

HIGH-DIMENSIONAL REGRESSION WITH NOISY AND MISSING DATA: PROVABLE GUARANTEES WITH NONCONVEXITY

BY PO-LING LOH^{1,2} AND MARTIN J. WAINWRIGHT²

University of California, Berkeley

Although the standard formulations of prediction problems involve fully-observed and noiseless data drawn in an i.i.d. manner, many applications involve noisy and/or missing data, possibly involving dependence, as well. We study these issues in the context of high-dimensional sparse linear regression, and propose novel estimators for the cases of noisy, missing and/or dependent data. Many standard approaches to noisy or missing data, such as those using the EM algorithm, lead to optimization problems that are inherently nonconvex, and it is difficult to establish theoretical guarantees on practical algorithms. While our approach also involves optimizing nonconvex programs, we are able to both analyze the statistical error associated with any global optimum, and more surprisingly, to prove that a simple algorithm based on projected gradient descent will converge in polynomial time to a small neighborhood of the set of all global minimizers. On the statistical side, we provide nonasymptotic bounds that hold with high probability for the cases of noisy, missing and/or dependent data. On the computational side, we prove that under the same types of conditions required for statistical consistency, the projected gradient descent algorithm is guaranteed to converge at a geometric rate to a near-global minimizer. We illustrate these theoretical predictions with simulations, showing close agreement with the predicted scalings.

1. Introduction. In standard formulations of prediction problems, it is assumed that the covariates are fully-observed and sampled independently from some underlying distribution. However, these assumptions are not realistic for many applications, in which covariates may be observed only partially, observed subject to corruption or exhibit some type of dependency. Consider the problem of modeling the voting behavior of politicians: in this setting, votes may be missing due to abstentions, and temporally dependent due to collusion or “tit-for-tat” behavior. Similarly, surveys often suffer from the missing data problem, since users fail to respond to all questions. Sensor network data also tends to be both noisy due to measurement error, and partially missing due to failures or drop-outs of sensors.

Received September 2011; revised May 2012.

¹Supported in part by a Hertz Foundation Fellowship and the Department of Defense (DoD) through an NDSEG Fellowship.

²Supported in part by NSF Grant DMS-09-07632 and Air Force Office of Scientific Research Grant AFOSR-09NL184.

MSC2010 subject classifications. Primary 62F12; secondary 68W25.

Key words and phrases. High-dimensional statistics, missing data, nonconvexity, regularization, sparse linear regression, M -estimation.

There are a variety of methods for dealing with noisy and/or missing data, including various heuristic methods, as well as likelihood-based methods involving the expectation–maximization (EM) algorithm (e.g., see the book [8] and references therein). A challenge in this context is the possible nonconvexity of associated optimization problems. For instance, in applications of EM, problems in which the negative likelihood is a convex function often become nonconvex with missing or noisy data. Consequently, although the EM algorithm will converge to a local minimum, it is difficult to guarantee that the local optimum is close to a global minimum.

In this paper, we study these issues in the context of high-dimensional sparse linear regression—in particular, in the case when the predictors or covariates are noisy, missing, and/or dependent. Our main contribution is to develop and study simple methods for handling these issues, and to prove theoretical results about both the associated statistical error and the optimization error. Like EM-based approaches, our estimators are based on solving optimization problems that may be nonconvex; however, despite this nonconvexity, we are still able to prove that a simple form of projected gradient descent will produce an output that is “sufficiently close”—as small as the statistical error—to any global optimum. As a second result, we bound the statistical error, showing that it has the same scaling as the minimax rates for the classical cases of perfectly observed and independently sampled covariates. In this way, we obtain estimators for noisy, missing, and/or dependent data that have the same scaling behavior as the usual fully-observed and independent case. The resulting estimators allow us to solve the problem of high-dimensional Gaussian graphical model selection with missing data.

There is a large body of work on the problem of corrupted covariates or error-in-variables for regression problems (e.g., see the papers and books [3, 6, 7, 21], as well as references therein). Much of the earlier theoretical work is classical in nature, meaning that it requires that the sample size n diverges with the dimension p fixed. Most relevant to this paper is more recent work that has examined issues of corrupted and/or missing data in the context of high-dimensional sparse linear models, allowing for $n \ll p$. Städler and Bühlmann [18] developed an EM-based method for sparse inverse covariance matrix estimation in the missing data regime, and used this result to derive an algorithm for sparse linear regression with missing data. As mentioned above, however, it is difficult to guarantee that EM will converge to a point close to a global optimum of the likelihood, in contrast to the methods studied here. Rosenbaum and Tsybakov [14] studied the sparse linear model when the covariates are corrupted by noise, and proposed a modified form of the Dantzig selector (see the discussion following our main results for a detailed comparison to this past work, and also to concurrent work [15] by the same authors). For the particular case of multiplicative noise, the type of estimator that we consider here has been studied in past work [21]; however, this theoretical analysis is of the classical type, holding only for $n \gg p$, in contrast to the high-dimensional models that are of interest here.

The remainder of this paper is organized as follows. We begin in Section 2 with background and a precise description of the problem. We then introduce the class of estimators we will consider and the form of the projected gradient descent algorithm. Section 3 is devoted to a description of our main results, including a pair of general theorems on the statistical and optimization error, and then a series of corollaries applying our results to the cases of noisy, missing, and dependent data. In Section 4, we demonstrate simulations to confirm that our methods work in practice, and verify the theoretically-predicted scaling laws. Section 5 contains proofs of some of the main results, with the remaining proofs contained in the supplementary Appendix [9].

NOTATION. For a matrix M , we write $\|M\|_{\max} := \max_{i,j} |m_{ij}|$ to be the elementwise ℓ_{∞} -norm of M . Furthermore, $\|M\|_1$ denotes the induced ℓ_1 -operator norm (maximum absolute column sum) of M , and $\|M\|_{\text{op}}$ is the spectral norm of M . We write $\kappa(M) := \frac{\lambda_{\max}(M)}{\lambda_{\min}(M)}$, the condition number of M . For matrices M_1, M_2 , we write $M_1 \odot M_2$ to denote the componentwise Hadamard product, and write $M_1 \oslash M_2$ to denote componentwise division. For functions $f(n)$ and $g(n)$, we write $f(n) \lesssim g(n)$ to mean that $f(n) \leq cg(n)$ for a universal constant $c \in (0, \infty)$, and similarly, $f(n) \gtrsim g(n)$ when $f(n) \geq c'g(n)$ for some universal constant $c' \in (0, \infty)$. Finally, we write $f(n) \asymp g(n)$ when $f(n) \lesssim g(n)$ and $f(n) \gtrsim g(n)$ hold simultaneously.

2. Background and problem setup. In this section, we provide background and a precise description of the problem, and then motivate the class of estimators analyzed in this paper. We then discuss a simple class of projected gradient descent algorithms that can be used to obtain an estimator.

2.1. *Observation model and high-dimensional framework.* Suppose we observe a response variable $y_i \in \mathbb{R}$ linked to a covariate vector $x_i \in \mathbb{R}^p$ via the linear model

$$(2.1) \quad y_i = \langle x_i, \beta^* \rangle + \varepsilon_i \quad \text{for } i = 1, 2, \dots, n.$$

Here, the regression vector $\beta^* \in \mathbb{R}^p$ is unknown, and $\varepsilon_i \in \mathbb{R}$ is observation noise, independent of x_i . Rather than directly observing each $x_i \in \mathbb{R}^p$, we observe a vector $z_i \in \mathbb{R}^p$ linked to x_i via some conditional distribution, that is,

$$(2.2) \quad z_i \sim \mathbb{Q}(\cdot | x_i) \quad \text{for } i = 1, 2, \dots, n.$$

This setup applies to various disturbances to the covariates, including:

- (a) *Covariates with additive noise:* We observe $z_i = x_i + w_i$, where $w_i \in \mathbb{R}^p$ is a random vector independent of x_i , say zero-mean with known covariance matrix Σ_w .

(b) *Missing data*: For some fraction $\rho \in [0, 1)$, we observe a random vector $z_i \in \mathbb{R}^p$ such that for each component j , we independently observe $z_{ij} = x_{ij}$ with probability $1 - \rho$, and $z_{ij} = *$ with probability ρ . We can also consider the case when the entries in the j th column have a different probability ρ_j of being missing.

(c) *Covariates with multiplicative noise*: Generalizing the missing data problem, suppose we observe $z_i = x_i \odot u_i$, where $u_i \in \mathbb{R}^p$ is again a random vector independent of x_i , and \odot is the Hadamard product. The problem of missing data is a special case of multiplicative noise, where all u_{ij} 's are independent and $u_{ij} \sim \text{Bernoulli}(1 - \rho_j)$.

Our first set of results is deterministic, depending on specific instantiations of the observations $\{(y_i, z_i)\}_{i=1}^n$. However, we are also interested in results that hold with high probability when the x_i 's and z_i 's are drawn at random. We consider both the case when the x_i 's are drawn i.i.d. from a fixed distribution; and the case of dependent covariates, when the x_i 's are generated according to a stationary vector autoregressive (VAR) process.

We work within a high-dimensional framework that allows the number of predictors p to grow and possibly exceed the sample size n . Of course, consistent estimation when $n \ll p$ is impossible unless the model is endowed with additional structure—for instance, sparsity in the parameter vector β^* . Consequently, we study the class of models where β^* has at most k nonzero parameters, where k is also allowed to increase to infinity with p and n .

2.2. *M-estimators for noisy and missing covariates*. In order to motivate the class of estimators we will consider, let us begin by examining a simple deterministic problem. Let $\Sigma_x \succ 0$ be the covariance matrix of the covariates, and consider the ℓ_1 -constrained quadratic program

$$(2.3) \quad \hat{\beta} \in \arg \min_{\|\beta\|_1 \leq R} \left\{ \frac{1}{2} \beta^T \Sigma_x \beta - \langle \Sigma_x \beta^*, \beta \rangle \right\}.$$

As long as the constraint radius R is at least $\|\beta^*\|_1$, the unique solution to this convex program is $\hat{\beta} = \beta^*$. Of course, this program is an idealization, since in practice we may not know the covariance matrix Σ_x , and we certainly do not know $\Sigma_x \beta^*$ —after all, β^* is the quantity we are trying to estimate!

Nonetheless, this idealization still provides useful intuition, as it suggests various estimators based on the plug-in principle. Given a set of samples, it is natural to form estimates of the quantities Σ_x and $\Sigma_x \beta^*$, which we denote by $\hat{\Gamma} \in \mathbb{R}^{p \times p}$ and $\hat{\gamma} \in \mathbb{R}^p$, respectively, and to consider the modified program

$$(2.4) \quad \hat{\beta} \in \arg \min_{\|\beta\|_1 \leq R} \left\{ \frac{1}{2} \beta^T \hat{\Gamma} \beta - \langle \hat{\gamma}, \beta \rangle \right\},$$

or alternatively, the regularized version

$$(2.5) \quad \hat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ \frac{1}{2} \beta^T \hat{\Gamma} \beta - \langle \hat{\gamma}, \beta \rangle + \lambda_n \|\beta\|_1 \right\},$$

where $\lambda_n > 0$ is a user-defined regularization parameter. Note that the two problems are equivalent by Lagrangian duality when the objectives are convex, but not in the case of a nonconvex objective. The Lasso [4, 19] is a special case of these programs, obtained by setting

$$(2.6) \quad \widehat{\Gamma}_{\text{Las}} := \frac{1}{n} X^T X \quad \text{and} \quad \widehat{\gamma}_{\text{Las}} := \frac{1}{n} X^T y,$$

where we have introduced the shorthand $y = (y_1, \dots, y_n)^T \in \mathbb{R}^n$, and $X \in \mathbb{R}^{n \times p}$, with x_i^T as its i th row. A simple calculation shows that $(\widehat{\Gamma}_{\text{Las}}, \widehat{\gamma}_{\text{Las}})$ are unbiased estimators of the pair $(\Sigma_x, \Sigma_x \beta^*)$. This unbiasedness and additional concentration inequalities (to be described in the sequel) underlie the well-known analysis of the Lasso in the high-dimensional regime.

In this paper, we focus on more general instantiations of the programs (2.4) and (2.5), involving different choices of the pair $(\widehat{\Gamma}, \widehat{\gamma})$ that are adapted to the cases of noisy and/or missing data. Note that the matrix $\widehat{\Gamma}_{\text{Las}}$ is positive semidefinite, so the Lasso program is convex. In sharp contrast, for the case of noisy or missing data, the most natural choice of the matrix $\widehat{\Gamma}$ is *not positive semidefinite*, hence the quadratic losses appearing in the problems (2.4) and (2.5) are *nonconvex*. Furthermore, when $\widehat{\Gamma}$ has negative eigenvalues, the objective in equation (2.5) is unbounded from below. Hence, we make use of the following regularized estimator:

$$(2.7) \quad \widehat{\beta} \in \arg \min_{\|\beta\|_1 \leq b_0 \sqrt{k}} \left\{ \frac{1}{2} \beta^T \widehat{\Gamma} \beta - \langle \widehat{\gamma}, \beta \rangle + \lambda_n \|\beta\|_1 \right\}$$

for a suitable constant b_0 .

In the presence of nonconvexity, it is generally impossible to provide a polynomial-time algorithm that converges to a (near) global optimum, due to the presence of local minima. Remarkably, we are able to prove that this issue is not significant in our setting, and a simple projected gradient descent algorithm applied to the programs (2.4) or (2.7) converges with high probability to a vector extremely close to any global optimum.

Let us illustrate these ideas with some examples. Recall that $(\widehat{\Gamma}, \widehat{\gamma})$ serve as unbiased estimators for $(\Sigma_x, \Sigma_x \beta^*)$.

EXAMPLE 1 (Additive noise). Suppose we observe $Z = X + W$, where W is a random matrix independent of X , with rows w_i drawn i.i.d. from a zero-mean distribution with known covariance Σ_w . We consider the pair

$$(2.8) \quad \widehat{\Gamma}_{\text{add}} := \frac{1}{n} Z^T Z - \Sigma_w \quad \text{and} \quad \widehat{\gamma}_{\text{add}} := \frac{1}{n} Z^T y.$$

Note that when $\Sigma_w = 0$ (corresponding to the noiseless case), the estimators reduce to the standard Lasso. However, when $\Sigma_w \neq 0$, the matrix $\widehat{\Gamma}_{\text{add}}$ is *not* positive semidefinite in the high-dimensional regime ($n \ll p$). Indeed, since the matrix $\frac{1}{n} Z^T Z$ has rank at most n , the subtracted matrix Σ_w may cause $\widehat{\Gamma}_{\text{add}}$ to have a

large number of negative eigenvalues. For instance, if $\Sigma_w = \sigma_w^2 I$ for $\sigma_w^2 > 0$, then $\widehat{\Gamma}_{\text{add}}$ has $p - n$ eigenvalues equal to $-\sigma_w^2$.

EXAMPLE 2 (Missing data). We now consider the case where the entries of X are missing at random. Let us first describe an estimator for the special case where each entry is missing at random, independently with some constant probability $\rho \in [0, 1)$. (In Example 3 to follow, we will describe the extension to general missing probabilities.) Consequently, we observe the matrix $Z \in \mathbb{R}^{n \times p}$ with entries

$$Z_{ij} = \begin{cases} X_{ij}, & \text{with probability } 1 - \rho, \\ 0, & \text{otherwise.} \end{cases}$$

Given the observed matrix $Z \in \mathbb{R}^{n \times p}$, we use

$$(2.9) \quad \widehat{\Gamma}_{\text{mis}} := \frac{\widetilde{Z}^T \widetilde{Z}}{n} - \rho \operatorname{diag}\left(\frac{\widetilde{Z}^T \widetilde{Z}}{n}\right) \quad \text{and} \quad \widehat{\gamma}_{\text{mis}} := \frac{1}{n} \widetilde{Z}^T y,$$

where $\widetilde{Z}_{ij} = Z_{ij}/(1 - \rho)$. It is easy to see that the pair $(\widehat{\Gamma}_{\text{mis}}, \widehat{\gamma}_{\text{mis}})$ reduces to the pair $(\widehat{\Gamma}_{\text{Las}}, \widehat{\gamma}_{\text{Las}})$ for the standard Lasso when $\rho = 0$, corresponding to no missing data. In the more interesting case when $\rho \in (0, 1)$, the matrix $\frac{\widetilde{Z}^T \widetilde{Z}}{n}$ in equation (2.9) has rank at most n , so the subtracted diagonal matrix may cause the matrix $\widehat{\Gamma}_{\text{mis}}$ to have a large number of negative eigenvalues when $n \ll p$. As a consequence, the matrix $\widehat{\Gamma}_{\text{mis}}$ is not (in general) positive semidefinite, so the associated quadratic function is not convex.

EXAMPLE 3 (Multiplicative noise). As a generalization of the previous example, we now consider the case of multiplicative noise. In particular, suppose we observe the quantity $Z = X \odot U$, where U is a matrix of nonnegative noise variables. In many applications, it is natural to assume that the rows u_i of U are drawn in an i.i.d. manner, say from some distribution in which both the vector $\mathbb{E}[u_i]$ and the matrix $\mathbb{E}[u_i u_i^T]$ have strictly positive entries. This general family of multiplicative noise models arises in various applications; we refer the reader to the papers [3, 6, 7, 21] for more discussion and examples. A natural choice of the pair $(\widehat{\Gamma}, \widehat{\gamma})$ is given by the quantities

$$(2.10) \quad \widehat{\Gamma}_{\text{mul}} := \frac{1}{n} Z^T Z \oplus \mathbb{E}(u_1 u_1^T) \quad \text{and} \quad \widehat{\Gamma}_{\text{mul}} := \frac{1}{n} Z^T y \oplus \mathbb{E}(u_1),$$

where \oplus denotes elementwise division. A small calculation shows that these are unbiased estimators of Σ_x and $\Sigma_x \beta^*$, respectively. The estimators (2.10) have been studied in past work [21], but only under classical scaling ($n \gg p$).

As a special case of the estimators (2.10), suppose the entries u_{ij} of U are independent Bernoulli($1 - \rho_j$) random variables. Then the observed matrix $Z = X \odot U$ corresponds to a missing-data matrix, where each element of the j th column has probability ρ_j of being missing. In this case, the estimators (2.10) become

$$(2.11) \quad \widehat{\Gamma}_{\text{mis}} = \frac{Z^T Z}{n} \oplus M \quad \text{and} \quad \widehat{\gamma}_{\text{mis}} = \frac{1}{n} Z^T y \oplus (\mathbf{1} - \rho),$$

where $M := \mathbb{E}(u_1 u_1^T)$ satisfies

$$M_{ij} = \begin{cases} (1 - \rho_i)(1 - \rho_j), & \text{if } i \neq j, \\ 1 - \rho_i, & \text{if } i = j, \end{cases}$$

ρ is the parameter vector containing the ρ_j 's, and $\mathbf{1}$ is the vector of all 1's. In this way, we obtain a generalization of the estimator discussed in Example 2.

2.3. *Restricted eigenvalue conditions.* Given an estimate $\widehat{\beta}$, there are various ways to assess its closeness to β^* . In this paper, we focus on the ℓ_2 -norm $\|\widehat{\beta} - \beta^*\|_2$, as well as the closely related ℓ_1 -norm $\|\widehat{\beta} - \beta^*\|_1$. When the covariate matrix X is fully observed (so that the Lasso can be applied), it is now well understood that a sufficient condition for ℓ_2 -recovery is that the matrix $\widehat{\Gamma}_{\text{Las}} = \frac{1}{n} X^T X$ satisfy a certain type of restricted eigenvalue (RE) condition (e.g., [2, 20]). In this paper, we make use of the following condition.

DEFINITION 1 (Lower-RE condition). The matrix $\widehat{\Gamma}$ satisfies a lower restricted eigenvalue condition with curvature $\alpha_1 > 0$ and tolerance $\tau(n, p) > 0$ if

$$(2.12) \quad \theta^T \widehat{\Gamma} \theta \geq \alpha_1 \|\theta\|_2^2 - \tau(n, p) \|\theta\|_1^2 \quad \text{for all } \theta \in \mathbb{R}^p.$$

It can be shown that when the Lasso matrix $\widehat{\Gamma}_{\text{Las}} = \frac{1}{n} X^T X$ satisfies this RE condition (2.12), the Lasso estimate has low ℓ_2 -error for any vector β^* supported on any subset of size at most $k \lesssim \frac{1}{\tau(n, p)}$. In particular, bound (2.12) implies a sparse RE condition for all k of this magnitude, and conversely, Lemma 11 in the Appendix of [9] shows that a sparse RE condition implies bound (2.12). In this paper, we work with condition (2.12), since it is especially convenient for analyzing optimization algorithms.

In the standard setting (with uncorrupted and fully observed design matrices), it is known that for many choices of the design matrix X (with rows having covariance Σ), the Lasso matrix $\widehat{\Gamma}_{\text{Las}}$ will satisfy such an RE condition with high probability (e.g., [13, 17]) with $\alpha_1 = \frac{1}{2} \lambda_{\min}(\Sigma)$ and $\tau(n, p) \asymp \frac{\log p}{n}$. A significant portion of the analysis in this paper is devoted to proving that different choices of $\widehat{\Gamma}$, such as the matrices $\widehat{\Gamma}_{\text{add}}$ and $\widehat{\Gamma}_{\text{mis}}$ defined earlier, also satisfy condition (2.12) with high probability. This fact is by no means obvious, since as previously discussed, the matrices $\widehat{\Gamma}_{\text{add}}$ and $\widehat{\Gamma}_{\text{mis}}$ generally have large numbers of negative eigenvalues.

Finally, although such upper bounds are not necessary for statistical consistency, our algorithmic results make use of the analogous upper restricted eigenvalue condition, formalized in the following:

DEFINITION 2 (Upper-RE condition). The matrix $\widehat{\Gamma}$ satisfies an upper restricted eigenvalue condition with smoothness $\alpha_2 > 0$ and tolerance $\tau(n, p) > 0$ if

$$(2.13) \quad \theta^T \widehat{\Gamma} \theta \leq \alpha_2 \|\theta\|_2^2 + \tau(n, p) \|\theta\|_1^2 \quad \text{for all } \theta \in \mathbb{R}^p.$$

In recent work on high-dimensional projected gradient descent, Agarwal et al. [1] make use of a more general form of the lower and upper bounds (2.12) and (2.13), applicable to nonquadratic losses as well, which are referred to as the restricted strong convexity (RSC) and restricted smoothness (RSM) conditions, respectively. For various class of random design matrices, it can be shown that the Lasso matrix $\widehat{\Gamma}_{\text{Las}}$ satisfies the upper bound (2.13) with $\alpha_2 = 2\lambda_{\max}(\Sigma_x)$ and $\tau(n, p) \asymp \frac{\log p}{n}$; see Raskutti et al. [13] for the Gaussian case and Rudelson and Zhou [17] for the sub-Gaussian setting. We will establish similar scaling for our choices of $\widehat{\Gamma}$.

2.4. *Gradient descent algorithms.* In addition to proving results about the global minima of the (possibly nonconvex) programs (2.4) and (2.5), we are also interested in polynomial-time procedures for approximating such optima. In this paper, we analyze some simple algorithms for solving either the constrained program (2.4) or the Lagrangian version (2.7). Note that the gradient of the quadratic loss function takes the form $\nabla \mathcal{L}(\beta) = \widehat{\Gamma}\beta - \widehat{\gamma}$. In application to the constrained version, the method of projected gradient descent generates a sequence of iterates $\{\beta^t, t = 0, 1, 2, \dots\}$ by the recursion

$$(2.14) \quad \beta^{t+1} = \arg \min_{\|\beta\|_1 \leq R} \left\{ \mathcal{L}(\beta^t) + \langle \nabla \mathcal{L}(\beta^t), \beta - \beta^t \rangle + \frac{\eta}{2} \|\beta - \beta^t\|_2^2 \right\},$$

where $\eta > 0$ is a stepsize parameter. Equivalently, this update can be written as $\beta^{t+1} = \Pi(\beta^t - \frac{1}{\eta} \nabla \mathcal{L}(\beta^t))$, where Π denotes the ℓ_2 -projection onto the ℓ_1 -ball of radius R . This projection can be computed rapidly in $\mathcal{O}(p)$ time using a procedure due to Duchi et al. [5]. For the Lagrangian update, we use a slight variant of the projected gradient update (2.14), namely

$$(2.15) \quad \beta^{t+1} = \arg \min_{\|\beta\|_1 \leq R} \left\{ \mathcal{L}(\beta^t) + \langle \nabla \mathcal{L}(\beta^t), \beta - \beta^t \rangle + \frac{\eta}{2} \|\beta - \beta^t\|_2^2 + \lambda_n \|\beta\|_1 \right\}$$

with the only difference being the inclusion of the regularization term. This update can also be performed efficiently by performing two projections onto the ℓ_1 -ball; see the paper [1] for details.

When the objective function is convex (equivalently, $\widehat{\Gamma}$ is positive semidefinite), the iterates (2.14) or (2.15) are guaranteed to converge to a global minimum of the objective functions (2.4) and (2.7), respectively. In our setting, the matrix $\widehat{\Gamma}$ need not be positive semidefinite, so the best generic guarantee is that the iterates converge to a local optimum. However, our analysis shows that for the family of programs (2.4) or (2.7), under a reasonable set of conditions satisfied by various statistical models, the iterates actually converge to a point extremely close to any global optimum in both ℓ_1 -norm and ℓ_2 -norm; see Theorem 2 to follow for a more detailed statement.

3. Main results and consequences. We now state our main results and discuss their consequences for noisy, missing, and dependent data.

3.1. *General results.* We provide theoretical guarantees for both the constrained estimator (2.4) and the Lagrangian version (2.7). Note that we obtain different optimization problems as we vary the choice of the pair $(\widehat{\Gamma}, \widehat{\gamma}) \in \mathbb{R}^{p \times p} \times \mathbb{R}^p$. We begin by stating a pair of general results, applicable to any pair that satisfies certain conditions. Our first result (Theorem 1) provides bounds on the statistical error, namely the quantity $\|\widehat{\beta} - \beta^*\|_2$, as well as the corresponding ℓ_1 -error, where $\widehat{\beta}$ is any global optimum of the programs (2.4) or (2.7). Since the problem may be nonconvex in general, it is not immediately obvious that one can obtain a *provably good* approximation to any global optimum without resorting to costly search methods. In order to assuage this concern, our second result (Theorem 2) provides rigorous bounds on the optimization error, namely the differences $\|\beta^t - \widehat{\beta}\|_2$ and $\|\beta^t - \widehat{\beta}\|_1$ incurred by the iterate β^t after running t rounds of the projected gradient descent updates (2.14) or (2.15).

3.1.1. *Statistical error.* In controlling the statistical error, we assume that the matrix $\widehat{\Gamma}$ satisfies a lower-RE condition with curvature α_1 and tolerance $\tau(n, p)$, as previously defined (2.12). Recall that $\widehat{\Gamma}$ and $\widehat{\gamma}$ serve as surrogates to the deterministic quantities $\Sigma_x \in \mathbb{R}^{p \times p}$ and $\Sigma_x \beta^* \in \mathbb{R}^p$, respectively. Our results also involve a measure of deviation in these surrogates. In particular, we assume that there is some function $\varphi(\mathbb{Q}, \sigma_\varepsilon)$, depending on the two sources of noise in our problem: the standard deviation σ_ε of the observation noise vector ε from equation (2.1), and the conditional distribution \mathbb{Q} from equation (2.2) that links the covariates x_i to the observed versions z_i . With this notation, we consider the deviation condition

$$(3.1) \quad \|\widehat{\gamma} - \widehat{\Gamma} \beta^*\|_\infty \leq \varphi(\mathbb{Q}, \sigma_\varepsilon) \sqrt{\frac{\log p}{n}}.$$

To aid intuition, note that inequality (3.1) holds whenever the following two deviation conditions are satisfied:

$$(3.2) \quad \begin{aligned} \|\widehat{\gamma} - \Sigma_x \beta^*\|_\infty &\leq \varphi(\mathbb{Q}, \sigma_\varepsilon) \sqrt{\frac{\log p}{n}} \quad \text{and} \\ \|\widehat{\Gamma} - \Sigma_x\|_\infty &\leq \varphi(\mathbb{Q}, \sigma_\varepsilon) \sqrt{\frac{\log p}{n}}. \end{aligned}$$

The pair of inequalities (3.2) clearly measures the deviation of the estimators $(\widehat{\Gamma}, \widehat{\gamma})$ from their population versions, and they are sometimes easier to verify theoretically. However, inequality (3.1) may be used directly to derive tighter bounds (e.g., in the additive noise case). Indeed, the bounds established via inequalities (3.2) is not sharp in the limit of low noise on the covariates, due to the second

inequality. In the proofs of our corollaries to follow, we will verify the deviation conditions for various forms of noisy, missing, and dependent data, with the quantity $\varphi(\mathbb{Q}, \sigma_\varepsilon)$ changing depending on the model. We have the following result, which applies to any global optimum $\hat{\beta}$ of the regularized version (2.7) with $\lambda_n \geq 4\varphi(\mathbb{Q}, \sigma_\varepsilon)\sqrt{\frac{\log p}{n}}$:

THEOREM 1 (Statistical error). *Suppose the surrogates $(\hat{\Gamma}, \hat{\gamma})$ satisfy the deviation bound (3.1), and the matrix $\hat{\Gamma}$ satisfies the lower-RE condition (2.12) with parameters (α_1, τ) such that*

$$(3.3) \quad \sqrt{k}\tau(n, p) \leq \min \left\{ \frac{\alpha_1}{128\sqrt{k}}, \frac{\varphi(\mathbb{Q}, \sigma_\varepsilon)}{b_0} \sqrt{\frac{\log p}{n}} \right\}.$$

Then for any vector β^* with sparsity at most k , there is a universal positive constant c_0 such that any global optimum $\hat{\beta}$ of the Lagrangian program (2.7) with any $b_0 \geq \|\beta^*\|_2$ satisfies the bounds

$$(3.4a) \quad \|\hat{\beta} - \beta^*\|_2 \leq \frac{c_0\sqrt{k}}{\alpha_1} \max \left\{ \varphi(\mathbb{Q}, \sigma_\varepsilon)\sqrt{\frac{\log p}{n}}, \lambda_n \right\} \quad \text{and}$$

$$(3.4b) \quad \|\hat{\beta} - \beta^*\|_1 \leq \frac{8c_0k}{\alpha_1} \max \left\{ \varphi(\mathbb{Q}, \sigma_\varepsilon)\sqrt{\frac{\log p}{n}}, \lambda_n \right\}.$$

The same bounds (without λ_n) also apply to the constrained program (2.4) with radius choice $R = \|\beta^*\|_1$.

Remarks. To be clear, all the claims of Theorem 1 are deterministic. Probabilistic conditions will enter when we analyze specific statistical models and certify that the RE condition (3.3) and deviation conditions are satisfied by a random pair $(\hat{\Gamma}, \hat{\gamma})$ with high probability. We note that for the standard Lasso choice $(\hat{\Gamma}_{\text{Las}}, \hat{\gamma}_{\text{Las}})$ of this matrix–vector pair, bounds of the form (3.4) for sub-Gaussian noise are well known from past work (e.g., [2, 11, 12, 23]). The novelty of Theorem 1 is in allowing for general pairs of such surrogates, which—as shown by the examples discussed earlier—can lead to nonconvexity in the underlying M -estimator. Moreover, some interesting differences arise due to the term $\varphi(\mathbb{Q}, \sigma_\varepsilon)$, which changes depending on the nature of the model (missing, noisy, and/or dependent). As will be clarified in the sequel. Proving that the conditions of Theorem 1 are satisfied with high probability for noisy/missing data requires some nontrivial analysis involving both concentration inequalities and random matrix theory.

Note that in the presence of nonconvexity, it is possible in principle for the optimization problems (2.4) and (2.7) to have *many* global optima that are separated by large distances. Interestingly, Theorem 1 guarantees that this unpleasant feature does not arise under the stated conditions: given any two global optima $\hat{\beta}$ and $\tilde{\beta}$

of the program (2.4), Theorem 1 combined with the triangle inequality guarantees that

$$\|\widehat{\beta} - \widetilde{\beta}\|_2 \leq \|\widehat{\beta} - \beta^*\|_2 + \|\widetilde{\beta} - \beta^*\|_2 \leq 2c_0 \frac{\varphi(\mathbb{Q}, \sigma_\varepsilon)}{\alpha_1} \sqrt{\frac{k \log p}{n}}$$

[and similarly for the program (2.7)]. Consequently, under any scaling such that $\frac{k \log p}{n} = o(1)$, the set of all global optima must lie within an ℓ_2 -ball whose radius shrinks to zero.

In addition, it is worth observing that Theorem 1 makes a specific prediction for the scaling behavior of the ℓ_2 -error $\|\widehat{\beta} - \beta^*\|_2$. In order to study this scaling prediction, we performed simulations under the additive noise model described in Example 1, using the parameter setting $\Sigma_x = I$ and $\Sigma_w = \sigma_w^2 I$ with $\sigma_w = 0.2$. Panel (a) of Figure 1 provides plots³ of the error $\|\widehat{\beta} - \beta^*\|_2$ versus the sample size n , for problem dimensions $p \in \{128, 256, 512\}$. Note that for all three choices of dimensions, the error decreases to zero as the sample size n increases, showing consistency of the method. The curves also shift to the right as the dimension p increases, reflecting the natural intuition that larger problems are harder in a certain sense. Theorem 1 makes a specific prediction about this scaling behavior: in particular, if we plot the ℓ_2 -error versus the rescaled sample size $n/(k \log p)$, the curves should roughly align for different values of p . Panel (b) shows the same data re-plotted on these rescaled axes, thus verifying the predicted “stacking behavior.”

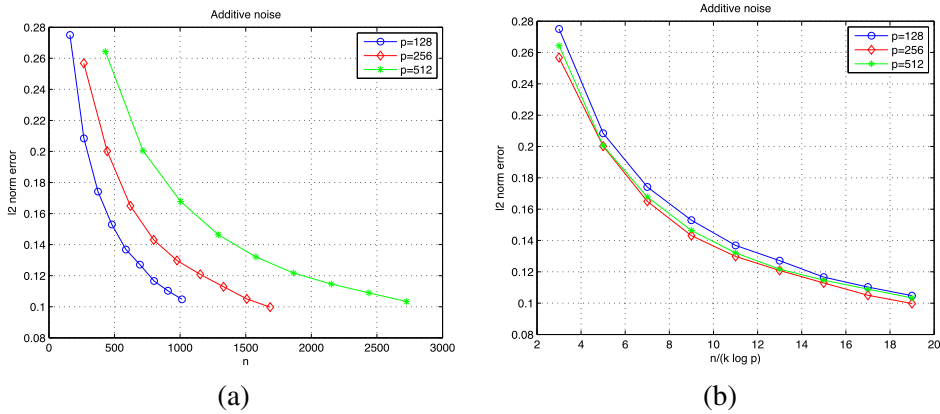


FIG. 1. Plots of the error $\|\widehat{\beta} - \beta^*\|_2$ after running projected gradient descent on the nonconvex objective, with sparsity $k \approx \sqrt{p}$. Plot (a) is an error plot for i.i.d. data with additive noise, and plot (b) shows ℓ_2 -error versus the rescaled sample size $\frac{n}{k \log p}$. As predicted by Theorem 1, the curves align for different values of p in the rescaled plot.

³Corollary 1, to be stated shortly, guarantees that the conditions of Theorem 1 are satisfied with high probability for the additive noise model. In addition, Theorem 2 to follow provides an efficient method of obtaining an accurate approximation of the global optimum.

Finally, as noted by a reviewer, the constraint $R = \|\beta^*\|_1$ in the program (2.4) is rather restrictive, since β^* is unknown. Theorem 1 merely establishes a heuristic for the scaling expected for this optimal radius. In this regard, the Lagrangian estimator (2.7) is more appealing, since it only requires choosing b_0 to be larger than $\|\beta^*\|_2$, and the conditions on the regularizer λ_n are the standard ones from past work on the Lasso.

3.1.2. *Optimization error.* Although Theorem 1 provides guarantees that hold uniformly for any global minimizer, it does not provide guidance on how to approximate such a global minimizer using a polynomial-time algorithm. Indeed, for nonconvex programs in general, gradient-type methods may become trapped in local minima, and it is impossible to guarantee that all such local minima are close to a global optimum. Nonetheless, we are able to show that for the family of programs (2.4), under reasonable conditions on $\widehat{\Gamma}$ satisfied in various settings, simple gradient methods will converge geometrically fast to a very good approximation of any global optimum. The following theorem supposes that we apply the projected gradient updates (2.14) to the constrained program (2.4), or the composite updates (2.15) to the Lagrangian program (2.7), with stepsize $\eta = 2\alpha_2$. In both cases, we assume that $n \gtrsim k \log p$, as is required for statistical consistency in Theorem 1.

THEOREM 2 (Optimization error). *Under the conditions of Theorem 1:*

(a) *For any global optimum $\widehat{\beta}$ of the constrained program (2.4), there are universal positive constants (c_1, c_2) and a contraction coefficient $\gamma \in (0, 1)$, independent of (n, p, k) , such that the gradient descent iterates (2.14) satisfy the bounds*

$$(3.5) \quad \|\beta^t - \widehat{\beta}\|_2^2 \leq \gamma^t \|\beta^0 - \widehat{\beta}\|_2^2 + c_1 \frac{\log p}{n} \|\widehat{\beta} - \beta^*\|_1^2 + c_2 \|\widehat{\beta} - \beta^*\|_2^2,$$

$$(3.6) \quad \|\beta^t - \widehat{\beta}\|_1 \leq 2\sqrt{k} \|\beta^t - \widehat{\beta}\|_2 + 2\sqrt{k} \|\widehat{\beta} - \beta^*\|_2 + 2\|\widehat{\beta} - \beta^*\|_1$$

for all $t \geq 0$.

(b) *Letting ϕ denote the objective function of Lagrangian program (2.7) with global optimum $\widehat{\beta}$, and applying composite gradient updates (2.15), there are universal positive constants (c_1, c_2) and a contraction coefficient $\gamma \in (0, 1)$, independent of (n, p, k) , such that*

$$(3.7) \quad \|\beta^t - \widehat{\beta}\|_2^2 \leq c_1 \underbrace{\|\widehat{\beta} - \beta^*\|_2^2}_{\delta^2} \quad \text{for all iterates } t \geq T,$$

where $T := c_2 \log \frac{(\phi(\beta^0) - \phi(\widehat{\beta}))}{\delta^2} / \log(1/\gamma)$.

Remarks. As with Theorem 1, these claims are deterministic in nature. Probabilistic conditions will enter into the corollaries, which involve proving that the surrogate matrices $\widehat{\Gamma}$ used for noisy, missing and/or dependent data satisfy the

lower- and upper-RE conditions with high probability. The proof of Theorem 2 itself is based on an extension of a result due to Agarwal et al. [1] on the convergence of projected gradient descent and composite gradient descent in high dimensions. Their result, as originally stated, imposed convexity of the loss function, but the proof can be modified so as to apply to the nonconvex loss functions of interest here. As noted following Theorem 1, all global minimizers of the nonconvex program (2.4) lie within a small ball. In addition, Theorem 2 guarantees that the local minimizers also lie within a ball of the same magnitude. Note that in order to show that Theorem 2 can be applied to the specific statistical models of interest in this paper, a considerable amount of technical analysis remains in order to establish that its conditions hold with high probability.

In order to understand the significance of the bounds (3.5) and (3.7), note that they provide upper bounds for the ℓ_2 -distance between the iterate β^t at time t , which is easily computed in polynomial-time, and any global optimum $\hat{\beta}$ of the program (2.4) or (2.7), which may be difficult to compute. Focusing on bound (3.5), since $\gamma \in (0, 1)$, the first term in the bound vanishes as t increases. The remaining terms involve the statistical errors $\|\hat{\beta} - \beta^*\|_q$, for $q = 1, 2$, which are controlled in Theorem 1. It can be verified that the two terms involving the statistical error on the right-hand side are bounded as $\mathcal{O}(\frac{k \log p}{n})$, so Theorem 2 guarantees that projected gradient descent produce an output that is essentially as good—in terms of statistical error—as any global optimum of the program (2.4). Bound (3.7) provides a similar guarantee for composite gradient descent applied to the Lagrangian version.

Experimentally, we have found that the predictions of Theorem 2 are borne out in simulations. Figure 2 shows the results of applying the projected gradient descent method to solve the optimization problem (2.4) in the case of additive noise




FIG. 2. Plots of the optimization error $\log(\|\beta^t - \hat{\beta}\|_2)$ and statistical error $\log(\|\beta^t - \beta^*\|_2)$ versus iteration number t , generated by running projected gradient descent on the nonconvex objective. Each plot shows the solution path for the same problem instance, using 10 different starting points. As predicted by Theorem 2, the optimization error decreases geometrically.

[panel (a)], and missing data [panel (b)]. In each case, we generated a random problem instance, and then applied the projected gradient descent method to compute an estimate $\hat{\beta}$. We then reapplied the projected gradient method to the same problem instance 10 times, each time with a random starting point, and measured the error $\|\beta^t - \hat{\beta}\|_2$ between the iterates and the first estimate (optimization error), and the error $\|\beta^t - \beta^*\|_2$ between the iterates and the truth (statistical error). Within each panel, the blue traces show the optimization error over 10 trials, and the red traces show the statistical error. On the logarithmic scale given, a geometric rate of convergence corresponds to a straight line. As predicted by Theorem 2, regardless of the starting point, the iterates $\{\beta^t\}$ exhibit geometric convergence to the same fixed point.⁴ The statistical error contracts geometrically up to a certain point, then flattens out.

3.2. Some consequences. As discussed previously, both Theorems 1 and 2 are deterministic results. Applying them to specific statistical models requires some additional work in order to establish that the stated conditions are met. We now turn to the statements of some consequences of these theorems for different cases of noisy, missing and dependent data. In all the corollaries below, the claims hold with probability greater than $1 - c_1 \exp(-c_2 \log p)$, where (c_1, c_2) are universal positive constants, independent of all other problem parameters. Note that in all corollaries, the triplet (n, p, k) is assumed to satisfy scaling of the form $n \gtrsim k \log p$, as is necessary for ℓ_2 -consistent estimation of k -sparse vectors in p dimensions.

DEFINITION 3. We say that a random matrix $X \in \mathbb{R}^{n \times p}$ is sub-Gaussian with parameters (Σ, σ^2) if:

- (a) each row $x_i^T \in \mathbb{R}^p$ is sampled independently from a zero-mean distribution with covariance Σ , and
- (b) for any unit vector $u \in \mathbb{R}^p$, the random variable $u^T x_i$ is sub-Gaussian with parameter at most σ .

For instance, if we form a random matrix by drawing each row independently from the distribution $N(0, \Sigma)$, then the resulting matrix $X \in \mathbb{R}^{n \times p}$ is a sub-Gaussian matrix with parameters $(\Sigma, \|\Sigma\|_{\text{op}})$.

3.2.1. Bounds for additive noise: i.i.d. case. We begin with the case of i.i.d. samples with additive noise, as described in Example 1.

COROLLARY 1. Suppose that we observe $Z = X + W$, where the random matrices $X, W \in \mathbb{R}^{n \times p}$ are sub-Gaussian with parameters (Σ_x, σ_x^2) , and let ε be

⁴To be precise, Theorem 2 states that the iterates will converge geometrically to a small neighborhood of all the global optima.

an i.i.d. sub-Gaussian vector with parameter σ_ε^2 . Let $\sigma_z^2 = \sigma_x^2 + \sigma_w^2$. Then under the scaling $n \gtrsim \max\{\frac{\sigma_z^4}{\lambda_{\min}^2(\Sigma_x)}, 1\}k \log p$, for the M -estimator based on the surrogates $(\widehat{\Gamma}_{\text{add}}, \widehat{\gamma}_{\text{add}})$, the results of Theorems 1 and 2 hold with parameters $\alpha_1 = \frac{1}{2}\lambda_{\min}(\Sigma_x)$ and $\varphi(\mathbb{Q}, \sigma_\varepsilon) = c_0\sigma_z(\sigma_w + \sigma_\varepsilon)\|\beta^*\|_2$, with probability at least $1 - c_1 \exp(-c_2 \log p)$.

Remarks. (a) Consequently, the ℓ_2 -error of any optimal solution $\widehat{\beta}$ satisfies the bound

$$\|\widehat{\beta} - \beta^*\|_2 \lesssim \frac{\sigma_z(\sigma_w + \sigma_\varepsilon)}{\lambda_{\min}(\Sigma_x)} \|\beta^*\|_2 \sqrt{\frac{k \log p}{n}}$$

with high probability. The prefactor in this bound has a natural interpretation as an inverse signal-to-noise ratio; for instance, when X and W are zero-mean Gaussian matrices with row covariances $\Sigma_x = \sigma_x^2 I$ and $\Sigma_w = \sigma_w^2 I$, respectively, we have $\lambda_{\min}(\Sigma_x) = \sigma_x^2$, so

$$\frac{(\sigma_w + \sigma_\varepsilon)\sqrt{\sigma_x^2 + \sigma_w^2}}{\lambda_{\min}(\Sigma_x)} = \frac{\sigma_w + \sigma_\varepsilon}{\sigma_x} \sqrt{1 + \frac{\sigma_w^2}{\sigma_x^2}}.$$

This quantity grows with the ratios σ_w/σ_x and $\sigma_\varepsilon/\sigma_x$, which measure the SNR of the observed covariates and predictors, respectively. Note that when $\sigma_w = 0$, corresponding to the case of uncorrupted covariates, the bound on ℓ_2 -error agrees with known results. See Section 4 for simulations and further discussions of the consequences of Corollary 1.

(b) We may also compare the results in (a) with bounds from past work on high-dimensional sparse regression with noisy covariates [15]. In this work, Rosenbaum and Tsybakov derive similar concentration bounds on sub-Gaussian matrices. The tolerance parameters are all $\mathcal{O}(\sqrt{\frac{\log p}{n}})$, with prefactors depending on the sub-Gaussian parameters of the matrices. In particular, in their notation,

$$v \asymp (\sigma_x\sigma_w + \sigma_w\sigma_\varepsilon + \sigma_w^2)\sqrt{\frac{\log p}{n}} \|\beta^*\|_1,$$

leading to the bound (cf. Theorem 2 of Rosenbaum and Tsybakov [15])

$$\|\widehat{\beta} - \beta^*\|_2 \lesssim \frac{v\sqrt{k}}{\lambda_{\min}(\Sigma_x)} \asymp \frac{\sigma^2}{\lambda_{\min}(\Sigma_x)} \sqrt{\frac{k \log p}{n}} \|\beta^*\|_1.$$

Extensions to unknown noise covariance. Situations may arise where the noise covariance Σ_w is unknown, and must be estimated from the data. One simple method is to assume that Σ_w is estimated from independent observations of the

noise. In this case, suppose we independently observe a matrix $W_0 \in \mathbb{R}^{n \times p}$ with n i.i.d. vectors of noise. Then we use $\widehat{\Sigma}_w = \frac{1}{n} W_0^T W_0$ as our estimate of Σ_w . A more sophisticated variant of this method (cf. Chapter 4 of Carroll et al. [3]) assumes that we observe k_i replicate measurements Z_{i1}, \dots, Z_{ik} for each x_i and form the estimator

$$(3.8) \quad \widehat{\Sigma}_w = \frac{\sum_{i=1}^n \sum_{j=1}^{k_i} (Z_{ij} - \bar{Z}_{i\cdot})(Z_{ij} - \bar{Z}_{i\cdot})^T}{\sum_{i=1}^n (k_i - 1)}.$$

Based on the estimator $\widehat{\Sigma}_w$, we form the pair $(\widetilde{\Gamma}, \widetilde{\gamma})$ such that $\widetilde{\gamma} = \frac{1}{n} Z^T y$ and $\widetilde{\Gamma} = \frac{Z^T Z}{n} - \widehat{\Sigma}_w$. In the proofs of Section 5, we will analyze the case where $\widehat{\Sigma}_w = \frac{1}{n} W_0^T W_0$ and show that the result of Corollary 1 still holds when Σ_w must be estimated from the data. Note that the estimator in equation (3.8) will also yield the same result, but the analysis is more complicated.

3.2.2. *Bounds for missing data: i.i.d. case.* Next, we turn to the case of i.i.d. samples with missing data, as discussed in Example 3. For a missing data parameter vector ρ , we define $\rho_{\max} := \max_j \rho_j$, and assume $\rho_{\max} < 1$.

COROLLARY 2. *Let $X \in \mathbb{R}^{n \times p}$ be sub-Gaussian with parameters (Σ_x, σ_x^2) , and Z the missing data matrix with parameter ρ . Let ε be an i.i.d. sub-Gaussian vector with parameter σ_ε^2 . If $n \gtrsim \max\left(\frac{1}{(1-\rho_{\max})^4} \frac{\sigma_x^4}{\lambda_{\min}^2(\Sigma_x)}, 1\right) k \log p$, then Theorems 1 and 2 hold with probability at least $1 - c_1 \exp(-c_2 \log p)$ for $\alpha_1 = \frac{1}{2} \lambda_{\min}(\Sigma_x)$ and $\varphi(\mathbb{Q}, \sigma_\varepsilon) = c_0 \frac{\sigma_x}{1-\rho_{\max}} \left(\sigma_\varepsilon + \frac{\sigma_x}{1-\rho_{\max}}\right) \|\beta^*\|_2$.*

Remarks. Suppose X is a Gaussian random matrix and $\rho_j = \rho$ for all j . In this case, the ratio $\frac{\sigma_x^2}{\lambda_{\min}(\Sigma_x)} = \frac{\lambda_{\max}(\Sigma_x)}{\lambda_{\min}(\Sigma_x)} = \kappa(\Sigma_x)$ is the condition number of Σ_x . Then

$$\frac{\varphi(\mathbb{Q}, \sigma_\varepsilon)}{\alpha} \asymp \left(\frac{1}{\lambda_{\min}(\Sigma_x)} \frac{\sigma_x \sigma_\varepsilon}{1-\rho} + \frac{\kappa(\Sigma_x)}{(1-\rho)^2} \right) \|\beta^*\|_2,$$

a quantity that depends on both the conditioning of Σ_x , and the fraction $\rho \in [0, 1)$ of missing data. We will consider the results of Corollary 2 applied to this example in the simulations of Section 4.

Extensions to unknown ρ . As in the additive noise case, we may wish to consider the case when the missing data parameters ρ are not observed and must be estimated from the data. For each $j = 1, 2, \dots, p$, we estimate ρ_j using $\widehat{\rho}_j$, the empirical average of the number of observed entries per column. Let $\widehat{\rho} \in \mathbb{R}^p$ denote the resulting estimator of ρ . Naturally, we use the pair of estimators $(\widetilde{\Gamma}, \widetilde{\gamma})$ defined by

$$(3.9) \quad \widetilde{\Gamma} = \frac{Z^T Z}{n} \oplus \widetilde{M} \quad \text{and} \quad \widetilde{\gamma} = \frac{1}{n} Z^T y \oplus (\mathbf{1} - \widehat{\rho}),$$

where

$$\widetilde{M}_{ij} = \begin{cases} (1 - \widehat{\rho}_i)(1 - \widehat{\rho}_j), & \text{if } i \neq j, \\ 1 - \widehat{\rho}_i, & \text{if } i = j. \end{cases}$$

We will show in Section 5 that Corollary 2 holds when ρ is estimated by $\widehat{\rho}$.

3.2.3. *Bounds for dependent data.* Turning to the case of dependent data, we consider the setting where the rows of X are drawn from a stationary vector autoregressive (VAR) process according to

$$(3.10) \quad x_{i+1} = Ax_i + v_i \quad \text{for } i = 1, 2, \dots, n - 1,$$

where $v_i \in \mathbb{R}^p$ is a zero-mean noise vector with covariance matrix Σ_v , and $A \in \mathbb{R}^{p \times p}$ is a driving matrix with spectral norm $\|A\|_2 < 1$. We assume the rows of X are drawn from a Gaussian distribution with covariance Σ_x , such that $\Sigma_x = A\Sigma_x A^T + \Sigma_v$. Hence, the rows of X are identically distributed but not independent, with the choice $A = 0$ giving rise to the i.i.d. scenario. Corollaries 3 and 4 correspond to the case of additive noise and missing data for a Gaussian VAR process.

COROLLARY 3. *Suppose the rows of X are drawn according to a Gaussian VAR process with driving matrix A . Suppose the additive noise matrix W is i.i.d. with Gaussian rows, and let ε be an i.i.d. sub-Gaussian vector with parameter σ_ε^2 . If $n \gtrsim \max(\frac{\zeta^4}{\lambda_{\min}^2(\Sigma_x)}, 1)k \log p$, with $\zeta^2 = \|\Sigma_w\|_{\text{op}} + \frac{2\|\Sigma_x\|_{\text{op}}}{1 - \|A\|_{\text{op}}}$, then Theorems 1 and 2 hold with probability at least $1 - c_1 \exp(-c_2 \log p)$ for $\alpha_1 = \frac{1}{2}\lambda_{\min}(\Sigma_x)$ and $\varphi(\mathbb{Q}, \sigma_\varepsilon) = c_0(\sigma_\varepsilon \zeta + \zeta^2)\|\beta^*\|_2$.*

COROLLARY 4. *Suppose the rows of X are drawn according to a Gaussian VAR process with driving matrix A , and Z is the observed matrix subject to missing data, with parameter ρ . Let ε be an i.i.d. sub-Gaussian vector with parameter σ_ε^2 . If $n \gtrsim \max(\frac{\zeta'^4}{\lambda_{\min}^2(\Sigma_x)}, 1)k \log p$, with $\zeta'^2 = \frac{1}{(1 - \rho_{\max})^2} \frac{2\|\Sigma_x\|_{\text{op}}}{1 - \|A\|_{\text{op}}}$, then Theorems 1 and 2 hold with probability at least $1 - c_1 \exp(-c_2 \log p)$ for $\alpha_1 = \frac{1}{2}\lambda_{\min}(\Sigma_x)$ and $\varphi(\mathbb{Q}, \sigma_\varepsilon) = c_0(\sigma_\varepsilon \zeta' + \zeta'^2)\|\beta^*\|_2$.*

REMARKS. Note that the scaling and the form of φ in Corollaries 2–4 are very similar, except with different effective variances $\sigma^2 = \frac{\sigma_x^2}{(1 - \rho_{\max})^2}$, ζ^2 or ζ'^2 , depending on the type of corruption in the data. As we will see in Section 5, the proofs involve verifying the deviation conditions (3.2) using similar techniques. On the other hand, the proof of Corollary 1 proceeds via deviation condition (3.1), which produces a tighter bound.

Note that we may extend the cases of dependent data to situations when Σ_w and ρ are unknown and must be estimated from the data. The proofs of these extensions are identical to the i.i.d case, so we will omit them.

3.3. *Application to graphical model inverse covariance estimation.* The problem of inverse covariance estimation for a Gaussian graphical model is also related to the Lasso. Meinshausen and Bühlmann [10] prescribed a way to recover the support of the precision matrix Θ when each column of Θ is k -sparse, via linear regression and the Lasso. More recently, Yuan [22] proposed a method for estimating Θ using the Dantzig selector, and obtained error bounds on $\|\hat{\Theta} - \Theta\|_1$ when the columns of Θ are bounded in ℓ_1 . Both of these results assume that X is fully-observed and has i.i.d. rows.

Suppose we are given a matrix $X \in \mathbb{R}^{n \times p}$ of samples from a multivariate Gaussian distribution, where each row is distributed according to $N(0, \Sigma)$. We assume the rows of X are either i.i.d. or sampled from a Gaussian VAR process. Based on the modified Lasso of the previous section, we devise a method to estimate Θ based on a corrupted observation matrix Z , when Θ is sparse. Our method bears similarity to the method of Yuan [22], but is valid in the case of corrupted data, and does not require an ℓ_1 column bound. Let X^j denote the j th column of X , and let X^{-j} denote the matrix X with j th column removed. By standard results on Gaussian graphical models, there exists a vector $\theta^j \in \mathbb{R}^{p-1}$ such that

$$(3.11) \quad X^j = X^{-j}\theta^j + \varepsilon^j,$$

where ε^j is a vector of i.i.d. Gaussians and $\varepsilon^j \perp\!\!\!\perp X^{-j}$ for each j . If we define $a_j := -(\Sigma_{jj} - \Sigma_{j,-j}\theta^j)^{-1}$, we can verify that $\Theta_{j,-j} = a_j\theta^j$. Our algorithm, described below, forms estimates $\hat{\theta}^j$ and \hat{a}_j for each j , then combines the estimates to obtain an estimate $\hat{\Theta}_{j,-j} = \hat{a}_j\hat{\theta}^j$.

In the additive noise case, we observe the matrix $Z = X + W$. From the equations (3.11), we obtain $Z^j = X^{-j}\theta^j + (\varepsilon^j + W^j)$. Note that $\delta^j = \varepsilon^j + W^j$ is a vector of i.i.d. Gaussians, and since $X \perp\!\!\!\perp W$, we have $\delta^j \perp\!\!\!\perp X^{-j}$. Hence, our results on covariates with additive noise allow us to recover θ^j from Z . We can verify that this reduces to solving the program (2.4) or (2.7) with the pair $(\hat{\Gamma}^{(j)}, \hat{\gamma}^{(j)}) = (\hat{\Sigma}_{-j,-j}, \frac{1}{n}Z^{-jT}Z^j)$, where $\hat{\Sigma} = \frac{1}{n}Z^T Z - \Sigma_w$.

When Z is a missing-data version of X , we similarly estimate the vectors θ^j via equation (3.11), using our results on the Lasso with missing covariates. Here, both covariates and responses are subject to missing data, but this makes no difference in our theoretical results. For each j , we use the pair

$$(\hat{\Gamma}^{(j)}, \hat{\gamma}^{(j)}) = \left(\hat{\Sigma}_{-j,-j}, \frac{1}{n}Z^{-jT}Z^j \ominus (\mathbf{1} - \rho^{-j})(1 - \rho_j) \right),$$

where $\hat{\Sigma} = \frac{1}{n}Z^T Z \ominus M$, and M is defined as in Example 3.

To obtain the estimate $\hat{\Theta}$, we therefore propose the following procedure, based on the estimators $\{(\hat{\Gamma}^{(j)}, \hat{\gamma}^{(j)})\}_{j=1}^p$ and $\hat{\Sigma}$.

ALGORITHM 3.1. (1) Perform p linear regressions of the variables Z^j upon the remaining variables Z^{-j} , using the program (2.4) or (2.7) with the estimators $(\hat{\Gamma}^{(j)}, \hat{\gamma}^{(j)})$, to obtain estimates $\hat{\theta}^j$ of θ^j .

(2) Estimate the scalars a_j using the quantity $\hat{a}_j := -(\hat{\Sigma}_{jj} - \hat{\Sigma}_{j,-j}\hat{\theta}^j)^{-1}$, based on the estimator $\hat{\Sigma}$. Form $\hat{\Theta}$ with $\hat{\Theta}_{j,-j} = \hat{a}_j\hat{\theta}^j$ and $\hat{\Theta}_{jj} = -\hat{a}_j$.

(3) Set $\hat{\Theta} = \arg \min_{\Theta \in S^p} \|\Theta - \hat{\Theta}\|_1$, where S^p is the set of symmetric matrices.

Note that the minimization in step (3) is a linear program, so is easily solved with standard methods. We have the following corollary about $\hat{\Theta}$:

COROLLARY 5. *Suppose the columns of the matrix Θ are k -sparse, and suppose the condition number $\kappa(\Theta)$ is nonzero and finite. Suppose we have*

$$(3.12) \quad \|\hat{\gamma}^{(j)} - \hat{\Gamma}^{(j)}\theta^j\|_\infty \leq \varphi(\mathbb{Q}, \sigma_\varepsilon)\sqrt{\frac{\log p}{n}} \quad \forall j,$$

and suppose we have the following additional deviation condition on $\hat{\Sigma}$:

$$(3.13) \quad \|\hat{\Sigma} - \Sigma\|_{\max} \leq c\varphi(\mathbb{Q}, \sigma_\varepsilon)\sqrt{\frac{\log p}{n}}.$$

Finally, suppose the lower-RE condition holds uniformly over the matrices $\hat{\Gamma}^{(j)}$ with the scaling (3.3). Then under the estimation procedure of Algorithm 3.1, there exists a universal constant c_0 such that

$$\|\hat{\Theta} - \Theta\|_{\text{op}} \leq \frac{c_0\kappa^2(\Sigma)}{\lambda_{\min}(\Sigma)} \left(\frac{\varphi(\mathbb{Q}, \sigma_\varepsilon)}{\lambda_{\min}(\Sigma)} + \frac{\varphi(\mathbb{Q}, \sigma_\varepsilon)}{\alpha_1} \right) k\sqrt{\frac{\log p}{n}}.$$

REMARKS. Note that Corollary 5 is again a deterministic result, with parallel structure to Theorem 1. Furthermore, the deviation bounds (3.12) and (3.13) hold for all scenarios considered in Section 3.2 above, using Corollaries 1–4 for the first two inequalities, and a similar bounding technique for $\|\hat{\Sigma} - \Sigma\|_{\max}$; and the lower-RE condition holds over all matrices $\hat{\Gamma}^{(j)}$ by the same technique used to establish the lower-RE condition for $\hat{\Gamma}$. The uniformity of the lower-RE bound over all sub-matrices holds because

$$0 < \lambda_{\min}(\Sigma) \leq \lambda_{\min}(\Sigma_{-j,-j}) \leq \lambda_{\max}(\Sigma_{-j,-j}) \leq \lambda_{\max}(\Sigma) < \infty.$$

Hence, the error bound in Corollary 5 holds with probability at least $1 - c_1 \exp(-c_2 \log p)$ when $n \gtrsim k \log p$, for the appropriate values of φ and α_1 .

4. Simulations. In this section, we report some additional simulation results to confirm that the scalings predicted by our theory are sharp. In Figure 1 following Theorem 1, we showed that the error curves align when plotted against a suitably rescaled sample size, in the case of additive noise perturbations. Panel (a) of Figure 3 shows these same types of rescaled curves for the case of missing data, with sparsity $k \approx \sqrt{p}$, covariate matrix $\Sigma_x = I$, and missing fraction $\rho = 0.2$, whereas panel (b) shows the rescaled plots for the vector autoregressive case with additive

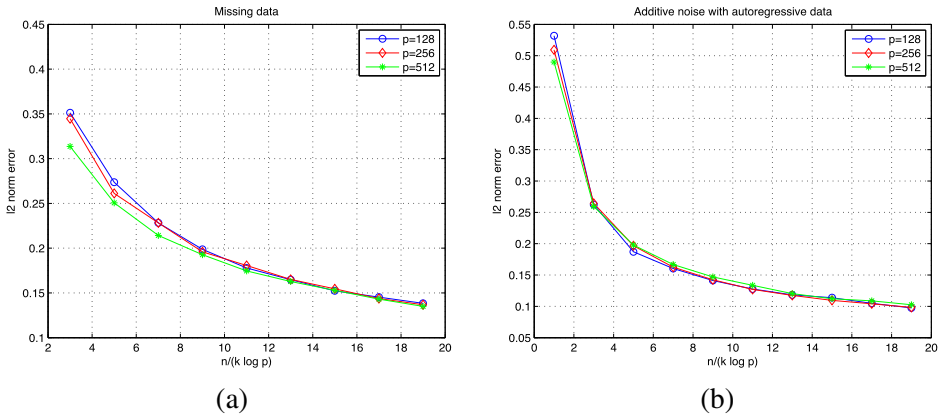


FIG. 3. Plots of the error $\|\hat{\beta} - \beta^*\|_2$ after running projected gradient descent on the nonconvex objective, with sparsity $k \approx \sqrt{p}$. In all cases, we plotted the error versus the rescaled sample size $\frac{n}{k \log p}$. As predicted by Theorems 1 and 2, the curves align for different values of p when plotted in this rescaled manner. (a) Missing data case with i.i.d. covariates. (b) Vector autoregressive data with additive noise. Each point represents an average over 100 trials.

noise perturbations, using a driving matrix A with $\|A\|_{\text{op}} = 0.2$. Each point corresponds to an average over 100 trials. Once again, we see excellent agreement with the scaling law provided by Theorem 1.

We also ran simulations to verify the form of the function $\varphi(\mathbb{Q}, \sigma_\varepsilon)$ appearing in Corollaries 1 and 2. In the additive noise setting for i.i.d. data, we set $\Sigma_x = I$ and ε equal to i.i.d. Gaussian noise with $\sigma_\varepsilon = 0.5$. For a fixed value of the parameters $p = 256$ and $k \approx \log p$, we ran the projected gradient descent algorithm for different values of $\sigma_w \in (0.1, 0.3)$, such that $\Sigma_w = \sigma_w^2 I$ and $n \approx 60(1 + \sigma_w^2)^2 k \log p$, with $\|\beta^*\|_2 = 1$. According to the theory, $\frac{\varphi(\mathbb{Q}, \sigma_\varepsilon)}{\alpha} \asymp (\sigma_w + 0.5)\sqrt{1 + \sigma_w^2}$, so that

$$\|\hat{\beta} - \beta^*\|_2 \lesssim (\sigma_w + 0.5)\sqrt{1 + \sigma_w^2} \sqrt{\frac{k \log p}{(1 + \sigma_w^2)^2 k \log p}} \asymp \frac{\sigma_w + 0.5}{\sqrt{1 + \sigma_w^2}}.$$

In order to verify this theoretical prediction, we plotted σ_w versus the rescaled error $\frac{\sqrt{1 + \sigma_w^2}}{\sigma_w + 0.5} \|\hat{\beta} - \beta^*\|_2$. As shown by Figure 4(a), the curve is roughly constant, as predicted by the theory.

Similarly, in the missing data setting for i.i.d. data, we set $\Sigma_x = I$ and ε equal to i.i.d. Gaussian noise with $\sigma_\varepsilon = 0.5$. For a fixed value of the parameters $p = 128$ and $k \approx \log p$, we ran simulations for different values of the missing data parameter $\rho \in (0, 0.3)$, such that $n \approx \frac{60}{(1 - \rho)^4} k \log p$. According to the theory, $\frac{\varphi(\mathbb{Q}, \sigma_\varepsilon)}{\alpha} \asymp \frac{\sigma_\varepsilon}{1 - \rho} + \frac{1}{(1 - \rho)^2}$. Consequently, with our specified scalings of (n, p, k) , we should expect a




FIG. 4. (a) Plot of the rescaled ℓ_2 -error $\frac{\sqrt{1+\sigma_w^2}}{\sigma_w+0.5} \|\hat{\beta} - \beta^*\|_2$ versus the additive noise standard deviation σ_w for the i.i.d. model with additive noise. (b) Plot of the rescaled ℓ_2 -error $\frac{\|\hat{\beta} - \beta^*\|_2}{1+0.5(1-\rho)}$ versus the missing fraction ρ for the i.i.d. model with missing data. Both curves are roughly constant, showing that our error bounds on $\|\hat{\beta} - \beta^*\|_2$ exhibit the proper scaling. Each point represents an average over 200 trials.

bound of the form

$$\|\hat{\beta} - \beta^*\|_2 \lesssim \frac{\varphi(\mathbb{Q}, \sigma_\varepsilon)}{\alpha} \sqrt{\frac{k \log p}{n}} \asymp 1 + 0.5(1 - \rho).$$

The plot of ρ versus the rescaled error $\frac{\|\hat{\beta} - \beta^*\|_2}{1+0.5(1-\rho)}$ is shown in Figure 4(b). The curve is again roughly constant, agreeing with theoretical results.

Finally, we studied the behavior of the inverse covariance matrix estimation algorithm on three types of Gaussian graphical models:

(a) *Chain-structured graphs.* In this case, all nodes of the graph are arranged in a linear chain. Hence, each node (except the two end nodes) has degree $k = 2$. The diagonal entries of Θ are set equal to 1, and all entries corresponding to links in the chain are set equal to 0.1. Then Θ is rescaled so $\|\Theta\|_{\text{op}} = 1$.

(b) *Star-structured graphs.* In this case, all nodes are connected to a central node, which has degree $k \approx 0.1p$. All other nodes have degree 1. The diagonal entries of Θ are set equal to 1, and all entries corresponding to edges in the graph are set equal to 0.1. Then Θ is rescaled so $\|\Theta\|_{\text{op}} = 1$.

(c) *Erdős–Renyi graphs.* This example comes from Rothman et al. [16]. For a sparsity parameter $k \approx \log p$, we randomly generate the matrix Θ by first generating the matrix B such that the diagonal entries are 0, and all other entries are independently equal to 0.5 with probability k/p , and 0 otherwise. Then δ is chosen so that $\Theta = B + \delta I$ has condition number p . Finally, Θ is rescaled so $\|\Theta\|_{\text{op}} = 1$.

After generating the matrix X of n i.i.d. samples from the appropriate graphical model, with covariance matrix $\Sigma_x = \Theta^{-1}$, we generated the corrupted matrix $Z = X + W$ with $\Sigma_w = (0.2)^2 I$ in the additive noise case, or the missing data matrix Z with $\rho = 0.2$ in the missing data case.

Panels (a) and (c) in Figure 5 show the rescaled ℓ_2 -error $\frac{1}{\sqrt{k}} \|\hat{\Theta} - \Theta\|_{\text{op}}$ plotted against the sample size n for a chain-structured graph. In panels (b) and (d), we have ℓ_2 -error plotted against the rescaled sample size, $n/(k \log p)$. Once again, we see good agreement with the theoretical predictions. We have obtained qualitatively similar results for the star and Erdős–Renyi graphs.

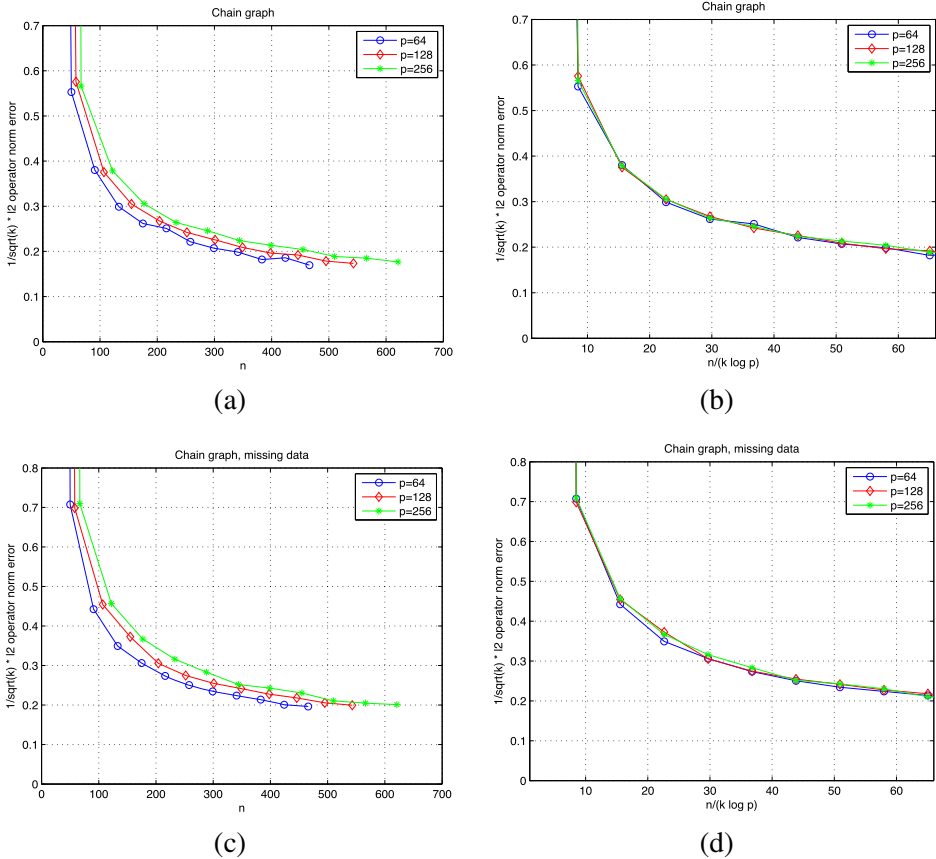


FIG. 5. (a) Plots of the error $\|\hat{\Theta} - \Theta\|_{\text{op}}$ after running projected gradient descent on the nonconvex objective for a chain-structured Gaussian graphical model with additive noise. As predicted by Theorems 1 and 2, all curves align when the error is rescaled by $\frac{1}{\sqrt{k}}$ and plotted against the ratio $\frac{n}{k \log p}$, as shown in (b). Plots (c) and (d) show the results of simulations on missing data sets. Each point represents the average over 50 trials.

5. Proofs. In this section, we prove our two main theorems. For the more technical proofs of the corollaries, see the supplementary Appendix [9].

5.1. *Proof of Theorem 1.* Let $\mathcal{L}(\beta) = \frac{1}{2}\beta^T \widehat{\Gamma} \beta - \langle \widehat{\gamma}, \beta \rangle + \lambda_n \|\beta\|_1$ denote the loss function to be minimized. This definition captures both the estimator (2.4) with $\lambda_n = 0$ and the estimator (2.7) with the choice of λ_n given in the theorem statement. For either estimator, we are guaranteed that β^* is feasible and $\widehat{\beta}$ is optimal for the program, so $\mathcal{L}(\widehat{\beta}) \leq \mathcal{L}(\beta^*)$. Indeed, in the regularized case, the k -sparsity of β^* implies that $\|\beta^*\|_1 \leq \sqrt{k}\|\beta^*\|_2 \leq b_0\sqrt{k}$. Defining the error vector $\widehat{v} := \widehat{\beta} - \beta^*$ and performing some algebra leads to the equivalent inequality

$$(5.1) \quad \frac{1}{2}\widehat{v}^T \widehat{\Gamma} \widehat{v} \leq \langle \widehat{v}, \widehat{\gamma} - \widehat{\Gamma} \beta^* \rangle + \lambda_n \{ \|\beta^*\|_1 - \|\beta^* + \widehat{v}\|_1 \}.$$

In the remainder of the proof, we first derive an upper bound for the right-hand side of this inequality. We then use this upper bound and the lower-RE condition to show that the error vector \widehat{v} must satisfy the inequality

$$(5.2) \quad \|\widehat{v}\|_1 \leq 8\sqrt{k}\|\widehat{v}\|_2.$$

Finally, we combine inequality (5.2) with the lower-RE condition to derive a lower bound on the left-hand side of the basic inequality (5.1). Combined with our earlier upper bound on the right-hand side, some algebra yields the claim.

Upper bound on right-hand side. We first upper-bound the right-hand side of inequality (5.1). Hölder’s inequality gives $\langle \widehat{v}, \widehat{\gamma} - \widehat{\Gamma} \beta^* \rangle \leq \|\widehat{v}\|_1 \|\widehat{\gamma} - \widehat{\Gamma} \beta^*\|_\infty$. By the triangle inequality, we have

$$\|\widehat{\gamma} - \widehat{\Gamma} \beta^*\|_\infty \leq \|\widehat{\gamma} - \Sigma_x \beta^*\|_\infty + \|(\Sigma_x - \widehat{\Gamma})\beta^*\|_\infty \stackrel{(i)}{\leq} 2\varphi(\mathbb{Q}, \sigma_\varepsilon) \sqrt{\frac{\log p}{n}},$$

where inequality (i) follows from the deviation conditions (3.2). Combining the pieces, we conclude that

$$(5.3) \quad \begin{aligned} \langle \widehat{v}, \widehat{\gamma} - \widehat{\Gamma} \beta^* \rangle &\leq 2\|\widehat{v}\|_1 \varphi(\mathbb{Q}, \sigma_\varepsilon) \sqrt{\frac{\log p}{n}} \\ &= (\|\widehat{v}_S\|_1 + \|\widehat{v}_{S^c}\|_1) 2\varphi(\mathbb{Q}, \sigma_\varepsilon) \sqrt{\frac{\log p}{n}}. \end{aligned}$$

On the other hand, we have

$$(5.4) \quad \begin{aligned} \|\beta^* + \widehat{v}\|_1 - \|\beta^*\|_1 &\geq \{ \|\beta_S^*\|_1 - \|\widehat{v}_S\|_1 \} + \|\widehat{v}_{S^c}\|_1 - \|\beta^*\|_1 \\ &= \|\widehat{v}_{S^c}\|_1 - \|\widehat{v}_S\|_1, \end{aligned}$$

where we have exploited the sparsity of β^* and applied the triangle inequality. Combining the pieces, we conclude that the right-hand side of inequality (5.1) is upper-bounded by

$$(5.5) \quad 2\varphi(\mathbb{Q}, \sigma_\varepsilon) \sqrt{\frac{\log p}{n}} (\|\widehat{v}_S\|_1 + \|\widehat{v}_{S^c}\|_1) + \lambda_n \{ \|\widehat{v}_S\|_1 - \|\widehat{v}_{S^c}\|_1 \},$$

a bound that holds for any nonnegative choice of λ_n .

Proof of inequality (5.2). We first consider the constrained program (2.4), with $R = \|\beta^*\|_1$, so $\|\widehat{\beta}\|_1 = \|\beta^* + \widehat{v}\|_1 \leq \|\beta^*\|_1$. Combined with inequality (5.4), we conclude that $\|\widehat{v}_{S^c}\|_1 \leq \|\widehat{v}_S\|_1$. Consequently, we have the inequality $\|\widehat{v}\|_1 \leq 2\|\widehat{v}_S\|_1 \leq 2\sqrt{k}\|\widehat{v}\|_2$, which is a slightly stronger form of the bound (5.2).

For the regularized estimator (2.7), we first note that our choice of λ_n guarantees that the term (5.5) is at most $\frac{3\lambda_n}{2}\|\widehat{v}_S\|_1 - \frac{\lambda_n}{2}\|\widehat{v}_{S^c}\|_1$. Returning to the basic inequality, we apply the lower-RE condition to lower-bound the left-hand side, thereby obtaining the inequality

$$-\frac{\tau}{2}\|\widehat{v}\|_1^2 \leq \frac{1}{2}(\alpha_1\|\widehat{v}\|_2^2 - \tau\|\widehat{v}\|_1^2) \leq \frac{3\lambda_n}{2}\|\widehat{v}_S\|_1 - \frac{\lambda_n}{2}\|\widehat{v}_{S^c}\|_1.$$

By the triangle inequality, we have $\|\widehat{v}\|_1 \leq \|\widehat{\beta}\|_1 + \|\beta^*\|_1 \leq 2b_0\sqrt{k}$. Since we have assumed $\sqrt{k}\tau(n, p) \leq \frac{\varphi(\mathbb{Q}, \sigma_\varepsilon)}{b_0} \sqrt{\frac{\log p}{n}}$, we are guaranteed that

$$\frac{\tau(n, p)}{2}\|\widehat{v}\|_1^2 \leq \varphi(\mathbb{Q}, \sigma_\varepsilon) \sqrt{\frac{\log p}{n}} \|\widehat{v}\|_1 \leq \frac{\lambda_n}{4}\|\widehat{v}\|_1$$

by our choice of λ_n . Combining the pieces, we conclude that

$$0 \leq \frac{3\lambda_n}{2}\|\widehat{v}_S\|_1 - \frac{\lambda_n}{2}\|\widehat{v}_{S^c}\|_1 + \frac{\lambda_n}{4}(\|\widehat{v}_S\|_1 + \|\widehat{v}_{S^c}\|_1) = \frac{7\lambda_n}{4}\|\widehat{v}_S\|_1 - \frac{\lambda_n}{4}\|\widehat{v}_{S^c}\|_1$$

and rearranging implies $\|\widehat{v}_{S^c}\|_1 \leq 7\|\widehat{v}_S\|_1$, from which we conclude that $\|\widehat{v}\|_1 \leq 8\sqrt{k}\|\widehat{v}\|_2$, as claimed.

Lower bound on left-hand side. We now derive a lower bound on the left-hand side of inequality (5.1). Combining inequality (5.2) with the RE condition (2.12) gives

$$(5.6) \quad \widehat{v}^T \widehat{\Gamma} \widehat{v} \geq \alpha_1 \|\widehat{v}\|_2^2 - \tau(n, p) \|\widehat{v}\|_1^2 \geq \{\alpha_1 - 64k\tau(n, p)\} \|\widehat{v}\|_2^2 \geq \frac{\alpha_1}{2} \|\widehat{v}\|_2^2,$$

where the final step uses our assumption that $k\tau(n, p) \leq \frac{\alpha_1}{128}$.

Finally, combining bounds (5.5), (5.2) and (5.6) yields

$$\begin{aligned} \frac{\alpha_1}{4} \|\widehat{v}\|_2^2 &\leq 2 \max \left\{ 2\varphi(\mathbb{Q}, \sigma_\varepsilon) \sqrt{\frac{\log p}{n}}, \lambda_n \right\} \|\widehat{v}\|_1 \\ &\leq 32\sqrt{k} \max \left\{ \varphi(\mathbb{Q}, \sigma_\varepsilon) \sqrt{\frac{\log p}{n}}, \lambda_n \right\} \|\widehat{v}\|_2, \end{aligned}$$

giving inequality (3.4a). Using inequality (5.2) again gives inequality (3.4b).

5.2. *Proof of Theorem 2.* We begin by proving the claims for the constrained problem, and projected gradient descent. For the ℓ_2 -error bound, we make use of Theorem 1 in the pre-print of Agarwal et al. [1]. Their theory, as originally stated, requires that the loss function be convex, but a careful examination of their proof shows that their arguments hinge on restricted strong convexity and smoothness assumptions, corresponding to a more general version of the lower- and upper-RE conditions given here. Apart from these conditions, the proof exploits the fact that the sub-problems defining the gradient updates (2.14) and (2.15) are convex. Since the loss function itself appears only in a linear term, their theory still applies.

In order to apply Theorem 1 in their paper, we first need to compute the tolerance parameter ε^2 defined there; since β^* is supported on the set S with $|S| = k$ and the RE conditions hold with $\tau \asymp \frac{\log p}{n}$, we find that

$$\begin{aligned} \varepsilon^2 &\leq c \frac{\log p}{\alpha_2 n} (\sqrt{k} \|\widehat{\beta} - \beta^*\|_2 + 2\|\widehat{\beta} - \beta^*\|_1)^2 \\ &\leq c'_2 \frac{k \log p}{\alpha_2 n} \|\widehat{\beta} - \beta^*\|_2^2 + c_1 \frac{\log p}{\alpha_2 n} \|\widehat{\beta} - \beta^*\|_1^2 \\ &\leq c_2 \|\widehat{\beta} - \beta^*\|_2^2 + c_1 \frac{\log p}{\alpha_2 n} \|\widehat{\beta} - \beta^*\|_1^2, \end{aligned}$$

where the final inequality makes use of the assumption that $n \gtrsim k \log p$. Similarly, we may compute the contraction coefficient to be

$$(5.7) \quad \gamma = \left(1 - \frac{\alpha_1}{\alpha_2} + \frac{c_1 k \log p}{\alpha_2 n}\right) \left(1 - \frac{c_2 k \log p}{\alpha_2 n}\right)^{-1},$$

so $\gamma \in (0, 1)$ for $n \gtrsim k \log p$.

We now establish the ℓ_1 -error bound. First, let $\Delta^t := \beta^t - \beta^*$. Since β^t is feasible and $\widehat{\beta}$ is optimal with an active constraint, we have $\|\beta^t\|_1 \leq \|\widehat{\beta}\|_1$. Applying the triangle inequality gives

$$\begin{aligned} \|\widehat{\beta}\|_1 &\leq \|\beta^*\|_1 + \|\widehat{\beta} - \beta^*\|_1 = \|\beta_S^*\|_1 + \|\widehat{\beta} - \beta^*\|_1, \\ \|\beta^t\|_1 &= \|\beta^* + \Delta^t\|_1 \geq \|\beta_S^* + \Delta_{Sc}^t\|_1 - \|\Delta_S^t\|_1 = \|\beta_S^*\|_1 + \|\Delta_{Sc}^t\|_1 - \|\Delta_S^t\|_1; \end{aligned}$$

combining the bounds yields $\|\Delta_{Sc}^t\|_1 \leq \|\Delta_S^t\|_1 + \|\widehat{\beta} - \beta^*\|_1$. Then

$$\|\Delta^t\|_1 \leq 2\|\Delta_S^t\|_1 + \|\widehat{\beta} - \beta^*\|_1 \leq 2\sqrt{k}\|\Delta^t\|_2 + \|\widehat{\beta} - \beta^*\|_1,$$

so

$$\|\beta^t - \widehat{\beta}\|_1 \leq \|\widehat{\beta} - \beta^*\|_1 + \|\Delta^t\|_1 \leq 2\sqrt{k}(\|\beta^t - \widehat{\beta}\|_2 + \|\widehat{\beta} - \beta^*\|_2) + 2\|\widehat{\beta} - \beta^*\|_1.$$

Turning to the Lagrangian version, we exploit Theorem 2 in Agarwal et al. [1], with \mathcal{M} corresponding to the subspace of all vectors with support contained within

the support set of β^* . With this choice, we have $\psi(\mathcal{M}) = \sqrt{k}$, and the contraction coefficient γ takes the previous form (5.7), so that the assumption $n \gtrsim k \log p$ guarantees that $\gamma \in (0, 1)$. It remains to verify that the requirements are satisfied. From the conditions in our Theorem 2 and using the notation of Agarwal et al. [1], we have $\beta(\mathcal{M}) = \mathcal{O}(\frac{\log p}{n})$ and $\bar{\rho} = \sqrt{k}$, and the condition $n \gtrsim k \log p$ implies that $\xi(\mathcal{M}) = \mathcal{O}(1)$. Putting together the pieces, we find that the compound tolerance parameter ε^2 satisfies the bound $\varepsilon^2 = \mathcal{O}(\frac{k \log p}{n} \|\hat{\beta} - \beta^*\|_2^2) = \mathcal{O}(\|\hat{\beta} - \beta^*\|_2^2)$, so the claim follows.

6. Discussion. In this paper, we formulated an ℓ_1 -constrained minimization problem for sparse linear regression on corrupted data. The source of corruption may be additive noise or missing data, and although the resulting objective is not generally convex, we showed that projected gradient descent is guaranteed to converge to a point within statistical precision of the optimum. In addition, we established ℓ_1 - and ℓ_2 -error bounds that hold with high probability when the data are drawn i.i.d. from a sub-Gaussian distribution, or drawn from a Gaussian vector autoregressive process. Finally, we applied our methods to sparse inverse covariance estimation for a Gaussian graphical model with corruptions, and obtained spectral norm rates of the same order as existing rates for uncorrupted, i.i.d. data.

Future directions of research include studying more general types of dependencies or corruption in the covariates of regression, such as more general types of multiplicative noise, and performing sparse linear regression for corrupted data with additive noise when the noise covariance is unknown and replicates of the data may be unavailable. As pointed out by a reviewer, it would also be interesting to study the performance of our algorithms on data that are not sub-Gaussian, or even under model mismatch. In addition, one might consider other loss functions, where it is more difficult to correct the objective for corrupted covariates. Finally, it remains to be seen whether or not our techniques—used to show that certain non-convex problems can be solved to statistical precision—can be applied more broadly.

Acknowledgments. The authors thank Alekh Agarwal, Sahand Negahban, John Duchi and Alexandre Tsybakov for useful discussions and guidance. They are also grateful to the Associate Editor and anonymous referees for improvements on the paper.

SUPPLEMENTARY MATERIAL

Supplementary material for: High-dimensional regression with noisy and missing data: Provable guarantees with nonconvexity (DOI: [10.1214/12-AOS1018SUPP](https://doi.org/10.1214/12-AOS1018SUPP); .pdf). Due to space constraints, we have relegated technical details of the remaining proofs to the supplement [9].

REFERENCES

- [1] AGARWAL, A., NEGAHBAN, S. and WAINWRIGHT, M. J. (2012). Fast global convergence of gradient methods for high-dimensional statistical recovery. Available at <http://arxiv.org/abs/1104.4824>.
- [2] BICKEL, P. J., RITOV, Y. and TSYBAKOV, A. B. (2009). Simultaneous analysis of Lasso and Dantzig selector. *Ann. Statist.* **37** 1705–1732. [MR2533469](#)
- [3] CARROLL, R. J., RUPPERT, D. and STEFANSKI, L. A. (1995). *Measurement Error in Non-linear Models. Monographs on Statistics and Applied Probability* **63**. Chapman & Hall, London. [MR1630517](#)
- [4] CHEN, S. S., DONOHO, D. L. and SAUNDERS, M. A. (1998). Atomic decomposition by basis pursuit. *SIAM J. Sci. Comput.* **20** 33–61. [MR1639094](#)
- [5] DUCHI, J., SHALEV-SHWARTZ, S., SINGER, Y. and CHANDRA, T. (2008). Efficient projections onto the ℓ_1 -ball for learning in high dimensions. In *International Conference on Machine Learning* 272–279. ACM, New York, NY.
- [6] HWANG, J. T. (1986). Multiplicative errors-in-variables models with applications to recent data released by the U.S. Department of Energy. *J. Amer. Statist. Assoc.* **81** 680–688. [MR0860499](#)
- [7] ITURRIA, S. J., CARROLL, R. J. and FIRTH, D. (1999). Polynomial regression and estimating functions in the presence of multiplicative measurement error. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **61** 547–561. [MR1707860](#)
- [8] LITTLE, R. J. A. and RUBIN, D. B. (1987). *Statistical Analysis with Missing Data*. Wiley, New York. [MR0890519](#)
- [9] LOH, P. and WAINWRIGHT, M. J. (2012). Supplement to “High-dimensional regression with noisy and missing data: Provable guarantees with nonconvexity.” DOI:10.1214/12-AOS1018SUPP.
- [10] MEINSHAUSEN, N. and BÜHLMANN, P. (2006). High-dimensional graphs and variable selection with the Lasso. *Ann. Statist.* **34** 1436–1462. [MR2278363](#)
- [11] MEINSHAUSEN, N. and YU, B. (2009). Lasso-type recovery of sparse representations for high-dimensional data. *Ann. Statist.* **37** 246–270. [MR2488351](#)
- [12] NEGAHBAN, S., RAVIKUMAR, P., WAINWRIGHT, M. J. and YU, B. (2009). A unified framework for the analysis of regularized M -estimators. In *Advances in Neural Information Processing Systems*. Curran Associates, Red Hook, NY.
- [13] RASKUTTI, G., WAINWRIGHT, M. J. and YU, B. (2010). Restricted eigenvalue properties for correlated Gaussian designs. *J. Mach. Learn. Res.* **11** 2241–2259. [MR2719855](#)
- [14] ROSENBAUM, M. and TSYBAKOV, A. B. (2010). Sparse recovery under matrix uncertainty. *Ann. Statist.* **38** 2620–2651. [MR2722451](#)
- [15] ROSENBAUM, M. and TSYBAKOV, A. B. (2011). Improved matrix uncertainty selector. Technical report. Available at <http://arxiv.org/abs/1112.4413>.
- [16] ROTHMAN, A. J., BICKEL, P. J., LEVINA, E. and ZHU, J. (2008). Sparse permutation invariant covariance estimation. *Electron. J. Stat.* **2** 494–515. [MR2417391](#)
- [17] RUDELSON, M. and ZHOU, S. (2011). Reconstruction from anisotropic random measurements. Technical report, Univ. Michigan.
- [18] STÄDLER, N. and BÜHLMANN, P. (2012). Missing values: Sparse inverse covariance estimation and an extension to sparse regression. *Statist. Comput.* **22** 219–235.
- [19] TIBSHIRANI, R. (1996). Regression shrinkage and selection via the Lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **58** 267–288. [MR1379242](#)
- [20] VAN DE GEER, S. A. and BÜHLMANN, P. (2009). On the conditions used to prove oracle results for the Lasso. *Electron. J. Stat.* **3** 1360–1392. [MR2576316](#)
- [21] XU, Q. and YOU, J. (2007). Covariate selection for linear errors-in-variables regression models. *Comm. Statist. Theory Methods* **36** 375–386. [MR2391878](#)

- [22] YUAN, M. (2010). High dimensional inverse covariance matrix estimation via linear programming. *J. Mach. Learn. Res.* **11** 2261–2286. [MR2719856](#)
- [23] ZHANG, C.-H. and HUANG, J. (2008). The sparsity and bias of the LASSO selection in high-dimensional linear regression. *Ann. Statist.* **36** 1567–1594. [MR2435448](#)

DEPARTMENT OF STATISTICS
UNIVERSITY OF CALIFORNIA, BERKELEY
BERKELEY, CALIFORNIA 94720
USA
E-MAIL: ploh@berkeley.edu
wainwrig@stat.berkeley.edu