The LASSO risk: asymptotic results and real world examples

Mohsen Bayati
Stanford University
bayati@stanford.edu

José Bento
Stanford University
jbento@stanford.edu

Andrea Montanari
Stanford University
montanar@stanford.edu

Abstract

We consider the problem of learning a coefficient vector $x_0 \in \mathbb{R}^N$ from noisy linear observation $y = Ax_0 + w \in \mathbb{R}^n$. In many contexts (ranging from model selection to image processing) it is desirable to construct a sparse estimator $\widehat{x}$. In this case, a popular approach consists in solving an ℓ_1-penalized least squares problem known as the LASSO or Basis Pursuit DeNoising (BPDN).

For sequences of matrices A of increasing dimensions, with independent gaussian entries, we prove that the normalized risk of the LASSO converges to a limit, and we obtain an explicit expression for this limit. Our result is the first rigorous derivation of an explicit formula for the asymptotic mean square error of the LASSO for random instances. The proof technique is based on the analysis of AMP, a recently developed efficient algorithm, that is inspired from graphical models ideas.

Through simulations on real data matrices (gene expression data and hospital medical records) we observe that these results can be relevant in a broad array of practical applications.

1 Introduction

Let $x_0 \in \mathbb{R}^N$ be an unknown vector, and assume that a vector $y \in \mathbb{R}^n$ of noisy linear measurements of x_0 is available. The problem of reconstructing x_0 from such measurements arises in a number of disciplines, ranging from statistical learning to signal processing. In many contexts the measurements are modeled by

$$y = Ax_0 + w\,, \tag{1.1}$$

where $A \in \mathbb{R}^{n \times N}$ is a known measurement matrix, and w is a noise vector.

The LASSO or Basis Pursuit Denoising (BPDN) is a method for reconstructing the unknown vector x_0 given y, A, and is particularly useful when one seeks sparse solutions. For given A, y, one considers the cost functions $\mathcal{C}_{A,y} : \mathbb{R}^N \to \mathbb{R}$ defined by

$$\mathcal{C}_{A,y}(x) = \frac{1}{2}\,\|y - Ax\|^2 + \lambda\|x\|_1\,, \tag{1.2}$$

with $\lambda > 0$. The original signal is estimated by

$$\widehat{x}(\lambda; A, y) = \operatorname{argmin}_x \mathcal{C}_{A,y}(x)\,. \tag{1.3}$$

In what follows we shall often omit the arguments A, y (and occasionally λ) from the above notations. We will also use $\widehat{x}(\lambda; N)$ to emphasize the N-dependence. Further $\|v\|_p \equiv (\sum_{i=1}^m v_i^p)^{1/p}$ denotes the ℓ_p-norm of a vector $v \in \mathbb{R}^p$ (the subscript p will often be omitted if $p = 2$).

A large and rapidly growing literature is devoted to (i) Developing fast algorithms for solving the optimization problem (1.3); (ii) Characterizing the performances and optimality of the estimator $\widehat{x}$. We refer to Section 1.3 for an unavoidably incomplete overview.

1

Despite such substantial effort, and many remarkable achievements, our understanding of (1.3) is not even comparable to the one we have of more classical topics in statistics and estimation theory. For instance, the best bound on the mean square error (MSE) of the estimator (1.3), i.e. on the quantity $N^{-1}\|\widehat{x} - x_0\|^2$, was proved by Candes, Romberg and Tao [CRT06] (who in fact did not consider the LASSO but a related optimization problem). Their result estimates the mean square error only up to an unknown numerical multiplicative factor. Work by Candes and Tao [CT07] on the analogous *Dantzig selector*, upper bounds the mean square error up to a factor $C \log N$, under somewhat different assumptions.

The objective of this paper is to complement this type of 'rough but robust' bounds by proving *asymptotically exact* expressions for the mean square error. Our asymptotic result holds almost surely for sequences of random matrices A with fixed aspect ratio and independent gaussian entries. While this setting is admittedly specific, the careful study of such matrix ensembles has a long tradition both in statistics and communications theory and has spurred many insights [Joh06, Tel99]. Further, our main result provides asymptotically exact expressions for other operating characteristics of the LASSO as well (e.g., False Positive Rate and True positive Rate). We carried out simulations on real data matrices with continuous entries (gene expression data) and binary feature matrices (hospital medical records). The results appear to be quite encouraging.

Although our rigorous results are asymptotic in the problem dimensions, numerical simulations have shown that they are accurate already on problems with a few hundreds of variables. Further, they seem to enjoy a remarkable *universality* property and to hold for a fairly broad family of matrices [DMM10]. Both these phenomena are analogous to ones in random matrix theory, where delicate asymptotic properties of gaussian ensembles were subsequently proved to hold for much broader classes of random matrices. Also, asymptotic statements in random matrix theory have been re-placed over time by concrete probability bounds in finite dimensions. Of course the optimization problem (1.2) is not immediately related to spectral properties of the random matrix A. As a consequence, universality and non-asymptotic results in random matrix theory cannot be directly exported to the present problem. Nevertheless, we expect such developments to be foreseeable.

Our proof is based on the analysis of an efficient iterative algorithm first proposed by [DMM09], and called AMP, for approximate message passing. The algorithm is inspired by belief-propagation on graphical models, although the resulting iteration is significantly simpler (and scales linearly in the number of nodes). Extensive simulations [DMM10] showed that, in a number of settings, AMP performances are statistically indistinguishable to the ones of LASSO, while its complexity is essentially as low as the one of the simplest greedy algorithms.

The proof technique just described is new. Earlier literature analyzes the convex optimization problem (1.3) –or similar problems– by a clever construction of an approximate optimum, or of a dual witness. Such constructions are largely explicit. Here instead we prove an asymptotically exact characterization of a rather non-trivial iterative algorithm. The algorithm is then proved to converge to the exact optimum. Due to limited space in this paper we only state the main steps of the proof. More details are available in [BM10b]

1.1 Definitions

In order to define the AMP algorithm, we denote by $\eta : \mathbb{R} \times \mathbb{R}_+ \to \mathbb{R}$ the soft thresholding function

$$\eta(x; \theta) = \begin{cases} x - \theta & \text{if } x > \theta, \\ 0 & \text{if } -\theta \leq x \leq \theta, \\ x + \theta & \text{otherwise.} \end{cases} \tag{1.4}$$

The algorithm constructs a sequence of estimates $x^t \in \mathbb{R}^N$, and residuals $z^t \in \mathbb{R}^n$, according to the iteration

$$x^{t+1} = \eta(A^* z^t + x^t; \theta_t), \tag{1.5}$$

$$z^t = y - Ax^t + \frac{\|x^t\|_0}{n} z^{t-1},$$

initialized with $x^0 = 0$. Here A^* denotes the transpose of matrix A, and $\|x^t\|_0$ is number of non-zero entries of x^t. Given a scalar function f and a vector $u \in \mathbb{R}^m$, we let $f(u)$ denote the vector $(f(u_1), \ldots, f(u_m)) \in \mathbb{R}^m$ obtained by applying f componentwise. Finally $\langle u \rangle \equiv m^{-1} \sum_{i=1}^m u_i$ is the average of the vector $u \in \mathbb{R}^m$.

2

As already mentioned, we will consider sequences of instances of increasing sizes, along which the LASSO behavior has a non-trivial limit.

Definition 1. *The sequence of instances* $\{x_0(N), w(N), A(N)\}_{N \in \mathbb{N}}$ *indexed by* N *is said to be a converging sequence if* $x_0(N) \in \mathbb{R}^N$, $w(N) \in \mathbb{R}^n$, $A(N) \in \mathbb{R}^{n \times N}$ *with* $n = n(N)$ *such that* $n/N \to \delta \in (0, \infty)$, *and in addition the following conditions hold:*

(a) The empirical[1] distribution of the entries of $x_0(N)$ *converges weakly to a probability measure* p_{X_0} *on* $\mathbb{R}$ *with bounded second moment. Further* $N^{-1} \sum_{i=1}^{N} x_{0,i}(N)^2 \to \mathbb{E}_{p_{X_0}}\{X_0^2\}$.

(b) The empirical distribution of the entries of $w(N)$ *converges weakly to a probability measure* p_W *on* $\mathbb{R}$ *with bounded second moment. Further* $n^{-1} \sum_{i=1}^{n} w_i(N)^2 \to \mathbb{E}_{p_W}\{W^2\}$.

(c) If $\{e_i\}_{1 \le i \le N}$, $e_i \in \mathbb{R}^N$ *denotes the standard basis, then* $\max_{i \in [N]} \|A(N)e_i\|_2$, $\min_{i \in [N]} \|A(N)e_i\|_2 \to 1$, *as* $N \to \infty$ *where* $[N] \equiv \{1, 2, \ldots, N\}$.

For a converging sequence of instances, and an arbitrary sequence of thresholds $\{\theta_t\}_{t \ge 0}$ (independent of N), the asymptotic behavior of the recursion (1.5) can be characterized as follows.

Define the sequence $\{\tau_t^2\}_{t \ge 0}$ by setting $\tau_0^2 = \sigma^2 + \mathbb{E}\{X_0^2\}/\delta$ (for $X_0 \sim p_{X_0}$ and $\sigma^2 \equiv \mathbb{E}\{W^2\}$, $W \sim p_W$) and letting, for all $t \ge 0$: $\tau_{t+1}^2 = \mathsf{F}(\tau_t^2, \theta_t)$ with

$$\mathsf{F}(\tau^2, \theta) \equiv \sigma^2 + \frac{1}{\delta} \mathbb{E}\left\{ [\eta(X_0 + \tau Z; \theta) - X_0]^2 \right\},$$

where $Z \sim \mathsf{N}(0, 1)$ is independent of X_0. Notice that the function F depends on the law p_{X_0}.

We say a function $\psi : \mathbb{R}^2 \to \mathbb{R}$ is *pseudo-Lipschitz* if there exist a constant $L > 0$ such that for all $x, y \in \mathbb{R}^2$: $|\psi(x) - \psi(y)| \le L(1 + \|x\|_2 + \|y\|_2)\|x - y\|_2$. (This is a special case of the definition used in [BM10a] where such a function is called pseudo-Lipschitz *of order 2*.)

Our next proposition that was conjectured in [DMM09] and proved in [BM10a]. It shows that the behavior of AMP can be tracked by the above one dimensional recursion. We often refer to this prediction by *state evolution*.

Theorem 1 ([BM10a]). *Let* $\{x_0(N), w(N), A(N)\}_{N \in \mathbb{N}}$ *be a converging sequence of instances with the entries of* $A(N)$ *iid normal with mean* 0 *and variance* $1/n$ *and let* $\psi : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ *be a pseudo-Lipschitz function. Then, almost surely*

$$\lim_{N \to \infty} \frac{1}{N} \sum_{i=1}^{N} \psi(x_i^{t+1}, x_{0,i}) = \mathbb{E}\left\{ \psi\big(\eta(X_0 + \tau_t Z; \theta_t), X_0\big) \right\}, \tag{1.6}$$

where $Z \sim \mathsf{N}(0, 1)$ *is independent of* $X_0 \sim p_{X_0}$.

In order to establish the connection with the LASSO, a specific policy has to be chosen for the thresholds $\{\theta_t\}_{t \ge 0}$. Throughout this paper we will take $\theta_t = \alpha \tau_t$ with α is fixed. In other words, the sequence $\{\tau_t\}_{t \ge 0}$ is given by the recursion $\tau_{t+1}^2 = \mathsf{F}(\tau_t^2, \alpha \tau_t)$. This choice enjoys several convenient properties [DMM09].

1.2 Main result

Before stating our results, we have to describe a *calibration* mapping between α and λ that was introduced in [DMM10] (Propositions 2, 3 and Corollary 4). Their proofs are presented in [BM10b].

Let us start by stating some convenient properties of the state evolution recursion.

Proposition 2 ([DMM09]). *Let* $\alpha_{\min} = \alpha_{\min}(\delta)$ *be the unique non-negative solution of the equation* $(1 + \alpha^2)\Phi(-\alpha) - \alpha\phi(\alpha) = \frac{\delta}{2}$, *with* $\phi(z) \equiv e^{-z^2/2}/\sqrt{2\pi}$ *the standard gaussian density and* $\Phi(z) \equiv \int_{-\infty}^{z} \phi(x) \, dx$.

For any $\sigma^2 > 0$, $\alpha > \alpha_{\min}(\delta)$, *the fixed point equation* $\tau^2 = \mathsf{F}(\tau^2, \alpha\tau)$ *admits a unique solution. Denoting by* $\tau_* = \tau_*(\alpha)$ *this solution, we have* $\lim_{t \to \infty} \tau_t = \tau_*(\alpha)$. *Further the convergence takes place for any initial condition and is monotone. Finally* $\left| \frac{d\mathsf{F}}{d\tau^2}(\tau^2, \alpha\tau) \right| < 1$ *at* $\tau = \tau_*$.

---

[1] The probability distribution that puts a point mass $1/N$ at each of the N entries of the vector.

3

We then define the function $\alpha \mapsto \lambda(\alpha)$ on $(\alpha_{\min}(\delta), \infty)$, by $\lambda(\alpha) \equiv \alpha\tau_*[1 - \frac{1}{\delta}\mathbb{P}\{|X_0 + \tau_*Z| \geq \alpha\tau_*\}]$.

This function defines a correspondence (calibration) between the sequence of thresholds $\{\theta_t\}_{t \geq 0}$ and the regularization parameter λ. It should be intuitively clear that larger λ corresponds to larger thresholds and hence larger α since both cases yield smaller estimates of x_0.

In the following we will need to invert this function. We thus define $\alpha : (0, \infty) \to (\alpha_{\min}, \infty)$ in such a way that $\alpha(\lambda) \in \big\{ a \in (\alpha_{\min}, \infty) : \lambda(a) = \lambda \big\}$.

The next result implies that the set on the right-hand side is non-empty and therefore the function $\lambda \mapsto \alpha(\lambda)$ is well defined.

Proposition 3 ([DMM10]). *The function $\alpha \mapsto \lambda(\alpha)$ is continuous on the interval $(\alpha_{\min}, \infty)$ with $\lambda(\alpha_{\min}+) = -\infty$ and $\lim_{\alpha \to \infty} \lambda(\alpha) = \infty$.*

Therefore the function $\lambda \mapsto \alpha(\lambda)$ satisfying $\alpha(\lambda) \in \big\{ a \in (\alpha_{\min}, \infty) : \lambda(a) = \lambda \big\}$ exists.

We will denote by $\mathcal{A} = \alpha((0, \infty))$ the image of the function α. Notice that the definition of α is *a priori* not unique. We will see that uniqueness follows from our main theorem.

Examples of the mappings $\tau^2 \mapsto \mathsf{F}(\tau^2, \alpha\tau)$, $\alpha \mapsto \tau_*(\alpha)$ and $\alpha \mapsto \lambda(\alpha)$ are presented in [BM10b].

We can now state our main result.

Theorem 2. *Let $\{x_0(N), w(N), A(N)\}_{N \in \mathbb{N}}$ be a converging sequence of instances with the entries of $A(N)$ iid normal with mean 0 and variance $1/n$. Denote by $\widehat{x}(\lambda; N)$ the* LASSO *estimator for instance $(x_0(N), w(N), A(N))$, with $\sigma^2, \lambda > 0$, $\mathbb{P}\{X_0 \neq 0\}$ and let $\psi : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ be a pseudo-Lipschitz function. Then, almost surely*

$$\lim_{N \to \infty} \frac{1}{N} \sum_{i=1}^{N} \psi\big(\widehat{x}_i, x_{0,i}\big) = \mathbb{E}\Big\{ \psi\big(\eta(X_0 + \tau_*Z; \theta_*), X_0\big) \Big\}, \tag{1.7}$$

where $Z \sim \mathsf{N}(0, 1)$ is independent of $X_0 \sim p_{X_0}$, $\tau_* = \tau_*(\alpha(\lambda))$ and $\theta_* = \alpha(\lambda)\tau_*(\alpha(\lambda))$.

As a corollary, the function $\lambda \mapsto \alpha(\lambda)$ is indeed uniquely defined.

Corollary 4. *For any $\lambda, \sigma^2 > 0$ there exists a unique $\alpha > \alpha_{\min}$ such that $\lambda(\alpha) = \lambda$ (with the function $\alpha \to \lambda(\alpha)$ defined by $\lambda(\alpha) = \alpha\tau_*[1 - \frac{1}{\delta}\mathbb{P}\{|X_0 + \tau_*Z| \geq \alpha\tau_*\}]$.*

Hence the function $\lambda \mapsto \alpha(\lambda)$ is continuous non-decreasing with $\alpha((0, \infty)) \equiv \mathcal{A} = (\alpha_0, \infty)$.

The assumption of a converging problem-sequence is important for the result to hold, while the hypothesis of gaussian measurement matrices $A(N)$ is necessary for the proof technique to be correct. On the other hand, the restrictions $\lambda, \sigma^2 > 0$, and $\mathbb{P}\{X_0 \neq 0\} > 0$ (whence $\tau_* \neq 0$ using $\lambda(\alpha) = \alpha\tau_*[1 - \frac{1}{\delta}\mathbb{P}\{|X_0 + \tau_*Z| \geq \alpha\tau_*\}]$) are made in order to avoid technical complications due to degenerate cases. Such cases can be resolved by continuity arguments.

1.3 Related work

The LASSO was introduced in [Tib96, CD95]. Several papers provide performance guarantees for the LASSO or similar convex optimization methods [CRT06, CT07], by proving upper bounds on the resulting mean square error. These works assume an appropriate 'isometry' condition to hold for A. While such condition hold with high probability for some random matrices, it is often difficult to verify them explicitly. Further, it is only applicable to very sparse vectors x_0. These restrictions are intrinsic to the worst-case point of view developed in [CRT06, CT07].

Guarantees have been proved for correct support recovery in [ZY06], under an appropriate 'irrepresentibility' assumption on A. While support recovery is an interesting conceptualization for some applications (e.g. model selection), the metric considered in the present paper (mean square error) provides complementary information and is quite standard in many different fields.

Closer to the spirit of this paper [RFG09] derived expressions for the mean square error under the same model considered here. Similar results were presented recently in [KWT09, GBS09]. These papers argue that a sharp asymptotic characterization of the LASSO risk can provide valuable

4

guidance in practical applications. For instance, it can be used to evaluate competing optimization methods on large scale applications, or to tune the regularization parameter λ.

Unfortunately, these results were non-rigorous and were obtained through the famously powerful 'replica method' from statistical physics [MM09].

Let us emphasize that the present paper offers two advantages over these recent developments: (i) It is completely *rigorous*, thus putting on a firmer basis this line of research; (ii) It is *algorithmic* in that the LASSO mean square error is shown to be equivalent to the one achieved by a low-complexity message passing algorithm.

2   Numerical illustrations

Theorem 2 assumes that the entries of matrix A are iid gaussians. We expect however that our predictions to be robust and hold for much larger family of matrices. Rigorous evidence in this direction is presented in [KM10] where the normalized cost $\mathcal{C}(\widehat{x})/N$ is shown to have a limit as $N \to \infty$ which is universal with respect to random matrices A with iid entries. (More precisely, it is universal if $\mathbb{E}\{A_{ij}\} = 0$, $\mathbb{E}\{A_{ij}^2\} = 1/n$ and $\mathbb{E}\{A_{ij}^6\} \leq C/n^3$ for a uniform constant C.)

Further, our result is asymptotic, while and one might wonder how accurate it is for instances of moderate dimensions.

Numerical simulations were carried out in [DMM10] and suggest that the result is robust and relevant already for N of the order of a few hundreds. As an illustration, we present in Figures 1-3 the outcome of such simulations for four types of real data and random matrices. We generated the signal vector randomly with entries in $\{+1, 0, -1\}$ and $\mathbb{P}(x_{0,i} = +1) = \mathbb{P}(x_{0,i} = -1) = 0.05$. The noise vector w was generated by using i.i.d. $\mathsf{N}(0, 0.2)$ entries.

We obtained the optimum estimator $\widehat{x}$ using `OWLQN` and `l1_ls`, packages for solving large-scale l1-regularized regressions [KKL$^+$07], [AJ07]. We used 40 values of λ between .05 and 2 and N equal to 500, 1000, and 2000. For each case, the point (λ, MSE) was plotted and the results are shown in the figures. Continuous lines corresponds to the asymptotic prediction by Theorem 2 for $\psi(a, b) = (a - b)^2$, namely $\mathrm{MSE} = \lim_{N \to \infty} N^{-1}\|\widehat{x} - x_0\|^2 = \mathbb{E}\big\{\big[\eta(X_0 + \tau_* Z; \theta_*) - X_0\big]^2\big\} = \delta(\tau_*^2 - \sigma^2)$.

The agreement is remarkably good already for N, n of the order of a few hundreds, and deviations are consistent with statistical fluctuations.

The four figures correspond to measurement matrices A:

Figure 1(a): Data consist of 2253 measurements of expression level of 7077 genes (this data is provided to us by Broad Institute). From this matrix we took sub-matrices A of aspect ratio δ for each N. The entries were continuous variables. We standardized all columns of A to have mean 0 and variance 1.

Figure 1(b): From a data set of 1932 patient records we extracted 4833 binary features describing demographic information, medical history, lab results, medications etc. The 0-1 matrix was sparse (with only 3.1% non-zero entries). Similar to genes data, for each N, the sub-matrices A with aspect ratio δ were selected and standardized.

Figure 2(a): Random ± 1 matrices with aspect ratio δ. Each entry is independently equal to $+1/\sqrt{n}$ or $-1/\sqrt{n}$ with equal probability.

Figure 2(b): Random gaussian matrices with aspect ratio δ and iid $\mathsf{N}(0, 1/n)$ entries (as in Theorem 2).

Notice the behavior appears to be essentially indistinguishable. Also the asymptotic prediction has a minimum as a function of λ. The location of this minimum can be used to select the regularization parameter. Further empirical analysis is presented in [BBM10].

For the second data set –patient records– we repeated the simulation 20 times (each time with fresh x_0 and w) and obtained the average and standard error for MSE, False Positive Rate (FPR) and True Positive Rate (TPR). The results with error bars are shown in Figure 3. The length of each error bar

5

Figure 1: Mean square error (MSE) as a function of the regularization parameter λ compared to the asymptotic prediction for $\delta = .5$ and $\sigma^2 = .2$. In plot (a) the measurement matrix A is a real valued (standardized) matrix of gene expression data and in plot (b) A is a (standardized) 0-1 feature matrix of hospital records. Each point in these plots is generated by finding the LASSO predictor $\widehat{x}$ using a measurement vector $y = Ax_0 + w$ for an independent signal vector x_0 and an independent noise vector w.
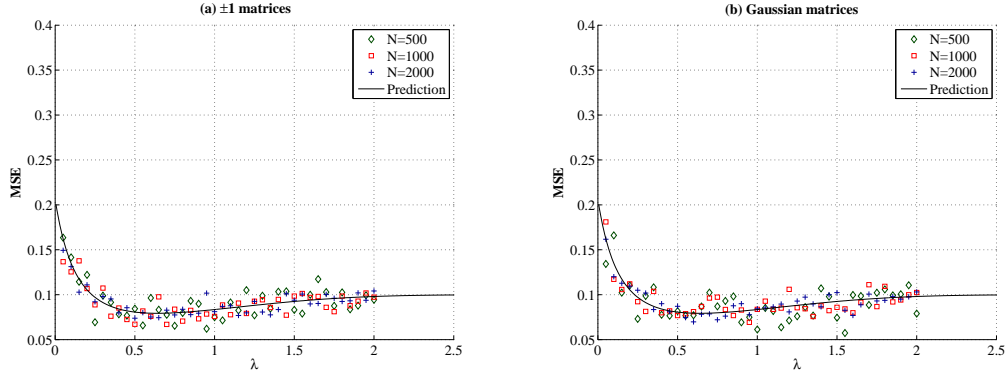
Figure 2: As in Figure 1, but the measurement matrix A has iid entries that are equal to $\pm 1/\sqrt{n}$ with equal probabilities in plot (a), and has iid $\mathsf{N}(0, 1/n)$ entries in plot (b). Additionally, each point in these plots uses an independent matrix A.

is equal to twice the standard error (in each direction). FPR and TPR are calculated using

$$\text{FPR} \equiv \frac{\sum_{i=1}^N \mathbb{I}_{\{\widehat{x}_i \neq 0\}}\mathbb{I}_{\{x_{i,0}=0\}}}{\sum_{i=1}^N \mathbb{I}_{\{x_{i,0}=0\}}}, \quad \text{TPR} \equiv \frac{\sum_{i=1}^N \mathbb{I}_{\{\widehat{x}_i \neq 0\}}\mathbb{I}_{\{x_{i,0}\neq 0\}}}{\sum_{i=1}^N \mathbb{I}_{\{x_{i,0}\neq 0\}}}, \quad (2.1)$$

where $\mathbb{I}_{\{\mathcal{S}\}} = 1$ if statement $\mathcal{S}$ holds and $\mathbb{I}_{\{\mathcal{S}\}} = 0$ otherwise. The predictions for FPR and TPR are obtained by applying Theorem 2 to $\psi_{fpr}(a, b) \equiv \mathbb{I}_{\{a\neq 0\}}\mathbb{I}_{\{b=0\}}$ and $\psi_{tpr}(a, b) = \mathbb{I}_{\{a\neq 0\}}\mathbb{I}_{\{b\neq 0\}}$, which yields

$$\lim_{N\to\infty} \text{FPR} = 2\Phi(-\alpha), \quad \lim_{N\to\infty} \text{TPR} = \Phi(-\alpha + \frac{1}{\tau_*}) + \Phi(-\alpha - \frac{1}{\tau_*}) \quad (2.2)$$

where Φ is defined in Proposition 2. Note that functions $\psi_{fpr}(a, b)$ and $\psi_{tpr}(a, b)$ are not pseudo-Lipschitz but the limits (2.2) follow from Theorem 2 via standard weak-convergence arguments.

3 A structural property and proof of the main theorem

We will prove the following theorem which implies our main result, Theorem 2.

Theorem 3. *Assume the hypotheses of Theorem 2. Denote by $\{x^t(N)\}_{t\geq 0}$ the sequence of estimates produced by* AMP. *Then* $\lim_{t\to\infty} \lim_{N\to\infty} N^{-1}\|x^t(N) - \widehat{x}(\lambda; N)\|_2^2 = 0$*, almost surely.*

6

Figure 3: Average of MSE, FPR and TPR versus λ for medical data, using 20 samples per λ and N. All parameters are similar to Figure 1(b). Error bars are twice the standard errors (in each direction).

The rest of the paper is devoted to the proof of this theorem. Section 3.1 proves a structural property that is the key tool in this proof. Section 3.2 uses this property together with a few lemmas to prove Theorem 3. Proofs of lemmas and more details can be found in [BM10b].

The proof of Theorem 2 follows immediately. Since when ψ is Lipschitz there is a constant B where $|\frac{1}{N}\sum_{i=1}^{N}\psi(x_i^{t+1},x_{0,i}) - \frac{1}{N}\sum_{i=1}^{N}\psi(\widehat{x}_i,x_{0,i})| \leq \mathsf{B}\|x^{t+1} - \widehat{x}\|_2$. We then obtain

$$\lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N}\psi(\widehat{x}_i,x_{0,i}) = \lim_{t\to\infty}\lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N}\psi(x_i^{t+1},x_{0,i}) = \mathbb{E}\{\psi(\eta(X_0 + \tau_* Z;\theta_*),X_0)\},$$

where we used Theorem 1 and Proposition 2. The case of pseudo-Lipschitz ψ is a straightforward generalization.

Some notations. For any non-empty subset S of $[m]$ and any $k \times m$ matrix M we refer by M_S to the k by $|S|$ sub-matrix of M that contains only the columns of M corresponding to S. Also define the scalar product $\langle u,v\rangle \equiv \frac{1}{m}\sum_{i=1}^{m}u_i v_i$ for $u,v \in \mathbb{R}^m$. Finally, the subgradient of a convex function $f : \mathbb{R}^m \to \mathbb{R}$ at point $x \in \mathbb{R}^m$ is denoted by $\partial f(x)$. In particular, remember that the subgradient of the ℓ_1 norm, $x \mapsto \|x\|_1$ is given by

$$\partial\|x\|_1 = \left\{v \in \mathbb{R}^m \text{ such that } |v_i| \leq 1 \,\forall i \text{ and } x_i \neq 0 \Rightarrow v_i = \mathrm{sign}(x_i)\right\}. \tag{3.1}$$

3.1 A structural property of the LASSO cost function

One main challenge in the proof of Theorem 2 lies in the fact that the function $x \mapsto \mathcal{C}_{A,y}(x)$ is not –in general– strictly convex. Hence there can be, in principle, vectors x of cost very close to the optimum and nevertheless far from the optimum. The following Lemma provides conditions under which this does not happen.

Lemma 1. *There exists a function $\xi(\varepsilon,c_1,\ldots,c_5)$ such that the following happens. If $x,r \in \mathbb{R}^N$ satisfy the following conditions:*

7

(1) $\|r\|_2 \leq c_1 \sqrt{N}$; *(2)* $\mathcal{C}(x+r) \leq \mathcal{C}(x)$;

(3) There exists a subgradient $\mathrm{sg}(\mathcal{C}, x) \in \partial \mathcal{C}(x)$ with $\|\mathrm{sg}(\mathcal{C}, x)\|_2 \leq \sqrt{N}\,\varepsilon$;

(4) Let $v \equiv (1/\lambda)[A^*(y - Ax) + \mathrm{sg}(\mathcal{C}, x)] \in \partial \|x\|_1$, and $S(c_2) \equiv \{i \in [N] : |v_i| \geq 1 - c_2\}$. Then, for any $S' \subseteq [N]$, $|S'| \leq c_3 N$, we have $\sigma_{\min}(A_{S(c_2) \cup S'}) \geq c_4$

(5) The maximum and minimum non-zero singular value of A satisfy $c_5^{-1} \leq \sigma_{\min}(A)^2 \leq \sigma_{\max}(A)^2 \leq c_5$.

Then $\|r\|_2 \leq \sqrt{N}\,\xi(\varepsilon, c_1, \ldots, c_5)$. Further for any $c_1, \ldots, c_5 > 0$, $\xi(\varepsilon, c_1, \ldots, c_5) \to 0$ as $\varepsilon \to 0$.

Further, if $\ker(A) = \{0\}$, the same conclusion holds under conditions 1, 2, 3, and 5.

3.2   Proof of Theorem 3

The proof is based on a series of Lemmas that are used to check the assumptions of Lemma 1

The next lemma implies that submatrices of A constructed using the first t iterations of the AMP algorithm are non-singular (more precisely, have singular values bounded away from 0).

Lemma 2. Let $S \subseteq [N]$ be measurable on the σ-algebra $\mathfrak{S}_t$ generated by $\{z^0, \ldots, z^{t-1}\}$ and $\{x^0 + A^* z^0, \ldots, x^{t-1} + A^* z^{t-1}\}$ and assume $|S| \leq N(\delta - c)$ for some $c > 0$. Then there exists $a_1 = a_1(c) > 0$ (independent of t) and $a_2 = a_2(c, t) > 0$ (depending on t and c) such that $\min_{S'} \{\sigma_{\min}(A_{S \cup S'}) : S' \subseteq [N], |S'| \leq a_1 N\} \geq a_2$, with probability converging to 1 as $N \to \infty$.

We will apply this lemma to a specific choice of the set S. Namely, defining

$$v^t \equiv \frac{1}{\theta_{t-1}} (x^{t-1} + A^* z^{t-1} - x^t), \tag{3.2}$$

our last lemma shows convergence of a particular sequence of sets provided by v^t.

Lemma 3. Fix $\gamma \in (0, 1)$ and let the sequence $\{S_t(\gamma)\}_{t \geq 0}$ be defined by $S_t(\gamma) \equiv \{i \in [N] : |v_i^t| \geq 1 - \gamma\}$. For any $\xi > 0$ there exists $t_* = t_*(\xi, \gamma) < \infty$ such that, for all $t_2 \geq t_1 \geq t_*$: $\lim_{N \to \infty} \mathbb{P}\{|S_{t_2}(\gamma) \setminus S_{t_1}(\gamma)| \geq N\xi\} = 0$.

The last two lemmas imply the following.

Proposition 5. There exist constants $\gamma_1 \in (0, 1)$, $\gamma_2, \gamma_3 > 0$ and $t_{\min} < \infty$ such that, for any $t \geq t_{\min}$, $\min\{\sigma_{\min}(A_{S_t(\gamma_1) \cup S'}) : S' \subseteq [N], |S'| \leq \gamma_2 N\} \geq \gamma_3$ with probability converging to 1 as $N \to \infty$.

Proof of Theorem 3. We apply Lemma 1 to $x = x^t$, the AMP estimate and $r = \widehat{x} - x^t$ the distance from the LASSO optimum. The thesis follows by checking conditions 1–5. Namely we need to show that there exists constants $c_1, \ldots, c_5 > 0$ and, for each $\varepsilon > 0$ some $t = t(\varepsilon)$ such that 1–5 hold with probability going to 1 as $N \to \infty$.

Condition 1 holds since $\lim_{N \to \infty} \langle \widehat{x}, \widehat{x} \rangle$ and $\lim_{N \to \infty} \langle x^t, x^t \rangle$ for all t are finite.

Condition 2 is immediate since $x + r = \widehat{x}$ minimizes $\mathcal{C}(\,\cdot\,)$.

Conditions 3-4. Take $v = v^t$ as defined in Eq. (3.2). Using the definition (1.5), it is easy to check that $v^t \in \partial \|x\|_1$. Further it can be shown that $v^t = (1/\lambda)[A^*(y - Ax^t) + \mathrm{sg}(\mathcal{C}, x^t)]$, with $\mathrm{sg}(\mathcal{C}, x^t)$ a subgradient satisfying $\lim_{t \to \infty} \lim_{N \to \infty} N^{-1} \|\mathrm{sg}(\mathcal{C}, x^t)\|^2 = 0$. This proves condition 3 and condition 4 holds by Proposition 5.

Condition 5 follows from standard limit theorems on the singular values of Wishart matrices. $\quad\square$

Acknowledgement

This work was partially supported by a Terman fellowship, the NSF CAREER award CCF-0743978, the NSF grant DMS-0806211 and a Portuguese Doctoral FCT fellowship.

8

References

[AJ07]   G. Andrew and G. Jianfeng, *Scalable training of l^1-regularized log-linear models*, Proceedings of the 24th international conference on Machine learning, 2007, pp. 33–40.

[BBM10]   M. Bayati, J .A. Bento, and A. Montanari, *The LASSO risk: asymptotic results and real world examples*, Long version (in preparation), 2010.

[BM10a]   M. Bayati and A. Montanari, *The dynamics of message passing on dense graphs, with applications to compressed sensing*, Proceedings of IEEE International Symposium on Inform. Theory (ISIT), 2010, Longer version in http://arxiv.org/abs/1001.3448.

[BM10b]   ______ , *The LASSO risk for gaussian matrices*, 2010, preprint available in http://arxiv.org/abs/1008.2581.

[CD95]   S.S. Chen and D.L. Donoho, *Examples of basis pursuit*, Proceedings of Wavelet Applications in Signal and Image Processing III (San Diego, CA), 1995.

[CRT06]   E. Candes, J. K. Romberg, and T. Tao, *Stable signal recovery from incomplete and inaccurate measurements*, Communications on Pure and Applied Mathematics **59** (2006), 1207–1223.

[CT07]   E. Candes and T. Tao, *The Dantzig selector: statistical estimation when p is much larger than n*, Annals of Statistics **35** (2007), 2313–2351.

[DMM09]   D. L. Donoho, A. Maleki, and A. Montanari, *Message Passing Algorithms for Compressed Sensing*, Proceedings of the National Academy of Sciences **106** (2009), 18914–18919.

[DMM10]   D.L. Donoho, A. Maleki, and A. Montanari, *The Noise Sensitivity Phase Transition in Compressed Sensing*, Preprint, 2010.

[GBS09]   D. Guo, D. Baron, and S. Shamai, *A single-letter characterization of optimal noisy compressed sensing*, 47th Annual Allerton Conference (Monticello, IL), September 2009.

[Joh06]   I. Johnstone, *High Dimensional Statistical Inference and Random Matrices*, Proc. International Congress of Mathematicians (Madrid), 2006.

[KKL+07] S. J. Kim, K. Koh, M. Lustig, S. Boyd, and D. Gorinevsky, *A method for large-scale ℓ_1-regularized least squares.*, IEEE Journal on Selected Topics in Signal Processing **1** (2007), 606–617.

[KM10]   S. Korada and A. Montanari, *Applications of Lindeberg Principle in Communications and Statistical Learning*, preprint available in http://arxiv.org/abs/1004.0557, 2010.

[KWT09]   Y. Kabashima, T. Wadayama, and T. Tanaka, *A typical reconstruction limit for compressed sensing based on lp-norm minimization*, J.Stat. Mech. (2009), L09003.

[MM09]   M. Mézard and A. Montanari, *Information, Physics and Computation*, Oxford University Press, Oxford, 2009.

[RFG09]   S. Rangan, A. K. Fletcher, and V. K. Goyal, *Asymptotic analysis of map estimation via the replica method and applications to compressed sensing*, PUT NIPS REF, 2009.

[Tel99]   E. Telatar, *Capacity of Multi-antenna Gaussian Channels*, European Transactions on Telecommunications **10** (1999), 585–595.

[Tib96]   R. Tibshirani, *Regression shrinkage and selection with the lasso*, J. Royal. Statist. Soc B **58** (1996), 267–288.

[ZY06]   P. Zhao and B. Yu, *On model selection consistency of Lasso*, The Journal of Machine Learning Research **7** (2006), 2541–2563.

9