
High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity

Po-Ling Loh
Department of Statistics
University of California, Berkeley
Berkeley, CA 94720
ploh@berkeley.edu

Martin J. Wainwright
Departments of Statistics and EECS
University of California, Berkeley
Berkeley, CA 94720
wainwrig@stat.berkeley.edu

Abstract

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 dependencies. 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 non-convex, and it is difficult to establish theoretical guarantees on practical algorithms. While our approach also involves optimizing non-convex programs, we are able to both analyze the statistical error associated with any global optimum, and prove that a simple projected gradient descent algorithm will converge in polynomial time to a small neighborhood of the set of global minimizers. On the statistical side, we provide non-asymptotic 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 will converge at geometric rates to a near-global minimizer. We illustrate these theoretical predictions with simulations, showing 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 with corruption, or exhibit dependencies. 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.

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 [1] and references therein). A challenge in this context is the possible non-convexity of associated optimization problems. For instance, in applications of EM, problems in which the negative likelihood is a convex function often become non-convex 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 the case when the predictors or covariates are noisy, missing, and/or dependent. Our main contribution is to develop and study some 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 non-convex; however, despite this non-convexity, we are still able to prove that a simple form of projected gradient descent will produce an output that is “sufficiently close”—meaning as

small as the statistical error—to any global optimum. As a second result, we bound the size of this 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 with guarantees similar to 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 errors-in-variables for regression problems (see the papers and books [2, 3, 4, 5] and references therein). Much of the earlier theoretical work is classical in nature, where the sample size n diverges with the dimension p held 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 [6] 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 [7] studied the sparse linear model when the covariates are corrupted by noise, and proposed a modified form of the Dantzig selector, involving a convex program. This convexity produces a computationally attractive method, but the statistical error bounds that they establish scale proportionally with the size of the additive perturbation, hence are often weaker than the bounds that can be proved using our methods.

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. For detailed proofs, we refer the reader to the technical report [8].

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 induced ℓ_2 -operator norm (spectral norm) of M . We write $\kappa(M) := \frac{\lambda_{\max}(M)}{\lambda_{\min}(M)}$, the condition number of M .

2 Background and problem set-up

In this section, we provide a formal description of the problem and motivate the class of estimators studied in the paper. We then describe a class of projected gradient descent algorithms to be used in the sequel.

2.1 Observation model and high-dimensional framework

Suppose we observe a response variable $y_i \in \mathbb{R}$ that is linked to a covariate vector $x_i \in \mathbb{R}^p$ via the linear model

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

Here, the regression vector $\beta^* \in \mathbb{R}^p$ is unknown, and $\epsilon_i \in \mathbb{R}$ 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:

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

This setup allows us to model various types of disturbances to the covariates, including

- (a) *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 a fraction $\rho \in [0, 1)$, we observe a random vector $z_i \in \mathbb{R}^p$ such that independently for each component j , we observe $z_{ij} = x_{ij}$ with probability $1 - \rho$, and $z_{ij} = *$ with probability ρ . This model can also be generalized to allow for different missing probabilities for each covariate.

Our first set of results is deterministic, depending on specific instantiations of the observed variables $\{(y_i, z_i)\}_{i=1}^n$. However, we are also interested in proving results that hold with high probability when the x_i 's and z_i 's are drawn at random from some distribution. We develop results for both the i.i.d. setting and the case of dependent covariates, where the x_i 's are generated according to a stationary vector autoregressive (VAR) process. Furthermore, we work within a high-dimensional framework where $n \ll p$, and assume β^* has at most k non-zero parameters, where the sparsity k is also allowed to increase to infinity with the sample size n . We assume the scaling $\|\beta^*\|_2 = \mathcal{O}(1)$, which is reasonable in order to have a non-diverging signal-to-noise ratio.

2.2 M -estimators for noisy and missing covariates

We begin by examining a simple deterministic problem. Let $\text{Cov}(X) = \Sigma_x \succ 0$, and consider the program

$$\hat{\beta} \in \arg \min_{\|\beta\|_1 \leq R} \left\{ \frac{1}{2} \beta^T \Sigma_x \beta - \langle \Sigma_x \beta^*, \beta \rangle \right\}. \quad (3)$$

As long as the constraint radius R is at least $\|\beta^*\|_1$, the unique solution to this convex program is $\hat{\beta} = \beta^*$. This idealization suggests various estimators based on the plug-in principle. We form unbiased estimates of Σ_x and $\Sigma_x \beta^*$, denoted by $\hat{\Gamma}$ and $\hat{\gamma}$, respectively, and consider the modified program and its regularized version:

$$\hat{\beta} \in \arg \min_{\|\beta\|_1 \leq R} \left\{ \frac{1}{2} \beta^T \hat{\Gamma} \beta - \langle \hat{\gamma}, \beta \rangle \right\}, \quad (4)$$

$$\hat{\beta} \in \arg \min_{\beta \in \mathbb{R}^p} \left\{ \frac{1}{2} \beta^T \hat{\Gamma} \beta - \langle \hat{\gamma}, \beta \rangle + \lambda_n \|\beta\|_1 \right\}, \quad (5)$$

where $\lambda_n > 0$ is the regularization parameter. The Lasso [9, 10] is a special case of these programs, where

$$\hat{\Gamma}_{\text{Las}} := \frac{1}{n} X^T X \quad \text{and} \quad \hat{\gamma}_{\text{Las}} := \frac{1}{n} X^T y, \quad (6)$$

and 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. In this paper, we focus on more general instantiations of the programs (4) and (5), involving different choices of the pair $(\hat{\Gamma}, \hat{\gamma})$ that are adapted to the cases of noisy and/or missing data. Note that the matrix $\hat{\Gamma}_{\text{Las}}$ is positive semidefinite, so the Lasso program is convex. In sharp contrast, for the cases of noisy or missing data, the most natural choice of the matrix $\hat{\Gamma}$ is *not positive semidefinite*, hence the loss functions appearing in the problems (4) and (5) are *non-convex*. It is generally impossible to provide a polynomial-time algorithm that converges to a (near) global optimum of a non-convex problem. Remarkably, we prove that a simple projected gradient descent algorithm still converges with high probability to a vector close to any global optimum in our setting.

Let us illustrate these ideas with some examples:

Example 1 (Additive noise). Suppose we observe the $n \times p$ matrix $Z = X + W$, 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 . Consider the pair

$$\hat{\Gamma}_{\text{add}} := \frac{1}{n} Z^T Z - \Sigma_w \quad \text{and} \quad \hat{\gamma}_{\text{add}} := \frac{1}{n} Z^T y, \quad (7)$$

which correspond to unbiased estimators of Σ_x and $\Sigma_x \beta^*$, respectively. 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 $\hat{\Gamma}_{\text{add}}$ is *not* positive semidefinite in the high-dimensional regime ($n \ll p$) of interest. Indeed, since the matrix $\frac{1}{n} Z^T Z$ has rank at most n , the subtracted matrix Σ_w may cause $\hat{\Gamma}_{\text{add}}$ to have a large number of negative eigenvalues.

Example 2 (Missing data). Suppose each entry of X is missing independently with probability $\rho \in [0, 1)$, and 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 $Z \in \mathbb{R}^{n \times p}$, consider an estimator of the general form (4), based on the choices

$$\hat{\Gamma}_{\text{mis}} := \frac{\tilde{Z}^T \tilde{Z}}{n} - \rho \text{diag} \left(\frac{\tilde{Z}^T \tilde{Z}}{n} \right) \quad \text{and} \quad \hat{\gamma}_{\text{mis}} := \frac{1}{n} \tilde{Z}^T y, \quad (8)$$

where $\tilde{Z}_{ij} = Z_{ij}/(1 - \rho)$. It is easy to see that the pair $(\hat{\Gamma}_{\text{mis}}, \hat{\gamma}_{\text{mis}})$ reduces to the pair $(\hat{\Gamma}_{\text{Las}}, \hat{\gamma}_{\text{Las}})$ for the standard Lasso when $\rho = 0$, corresponding to no missing data. In the more interesting case when $\rho \in (0, 1)$, the matrix $\frac{\tilde{Z}^T \tilde{Z}}{n}$ in equation (8) has rank at most n , so the subtracted diagonal matrix may cause the matrix $\hat{\Gamma}_{\text{mis}}$ to have a large number of negative eigenvalues when $n \ll p$, and the associated quadratic function is not convex.

2.3 Restricted eigenvalue conditions

Given an estimate $\hat{\beta}$, there are various ways to assess its closeness to β^* . We focus on the ℓ_2 -norm $\|\hat{\beta} - \beta^*\|_2$, as well as the closely related ℓ_1 -norm $\|\hat{\beta} - \beta^*\|_1$. When the covariate matrix X is fully observed (so that the Lasso can be applied), it is well understood that a sufficient condition for ℓ_2 -recovery is that the matrix $\hat{\Gamma}_{\text{Las}} = \frac{1}{n} X^T X$ satisfy a restricted eigenvalue (RE) condition (e.g., [11, 12, 13]). In this paper, we use the following condition:

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

$$\theta^T \widehat{\Gamma} \theta \geq \alpha_\ell \|\theta\|_2^2 - \tau_\ell(n, p) \|\theta\|_1^2 \quad \text{for all } \theta \in \mathbb{R}^p. \quad (9)$$

It can be shown that when the Lasso matrix $\widehat{\Gamma}_{\text{Las}} = \frac{1}{n} X^T X$ satisfies this RE condition (9), the Lasso estimate has low ℓ_2 -error for any vector β^* supported on any subset of size at most $k \lesssim \frac{1}{\tau_\ell(n, p)}$. Moreover, it is known that for various random choices of the design matrix X , the Lasso matrix $\widehat{\Gamma}_{\text{Las}}$ will satisfy such an RE condition with high probability (e.g., [14]). We also make use of the analogous upper restricted eigenvalue condition:

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

$$\theta^T \widehat{\Gamma} \theta \leq \alpha_u \|\theta\|_2^2 + \tau_u(n, p) \|\theta\|_1^2 \quad \text{for all } \theta \in \mathbb{R}^p. \quad (10)$$

In recent work on high-dimensional projected gradient descent, Agarwal et al. [15] use a more general form of bounds (9) and (10), called the restricted strong convexity (RSC) and restricted smoothness (RSM) conditions.

2.4 Projected gradient descent

In addition to proving results about the global minima of programs (4) and (5), we are also interested in polynomial-time procedures for approximating such optima. We show that the simple projected gradient descent algorithm can be used to solve the program (4). The algorithm generates a sequence of iterates β^t according to

$$\beta^{t+1} = \Pi\left(\beta^t - \frac{1}{\eta}(\widehat{\Gamma}\beta^t - \widehat{\gamma})\right), \quad (11)$$

where $\eta > 0$ is a stepsize parameter, and 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, for instance using a procedure due to Duchi et al. [16]. Our analysis shows that under a reasonable set of conditions, the iterates for the family of programs (4) converges to a point extremely close to any global optimum in both ℓ_1 -norm and ℓ_2 -norm, even for the non-convex program.

3 Main results and consequences

We provide theoretical guarantees for both the constrained estimator (4) and the regularized variant

$$\widehat{\beta} \in \arg \min_{\|\beta\|_1 \leq b_0 \sqrt{k}} \left\{ \frac{1}{2} \beta^T \widehat{\Gamma} \beta - \langle \widehat{\gamma}, \beta \rangle + \lambda_n \|\beta\|_1 \right\}, \quad (12)$$

for a constant $b_0 \geq \|\beta^*\|_2$, which is a hybrid between the constrained (4) and regularized (5) programs.

3.1 Statistical error

In controlling the statistical error, we assume that the matrix $\widehat{\Gamma}$ satisfies a lower-RE condition with curvature α_ℓ and tolerance $\tau_\ell(n, p)$, as previously defined (9). In addition, recall that the matrix $\widehat{\Gamma}$ and vector $\widehat{\gamma}$ serve as surrogates to the deterministic quantities $\Sigma_x \in \mathbb{R}^{p \times p}$ and $\Sigma_x \beta^* \in \mathbb{R}^p$, respectively. We assume there is a function $\varphi(\mathbb{Q}, \sigma_\epsilon)$, depending on the standard deviation σ_ϵ of the observation noise vector ϵ from equation (1) and the conditional distribution \mathbb{Q} from equation (2), such that the following deviation conditions are satisfied:

$$\|\widehat{\gamma} - \Sigma_x \beta^*\|_\infty \leq \varphi(\mathbb{Q}, \sigma_\epsilon) \sqrt{\frac{\log p}{n}} \quad \text{and} \quad \|(\widehat{\Gamma} - \Sigma_x) \beta^*\|_\infty \leq \varphi(\mathbb{Q}, \sigma_\epsilon) \sqrt{\frac{\log p}{n}}. \quad (13)$$

The following result applies to any global optimum $\widehat{\beta}$ of the program (12) with $\lambda_n \geq 4 \varphi(\mathbb{Q}, \sigma_\epsilon) \sqrt{\frac{\log p}{n}}$:

Theorem 1 (Statistical error). *Suppose the surrogates $(\widehat{\Gamma}, \widehat{\gamma})$ satisfy the deviation bounds (13), and the matrix $\widehat{\Gamma}$ satisfies the lower-RE condition (9) with parameters (α_ℓ, τ_ℓ) such that*

$$\sqrt{k} \tau_\ell(n, p) \leq \min \left\{ \frac{\alpha_\ell}{128 \sqrt{k}}, \frac{\varphi(\mathbb{Q}, \sigma_\epsilon)}{2 b_0} \sqrt{\frac{\log p}{n}} \right\}. \quad (14)$$

Then for any vector β^* with sparsity at most k , there is a universal positive constant c_0 such that any global optimum $\hat{\beta}$ satisfies the bounds

$$\|\hat{\beta} - \beta^*\|_2 \leq \frac{c_0 \sqrt{k}}{\alpha_\ell} \max\{\varphi(\mathbb{Q}, \sigma_\epsilon) \sqrt{\frac{\log p}{n}}, \lambda_n\}, \quad \text{and} \quad (15a)$$

$$\|\hat{\beta} - \beta^*\|_1 \leq \frac{8 c_0 k}{\alpha_\ell} \max\{\varphi(\mathbb{Q}, \sigma_\epsilon) \sqrt{\frac{\log p}{n}}, \lambda_n\}. \quad (15b)$$

The same bounds (without λ_n) also apply to the constrained program (4) with radius choice $R = \|\beta^*\|_1$.

Remarks: Note that for the standard Lasso pair $(\hat{\Gamma}_{\text{Las}}, \hat{\gamma}_{\text{Las}})$, bounds of the form (15) for sub-Gaussian noise are well-known from past work (e.g., [12, 17, 18, 19]). The novelty of Theorem 1 is in allowing for general pairs of such surrogates, which can lead to non-convexity in the underlying M -estimator.

3.2 Optimization error

Although Theorem 1 provides guarantees that hold uniformly for any choice of global minimizer, it does not provide any guidance on how to approximate such a global minimizer using a polynomial-time algorithm. Nonetheless, we are able to show that for the family of programs (4), under reasonable conditions on $\hat{\Gamma}$ satisfied in various settings, a simple projected gradient method will converge geometrically fast to a very good approximation of any global optimum.

Theorem 2 (Optimization error). *Consider the program (4) with any choice of radius R for which the constraint is active. Suppose that the surrogate matrix $\hat{\Gamma}$ satisfies the lower-RE (9) and upper-RE (10) conditions with $\tau_u, \tau_l \asymp \frac{\log p}{n}$, and that we apply projected gradient descent (11) with constant stepsize $\eta = 2\alpha_u$. Then as long as $n \gtrsim k \log p$, there is a contraction coefficient $\gamma \in (0, 1)$ independent of (n, p, k) and universal positive constants (c_1, c_2) such that for any global optimum $\hat{\beta}$, the gradient descent iterates $\{\beta^t\}_{t=0}^\infty$ satisfy the bound*

$$\|\beta^t - \hat{\beta}\|_2^2 \leq \gamma^t \|\beta^0 - \hat{\beta}\|_2^2 + c_1 \frac{\log p}{n} \|\hat{\beta} - \beta^*\|_1^2 + c_2 \|\hat{\beta} - \beta^*\|_2^2 \quad \text{for all } t = 0, 1, 2, \dots \quad (16)$$

In addition, we have the ℓ_1 -bound

$$\|\beta^t - \hat{\beta}\|_1 \leq 2\sqrt{k} \|\beta^t - \hat{\beta}\|_2 + 2\sqrt{k} \|\hat{\beta} - \beta^*\|_2 + 2\|\hat{\beta} - \beta^*\|_1 \quad \text{for all } t = 0, 1, 2, \dots \quad (17)$$

Note that the bound (16) controls 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 (4), which may be difficult to compute. Since $\gamma \in (0, 1)$, the first term in the bound vanishes as t increases. Together with Theorem 1, equations (16) and (17) imply that the ℓ_2 - and ℓ_1 -optimization error are bounded as $\mathcal{O}(\frac{k \log p}{n})$ and $\mathcal{O}(k\sqrt{\frac{\log p}{n}})$, respectively.

3.3 Some consequences

Both Theorems 1 and 2 are deterministic results; applying them to specific models requires additional work to establish the stated conditions. We turn to the statements of some consequences of these theorems for different cases of noisy, missing, and dependent data. A zero-mean random variable Z is sub-Gaussian with parameter $\sigma > 0$ if $\mathbb{E}(e^{\lambda Z}) \leq \exp(\lambda^2 \sigma^2 / 2)$ for all $\lambda \in \mathbb{R}$. We say that a random matrix $X \in \mathbb{R}^{n \times p}$ is sub-Gaussian with parameters (Σ, σ^2) if each row $x_i^T \in \mathbb{R}^p$ is sampled independently from a zero-mean distribution with covariance Σ , and for any unit vector $u \in \mathbb{R}^p$, the random variable $u^T x_i$ is sub-Gaussian with parameter at most σ .

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

Corollary 1. *Suppose 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 (Σ_w, σ_w^2) , respectively, and the sample size is lower-bounded as $n \gtrsim \max\left\{\left(\frac{\sigma_x^2 + \sigma_w^2}{\lambda_{\min}(\Sigma_x)}\right)^2, 1\right\} k \log p$. Then for the M -estimator based on the surrogates $(\hat{\Gamma}_{\text{add}}, \hat{\gamma}_{\text{add}})$, the results of Theorems 1 and 2 hold with parameters*

$$\alpha_\ell = \frac{1}{2} \lambda_{\min}(\Sigma_x) \quad \text{and} \quad \varphi(\mathbb{Q}, \sigma_\epsilon) = c_0 \{\sigma_x^2 + \sigma_w^2 + \sigma_\epsilon \sqrt{\sigma_x^2 + \sigma_w^2}\},$$

with probability at least $1 - c_1 \exp(-c_2 \log p)$.

For i.i.d. samples with missing data, we have the following:

Corollary 2. *Suppose $X \in \mathbb{R}^{n \times p}$ is a sub-Gaussian matrix with parameters (Σ_x, σ_x^2) , and Z is the missing data matrix with parameter ρ . If $n \gtrsim \max\left(\frac{1}{(1-\rho)^4}, \frac{\sigma_x^4}{\lambda_{\min}^2(\Sigma_x)}, 1\right)k \log p$, then Theorems 1 and 2 hold with probability at least $1 - c_1 \exp(-c_2 \log p)$ for $\alpha_\ell = \frac{1}{2} \lambda_{\min}(\Sigma_x)$ and*

$$\varphi(\mathbb{Q}, \sigma_\epsilon) = c_0 \frac{\sigma_x}{1-\rho} \left(\sigma_\epsilon + \frac{\sigma_x}{1-\rho} \right).$$

Now consider the case where the rows of X are drawn from a vector autoregressive (VAR) process according to

$$x_{i+1} = Ax_i + v_i, \quad \text{for } i = 1, 2, \dots, n-1, \quad (18)$$

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 $\Sigma_x = A\Sigma_x A^T + \Sigma_v$, so the process is stationary. Corollary 3 corresponds to the case of additive noise for a Gaussian VAR process. A similar result can be derived in the missing data setting.

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. If $n \gtrsim \max\left(\frac{\zeta^4}{\lambda_{\min}^2(\Sigma_x)}, 1\right)k \log p$, with*

$$\zeta^2 = \|\Sigma_w\|_{op} + \frac{2\|\Sigma_x\|_{op}}{1 - \|A\|_{op}},$$

then Theorems 1 and 2 hold with probability at least $1 - c_1 \exp(-c_2 \log p)$ for $\alpha_\ell = \frac{1}{2} \lambda_{\min}(\Sigma_x)$ and

$$\varphi(\mathbb{Q}, \sigma_\epsilon) = c_0(\sigma_\epsilon \zeta + \zeta^2).$$

3.4 Application to graphical model inverse covariance estimation

The problem of inverse covariance estimation for a Gaussian graphical model is closely related to the Lasso. Meinshausen and Bühlmann [20] 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 [21] proposed a method for estimating Θ using linear regression and 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 the rows of X are observed noiselessly and independently.

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 (18). Based on the modified Lasso, we devise a method to estimate Θ based on a corrupted observation matrix Z . 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

$$X^j = X^{-j} \theta^j + \epsilon^j, \quad (19)$$

where ϵ^j is a vector of i.i.d. Gaussians and $\epsilon^j \perp\!\!\!\perp X^{-j}$. Defining $a_j := -(\Sigma_{jj} - \Sigma_{j,-j} \theta^j)^{-1}$, we have $\Theta_{j,-j} = a_j \theta^j$. Our algorithm estimates $\hat{\theta}^j$ and \hat{a}_j for each j and combines the estimates to obtain $\hat{\Theta}_{j,-j} = \hat{a}_j \hat{\theta}^j$.

In the additive noise case, we observe $Z = X + W$. The equations (19) yield $Z^j = X^{-j} \theta^j + (\epsilon^j + W^j)$. Note that $\delta^j = \epsilon^j + W^j$ 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 produce an estimate of θ^j by solving the program (4) or (12) 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 with suitable corrections. We arrive at the following algorithm:

Algorithm 3.1. (1) Perform p linear regressions of the variables Z^j upon the remaining variables Z^{-j} , using the modified Lasso program (4) or (12) with the estimators $(\hat{\Gamma}^{(j)}, \hat{\gamma}^{(j)})$, to obtain estimates $\hat{\theta}^j$.

(2) Estimate the scalars a_j using $\hat{a}_j := -(\hat{\Sigma}_{jj} - \hat{\Sigma}_{j,-j} \hat{\theta}^j)^{-1}$. Set $\tilde{\Theta}_{j,-j} = \hat{a}_j \hat{\theta}^j$ and $\tilde{\Theta}_{jj} = -\hat{a}_j$.

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

Note that the minimization in step (3) is a linear program, so is easily solved with standard methods. We have:

Corollary 4. Suppose the columns of the matrix Θ are k -sparse, and suppose the condition number $\kappa(\Theta)$ is nonzero and finite. Suppose the deviation conditions

$$\|\widehat{\gamma}^{(j)} - \Sigma_{-j,-j}\theta^j\|_\infty \leq \varphi(\mathbb{Q}, \sigma_\epsilon) \sqrt{\frac{\log p}{n}} \quad \text{and} \quad \|(\widehat{\Gamma}^{(j)} - \Sigma_{-j,-j})\theta^j\|_\infty \leq \varphi(\mathbb{Q}, \sigma_\epsilon) \sqrt{\frac{\log p}{n}} \quad (20)$$

hold for all j , and suppose we have the following additional deviation condition on $\widehat{\Sigma}$:

$$\|\widehat{\Sigma} - \Sigma\|_{\max} \leq c\varphi(\mathbb{Q}, \sigma_\epsilon) \sqrt{\frac{\log p}{n}}. \quad (21)$$

Finally, suppose the lower-RE condition holds uniformly over the matrices $\widehat{\Gamma}^{(j)}$ with the scaling (14). Then under the estimation procedure of Algorithm 3.1, there exists a universal constant c_0 such that

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

4 Simulations

In this section, we provide simulation results to confirm that the scalings predicted by our theory are sharp. In Figure 1, we plot the results of simulations under the additive noise model described in Example 1, using $\Sigma_x = I$ and $\Sigma_w = \sigma_w^2 I$ with $\sigma_w = 0.2$. Panel (a) provides plots of ℓ_2 -error versus the sample size n , for $p \in \{128, 256, 512\}$. For all three choices of dimensions, the error decreases to zero as the sample size n increases, showing consistency of the method. If we plot the ℓ_2 -error versus the rescaled sample size $n/(k \log p)$, as depicted in panel (b), the curves roughly align for different values of p , agreeing with Theorem 1. Panel (c) shows analogous curves for VAR data with additive noise, using a driving matrix A with $\|A\|_{op} = 0.2$.




Figure 1. Plots of the error $\|\widehat{\beta} - \beta^*\|_2$ after running projected gradient descent on the non-convex 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 $\frac{n}{k \log p}$. Plot (c) depicts a similar (rescaled) plot for VAR data with additive noise. As predicted by Theorem 1, the curves align for different values of p in the rescaled plot.

We also verified the results of Theorem 2 empirically. Figure 2 shows the results of applying projected gradient descent to solve the optimization problem (4) in the cases of additive noise and missing data. We first applied projected gradient to obtain an initial estimate $\widehat{\beta}$, then reapplied projected gradient descent 10 times, tracking the optimization error $\|\beta^t - \widehat{\beta}\|_2$ (in blue) and statistical error $\|\beta^t - \beta^*\|_2$ (in red). As predicted by Theorem 2, the iterates exhibit geometric convergence to roughly the same fixed point, regardless of starting point.

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

- Chain-structured graphs.* In this case, all nodes are arranged in a line. The diagonal entries of Θ equal 1, and entries corresponding to links in the chain equal 0.1. Then Θ is rescaled so $\|\Theta\|_{op} = 1$.
- 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 $\|\Theta\|_{op} = 1$.
- Erdős-Renyi graphs.* As in Rothman et al. [22], we first generate a matrix B with diagonal entries 0, and all other entries independently equal to 0.5 with probability k/p , and 0 otherwise. Then δ is chosen so $\Theta = B + \delta I$ has condition number p , and Θ is rescaled so $\|\Theta\|_{op} = 1$.





Figure 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 non-convex objective. As predicted by Theorem 2, the optimization error decreases geometrically.

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 $\Sigma_w = (0.2)^2 I$. Figure 3 shows 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, with panel (a) showing the original plot and panel (b) plotting against the rescaled sample size. We obtained qualitatively similar results for the star and Erdős-Rényi graphs, in the presence of missing and/or dependent data.



(a) ℓ_2 error plot for chain graph, additive noise

(b) rescaled plot

Figure 3. (a) Plots of the rescaled error $\frac{1}{\sqrt{k}} \|\hat{\Theta} - \Theta\|_{\text{op}}$ after running projected gradient descent for a chain-structured Gaussian graphical model with additive noise. As predicted by Theorems 1 and 2, all curves align when the rescaled error is plotted against the ratio $\frac{n}{k \log p}$, as shown in (b). Each point represents the average over 50 trials.

5 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 VAR process. Finally, we used our results on linear regression to perform sparse inverse covariance estimation for a Gaussian graphical model, based on corrupted data. The bounds we obtain for the spectral norm of the error are of the same order as existing bounds for inverse covariance matrix estimation with uncorrupted, i.i.d. data.

Acknowledgments

PL acknowledges support from a Hertz Foundation Fellowship and an NDSEG Fellowship; MJW and PL were also partially supported by grants NSF-DMS-0907632 and AFOSR-09NL184. The authors thank Alekh Agarwal, Sahand Negahban, and John Duchi for discussions and guidance.

References

- [1] R. Little and D. B. Rubin. *Statistical analysis with missing data*. Wiley, New York, 1987.
- [2] J. T. Hwang. Multiplicative errors-in-variables models with applications to recent data released by the U.S. Department of Energy. *Journal of the American Statistical Association*, 81(395):pp. 680–688, 1986.
- [3] R. J. Carroll, D. Ruppert, and L. A. Stefanski. *Measurement Error in Nonlinear Models*. Chapman and Hall, 1995.
- [4] S. J. Iturria, R. J. Carroll, and D. Firth. Polynomial regression and estimating functions in the presence of multiplicative measurement error. *Journal of the Royal Statistical Society Series B - Statistical Methodology*, 61:547–561, 1999.
- [5] Q. Xu and J. You. Covariate selection for linear errors-in-variables regression models. *Communications in Statistics - Theory and Methods*, 36(2):375–386, 2007.
- [6] N. Städler and P. Bühlmann. Missing values: Sparse inverse covariance estimation and an extension to sparse regression. *Statistics and Computing*, pages 1–17, 2010.
- [7] M. Rosenbaum and A. B. Tsybakov. Sparse recovery under matrix uncertainty. *Annals of Statistics*, 38:2620–2651, 2010.
- [8] P. Loh and M.J. Wainwright. High-dimensional regression with noisy and missing data: Provable guarantees with non-convexity. Technical report, UC Berkeley, September 2011. Available at <http://arxiv.org/abs/1109.3714>.
- [9] R. Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society, Series B*, 58(1):267–288, 1996.
- [10] S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM Journal on Scientific Computing*, 20(1):33–61, 1998.
- [11] S. van de Geer. The deterministic Lasso. In *Proceedings of Joint Statistical Meeting*, 2007.
- [12] P. J. Bickel, Y. Ritov, and A. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *Annals of Statistics*, 37(4):1705–1732, 2009.
- [13] 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.
- [14] G. Raskutti, M. J. Wainwright, and B. Yu. Restricted eigenvalue properties for correlated Gaussian designs. *Journal of Machine Learning Research*, 11:2241–2259, 2010.
- [15] A. Agarwal, S. Negahban, and M.J. Wainwright. Fast global convergence of gradient methods for high-dimensional statistical recovery. Technical report, UC Berkeley, April 2011. Available at <http://arxiv.org/abs/1104.4824>.
- [16] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra. Efficient projections onto the ℓ_1 -ball for learning in high dimensions. In *International Conference on Machine Learning*, pages 272–279, 2008.
- [17] 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.
- [18] N. Meinshausen and B. Yu. Lasso-type recovery of sparse representations for high-dimensional data. *Annals of Statistics*, 37(1):246–270, 2009.
- [19] S. Negahban, P. Ravikumar, M. J. Wainwright, and B. Yu. A unified framework for the analysis of regularized M -estimators. In *Advances in Neural Information Processing Systems*, 2009.
- [20] N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the Lasso. *Annals of Statistics*, 34:1436–1462, 2006.
- [21] M. Yuan. High-dimensional inverse covariance matrix estimation via linear programming. *Journal of Machine Learning Research*, 99:2261–2286, August 2010.
- [22] A. J. Rothman, P. J. Bickel, E. Levina, and J. Zhu. Sparse permutation invariant covariance estimation. *Electronic Journal of Statistics*, 2:494–515, 2008.