma \in \overline{\mathcal{M}}^\perp$, we have $\Theta^T \Gamma = 0$ by definition, which implies that $\langle \Theta, \Gamma \rangle = \text{trace}(\Theta^T \Gamma) = 0$. Since $\Gamma \in \overline{\mathcal{M}}^\perp$ was arbitrary, we have shown that Θ is orthogonal to the space $\overline{\mathcal{M}}^\perp$, meaning that it must belong to $\overline{\mathcal{M}}$.

We claim that this upper bound actually holds with equality. In particular, letting j be any index for which $|v_j|$ achieves the maximum $\|v\|_\infty = \max_{k=1,\dots,p} |v_k|$, suppose that we form a vector $\bar{u} \in \mathbb{R}^p$ with $\bar{u}_j = \text{sign}(v_j)$, and $\bar{u}_k = 0$ for all $k \neq j$. With this choice, we have $\|\bar{u}\|_1 \leq 1$, and hence

$$\sup_{\|u\|_1 \leq 1} \langle u, v \rangle \geq \sum_{k=1}^p \bar{u}_k v_k = \|v\|_\infty,$$

showing that the dual of the ℓ_1 -norm is the ℓ_∞ -norm.

Dual of group norm: Now recall the group norm from Example 2, specified in terms of a vector $\vec{\alpha} \in [2, \infty]^{N_G}$. A similar calculation shows that its dual norm, again with respect to the Euclidean norm on \mathbb{R}^p , is given by

$$(15) \quad \|v\|_{\mathcal{G}, \vec{\alpha}^*} = \max_{t=1,\dots,N_G} \|v\|_{\alpha_t^*} \quad \text{where } \frac{1}{\alpha_t} + \frac{1}{\alpha_t^*} = 1 \text{ are dual exponents.}$$

As special cases of this general duality relation, the block $(1, 2)$ norm that underlies the usual group Lasso leads to a block $(\infty, 2)$ norm as the dual, whereas the the block $(1, \infty)$ norm leads to a block $(\infty, 1)$ norm as the dual.

Dual of nuclear norm: For the nuclear norm, the dual is defined with respect to the trace inner product on the space of matrices. For any matrix $N \in \mathbb{R}^{p_1 \times p_2}$, it can be shown that

$$\mathcal{R}^*(N) = \sup_{\|M\|_{\text{nuc}} \leq 1} \langle M, N \rangle = \|N\|_{\text{op}} = \max_{j=1,\dots,\min\{p_1,p_2\}} \sigma_j(N),$$

corresponding to the ℓ_∞ -norm applied to the vector $\sigma(N)$ of singular values. In the special case of diagonal matrices, this fact reduces to the dual relationship between the vector ℓ_1 and ℓ_∞ norms.

The dual norm plays a key role in our general theory, in particular by specifying a suitable choice of the regularization weight λ_n . We summarize in the following:

LEMMA 1. *Suppose that \mathcal{L} is a convex and differentiable function, and consider any optimal solution $\hat{\theta}$ to the optimization problem (1) with a strictly positive regularization parameter satisfying*

$$(16) \quad \lambda_n \geq 2\mathcal{R}^*(\nabla\mathcal{L}(\theta^*; Z_1^n)).$$

Then for any pair $(\mathcal{M}, \bar{\mathcal{M}}^\perp)$ over which \mathcal{R} is decomposable, the error $\hat{\Delta} = \hat{\theta}_{\lambda_n} - \theta^$ belongs to the set*

$$(17) \quad \mathbb{C}(\mathcal{M}, \bar{\mathcal{M}}^\perp; \theta^*) := \{\Delta \in \mathbb{R}^p \mid \mathcal{R}(\Delta_{\bar{\mathcal{M}}^\perp}) \leq 3\mathcal{R}(\Delta_{\bar{\mathcal{M}}}) + 4\mathcal{R}(\theta^*_{\bar{\mathcal{M}}^\perp})\}.$$

We prove this result in the supplementary appendix [51]. It has the following important consequence: for any decomposable regularizer and an appropriate choice (16) of regularization parameter, we are guaranteed that the error vector $\hat{\Delta}$ belongs to a very specific set, depending on the unknown vector θ^* . As illustrated in Figure 1, the geometry of the set \mathbb{C} depends on the relation between θ^* and the model subspace \mathcal{M} . When $\theta^* \in \mathcal{M}$, then we are guaranteed that $\mathcal{R}(\theta^*_{\bar{\mathcal{M}}^\perp}) = 0$. In this case, the constraint (17) reduces to $\mathcal{R}(\Delta_{\bar{\mathcal{M}}^\perp}) \leq 3\mathcal{R}(\Delta_{\bar{\mathcal{M}}})$, so that \mathbb{C} is a cone, as illustrated in panel (a). In the more general case when $\theta^* \notin \mathcal{M}$ so that $\mathcal{R}(\theta^*_{\bar{\mathcal{M}}^\perp}) \neq 0$, the set \mathbb{C} is *not* a cone, but rather a star-shaped set (panel (b)). As will be clarified in the sequel, the case $\theta^* \notin \mathcal{M}$ requires a more delicate treatment.

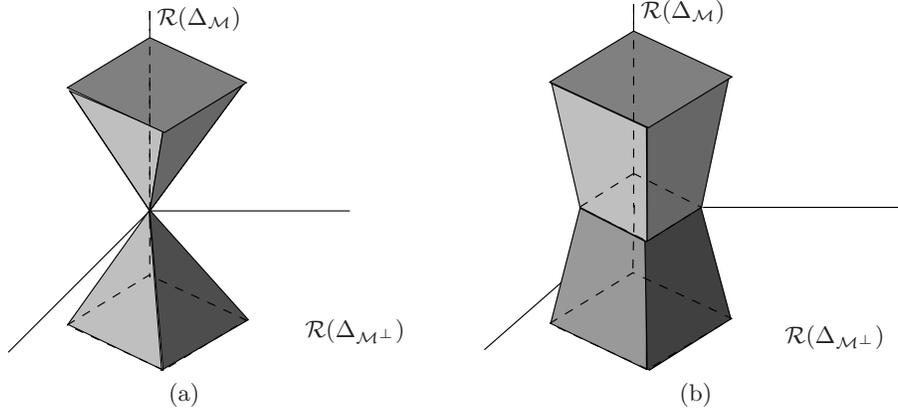


Fig 1. Illustration of the set $\mathbb{C}(\mathcal{M}, \mathcal{M}^\perp; \theta^*)$ in the special case $\Delta = (\Delta_1, \Delta_2, \Delta_3) \in \mathbb{R}^3$ and regularizer $\mathcal{R}(\Delta) = \|\Delta\|_1$, relevant for sparse vectors (Example 1). This picture shows the case $S = \{3\}$, so that the model subspace is $\mathcal{M}(S) = \{\Delta \in \mathbb{R}^3 \mid \Delta_1 = \Delta_2 = 0\}$, and its orthogonal complement is given by $\mathcal{M}^\perp(S) = \{\Delta \in \mathbb{R}^3 \mid \Delta_3 = 0\}$. (a) In the special case when $\theta_1^* = \theta_2^* = 0$, so that $\theta^* \in \mathcal{M}$, the set $\mathbb{C}(\mathcal{M}, \mathcal{M}^\perp; \theta^*)$ is a cone. (b) When θ^* does not belong to \mathcal{M} , the set $\mathbb{C}(\mathcal{M}, \mathcal{M}^\perp; \theta^*)$ is enlarged in the co-ordinates (Δ_1, Δ_2) that span \mathcal{M}^\perp . It is no longer a cone, but is still a star-shaped set.

2.4 Restricted strong convexity

We now turn to an important requirement of the loss function, and its interaction with the statistical model. Recall that $\widehat{\Delta} = \widehat{\theta}_{\lambda_n} - \theta^*$ is the difference between an optimal solution $\widehat{\theta}_{\lambda_n}$ and the true parameter, and consider the loss difference³ $\mathcal{L}(\widehat{\theta}_{\lambda_n}) - \mathcal{L}(\theta^*)$. In the classical setting, under fairly mild conditions, one expects that the loss difference should converge to zero as the sample size n increases. It is important to note, however, that such convergence on its own is *not sufficient* to guarantee that $\widehat{\theta}_{\lambda_n}$ and θ^* are close, or equivalently that $\widehat{\Delta}$ is small. Rather, the closeness depends on the curvature of the loss function, as illustrated in Figure 2. In a desirable setting (panel (a)), the loss function is sharply curved around its optimum $\widehat{\theta}_{\lambda_n}$, so that having a small loss difference $|\mathcal{L}(\theta^*) - \mathcal{L}(\widehat{\theta}_{\lambda_n})|$ translates to a small error $\widehat{\Delta} = \widehat{\theta}_{\lambda_n} - \theta^*$. Panel (b) illustrates a less desirable setting, in which the loss function is relatively flat, so that the loss difference can be small while the error $\widehat{\Delta}$ is relatively large.

The standard way to ensure that a function is “not too flat” is via the notion of strong convexity. Since \mathcal{L} is differentiable by assumption, we may perform a first-order Taylor series expansion at θ^* , and in some direction Δ ; the error in this Taylor series is given by

$$(18) \quad \delta\mathcal{L}(\Delta, \theta^*) := \mathcal{L}(\theta^* + \Delta) - \mathcal{L}(\theta^*) - \langle \nabla\mathcal{L}(\theta^*), \Delta \rangle.$$

One way in which to enforce that \mathcal{L} is strongly convex is to require the existence of some positive constant $\kappa > 0$ such that $\delta\mathcal{L}(\Delta, \theta^*) \geq \kappa\|\Delta\|^2$ for all $\Delta \in \mathbb{R}^p$ in a neighborhood of θ^* . When the loss function is twice differentiable, strong convexity amounts to lower bound on the eigenvalues of the Hessian $\nabla^2\mathcal{L}(\theta)$, holding uniformly for all θ in a neighborhood of θ^* .

³To simplify notation, we frequently write $\mathcal{L}(\theta)$ as shorthand for $\mathcal{L}(\theta; Z_1^n)$ when the underlying data Z_1^n is clear from context.

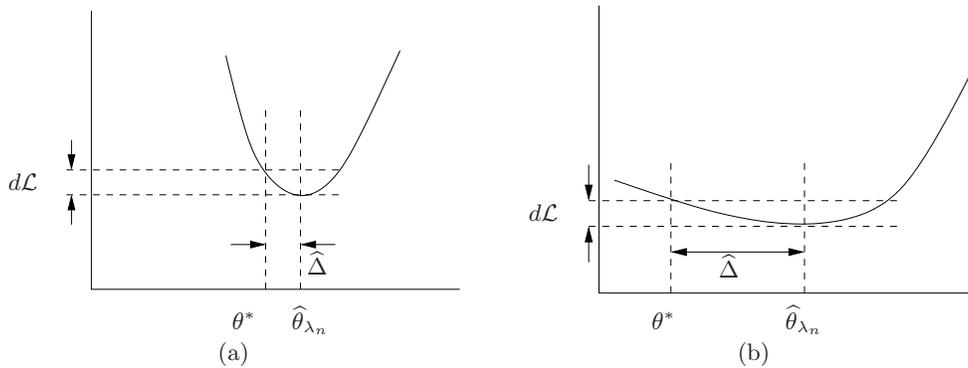


Fig 2. Role of curvature in distinguishing parameters. (a) Loss function has high curvature around $\widehat{\Delta}$. A small excess loss $d\mathcal{L} = |\mathcal{L}(\widehat{\theta}_{\lambda_n}) - \mathcal{L}(\theta^*)|$ guarantees that the parameter error $\widehat{\Delta} = \widehat{\theta}_{\lambda_n} - \theta^*$ is also small. (b) A less desirable setting, in which the loss function has relatively low curvature around the optimum.

Under classical “fixed p , large n ” scaling, the loss function will be strongly convex under mild conditions. For instance, suppose that population risk $\bar{\mathcal{L}}$ is strongly convex, or equivalently, that the Hessian $\nabla^2 \bar{\mathcal{L}}(\theta)$ is strictly positive definite in a neighborhood of θ^* . As a concrete example, when the loss function \mathcal{L} is defined based on negative log likelihood of a statistical model, then the Hessian $\nabla^2 \bar{\mathcal{L}}(\theta)$ corresponds to the Fisher information matrix, a quantity which arises naturally in asymptotic statistics. If the dimension p is fixed while the sample size n goes to infinity, standard arguments can be used to show that (under mild regularity conditions) the random Hessian $\nabla^2 \mathcal{L}(\theta)$ converges to $\nabla^2 \bar{\mathcal{L}}(\theta)$ uniformly for all θ in an open neighborhood of θ^* . In contrast, whenever the pair (n, p) both increase in such a way that $p > n$, the situation is drastically different: the Hessian matrix $\nabla^2 \mathcal{L}(\theta)$ is often singular. As a concrete example, consider linear regression based on samples $Z_i = (y_i, x_i) \in \mathbb{R} \times \mathbb{R}^p$, for $i = 1, 2, \dots, n$. Using the least-squares loss $\mathcal{L}(\theta) = \frac{1}{2n} \|y - X\theta\|_2^2$, the $p \times p$ Hessian matrix $\nabla^2 \mathcal{L}(\theta) = \frac{1}{n} X^T X$ has rank at most n , meaning that the loss cannot be strongly convex when $p > n$. Consequently, it is impossible to guarantee global strong convexity, so that we need to restrict the set of directions Δ in which we require a curvature condition.

Ultimately, the only direction of interest is given by the error vector $\widehat{\Delta} = \widehat{\theta}_{\lambda_n} - \theta^*$. Recall that Lemma 1 guarantees that, for suitable choices of the regularization parameter λ_n , this error vector must belong to the set $\mathbb{C}(\mathcal{M}, \bar{\mathcal{M}}^\perp; \theta^*)$, as previously defined (17). Consequently, it suffices to ensure that the function is strongly convex over this set, as formalized in the following:

DEFINITION 2. The loss function satisfies a **restricted strong convexity** (RSC) condition with *curvature* $\kappa_{\mathcal{L}} > 0$ and *tolerance function* $\tau_{\mathcal{L}}$ if

$$(19) \quad \delta\mathcal{L}(\Delta, \theta^*) \geq \kappa_{\mathcal{L}} \|\Delta\|^2 - \tau_{\mathcal{L}}^2(\theta^*) \quad \text{for all } \Delta \in \mathbb{C}(\mathcal{M}, \bar{\mathcal{M}}^\perp; \theta^*).$$

In the simplest of cases—in particular, when $\theta^* \in \mathcal{M}$ —there are many statistical models for which this RSC condition holds with tolerance $\tau_{\mathcal{L}}(\theta^*) = 0$. In the more general setting, it can hold only with a non-zero tolerance term, as illustrated in Figure 3(b). As our proofs will clarify, we in fact require only the lower

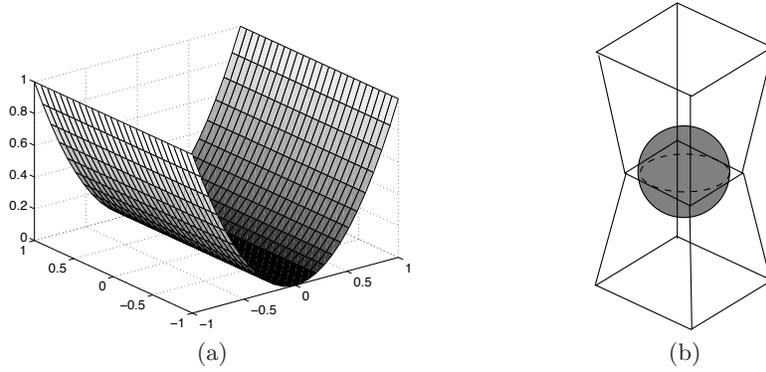


Fig 3. (a) Illustration of a generic loss function in the high-dimensional $p > n$ setting: it is curved in certain directions, but completely flat in others. (b) When $\theta^* \notin \mathcal{M}$, the set $\mathbb{C}(\mathcal{M}, \mathcal{M}^\perp; \theta^*)$ contains a ball centered at the origin, which necessitates a tolerance term $\tau_{\mathcal{L}}(\theta^*) > 0$ in the definition of restricted strong convexity.

bound (19) to hold for the intersection of \mathbb{C} with a local ball $\{\|\Delta\| \leq R\}$ of some radius centered at zero. As will be clarified later, this restriction is not necessary for the least-squares loss, but is essential for more general loss functions, such as those that arise in generalized linear models.

We will see in the sequel that for many loss functions, it is possible to prove that with high probability the first-order Taylor series error satisfies a lower bound of the form

$$(20) \quad \delta\mathcal{L}(\Delta, \theta^*) \geq \kappa_1 \|\Delta\|^2 - \kappa_2 g(n, p) \mathcal{R}^2(\Delta) \quad \text{for all } \|\Delta\| \leq 1,$$

where κ_1, κ_2 are positive constants, and $g(n, p)$ is a function of the sample size n and ambient dimension p , decreasing in the sample size. For instance, in the case of ℓ_1 -regularization, for covariates with suitably controlled tails, this type of bound holds for the least squares loss with the function $g(n, p) = \frac{\log p}{n}$; see equation (31) to follow. For generalized linear models and the ℓ_1 -norm, a similar type of bound is given in equation (43). We also provide a bound of this form for the least-squares loss group-structured norms in equation (46), with a different choice of the function g depending on the group structure.

A bound of the form (20) implies a form of restricted strong convexity as long as $\mathcal{R}(\Delta)$ is not “too large” relative to $\|\Delta\|$. In order to formalize this notion, we define a quantity that relates the error norm and the regularizer:

DEFINITION 3 (Subspace compatibility constant). For any subspace \mathcal{M} of \mathbb{R}^p , the *subspace compatibility constant* with respect to the pair $(\mathcal{R}, \|\cdot\|)$ is given by

$$(21) \quad \Psi(\mathcal{M}) := \sup_{u \in \mathcal{M} \setminus \{0\}} \frac{\mathcal{R}(u)}{\|u\|}.$$

This quantity reflects the degree of compatibility between the regularizer and the error norm over the subspace \mathcal{M} . In alternative terms, it is the Lipschitz constant of the regularizer with respect to the error norm, restricted to the subspace \mathcal{M} .

As a simple example, if \mathcal{M} is a s -dimensional co-ordinate subspace, with regularizer $\mathcal{R}(u) = \|u\|_1$ and error norm $\|u\| = \|u\|_2$, then we have $\Psi(\mathcal{M}) = \sqrt{s}$.

This compatibility constant appears explicitly in the bounds of our main theorem, and also arises in establishing restricted strong convexity. Let us now illustrate how it can be used to show that the condition (20) implies a form of restricted strong convexity. To be concrete, let us suppose that θ^* belongs to a subspace \mathcal{M} ; in this case, membership of Δ in the set $\mathbb{C}(\mathcal{M}, \overline{\mathcal{M}}^\perp; \theta^*)$ implies that $\mathcal{R}(\Delta_{\overline{\mathcal{M}}^\perp}) \leq 3\mathcal{R}(\Delta_{\overline{\mathcal{M}}})$. Consequently, by triangle inequality and the definition (21), we have

$$\mathcal{R}(\Delta) \leq \mathcal{R}(\Delta_{\overline{\mathcal{M}}^\perp}) + \mathcal{R}(\Delta_{\overline{\mathcal{M}}}) \leq 4\mathcal{R}(\Delta_{\overline{\mathcal{M}}}) \leq 4\Psi(\overline{\mathcal{M}})\|\Delta\|.$$

Therefore, whenever a bound of the form (20) holds and $\theta^* \in \mathcal{M}$, we are guaranteed that

$$\delta\mathcal{L}(\Delta, \theta^*) \geq \{\kappa_1 - 16\kappa_2\Psi^2(\overline{\mathcal{M}})g(n, p)\}\|\Delta\|^2 \quad \text{for all } \|\Delta\| \leq 1.$$

Consequently, as long as the sample size is large enough that $16\kappa_2\Psi^2(\overline{\mathcal{M}})g(n, p) < \frac{\kappa_1}{2}$, the restricted strong convexity condition will hold with $\kappa_{\mathcal{L}} = \frac{\kappa_1}{2}$ and $\tau_{\mathcal{L}}(\theta^*) = 0$. We make use of arguments of this flavor throughout this paper.

3. BOUNDS FOR GENERAL M -ESTIMATORS

We are now ready to state a general result that provides bounds and hence convergence rates for the error $\|\hat{\theta}_{\lambda_n} - \theta^*\|$, where $\hat{\theta}_{\lambda_n}$ is any optimal solution of the convex program (1). Although it may appear somewhat abstract at first sight, this result has a number of concrete and useful consequences for specific models. In particular, we recover as an immediate corollary the best known results about estimation in sparse linear models with general designs [9, 48], as well as a number of new results, including minimax-optimal rates for estimation under ℓ_q -sparsity constraints and estimation of block-structured sparse matrices. In results that we report elsewhere, we also apply these theorems to establishing results for sparse generalized linear models [50], estimation of low-rank matrices [54, 52], matrix decomposition problems [2], and sparse non-parametric regression models [59].

Let us recall our running assumptions on the structure of the convex program (1).

- (G1) The regularizer \mathcal{R} is a norm, and is decomposable with respect to the subspace pair $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$, where $\mathcal{M} \subseteq \overline{\mathcal{M}}$.
- (G2) The loss function \mathcal{L} is convex and differentiable, and satisfies restricted strong convexity with curvature $\kappa_{\mathcal{L}}$ and tolerance $\tau_{\mathcal{L}}$.

The reader should also recall the definition (21) of the subspace compatibility constant. With this notation, we can now state the main result of this paper:

THEOREM 1 (Bounds for general models). *Under conditions (G1) and (G2), consider the problem (1) based on a strictly positive regularization constant $\lambda_n \geq 2\mathcal{R}^*(\nabla\mathcal{L}(\theta^*))$. Then any optimal solution $\hat{\theta}_{\lambda_n}$ to the convex program (1) satisfies the bound*

$$(22) \quad \|\hat{\theta}_{\lambda_n} - \theta^*\|^2 \leq 9\frac{\lambda_n^2}{\kappa_{\mathcal{L}}^2}\Psi^2(\overline{\mathcal{M}}) + \frac{\lambda_n}{\kappa_{\mathcal{L}}}\{2\tau_{\mathcal{L}}^2(\theta^*) + 4\mathcal{R}(\theta_{\mathcal{M}^\perp}^*)\},$$

Remarks: Let us consider in more detail some different features of this result.

- (a) It should be noted that Theorem 1 is actually a *deterministic* statement about the set of optimizers of the convex program (1) for a fixed choice of λ_n . Although the program is convex, it need not be strictly convex, so that the global optimum might be attained at more than one point $\widehat{\theta}_{\lambda_n}$. The stated bound holds for any of these optima. Probabilistic analysis is required when Theorem 1 is applied to particular statistical models, and we need to verify that the regularizer satisfies the condition

$$(23) \quad \lambda_n \geq 2\mathcal{R}^*(\nabla\mathcal{L}(\theta^*)),$$

and that the loss satisfies the RSC condition. A challenge here is that since θ^* is unknown, it is usually impossible to compute the right-hand side of the condition (23). Instead, when we derive consequences of Theorem 1 for different statistical models, we use concentration inequalities in order to provide bounds that hold with high probability over the data.

- (b) Second, note that Theorem 1 actually provides a *family of bounds*, one for each pair $(\mathcal{M}, \overline{\mathcal{M}}^\perp)$ of subspaces for which the regularizer is decomposable. Ignoring the term involving $\tau_{\mathcal{L}}$ for the moment, for any given pair, the error bound is the sum of two terms, corresponding to estimation error \mathcal{E}_{err} and approximation error \mathcal{E}_{app} , given by (respectively)

$$(24) \quad \mathcal{E}_{\text{err}} := 9 \frac{\lambda_n^2}{\kappa_{\mathcal{L}}^2} \Psi^2(\overline{\mathcal{M}}), \quad \text{and} \quad \mathcal{E}_{\text{app}} := 4 \frac{\lambda_n}{\kappa_{\mathcal{L}}} \mathcal{R}(\theta^*_{\mathcal{M}^\perp}).$$

As the dimension of the subspace \mathcal{M} increases (so that the dimension of \mathcal{M}^\perp decreases), the approximation error tends to zero. But since $\mathcal{M} \subseteq \overline{\mathcal{M}}$, the estimation error is increasing at the same time. Thus, in the usual way, optimal rates are obtained by choosing \mathcal{M} and $\overline{\mathcal{M}}$ so as to balance these two contributions to the error. We illustrate such choices for various specific models to follow.

- (c) As will be clarified in the sequel, many high-dimensional statistical models have an unidentifiable component, and the tolerance term $\tau_{\mathcal{L}}$ reflects the degree of this non-identifiability.

A large body of past work on sparse linear regression has focused on the case of exactly sparse regression models for which the unknown regression vector θ^* is s -sparse. For this special case, recall from Example 1 in Section 2.2 that we can define an s -dimensional subspace \mathcal{M} that contains θ^* . Consequently, the associated set $\mathbb{C}(\mathcal{M}, \mathcal{M}^\perp; \theta^*)$ is a cone (see Figure 1(a)), and it is thus possible to establish that restricted strong convexity (RSC) holds with tolerance parameter $\tau_{\mathcal{L}}(\theta^*) = 0$. This same reasoning applies to other statistical models, among them group-sparse regression, in which a small subset of groups are active, as well as low-rank matrix estimation. The following corollary provides a simply stated bound that covers all of these models:

COROLLARY 1. *Suppose that, in addition to the conditions of Theorem 1, the unknown θ^* belongs to \mathcal{M} and the RSC condition holds over $\mathbb{C}(\mathcal{M}, \overline{\mathcal{M}}, \theta^*)$ with*

$\tau_{\mathcal{L}}(\theta^*) = 0$. Then any optimal solution $\widehat{\theta}_{\lambda_n}$ to the convex program (1) satisfies the bounds

$$(25a) \quad \|\widehat{\theta}_{\lambda_n} - \theta^*\| \leq 9 \frac{\lambda_n^2}{\kappa_{\mathcal{L}}} \Psi^2(\overline{\mathcal{M}}), \quad \text{and}$$

$$(25b) \quad \mathcal{R}(\widehat{\theta}_{\lambda_n} - \theta^*) \leq 12 \frac{\lambda_n}{\kappa_{\mathcal{L}}} \Psi^2(\overline{\mathcal{M}}).$$

Focusing first on the bound (25a), it consists of three terms, each of which has a natural interpretation. First, it is inversely proportional to the RSC constant $\kappa_{\mathcal{L}}$, so that higher curvature guarantees lower error, as is to be expected. The error bound grows proportionally with the subspace compatibility constant $\Psi(\overline{\mathcal{M}})$, which measures the compatibility between the regularizer \mathcal{R} and error norm $\|\cdot\|$ over the subspace $\overline{\mathcal{M}}$ (see Definition 3). This term increases with the size of subspace $\overline{\mathcal{M}}$, which contains the model subspace \mathcal{M} . Third, the bound also scales linearly with the regularization parameter λ_n , which must be strictly positive and satisfy the lower bound (23). The bound (25b) on the error measured in the regularizer norm is similar, except that it scales quadratically with the subspace compatibility constant. As the proof clarifies, this additional dependence arises since the regularizer over the subspace $\overline{\mathcal{M}}$ is larger than the norm $\|\cdot\|$ by a factor of at most $\Psi(\overline{\mathcal{M}})$ (see Definition 3).

Obtaining concrete rates using Corollary 1 requires some work in order to verify the conditions of Theorem 1, and to provide control on the three quantities in the bounds (25a) and (25b), as illustrated in the examples to follow.

4. CONVERGENCE RATES FOR SPARSE REGRESSION

As an illustration, we begin with one of the simplest statistical models, namely the standard linear model. It is based on n observations $Z_i = (x_i, y_i) \in \mathbb{R}^p \times \mathbb{R}$ of covariate-response pairs. Let $y \in \mathbb{R}^n$ denote a vector of the responses, and let $X \in \mathbb{R}^{n \times p}$ be the design matrix, where $x_i \in \mathbb{R}^p$ is the i^{th} row. This pair is linked via the linear model

$$(26) \quad y = X\theta^* + w,$$

where $\theta^* \in \mathbb{R}^p$ is the unknown regression vector, and $w \in \mathbb{R}^n$ is a noise vector. To begin, we focus on this simple linear set-up, and describe extensions to generalized models in Section 4.4.

Given the data set $Z_1^n = (y, X) \in \mathbb{R}^n \times \mathbb{R}^{n \times p}$, our goal is to obtain a “good” estimate $\widehat{\theta}$ of the regression vector θ^* , assessed either in terms of its ℓ_2 -error $\|\widehat{\theta} - \theta^*\|_2$ or its ℓ_1 -error $\|\widehat{\theta} - \theta^*\|_1$. It is worth noting that whenever $p > n$, the standard linear model (26) is unidentifiable in a certain sense, since the rectangular matrix $X \in \mathbb{R}^{n \times p}$ has a nullspace of dimension at least $p - n$. Consequently, in order to obtain an identifiable model—or at the very least, to bound the degree of non-identifiability—it is essential to impose additional constraints on the regression vector θ^* . One natural constraint is some type of sparsity in the regression vector; for instance, one might assume that θ^* has at most s non-zero coefficients, as discussed at more length in Section 4.2. More generally, one might assume that although θ^* is not exactly sparse, it can be well-approximated by a

sparse vector, in which case one might say that θ^* is “weakly sparse”, “sparsifiable” or “compressible”. Section 4.3 is devoted to a more detailed discussion of this weakly sparse case.

A natural M -estimator for this problem is the Lasso [20, 69], obtained by solving the ℓ_1 -penalized quadratic program

$$(27) \quad \hat{\theta}_{\lambda_n} \in \arg \min_{\theta \in \mathbb{R}^p} \left\{ \frac{1}{2n} \|y - X\theta\|_2^2 + \lambda_n \|\theta\|_1 \right\},$$

for some choice $\lambda_n > 0$ of regularization parameter. Note that this Lasso estimator is a particular case of the general M -estimator (1), based on the loss function and regularization pair $\mathcal{L}(\theta; Z_1^n) = \frac{1}{2n} \|y - X\theta\|_2^2$ and $\mathcal{R}(\theta) = \sum_{j=1}^p |\theta_j| = \|\theta\|_1$. We now show how Theorem 1 can be specialized to obtain bounds on the error $\hat{\theta}_{\lambda_n} - \theta^*$ for the Lasso estimate.

4.1 Restricted eigenvalues for sparse linear regression

For the least-squares loss function that underlies the Lasso, the first-order Taylor series expansion from Definition 2 is exact, so that

$$\delta\mathcal{L}(\Delta, \theta^*) = \langle \Delta, \frac{1}{n} X^T X \Delta \rangle = \frac{1}{n} \|X\Delta\|_2^2.$$

Thus, in this special case, the Taylor series error is independent of θ^* , a fact which allows for substantial theoretical simplification. More precisely, in order to establish restricted strong convexity, it suffices to establish a lower bound on $\|X\Delta\|_2^2/n$ that holds uniformly for an appropriately restricted subset of p -dimensional vectors Δ .

As previously discussed in Example 1, for any subset $S \subseteq \{1, 2, \dots, p\}$, the ℓ_1 -norm is decomposable with respect to the subspace $\mathcal{M}(S) = \{\theta \in \mathbb{R}^p \mid \theta_{S^c} = 0\}$ and its orthogonal complement. When the unknown regression vector $\theta^* \in \mathbb{R}^p$ is exactly sparse, it is natural to choose S equal to the support set of θ^* . By appropriately specializing the definition (17) of \mathbb{C} , we are led to consider the cone

$$(28) \quad \mathbb{C}(S) := \left\{ \Delta \in \mathbb{R}^p \mid \|\Delta_{S^c}\|_1 \leq 3\|\Delta_S\|_1 \right\}.$$

See Figure 1(a) for an illustration of this set in three dimensions. With this choice, restricted strong convexity with respect to the ℓ_2 -norm is equivalent to requiring that the design matrix X satisfy the condition

$$(29) \quad \frac{\|X\theta\|_2^2}{n} \geq \kappa_{\mathcal{L}} \|\theta\|_2^2 \quad \text{for all } \theta \in \mathbb{C}(S).$$

This lower bound is a type of *restricted eigenvalue* (RE) condition, and has been studied in past work on basis pursuit and the Lasso (e.g., [9, 48, 58, 73]). One could also require that a related condition hold with respect to the ℓ_1 -norm—viz.

$$(30) \quad \frac{\|X\theta\|_2^2}{n} \geq \kappa'_{\mathcal{L}} \frac{\|\theta\|_1^2}{|S|} \quad \text{for all } \theta \in \mathbb{C}(S).$$

This type of ℓ_1 -based RE condition is less restrictive than the corresponding ℓ_2 -version (29). We refer the reader to the paper by van de Geer and Bühlmann [73]

for an extensive discussion of different types of restricted eigenvalue or compatibility conditions.

It is natural to ask whether there are many matrices that satisfy these types of RE conditions. If X has i.i.d. entries following a sub-Gaussian distribution (including Gaussian and Bernoulli variables as special cases), then known results in random matrix theory imply that the restricted isometry property [16] holds with high probability, which in turn implies that the RE condition holds [9, 73]. Since statistical applications involve design matrices with substantial dependency, it is natural to ask whether an RE condition also holds for more general random designs. This question was addressed by Raskutti et al. [58, 57], who showed that if the design matrix $X \in \mathbb{R}^{n \times p}$ is formed by independently sampling each row $X_i \sim N(0, \Sigma)$, referred to as the Σ -Gaussian ensemble, then there are strictly positive constants (κ_1, κ_2) , depending only on the positive definite matrix Σ , such that

$$(31) \quad \frac{\|X\theta\|_2^2}{n} \geq \kappa_1 \|\theta\|_2^2 - \kappa_2 \frac{\log p}{n} \|\theta\|_1^2 \quad \text{for all } \theta \in \mathbb{R}^p$$

with probability greater than $1 - c_1 \exp(-c_2 n)$. The bound (31) has an important consequence: it guarantees that the RE property (29) holds⁴ with $\kappa_{\mathcal{L}} = \frac{\kappa_1}{2} > 0$ as long as $n > 64(\kappa_2/\kappa_1) s \log p$. Therefore, not only do there exist matrices satisfying the RE property (29), but any matrix sampled from a Σ -Gaussian ensemble will satisfy it with high probability. Related analysis by Rudelson and Zhou [67] extends these types of guarantees to the case of sub-Gaussian designs, also allowing for substantial dependencies among the covariates.

4.2 Lasso estimates with exact sparsity

We now show how Corollary 1 can be used to derive convergence rates for the error of the Lasso estimate when the unknown regression vector θ^* is s -sparse. In order to state these results, we require some additional notation. Using $X_j \in \mathbb{R}^n$ to denote the j^{th} column of X , we say that X is *column-normalized* if

$$(32) \quad \frac{\|X_j\|_2}{\sqrt{n}} \leq 1 \quad \text{for all } j = 1, 2, \dots, p.$$

Here we have set the upper bound to one in order to simplify notation. This particular choice entails no loss of generality, since we can always rescale the linear model appropriately (including the observation noise variance) so that it holds.

In addition, we assume that the noise vector $w \in \mathbb{R}^n$ is zero-mean and has *sub-Gaussian tails*, meaning that there is a constant $\sigma > 0$ such that for any fixed $\|v\|_2 = 1$,

$$(33) \quad \mathbb{P}[|\langle v, w \rangle| \geq t] \leq 2 \exp\left(-\frac{\delta^2}{2\sigma^2}\right) \quad \text{for all } \delta > 0.$$

For instance, this condition holds when the noise vector w has i.i.d. $N(0, 1)$ entries, or consists of independent bounded random variables. Under these conditions, we recover as a corollary of Theorem 1 the following result:

⁴To see this fact, note that for any $\theta \in \mathbb{C}(S)$, we have $\|\theta\|_1 \leq 4\|\theta_S\|_1 \leq 4\sqrt{s}\|\theta_S\|_2$. Given the lower bound (31), for any $\theta \in \mathbb{C}(S)$, we have the lower bound $\frac{\|X\theta\|_2}{\sqrt{n}} \geq \{\kappa_1 - 4\kappa_2\sqrt{\frac{s \log p}{n}}\} \|\theta\|_2 \geq \frac{\kappa_1}{2} \|\theta\|_2$, where final inequality follows as long as $n > 64(\kappa_2/\kappa_1)^2 s \log p$.

COROLLARY 2. Consider an s -sparse instance of the linear regression model (26) such that X satisfies the RE condition (29), and the column normalization condition (32). Given the Lasso program (27) with regularization parameter $\lambda_n = 4\sigma\sqrt{\frac{\log p}{n}}$, then with probability at least $1 - c_1 \exp(-c_2 n \lambda_n^2)$, any optimal solution $\hat{\theta}_{\lambda_n}$ satisfies the bounds

$$(34) \quad \|\hat{\theta}_{\lambda_n} - \theta^*\|_2^2 \leq \frac{64\sigma^2}{\kappa_{\mathcal{L}}^2} \frac{s \log p}{n}, \quad \text{and} \quad \|\hat{\theta}_{\lambda_n} - \theta^*\|_1 \leq \frac{24\sigma}{\kappa_{\mathcal{L}}} s \sqrt{\frac{\log p}{n}}.$$

Although error bounds of this form are known from past work (e.g., [9, 16, 48]), our proof illuminates the underlying structure that leads to the different terms in the bound—in particular, see equations (25a) and (25b) in the statement of Corollary 1.

PROOF. We first note that the RE condition (30) implies that RSC holds with respect to the subspace $\mathcal{M}(S)$. As discussed in Example 1, the ℓ_1 -norm is decomposable with respect to $\mathcal{M}(S)$ and its orthogonal complement, so that we may set $\overline{\mathcal{M}}(S) = \mathcal{M}(S)$. Since any vector $\theta \in \mathcal{M}(S)$ has at most s non-zero entries, the subspace compatibility constant is given by $\Psi(\mathcal{M}(S)) = \sup_{\theta \in \mathcal{M}(S) \setminus \{0\}} \frac{\|\theta\|_1}{\|\theta\|_2} = \sqrt{s}$.

The final step is to compute an appropriate choice of the regularization parameter. The gradient of the quadratic loss is given by $\nabla \mathcal{L}(\theta; (y, X)) = \frac{1}{n} X^T w$, whereas the dual norm of the ℓ_1 -norm is the ℓ_∞ -norm. Consequently, we need to specify a choice of $\lambda_n > 0$ such that

$$\lambda_n \geq 2 \mathcal{R}^*(\nabla \mathcal{L}(\theta^*)) = 2 \left\| \frac{1}{n} X^T w \right\|_\infty$$

with high probability. Using the column normalization (32) and sub-Gaussian (33) conditions, for each $j = 1, \dots, p$, we have the tail bound $\mathbb{P}[|\langle X_j, w \rangle / n| \geq t] \leq 2 \exp(-\frac{nt^2}{2\sigma^2})$. Consequently, by union bound, we conclude that $\mathbb{P}[\|X^T w / n\|_\infty \geq t] \leq 2 \exp(-\frac{nt^2}{2\sigma^2} + \log p)$. Setting $t^2 = \frac{4\sigma^2 \log p}{n}$, we see that the choice of λ_n given in the statement is valid with probability at least $1 - c_1 \exp(-c_2 n \lambda_n^2)$. Consequently, the claims (34) follow from the bounds (25a) and (25b) in Corollary 1. \square

4.3 Lasso estimates with weakly sparse models

We now consider regression models for which θ^* is not exactly sparse, but rather can be approximated well by a sparse vector. One way in which to formalize this notion is by considering the ℓ_q “ball” of radius R_q , given by

$$\mathbb{B}_q(R_q) := \left\{ \theta \in \mathbb{R}^p \mid \sum_{i=1}^p |\theta_i|^q \leq R_q \right\}, \quad \text{where } q \in [0, 1] \text{ is fixed.}$$

In the special case $q = 0$, this set corresponds to an exact sparsity constraint—that is, $\theta^* \in \mathbb{B}_0(R_0)$ if and only if θ^* has at most R_0 non-zero entries. More generally, for $q \in (0, 1]$, the set $\mathbb{B}_q(R_q)$ enforces a certain decay rate on the ordered absolute values of θ^* .

In the case of weakly sparse vectors, the constraint set \mathbb{C} takes the form

$$(35) \quad \mathbb{C}(\mathcal{M}, \overline{\mathcal{M}}; \theta^*) = \{\Delta \in \mathbb{R}^p \mid \|\Delta_{S^c}\|_1 \leq 3\|\Delta_S\|_1 + 4\|\theta_{S^c}^*\|_1\}.$$

In contrast to the case of exact sparsity, the set \mathbb{C} is no longer a cone, but rather contains a ball centered at the origin—compare panels (a) and (b) of Figure 1. As a consequence, it is *never* possible to ensure that $\|X\theta\|_2/\sqrt{n}$ is uniformly bounded from below for all vectors θ in the set (35), and so a strictly positive tolerance term $\tau_{\mathcal{L}}(\theta^*) > 0$ is required. The random matrix result (31), stated in the previous section, allows us to establish a form of RSC that is appropriate for the setting of ℓ_q -ball sparsity. We summarize our conclusions in the following:

COROLLARY 3. *Suppose that X satisfies the RE condition (31) as well as the column normalization condition (32), the noise w is sub-Gaussian (33), and θ^* belongs to $\mathbb{B}_q(R_q)$ for a radius R_q such that $\sqrt{R_q} \left(\frac{\log p}{n}\right)^{\frac{1}{2}-\frac{q}{4}} \leq 1$. Then if we solve the Lasso with regularization parameter $\lambda_n = 4\sigma\sqrt{\frac{\log p}{n}}$, there are universal positive constants (c_0, c_1, c_2) such that any optimal solution $\widehat{\theta}_{\lambda_n}$ satisfies*

$$(36) \quad \|\widehat{\theta}_{\lambda_n} - \theta^*\|_2^2 \leq c_0 R_q \left(\frac{\sigma^2 \log p}{\kappa_1^2 n}\right)^{1-\frac{q}{2}}$$

with probability at least $1 - c_1 \exp(-c_2 n \lambda_n^2)$.

Remarks: Note that this corollary is a strict generalization of Corollary 2, to which it reduces when $q = 0$. More generally, the parameter $q \in [0, 1]$ controls the relative “sparsifiability” of θ^* , with larger values corresponding to lesser sparsity. Naturally then, the rate slows down as q increases from 0 towards 1. In fact, Raskutti et al. [58] show that the rates (36) are minimax-optimal over the ℓ_q -balls—implying that not only are the consequences of Theorem 1 sharp for the Lasso, but more generally, no algorithm can achieve faster rates.

PROOF. Since the loss function \mathcal{L} is quadratic, the proof of Corollary 2 shows that the stated choice $\lambda_n = 4\sqrt{\frac{\sigma^2 \log p}{n}}$ is valid with probability at least $1 - c \exp(-c' n \lambda_n^2)$. Let us now show that the RSC condition holds. We do so via condition (31) applied to equation (35). For a threshold $\eta > 0$ to be chosen, define the thresholded subset

$$(37) \quad S_\eta := \{j \in \{1, 2, \dots, p\} \mid |\theta_j^*| > \eta\}.$$

Now recall the subspaces $\mathcal{M}(S_\eta)$ and $\mathcal{M}^\perp(S_\eta)$ previously defined in equations (5) and (6) of Example 1, where we set $S = S_\eta$. The following lemma, proved in the supplement [51], provides sufficient conditions for restricted strong convexity with respect to these subspace pairs:

LEMMA 2. *Suppose that the conditions of Corollary 3 hold, and $n > 9\kappa_2|S_\eta|\log p$. Then with the choice $\eta = \frac{\lambda_n}{\kappa_1}$, the RSC condition holds over $\mathbb{C}(\mathcal{M}(S_\eta), \mathcal{M}^\perp(S_\eta), \theta^*)$ with $\kappa_{\mathcal{L}} = \kappa_1/4$ and $\tau_{\mathcal{L}}^2 = 8\kappa_2\frac{\log p}{n}\|\theta_{S_\eta^c}^*\|_1^2$.*

Consequently, we may apply Theorem 1 with $\kappa_{\mathcal{L}} = \kappa_1/4$ and $\tau_{\mathcal{L}}^2(\theta^*) = 8\kappa_2 \frac{\log p}{n} \|\theta_{S_\eta^c}^*\|_1^2$ to conclude that

$$(38) \quad \|\widehat{\theta}_{\lambda_n} - \theta^*\|_2^2 \leq 144 \frac{\lambda_n^2}{\kappa_1^2} |S_\eta| + \frac{4\lambda_n}{\kappa_1} \left\{ 16\kappa_2 \frac{\log p}{n} \|\theta_{S_\eta^c}^*\|_1^2 + 4\|\theta_{S_\eta^c}^*\|_1 \right\},$$

where we have used the fact that $\Psi^2(S_\eta) = |S_\eta|$, as noted in the proof of Corollary 2.

It remains to upper bound the cardinality of S_η in terms of the threshold η and ℓ_q -ball radius R_q . Note that we have

$$(39) \quad R_q \geq \sum_{j=1}^p |\theta_j^*|^q \geq \sum_{j \in S_\eta} |\theta_j^*|^q \geq \eta^q |S_\eta|,$$

whence $|S_\eta| \leq \eta^{-q} R_q$ for any $\eta > 0$. Next we upper bound the approximation error $\|\theta_{S_\eta^c}^*\|_1$, using the fact that $\theta^* \in \mathbb{B}_q(R_q)$. Letting S_η^c denote the complementary set $S_\eta \setminus \{1, 2, \dots, p\}$, we have

$$(40) \quad \|\theta_{S_\eta^c}^*\|_1 = \sum_{j \in S_\eta^c} |\theta_j^*| = \sum_{j \in S_\eta^c} |\theta_j^*|^q |\theta_j^*|^{1-q} \leq R_q \eta^{1-q}.$$

Setting $\eta = \lambda_n/\kappa_1$ and then substituting the bounds (39) and (40) into the bound (38) yields

$$\|\widehat{\theta}_{\lambda_n} - \theta^*\|_2^2 \leq 160 \left(\frac{\lambda_n^2}{\kappa_1^2}\right)^{1-\frac{q}{2}} R_q + 64\kappa_2 \left\{ \left(\frac{\lambda_n^2}{\kappa_1^2}\right)^{1-\frac{q}{2}} R_q \right\}^2 \frac{(\log p)/n}{\lambda_n/\kappa_1}.$$

For any fixed noise variance, our choice of regularization parameter ensures that the ratio $\frac{(\log p)/n}{\lambda_n/\kappa_1}$ is of order one, so that the claim follows. \square

4.4 Extensions to generalized linear models

In this section, we briefly outline extensions of the preceding results to the family of generalized linear models (GLM). Suppose that conditioned on a vector $x \in \mathbb{R}^p$ of covariates, a response variable $y \in \mathcal{Y}$ has the distribution

$$(41) \quad \mathbb{P}_{\theta^*}(y \mid x) \propto \exp \left\{ \frac{y \langle \theta^*, x \rangle - \Phi(\langle \theta^*, x \rangle)}{c(\sigma)} \right\}.$$

Here the quantity $c(\sigma)$ is a fixed and known scale parameter, and the function $\Phi : \mathbb{R} \rightarrow \mathbb{R}$ is the link function, also known. The family (41) includes many well-known classes of regression models as special cases, including ordinary linear regression (obtained with $\mathcal{Y} = \mathbb{R}$, $\Phi(t) = t^2/2$ and $c(\sigma) = \sigma^2$), and logistic regression (obtained with $\mathcal{Y} = \{0, 1\}$, $c(\sigma) = 1$ and $\Phi(t) = \log(1 + \exp(t))$).

Given samples $Z_i = (x_i, y_i) \in \mathbb{R}^p \times \mathcal{Y}$, the goal is to estimate the unknown vector $\theta^* \in \mathbb{R}^p$. Under a sparsity assumption on θ^* , a natural estimator is based on minimizing the (negative) log likelihood, combined with an ℓ_1 -regularization term. This combination leads to the convex program

$$(42) \quad \widehat{\theta}_{\lambda_n} \in \arg \min_{\theta \in \mathbb{R}^p} \left\{ \underbrace{\frac{1}{n} \sum_{i=1}^n \left\{ -y_i \langle \theta, x_i \rangle + \Phi(\langle \theta, x_i \rangle) \right\}}_{\mathcal{L}(\theta; Z_1^n)} + \lambda_n \|\theta\|_1 \right\}.$$

In order to extend the error bounds from the previous section, a key ingredient is to establish that this GLM-based loss function satisfies a form of restricted strong convexity. Along these lines, Negahban et al. [50] proved the following result: suppose that the covariate vectors x_i are zero-mean with covariance matrix $\Sigma \succ 0$, and are drawn i.i.d. from a distribution with sub-Gaussian tails (see equation (33)). Then there are constants κ_1, κ_2 such that the first-order Taylor series error for the GLM-based loss (42) satisfies the lower bound

$$(43) \quad \delta\mathcal{L}(\Delta, \theta^*) \geq \kappa_1 \|\Delta\|_2^2 - \kappa_2 \frac{\log p}{n} \|\Delta\|_1^2 \quad \text{for all } \|\Delta\|_2 \leq 1.$$

As discussed following Definition 2, this type of lower bound implies that \mathcal{L} satisfies a form of RSC, as long as the sample size scales as $n = \Omega(s \log p)$, where s is the target sparsity. Consequently, this lower bound (43) allows us to recover analogous bounds on the error $\|\hat{\theta}_{\lambda_n} - \theta^*\|_2$ of the GLM-based estimator (42).

5. CONVERGENCE RATES FOR GROUP-STRUCTURED NORMS

The preceding two sections addressed M -estimators based on ℓ_1 -regularization, the simplest type of decomposable regularizer. We now turn to some extensions of our results to more complex regularizers that are also decomposable. Various researchers have proposed extensions of the Lasso based on regularizers that have more structure than the ℓ_1 norm (e.g., [72, 79, 81, 46, 6]). Such regularizers allow one to impose different types of block-sparsity constraints, in which groups of parameters are assumed to be active (or inactive) simultaneously. These norms arise in the context of multivariate regression, where the goal is to predict a multivariate output in \mathbb{R}^m on the basis of a set of p covariates. Here it is appropriate to assume that groups of covariates are useful for predicting the different elements of the m -dimensional output vector. We refer the reader to the papers [72, 79, 81, 46, 6] for further discussion of and motivation for the use of block-structured norms.

Given a collection $\mathcal{G} = \{G_1, \dots, G_{N_{\mathcal{G}}}\}$ of groups, recall from Example 2 in Section 2.2 the definition of the group norm $\|\cdot\|_{\mathcal{G}, \vec{\alpha}}$. In full generality, this group norm is based on a weight vector $\vec{\alpha} = (\alpha_1, \dots, \alpha_{N_{\mathcal{G}}}) \in [2, \infty]^{N_{\mathcal{G}}}$, one for each group. For simplicity, here we consider the case when $\alpha_t = \alpha$ for all $t = 1, 2, \dots, N_{\mathcal{G}}$, and we use $\|\cdot\|_{\mathcal{G}, \alpha}$ to denote the associated group norm. As a natural extension of the Lasso, we consider the *block Lasso* estimator

$$(44) \quad \hat{\theta} \in \arg \min_{\theta \in \mathbb{R}^p} \left\{ \frac{1}{n} \|y - X\theta\|_2^2 + \lambda_n \|\theta\|_{\mathcal{G}, \alpha} \right\},$$

where $\lambda_n > 0$ is a user-defined regularization parameter. Different choices of the parameter α yield different estimators, and in this section, we consider the range $\alpha \in [2, \infty]$. This range covers the two most commonly applied choices, $\alpha = 2$, often referred to as the group Lasso, as well as the choice $\alpha = +\infty$.

5.1 Restricted strong convexity for group sparsity

As a parallel to our analysis of ordinary sparse regression, our first step is to provide a condition sufficient to guarantee restricted strong convexity for the group-sparse setting. More specifically, we state the natural extension of condition (31) to the block-sparse setting, and prove that it holds with high probability for the class of Σ -Gaussian random designs. Recall from Theorem 1 that the

dual norm of the regularizer plays a central role. As discussed previously, for the block- $(1, \alpha)$ -regularizer, the associated dual norm is a block- (∞, α^*) norm, where (α, α^*) are conjugate exponents satisfying $\frac{1}{\alpha} + \frac{1}{\alpha^*} = 1$.

Letting $\varepsilon \sim N(0, I_{p \times p})$ be a standard normal vector, we consider the following condition. Suppose that there are strictly positive constants (κ_1, κ_2) such that, for all $\Delta \in \mathbb{R}^p$, we have

$$(45) \quad \frac{\|X\Delta\|_2^2}{n} \geq \kappa_1 \|\Delta\|_2^2 - \kappa_2 \rho_{\mathcal{G}}^2(\alpha^*) \|\Delta\|_{1,\alpha}^2,$$

where $\rho_{\mathcal{G}}(\alpha^*) := \mathbb{E} \left[\max_{t=1,2,\dots,N_{\mathcal{G}}} \frac{\|\varepsilon_{G_t}\|_{\alpha^*}}{\sqrt{n}} \right]$. To understand this condition, first consider the special case of $N_{\mathcal{G}} = p$ groups, each of size one, so that the group-sparse norm reduces to the ordinary ℓ_1 -norm, and its dual is the ℓ_{∞} -norm. Using $\alpha = 2$ for concreteness, we have $\rho_{\mathcal{G}}(2) = \mathbb{E}[\|\varepsilon\|_{\infty}]/\sqrt{n} \leq \sqrt{\frac{3 \log p}{n}}$ for all $p \geq 10$, using standard bounds on Gaussian maxima. Therefore, condition (45) reduces to the earlier condition (31) in this special case.

Let us consider a more general setting, say with $\alpha = 2$ and $N_{\mathcal{G}}$ groups each of size m , so that $p = N_{\mathcal{G}}m$. For this choice of groups and norm, we have

$$\rho_{\mathcal{G}}(2) = \mathbb{E} \left[\max_{t=1,\dots,N_{\mathcal{G}}} \frac{\|\varepsilon_{G_t}\|_2}{\sqrt{n}} \right],$$

where each sub-vector w_{G_t} is a standard Gaussian vector with m elements. Since $\mathbb{E}[\|\varepsilon_{G_t}\|_2] \leq \sqrt{m}$, tail bounds for χ^2 -variables yield $\rho_{\mathcal{G}}(2) \leq \sqrt{\frac{m}{n}} + \sqrt{\frac{3 \log N_{\mathcal{G}}}{n}}$, so that the condition (45) is equivalent to

$$(46) \quad \frac{\|X\Delta\|_2^2}{n} \geq \kappa_1 \|\Delta\|_2^2 - \kappa_2 \left[\sqrt{\frac{m}{n}} + \sqrt{\frac{3 \log N_{\mathcal{G}}}{n}} \right]^2 \|\Delta\|_{\mathcal{G},2}^2 \quad \text{for all } \Delta \in \mathbb{R}^p.$$

Thus far, we have seen the form that condition (45) takes for different choices of the groups and parameter α . It is natural to ask whether there are any matrices that satisfy the condition (45). As shown in the following result, the answer is affirmative—more strongly, almost every matrix satisfied from the Σ -Gaussian ensemble will satisfy this condition with high probability. (Here we recall that for a non-degenerate covariance matrix, a random design matrix $X \in \mathbb{R}^{n \times p}$ is drawn from the Σ -Gaussian ensemble if each row $x_i \sim N(0, \Sigma)$, i.i.d. for $i = 1, 2, \dots, n$.)

PROPOSITION 1. *For a design matrix $X \in \mathbb{R}^{n \times p}$ from the Σ -ensemble, there are constants (κ_1, κ_2) depending only Σ such that condition (45) holds with probability greater than $1 - c_1 \exp(-c_2 n)$.*

We provide the proof of this result in the supplement [51]. This condition can be used to show that appropriate forms of RSC hold, for both the cases of exactly group-sparse and weakly sparse vectors. As with ℓ_1 -regularization, these RSC conditions are milder than analogous group-based RIP conditions (e.g., [28, 68, 6]), which require that all sub-matrices up to a certain size are close to isometries.

5.2 Convergence rates

Apart from RSC, we impose one additional condition on the design matrix. For a given group G of size m , let us view the matrix $X_G \in \mathbb{R}^{n \times m}$ as an operator from $\ell_\alpha^m \rightarrow \ell_2^n$, and define the associated operator norm $\|X_G\|_{\alpha \rightarrow 2} := \max_{\|\theta\|_\alpha=1} \|X_G \theta\|_2$.

We then require that

$$(47) \quad \frac{\|X_{G_t}\|_{\alpha \rightarrow 2}}{\sqrt{n}} \leq 1 \quad \text{for all } t = 1, 2, \dots, N_G.$$

Note that this is a natural generalization of the column normalization condition (32), to which it reduces when we have $N_G = p$ groups, each of size one. As before, we may assume without loss of generality, rescaling X and the noise as necessary, that condition (47) holds with constant one. Finally, we define the maximum group size $m = \max_{t=1, \dots, N_G} |G_t|$. With this notation, we have the following novel result:

COROLLARY 4. *Suppose that the noise w is sub-Gaussian (33), and the design matrix X satisfies condition (45) and the block normalization condition (47). If we solve the group Lasso with*

$$(48) \quad \lambda_n \geq 2\sigma \left\{ \frac{m^{1-1/\alpha}}{\sqrt{n}} + \sqrt{\frac{\log N_G}{n}} \right\},$$

then with probability at least $1 - 2/N_G^2$, for any group subset $S_G \subseteq \{1, 2, \dots, N_G\}$ with cardinality $|S_G| = s_G$, any optimal solution $\hat{\theta}_{\lambda_n}$ satisfies

$$(49) \quad \|\hat{\theta}_{\lambda_n} - \theta^*\|_2^2 \leq \frac{4\lambda_n^2}{\kappa_{\mathcal{L}}^2} s_G + \frac{4\lambda_n}{\kappa_{\mathcal{L}}} \sum_{t \notin S_G} \|\theta_{G_t}^*\|_\alpha.$$

Remarks: Since the result applies to any $\alpha \in [2, \infty]$, we can observe how the choices of different group-sparse norms affect the convergence rates. So as to simplify this discussion, let us assume that the groups are all of equal size m , so that $p = mN_G$ is the ambient dimension of the problem.

Case $\alpha = 2$: The case $\alpha = 2$ corresponds to the block $(1, 2)$ norm, and the resulting estimator is frequently referred to as the group Lasso. For this case, we can set the regularization parameter as $\lambda_n = 2\sigma \left\{ \sqrt{\frac{m}{n}} + \sqrt{\frac{\log N_G}{n}} \right\}$. If we assume moreover that θ^* is exactly group-sparse, say supported on a group subset $S_G \subseteq \{1, 2, \dots, N_G\}$ of cardinality s_G , then the bound (49) takes the form

$$(50) \quad \|\hat{\theta} - \theta^*\|_2^2 \lesssim \frac{s_G m}{n} + \frac{s_G \log N_G}{n}.$$

Similar bounds were derived in independent work by Lounici et al. [41] and Huang and Zhang [28] for this special case of exact block sparsity. The analysis here shows how the different terms arise, in particular via the noise magnitude measured in the dual norm of the block regularizer.

In the more general setting of weak block sparsity, Corollary 4 yields a number of novel results. For instance, for a given set of groups \mathcal{G} , we can consider the block sparse analog of the ℓ_q -“ball”—namely the set

$$\mathbb{B}_q(R_q; \mathcal{G}, 2) := \left\{ \theta \in \mathbb{R}^p \mid \sum_{t=1}^{N_{\mathcal{G}}} \|\theta_{G_t}\|_2^q \leq R_q \right\}.$$

In this case, if we optimize the choice of S in the bound (49) so as to trade off the estimation and approximation errors, then we obtain

$$\|\hat{\theta} - \theta^*\|_2^2 \lesssim R_q \left(\frac{m}{n} + \frac{\log N_{\mathcal{G}}}{n} \right)^{1-\frac{q}{2}},$$

which is a novel result. This result is a generalization of our earlier Corollary 3, to which it reduces when we have $N_{\mathcal{G}} = p$ groups each of size $m = 1$.

Case $\alpha = +\infty$: Now consider the case of ℓ_1/ℓ_∞ regularization, as suggested in past work [72]. In this case, Corollary 4 implies that $\|\hat{\theta} - \theta^*\|_2^2 \lesssim \frac{sm^2}{n} + \frac{s \log N_{\mathcal{G}}}{n}$. Similar to the case $\alpha = 2$, this bound consists of an estimation term, and a search term. The estimation term $\frac{sm^2}{n}$ is larger by a factor of m , which corresponds to amount by which an ℓ_∞ -ball in m dimensions is larger than the corresponding ℓ_2 -ball.

We provide the proof of Corollary 4 in the supplementary appendix [51]. It is based on verifying the conditions of Theorem 1: more precisely, we use Proposition 1 in order to establish RSC, and we provide a lemma that shows that the regularization choice (48) is valid in the context of Theorem 1.

6. DISCUSSION

In this paper, we have presented a unified framework for deriving error bounds and convergence rates for a class of regularized M -estimators. The theory is high-dimensional and non-asymptotic in nature, meaning that it yields explicit bounds that hold with high probability for finite sample sizes, and reveals the dependence on dimension and other structural parameters of the model. Two properties of the M -estimator play a central role in our framework. We isolated the notion of a regularizer being *decomposable* with respect to a pair of subspaces, and showed how it constrains the error vector—meaning the difference between any solution and the nominal parameter—to lie within a very specific set. This fact is significant, because it allows for a fruitful notion of *restricted strong convexity* to be developed for the loss function. Since the usual form of strong convexity cannot hold under high-dimensional scaling, this interaction between the decomposable regularizer and the loss function is essential.

Our main result (Theorem 1) provides a deterministic bound on the error for a broad class of regularized M -estimators. By specializing this result to different statistical models, we derived various explicit convergence rates for different estimators, including some known results and a range of novel results. We derived convergence rates for sparse linear models, both under exact and approximate sparsity assumptions, and these results have been shown to be minimax optimal [58]. In the case of sparse group regularization, we established a novel upper

bound of the oracle type, with a separation between the approximation and estimation error terms. For matrix estimation, the framework described here has been used to derive bounds on Frobenius error that are known to be minimax-optimal, both for multitask regression and autoregressive estimation [52], as well as the matrix completion problem [54]. In recent work [2], this framework has also been applied to obtain minimax-optimal rates for noisy matrix decomposition, which involves using a combination of the nuclear norm and elementwise ℓ_1 -norm. Finally, in shown in the paper [50], these results may be applied to derive convergence rates for generalized linear models. Doing so requires leveraging that restricted strong convexity can also be shown to hold for these models, as stated in the bound (43).

There are a variety of interesting open questions associated with our work. In this paper, for simplicity of exposition, we have specified the regularization parameter in terms of the dual norm \mathcal{R}^* of the regularizer. In many cases, this choice leads to optimal convergence rates, including linear regression over ℓ_q -balls (Corollary 3) for sufficiently small radii, and various instances of low-rank matrix regression. In other cases, some refinements of our convergence rates are possible; for instance, for the special case of linear sparsity regression (i.e., an exactly sparse vector, with a constant fraction of non-zero elements), our rates can be sharpened by a more careful analysis of the noise term, which allows for a slightly smaller choice of the regularization parameter. Similarly, there are other non-parametric settings in which a more delicate choice of the regularization parameter is required [36, 59]. Last, we suspect that there are many other statistical models, not discussed in this paper, for which this framework can yield useful results. Some examples include different types of hierarchical regularizers and/or overlapping group regularizers [29, 30], as well as methods using combinations of decomposable regularizers, such as the fused Lasso [70].

Acknowledgements

All authors were partially supported by NSF grants DMS-0605165 and DMS-0907632. BY acknowledges additional support from NSF-grant SES-0835531 (CDI); MJW and SN acknowledge additional support from the NSF grant CDI-0941742, and AFOSR grant 09NL184; and PR acknowledges additional support from NSF grant IIS-101842. We thank a number of people, including Arash Amini, Francis Bach, Peter Buhlmann, Garvesh Raskutti, Alexandre Tsybakov, Sara van de Geer and Tong Zhang for helpful discussions.

REFERENCES

- [1] Large Synoptic Survey Telescope, 2003. URL: www.lsst.org.
- [2] A. Agarwal, S. Negahban, and M. J. Wainwright. Noisy matrix decomposition via convex relaxation: Optimal rates in high dimensions. *To appear in Annals of Statistics*, 2011. Appeared as <http://arxiv.org/abs/1102.4807>.
- [3] F. Bach. Consistency of the group Lasso and multiple kernel learning. *Journal of Machine Learning Research*, 9:1179–1225, 2008.
- [4] F. Bach. Consistency of trace norm minimization. *Journal of Machine Learning Research*, 9:1019–1048, June 2008.
- [5] F. Bach. Self-concordant analysis for logistic regression. *Electronic Journal of Statistics*, 4:384–414, 2010.
- [6] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde. Model-based compressive sensing. Technical report, Rice University, 2008. Available at [arxiv:0808.3572](http://arxiv.org/abs/0808.3572).
- [7] P. J. Bickel, J. B. Brown, H. Huang, and Q. Li. An overview of recent developments in genomics and associated statistical methods. *Phil. Trans. Royal Society A*, 367:4313–4337, 2009.
- [8] P. J. Bickel and E. Levina. Covariance regularization by thresholding. *Annals of Statistics*, 36(6):2577–2604, 2008.
- [9] P. J. Bickel, Y. Ritov, and A. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *Annals of Statistics*, 37(4):1705–1732, 2009.
- [10] F. Bunea. Honest variable selection in linear and logistic regression models via ℓ_1 and $\ell_1 + \ell_2$ penalization. *Electronic Journal of Statistics*, 2:1153–1194, 2008.
- [11] F. Bunea, Y. She, and M. Wegkamp. Adaptive rank penalized estimators in multivariate regression. Technical report, Florida State, 2010. available at [arXiv:1004.2995](http://arxiv.org/abs/1004.2995).
- [12] F. Bunea, A. Tsybakov, and M. Wegkamp. Aggregation for gaussian regression. *Annals of Statistics*, 35(4):1674–1697, 2007.
- [13] F. Bunea, A. Tsybakov, and M. Wegkamp. Sparsity oracle inequalities for the Lasso. *Electronic Journal of Statistics*, pages 169–194, 2007.
- [14] T. Cai and H. Zhou. Optimal rates of convergence for sparse covariance matrix estimation. Technical report, Wharton School of Business, University of Pennsylvania, 2010. available at <http://www-stat.wharton.upenn.edu/~tcai/paper/html/Sparse-Covariance-Matrix.html>.
- [15] E. Candes and T. Tao. Decoding by linear programming. *IEEE Trans. Info Theory*, 51(12):4203–4215, December 2005.
- [16] E. Candes and T. Tao. The Dantzig selector: Statistical estimation when p is much larger than n . *Annals of Statistics*, 35(6):2313–2351, 2007.
- [17] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Found. Comput. Math.*, 9(6):717–772, 2009.
- [18] E. J. Candès, Y. Ma X. Li, and J. Wright. Stable principal component pursuit. In *International Symposium on Information Theory*, June 2010.
- [19] V. Chandrasekaran, S. Sanghavi, P. A. Parrilo, and A. S. Willsky. Rank-sparsity incoherence for matrix decomposition. Technical report, MIT, June 2009. Available at [arXiv:0906.2220v1](http://arxiv.org/abs/0906.2220v1).
- [20] S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM J. Sci. Computing*, 20(1):33–61, 1998.
- [21] D. L. Donoho. Compressed sensing. *IEEE Trans. Info. Theory*, 52(4):1289–1306, April 2006.
- [22] D. L. Donoho and J. M. Tanner. Neighborliness of randomly-projected simplices in high dimensions. *Proceedings of the National Academy of Sciences*, 102(27):9452–9457, 2005.
- [23] M. Fazel. *Matrix Rank Minimization with Applications*. PhD thesis, Stanford, 2002. Available online: <http://faculty.washington.edu/mfazel/thesis-final.pdf>.
- [24] V. L. Girko. *Statistical analysis of observations of increasing dimension*. Kluwer Academic, New York, 1995.
- [25] E. Greenshtein and Y. Ritov. Persistency in high dimensional linear predictor-selection and the virtue of over-parametrization. *Bernoulli*, 10:971–988, 2004.
- [26] D. Hsu, S. M. Kakade, and T. Zhang. Robust matrix decomposition with sparse corruptions. Technical report, Univ. Pennsylvania, November 2010.
- [27] H. Hu, C. Caramanis, and S. Sanghavi. Robust PCA via outlier pursuit. Technical report, UT Austin, 2010.
- [28] J. Huang and T. Zhang. The benefit of group sparsity. *The Annals of Statistics*, 38(4):1978–

- 2004, 2010.
- [29] L. Jacob, G. Obozinski, and J. P. Vert. Group Lasso with Overlap and Graph Lasso. In *International Conference on Machine Learning (ICML)*, pages 433–440, 2009.
 - [30] R. Jenatton, J. Mairal, G. Obozinski, and F. Bach. Proximal methods for hierarchical sparse coding. Technical report, HAL-Inria, 2010. available at inria-00516723.
 - [31] S. M. Kakade, O. Shamir, K. Sridharan, and A. Tewari. Learning exponential families in high-dimensions: Strong convexity and sparsity. In *AISTATS*, 2010.
 - [32] N. El Karoui. Operator norm consistent estimation of large-dimensional sparse covariance matrices. *Annals of Statistics*, 36(6):2717–2756, 2008.
 - [33] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from noisy entries. Technical report, Stanford, June 2009. Preprint available at <http://arxiv.org/abs/0906.2027v1>.
 - [34] Y. Kim, J. Kim, and Y. Kim. Blockwise sparse regression. *Statistica Sinica*, 16(2), 2006.
 - [35] V. Koltchinskii and M. Yuan. Sparse recovery in large ensembles of kernel machines. In *Proceedings of COLT*, 2008.
 - [36] V. Koltchinskii and M. Yuan. Sparsity in multiple kernel learning. *Annals of Statistics*, 38:3660–3695, 2010.
 - [37] C. Lam and J. Fan. Sparsistency and rates of convergence in large covariance matrix estimation. *Annals of Statistics*, 37:4254–4278, 2009.
 - [38] D. Landgrebe. Hyperspectral image data analysis as a high-dimensional signal processing problem. *IEEE Signal Processing Magazine*, 19(1):17–28, January 2008.
 - [39] K. Lee and Y. Bresler. Guaranteed minimum rank approximation from linear observations by nuclear norm minimization with an ellipsoidal constraint. Technical report, UIUC, 2009. Available at arXiv:0903.4742.
 - [40] Z. Liu and L. Vandenberghe. Interior-point method for nuclear norm optimization with application to system identification. *SIAM Journal on Matrix Analysis and Applications*, 31(3):1235–1256, 2009.
 - [41] K. Lounici, M. Pontil, A. B. Tsybakov, and S. van de Geer. Taking advantage of sparsity in multi-task learning. Technical Report arXiv:0903.1468, ETH Zurich, March 2009.
 - [42] M. Lustig, D. Donoho, J. Santos, and J. Pauly. Compressed sensing MRI. *IEEE Signal Processing Magazine*, 27:72–82, March 2008.
 - [43] M. McCoy and J. Tropp. Two Proposals for Robust PCA using Semidefinite Programming. Technical report, California Institute of Technology, 2010.
 - [44] M. L. Mehta. *Random matrices*. Academic Press, New York, NY, 1991.
 - [45] L. Meier, S. van de Geer, and P. Bühlmann. High-dimensional additive modeling. *Annals of Statistics*, 37:3779–3821, 2009.
 - [46] N. Meinshausen. A note on the Lasso for graphical Gaussian model selection. *Statistics and Probability Letters*, 78(7):880–884, 2008.
 - [47] N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the Lasso. *Annals of Statistics*, 34:1436–1462, 2006.
 - [48] N. Meinshausen and B. Yu. Lasso-type recovery of sparse representations for high-dimensional data. *Annals of Statistics*, 37(1):246–270, 2009.
 - [49] Y. Nardi and A. Rinaldo. On the asymptotic properties of the group lasso estimator for linear models. *Electronic Journal of Statistics*, 2:605–633, 2008.
 - [50] S. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. A unified framework for high-dimensional analysis of M -estimators with decomposable regularizers. In *NIPS Conference*, 2009.
 - [51] S. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. Supplement to “a unified framework for high-dimensional analysis of M -estimators with decomposable regularizers”. 2012.
 - [52] S. Negahban and M. J. Wainwright. Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *Annals of Statistics*, 39(2):1069–1097, 2011.
 - [53] S. Negahban and M. J. Wainwright. Simultaneous support recovery in high-dimensional regression: Benefits and perils of $\ell_{1,\infty}$ -regularization. *IEEE Transactions on Information Theory*, 57(6):3481–3863, June 2011.
 - [54] S. Negahban and M. J. Wainwright. Restricted strong convexity and (weighted) matrix completion: Optimal bounds with noise. *Journal of Machine Learning Research*, 2012. Originally posted as arxiv:1009.2118.
 - [55] G. Obozinski, M. J. Wainwright, and M. I. Jordan. Union support recovery in high-dimensional multivariate regression. *Annals of Statistics*, 39(1):1–47, January 2011.

- [56] L. A. Pastur. On the spectrum of random matrices. *Theoretical and Mathematical Physics*, 10:67–74, 1972.
- [57] G. Raskutti, M. J. Wainwright, and B. Yu. Restricted eigenvalue conditions for correlated Gaussian designs. *Journal of Machine Learning Research*, 11:2241–2259, August 2010.
- [58] G. Raskutti, M. J. Wainwright, and B. Yu. Minimax rates of estimation for high-dimensional linear regression over ℓ_q -balls. *IEEE Trans. Information Theory*, 57(10):6976–6994, October 2011.
- [59] G. Raskutti, M. J. Wainwright, and B. Yu. Minimax-optimal rates for sparse additive models over kernel classes via convex programming. *Journal of Machine Learning Research*, 12:389–427, March 2012.
- [60] P. Ravikumar, H. Liu, J. Lafferty, and L. Wasserman. SpAM: sparse additive models. *Journal of the Royal Statistical Society, Series B*, 71(5):1009–1030, 2009.
- [61] P. Ravikumar, M. J. Wainwright, and J. Lafferty. High-dimensional Ising model selection using ℓ_1 -regularized logistic regression. *Annals of Statistics*, 38(3):1287–1319, 2010.
- [62] P. Ravikumar, M. J. Wainwright, G. Raskutti, and B. Yu. High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence. *Electron. J. Statist.*, 5:935–980, 2011.
- [63] B. Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12:3413–3430, 2011.
- [64] B. Recht, M. Fazel, and P. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.
- [65] A. Rohde and A. Tsybakov. Estimation of high-dimensional low-rank matrices. *Annals of Statistics*, 39(2):887–930, 2011.
- [66] A. J. Rothman, P. J. Bickel, E. Levina, and J. Zhu. Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics*, 2:494–515, 2008.
- [67] M. Rudelson and S. Zhou. Reconstruction from anisotropic random measurements. Technical report, University of Michigan, July 2011.
- [68] M. Stojnic, F. Parvaresh, and B. Hassibi. On the reconstruction of block-sparse signals with an optimal number of measurements. *IEEE Transactions on Signal Processing*, 57(8):3075–3085, 2009.
- [69] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58(1):267–288, 1996.
- [70] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight. Sparsity and smoothness via the fused Lasso. *J. R. Statistical Soc. B*, 67:91–108, 2005.
- [71] J. A. Tropp, A. C. Gilbert, and M. J. Strauss. Algorithms for simultaneous sparse approximation. *Signal Processing*, 86:572–602, April 2006. Special issue on "Sparse approximations in signal and image processing".
- [72] B. Turlach, W.N. Venables, and S.J. Wright. Simultaneous variable selection. *Technometrics*, 27:349–363, 2005.
- [73] S. van de Geer and P. Bühlmann. On the conditions used to prove oracle results for the lasso. *Electronic Journal of Statistics*, 3:1360–1392, 2009.
- [74] S. A. van de Geer. High-dimensional generalized linear models and the lasso. *The Annals of Statistics*, 36:614–645, 2008.
- [75] M. J. Wainwright. Information-theoretic bounds on sparsity recovery in the high-dimensional and noisy setting. *IEEE Trans. Info. Theory*, 55:5728–5741, December 2009.
- [76] M. J. Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (Lasso). *IEEE Trans. Information Theory*, 55:2183–2202, May 2009.
- [77] E. Wigner. Characteristic vectors of bordered matrices with infinite dimensions. *Annals of Mathematics*, 62:548–564, 1955.
- [78] M. Yuan, A. Ekici, Z. Lu, and R. Monteiro. Dimension reduction and coefficient estimation in multivariate linear regression. *Journal Of The Royal Statistical Society Series B*, 69(3):329–346, 2007.
- [79] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society B*, 1(68):49, 2006.
- [80] C. H. Zhang and J. Huang. The sparsity and bias of the lasso selection in high-dimensional linear regression. *Annals of Statistics*, 36(4):1567–1594, 2008.
- [81] P. Zhao, G. Rocha, and B. Yu. Grouped and hierarchical model selection through composite absolute penalties. *Annals of Statistics*, 37(6A):3468–3497, 2009.

- [82] P. Zhao and B. Yu. On model selection consistency of Lasso. *Journal of Machine Learning Research*, 7:2541–2567, 2006.
- [83] S. Zhou, J. Lafferty, and L. Wasserman. Time-varying undirected graphs. In *21st Annual Conference on Learning Theory (COLT)*, Helsinki, Finland, July 2008.