Estimating LASSO Risk and Noise Level

Mohsen Bayati
Stanford University
bayati@stanford.edu

Murat A. Erdogdu
Stanford University
erdogdu@stanford.edu

Andrea Montanari
Stanford University
montanar@stanford.edu

Abstract

We study the fundamental problems of variance and risk estimation in high dimensional statistical modeling. In particular, we consider the problem of learning a coefficient vector $\theta_0 \in \mathbb{R}^p$ from noisy linear observations $y = X\theta_0 + w \in \mathbb{R}^n$ ($p > n$) and the popular estimation procedure of solving the ℓ_1-penalized least squares objective known as the LASSO or Basis Pursuit DeNoising (BPDN). In this context, we develop new estimators for the ℓ_2 estimation risk $\|\widehat{\theta} - \theta_0\|_2$ and the variance of the noise when distributions of θ_0 and w are unknown. These can be used to select the regularization parameter optimally. Our approach combines Stein's unbiased risk estimate [Ste81] and the recent results of [BM12a][BM12b] on the analysis of approximate message passing and the risk of LASSO.

We establish high-dimensional consistency of our estimators for sequences of matrices X of increasing dimensions, with independent Gaussian entries. We establish validity for a broader class of Gaussian designs, conditional on a certain conjecture from statistical physics.

To the best of our knowledge, this result is the first that provides an asymptotically consistent risk estimator for the LASSO solely based on data. In addition, we demonstrate through simulations that our variance estimation outperforms several existing methods in the literature.

1 Introduction

In Gaussian random design model for the linear regression, we seek to reconstruct an unknown coefficient vector $\theta_0 \in \mathbb{R}^p$ from a vector of noisy linear measurements $y \in \mathbb{R}^n$:

$$y = X\theta_0 + w, \tag{1.1}$$

where $X \in \mathbb{R}^{n\times p}$ is a measurement (or feature) matrix with iid rows generated through a multivariate normal density. The noise vector, w, has iid entries with mean 0 and variance σ^2. While this problem is well understood in the low dimensional regime $p \ll n$, a growing corpus of research addresses the more challenging high-dimensional scenario in which $p > n$. The Basis Pursuit Denoising (BPDN) or LASSO [CD95, Tib96] is an extremely popular approach in this regime, that finds an estimate for θ_0 by minimizing the following cost function

$$\mathcal{C}_{X,y}(\lambda, \theta) \equiv (2n)^{-1}\, \|y - X\theta\|_2^2 + \lambda\|\theta\|_1\,, \tag{1.2}$$

with $\lambda > 0$. In particular, θ_0 is estimated by $\widehat{\theta}(\lambda; X, y) = \mathrm{argmin}_\theta\, \mathcal{C}_{X,y}(\lambda, \theta)$. This method is well suited for the ubiquitous case in which θ_0 is sparse, i.e. a small number of features effectively predict the outcome. Since this optimization problem is convex, it can be solved efficiently, and fast specialized algorithms have been developed for this purpose [BT09].

Research has established a number of important properties of LASSO estimator under suitable conditions on the design matrix X, and for sufficiently sparse vectors θ_0. Under irrepresentability conditions, the LASSO correctly recovers the support of θ_0 [ZY06, MB06, Wai09]. Under weaker conditions such as restricted isometry or compatibility properties the correct recovery of support fails however, the ℓ_2 estimation error $\|\widehat{\theta}-\theta_0\|_2$ is of the same order as the one achieved by an oracle estimator that knows the support [CRT06, CT07, BRT09, BdG11]. Finally, [DMM09, RFG09, BM12b] provided asymptotic formulas for MSE or other operating characteristics of $\widehat{\theta}$, for Gaussian design matrices X.

While the aforementioned research provides solid justification for using the LASSO estimator, it is of limited guidance to the practitioner. For instance, a crucial question is how to set the regularization parameter λ. This question becomes even more urgent for high-dimensional methods with multiple regularization terms. The oracle bounds of [CRT06, CT07, BRT09, BdG11] suggest to take $\lambda = c\,\sigma\sqrt{\log p}$ with c a dimension-independent constant (say $c = 1$ or 2). However, in practice a factor two in λ can make a substantial difference for statistical applications. Related to this issue is the question of estimating accurately the ℓ_2 error $\|\widehat{\theta}-\theta_0\|_2^2$. The above oracle bounds have the form $\|\widehat{\theta}-\theta_0\|_2^2 \le C\,k\lambda^2$, with $k = \|\theta_0\|_0$ the number of nonzero entries in θ_0, as long as $\lambda \ge c\sigma\sqrt{\log p}$. As a consequence, minimizing the bound does not yield a recipe for setting λ. Finally, estimating the noise level is necessary for applying these formulae, and this is in itself a challenging question.

The results of [DMM09, BM12b] provide exact asymptotic formulae for the risk, and its dependence on the regularization parameter λ. This might appear promising for choosing the optimal value of λ, but has one serious drawback. The formulae of [DMM09, BM12b] depend on the empirical distribution[1] of the entries of θ_0, which is of course unknown, as well as on the noise level[2]. A step towards the resolution of this problem was taken in [DMM11], which determined the least favorable noise level and distribution of entries, and hence suggested a prescription for λ, and a predicted risk in this case. While this settles the question (in an asymptotic sense) from a minimax point of view, it would be preferable to have a prescription that is adaptive to the distribution of the entries of θ_0 and to the noise level.

Our starting point is the asymptotic results of [DMM09, DMM11, BM12a, BM12b]. These provide a construction of an unbiased *pseudo-data* $\widehat{\theta}^u$ that is asymptotically Gaussian with mean θ_0. The LASSO estimator $\widehat{\theta}$ is obtained by applying a denoiser function to $\widehat{\theta}^u$. We then use Stein's Unbiased Risk Estimate (*SURE*) [Ste81] to derive an expression for the ℓ_2 risk (mean squared error) of this operation. What results is an expression for the mean squared error of the LASSO that only depends on the observed data y and X. Finally, by modifying this formula we obtain an estimator for the noise level.

We prove that these estimators are asymptotically consistent for sequences of design matrices X with converging aspect ratio and iid Gaussian entries. We expect that the consistency holds far beyond this case. In particular, for the case of general Gaussian design matrices, consistency holds conditionally on a conjectured formula stated in [JM13] on the basis of the "replica method" from statistical physics.

For the sake of concreteness, let us briefly describe our method in the case of standard Gaussian design that is when the design matrix X has iid Gaussian entries. We construct the unbiased pseudo-data vector by

$$\widehat{\theta}^u = \widehat{\theta} + X^T(y - X\widehat{\theta})/[n - \|\widehat{\theta}\|_0]\,. \tag{1.3}$$

Our estimator of the mean squared error is derived from applying SURE to unbiased pseudo-data. In particular, our estimator is $\widehat{R}(y, X, \lambda, \widehat{\tau})$ where

$$\widehat{R}(y, X, \lambda, \tau) \equiv \tau^2\left(2\|\widehat{\theta}\|_0/p - 1\right) + \|X^T(y - X\widehat{\theta})\|_2^2 \Big/ \left[p(n - \|\widehat{\theta}\|_0)^2\right] \tag{1.4}$$

Here $\widehat{\theta}(\lambda; X, y)$ is the LASSO estimator and $\widehat{\tau} = \|y - X\widehat{\theta}\|_2/[n - \|\widehat{\theta}\|_0]$.

Our estimator of the noise level is

$$\widehat{\sigma}^2/n = \widehat{\tau}^2 - \widehat{R}(y, X, \lambda, \widehat{\tau}\,)/\delta$$

where $\delta = n/p$. Although our rigorous results are asymptotic in the problem dimensions, we show through numerical simulations that they are accurate already on problems with a few thousands of

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

[2]Note that our definition of noise level σ corresponds to $\sigma\sqrt{n}$ in most of the compressed sensing literature.

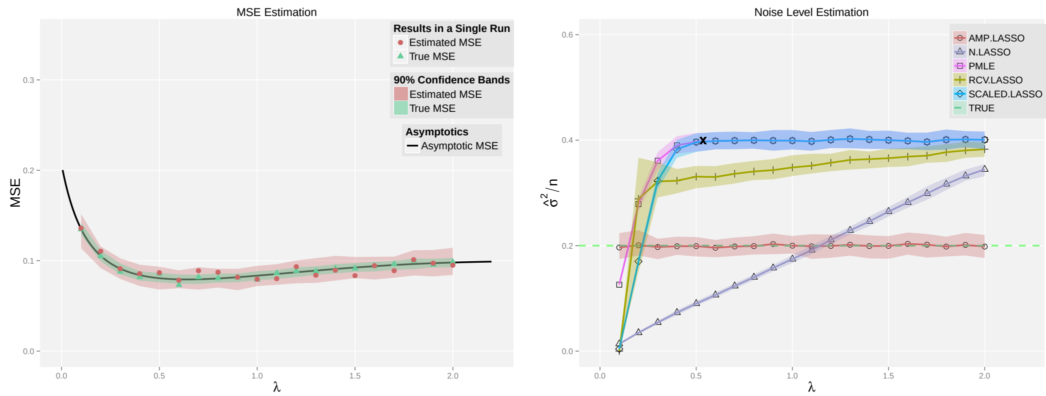

Figure 1: Red color represents the estimated values by our estimators and green color represents the true values to be estimated. *Left:* MSE versus regularization parameter λ. Here, $\delta = 0.5$, $\sigma^2/n = 0.2$, $X \in \mathbb{R}^{n \times p}$ with iid $\mathsf{N}_1(0,1)$ entries where $n = 4000$. *Right:* $\hat{\sigma}^2/n$ versus λ. Comparison of different estimators of σ^2 under the same model parameters. Scaled Lasso's prescribed choice of $(\lambda, \hat{\sigma}^2/n)$ is marked with a bold **x**.

variables. To the best of our knowledge, this is the first method for estimating the LASSO mean square error solely based on data. We compare our approach with earlier work on the estimation of the noise level. The authors of [NSvdG10] target this problem by using a ℓ_1-penalized maximum log-likelihood estimator (PMLE) and a related method called "Scaled Lasso" [SZ12] (also studied by [BC13]) considers an iterative algorithm to jointly estimate the noise level and θ_0. Moreover, authors of [FGH12] developed a refitted cross-validation (RCV) procedure for the same task. Under some conditions, the aforementioned studies provide consistency results for their noise level estimators. We compare our estimator with these methods through extensive numerical simulations.

The rest of the paper is organized as follows. In order to motivate our theoretical work, we start with numerical simulations in Section 2. The necessary background on SURE and asymptotic distributional characterization of the LASSO is presented in Section 3. Finally, our main theoretical results can be found in Section 4.

2 Simulation Results

In this section, we validate the accuracy of our estimators through numerical simulations. We also analyze the behavior of our variance estimator as λ varies, along with four other methods. Two of these methods rely on the minimization problem,

$$(\widehat{\theta}, \widehat{\sigma}) = \mathrm{argmin}_{\theta,\sigma} \left\{ \frac{\|y - X\theta\|_2^2}{2nh_1(\sigma)} + h_2(\sigma) + \lambda \, \frac{\|\theta\|_1}{2^3 \, h_3(\sigma)} \right\} ,$$

where for PMLE $h_1(\sigma) = \sigma^2$, $h_2(\sigma) = \log(\sigma)$, $h_3(\sigma) = \sigma$ and for the Scaled Lasso $h_1(\sigma) = \sigma$, $h_2(\sigma) = \sigma/2$, and $h_3(\sigma) = 1$. The third method is a *naïve* procedure that estimates the variance in two steps: (i) use the LASSO to determine the relevant variables; (ii) apply ordinary least squares on the selected variables to estimate the variance. The fourth method is *Refitted Cross-Validation* (RCV) by [FGH12] which also has two-stages. RCV requires *sure screening property* that is the model selected in its first stage includes all the relevant variables. Note that this requirement may not be satisfied for many values of λ. In our implementation of RCV, we used the LASSO for variable selection.

In our simulation studies, we used the LASSO solver `l1_ls` [SJKG07]. We simulated across 50 replications within each, we generated a new Gaussian design matrix X. We solved for LASSO over 20 equidistant λ's in the interval $[0.1, 2]$. For each λ, a new signal θ_0 and noise independent from X were generated.

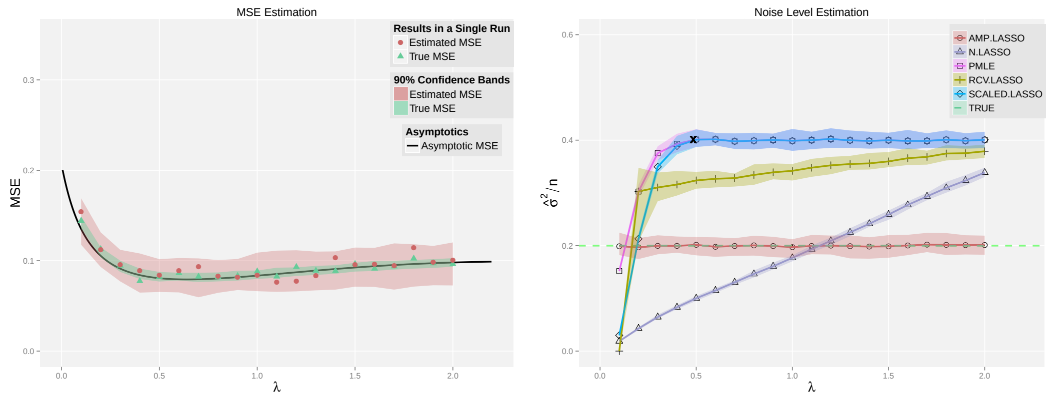

Figure 2: Red color represents the estimated values by our estimators and green color represents the true values to be estimated. *Left:* MSE versus regularization parameter λ. Here, $\delta = 0.5$, $\sigma^2/n = 0.2$, rows of $X \in \mathbb{R}^{n\times p}$ are iid from $\mathsf{N}_p(0,\Sigma)$ where $n = 5000$ and Σ has entries 1 on the main diagonal, 0.4 on above and below the main diagonal. *Right:* Comparison of different estimators of σ^2/n. Parameter values are the same as in Figure 1. Scaled Lasso's prescribed choice of $(\lambda, \hat{\sigma}^2/n)$ is marked with a bold **x**.

The results are demonstrated in Figures 1 and 2. Figure 1 is obtained using $n = 4000$, $\delta = 0.5$ and $\sigma^2/n = 0.2$. The coordinates of true signal independently get values 0, 1, -1 with probabilities 0.9, 0.05, 0.05 respectively. For each replication, we used a design matrix X where $X_{i,j} \overset{iid}{\sim} \mathsf{N}_1(0,1)$. Figure 2 is obtained with $n = 5000$ and same values of δ and σ^2 as in Figure 1. The coordinates of true signal independently get values 0, 1, -1 with probabilities 0.9, 0.05, 0.05 respectively. For each replication, we used a design matrix X where each row is independently generated through $\mathsf{N}_p(0,\Sigma)$ where Σ has 1 on the main diagonal and 0.4 above and below the diagonal.

As can be seen from the figures, the asymptotic theory applies quite well to the finite dimensional data. We refer reader to [BEM13] for a more detailed simulation analysis.

3 Background and Notations

3.1 Preliminaries and Definitions

First, we need to provide a brief introduction to *approximate message passing* (AMP) algorithm suggested by [DMM09] and its connection to LASSO (see [DMM09, BM12b] for more details).

For an appropriate sequence of non-linear denoisers $\{\eta_t\}_{t\ge 0}$, the AMP algorithm constructs a sequence of estimates $\{\theta^t\}_{t\ge 0}$, pseudo-data $\{y^t\}_{t\ge 0}$, and residuals $\{\epsilon^t\}_{t\ge 0}$ where $\theta^t, y_t \in \mathbb{R}^p$ and $\epsilon^t \in \mathbb{R}^n$. These sequences are generated according to the iteration

$$\theta^{t+1} = \eta_t(y^t)\,, \qquad y^t = \theta^t + X^T\epsilon^t/n\,, \qquad \epsilon^t = y - X\theta^t + \frac{1}{\delta}\epsilon^{t-1}\left\langle \eta_t'(y^{t-1})\right\rangle\,, \tag{3.1}$$

where $\delta \equiv n/p$ and the algorithm is initialized with $\theta^0 = \epsilon^0 = 0 \in \mathbb{R}^p$. In addition, each denoiser $\eta_t(\cdot)$ is a separable function and its derivative is denoted by $\eta_t'(\,\cdot\,)$. Given a scalar function f and a vector $u \in \mathbb{R}^m$, we let $f(u)$ denote the vector $(f(u_1),\dots,f(u_m)) \in \mathbb{R}^m$ obtained by applying f component-wise and $\langle u\rangle \equiv m^{-1}\sum_{i=1}^m u_i$ is the average of the vector $u \in \mathbb{R}^m$.

Next, consider the *state evolution* for the AMP algorithm. For the random variable $\Theta_0 \sim p_{\theta_0}$, a positive constant σ^2 and a given sequence of non-linear denoisers $\{\eta_t\}_{t\ge 0}$, define the sequence $\{\tau_t^2\}_{t\ge 0}$ iteratively by

$$\tau_{t+1}^2 = \mathsf{F}_t(\tau_t^2)\,, \qquad \mathsf{F}_t(\tau^2) \equiv \sigma^2 + \frac{1}{\delta}\,\mathbb{E}\{\,[\eta_t(\Theta_0 + \tau Z) - \Theta_0]^2\}\,, \tag{3.2}$$

where $\tau_0^2 = \sigma^2 + \mathbb{E}\{\Theta_0^2\}/\delta$ and $Z \sim \mathsf{N}_1(0,1)$ is independent of Θ_0. From Eq. 3.2, it is apparent that the function F_t depends on the distribution of Θ_0. It is shown in [BM12a] that the pseudo-data

y^t has the same asymptotic distribution as $\Theta_0 + \tau_t Z$. This result can be roughly interpreted as the pseudo-data generated by AMP is the summation of the true signal and a normally distributed noise which has zero mean. Its variance is determined by the state evolution. In other words, each iteration produces a pseudo-data that is distributed normally around the true signal, i.e. $y_i^t \approx \theta_{0,i} + \mathsf{N}_1(0, \tau_t^2)$. The importance of this result will appear later when we use Stein's method in order to obtain an estimator for the MSE and the variance of the noise.

We will use state evolution in order to describe the behavior of a specific type of converging sequence defined as the following:

Definition 1. *The sequence of instances $\{\theta_0(n), X(n), \sigma^2(n)\}_{n\in\mathbb{N}}$ indexed by n is said to be a* converging sequence *if $\theta_0(n) \in \mathbb{R}^p$, $X(n) \in \mathbb{R}^{n\times p}$, $\sigma^2(n) \in \mathbb{R}$ and $p = p(n)$ is such that $n/p \to \delta \in (0,\infty)$, $\sigma^2(n)/n \to \sigma_0^2$ for some $\sigma_0 \in \mathbb{R}$ and in addition the following conditions hold:*

(a) The empirical distribution of $\{\theta_{0,i}(n)\}_{i=1}^p$, converges in distribution to a probability measure p_{θ_0} on $\mathbb{R}$ with bounded 2^{nd} moment. Further, as $n \to \infty$, $p^{-1}\sum_{i=1}^p \theta_{0,i}(n)^2 \to \mathbb{E}_{p_{\theta_0}}\{\Theta_0^2\}$.

(b) If $\{e_i\}_{1\le i\le p} \subset \mathbb{R}^p$ denotes the standard basis, then $n^{-1/2}\max_{i\in[p]}\|X(n)e_i\|_2 \to 1$, $n^{-1/2}\min_{i\in[p]}\|X(n)e_i\|_2 \to 1$, as $n\to\infty$ with $[p] \equiv \{1,\dots,p\}$.

We provide rigorous results for the special class of converging sequences when entries of X are iid $\mathsf{N}_1(0,1)$ (i.e., *standard gaussian design model*). We also provide results (assuming Conjecture 4.4 is correct) when rows of X are iid multivariate normal $\mathsf{N}_p(0,\Sigma)$ (i.e., *general gaussian design model*).

In order to discuss the LASSO connection for the AMP algorithm, we need to use a specific class of denoisers and apply an appropriate calibration to the state evolution. Here, we provide briefly how this can be done and we refer the reader to [BEM13] for a detailed discussion.

Denote by $\eta : \mathbb{R}\times\mathbb{R}_+ \to \mathbb{R}$ the soft thresholding denoiser where

$$\eta(x;\xi) = \begin{cases} x-\xi & \text{if } x > \xi \\ 0 & \text{if } -\xi \le x \le \xi \\ x+\xi & \text{if } x < -\xi \end{cases}.$$

Also, denote by $\eta'(\,\cdot\,;\,\cdot\,)$, the derivative of the soft-thresholding function with respect to its first argument. We will use the AMP algorithm with the soft-thresholding denoiser $\eta_t(\,\cdot\,) = \eta(\,\cdot\,;\xi_t\,)$ along with a suitable sequence of thresholds $\{\xi_t\}_{t\ge 0}$ in order to obtain a connection to the LASSO.

Let $\alpha > 0$ be a constant and at every iteration t, choose the threshold $\xi_t = \alpha\tau_t$. It was shown in [DMM09] and [BM12b] that the state evolution has a unique fixed point $\tau_* = \lim_{t\to\infty}\tau_t$, and there exists a mapping $\alpha \mapsto \tau_*(\alpha)$, between those two parameters. Further, it was shown that a function $\alpha \mapsto \lambda(\alpha)$ with domain $(\alpha_{\min}(\delta),\infty)$ for some constant $\alpha_{\min}$, and given by

$$\lambda(\alpha) \equiv \alpha\tau_*\big(1 - \frac{1}{\delta}\mathbb{E}\big[\eta'(\Theta_0 + \tau_* Z; \alpha\tau_*)\big]\big)\,,$$

admits a well-defined continuous and non-decreasing inverse $\alpha : (0,\infty) \to (\alpha_{\min},\infty)$. In particular, the functions $\lambda \mapsto \alpha(\lambda)$ and $\alpha \mapsto \tau_*(\alpha)$ provide a calibration between the AMP algorithm and the LASSO where λ is the regularization parameter.

3.2 Distributional Results for the LASSO

We will proceed by stating a distributional result on LASSO which was established in [BM12b].

Theorem 3.1. *Let $\{\theta_0(n), X(n), \sigma^2(n)\}_{n\in\mathbb{N}}$ be a converging sequence of instances of the standard Gaussian design model. Denote the* LASSO *estimator of $\theta_0(n)$ by $\widehat{\theta}(n,\lambda)$ and the unbiased pseudo-data generated by* LASSO *by $\widehat{\theta}^u(n,\lambda) \equiv \widehat{\theta} + X^T(y - X\widehat{\theta})/[n - \|\widehat{\theta}\|_0]$.*
Then, as $n\to\infty$, the empirical distribution of $\{\widehat{\theta}_i^u, \theta_{0,i}\}_{i=1}^p$ converges weakly to the joint distribution of $(\Theta_0 + \tau_* Z, \Theta_0)$ where $\Theta_0 \sim p_{\theta_0}$, $\tau_* = \tau_*(\alpha(\lambda))$, $Z \sim \mathsf{N}_1(0,1)$ and Θ_0 and Z are independent random variables.

The above theorem combined with the stationarity condition of the LASSO implies that the empirical distribution of $\{\widehat{\theta}_i, \theta_{0,i}\}_{i=1}^p$ converges weakly to the joint distribution of $\big(\eta(\Theta_0 + \tau_* Z; \xi_*), \Theta_0\big)$

where $\xi_* = \alpha(\lambda)\tau_*(\alpha(\lambda))$. It is also important to emphasize a relation between the asymptotic MSE, τ_*^2 and the model variance. By Theorem 3.1 and the state evolution recursion, almost surely,

$$\lim_{p\to\infty} \|\widehat{\theta} - \theta_0\|_2^2/p = \mathbb{E}\Big[\,[\eta(\Theta_0 + \tau_* Z; \xi_*) - \Theta_0]^2\,\Big] = \delta(\tau_*^2 - \sigma_0^2)\,, \tag{3.3}$$

which will be helpful to get an estimator for the noise level.

3.3 Stein's Unbiased Risk Estimator

In [Ste81], Stein proposed a method to estimate the risk of an almost arbitrary estimator of the mean of a multivariate normal vector. A generalized form of his method can be stated as the following.

Proposition 3.2. *[Ste81]&[Joh12] Let $x, \mu \in \mathbb{R}^n$ and $\boldsymbol{V} \in \mathbb{R}^{n\times n}$ be such that $x \sim \mathsf{N}_n(\mu, \boldsymbol{V})$. Suppose that $\hat{\mu}(x) \in \mathbb{R}^n$ is an estimator of μ for which $\hat{\mu}(x) = x + g(x)$ and that $g : \mathbb{R}^n \to \mathbb{R}^n$ is weakly differentiable and that $\forall i, j \in [n]$, $\mathbb{E}_\nu[|x_i g_i(x)| + |x_j g_j(x)|] < \infty$ where ν is the measure corresponding to the multivariate Gaussian distribution $\mathsf{N}_n(\mu, \boldsymbol{V})$. Define the functional*

$$S(x, \hat{\mu}) \equiv \mathrm{Tr}(\boldsymbol{V}) + 2\mathrm{Tr}(\boldsymbol{V} Dg(x)) + \|g(x)\|_2^2\,,$$

where Dg is the vector derivative. $S(x, \hat{\mu})$ is an unbiased estimator of the risk, i.e. $\mathbb{E}_\nu \|\hat{\mu}(x) - \mu\|_2^2 = \mathbb{E}_\nu[S(x, \hat{\mu})]$.

In the literature of statistics, the above estimator is called "Stein's Unbiased Risk Estimator" or SURE. The following remark will be helpful to build intuition about our approach.

Remark 1. *If we consider the risk of soft thresholding estimator $\eta(x_i; \xi)$ for μ_i when $x_i \sim \mathsf{N}_1(\mu_i, \sigma^2)$ for $i \in [m]$, the above formula suggests the functional*

$$\frac{S(x, \eta(\,\cdot\,;\xi))}{m} = \sigma^2 - \frac{2\sigma^2}{m} \sum_{i=1}^{m} \mathbb{1}_{\{|x_i| \le \xi\}} + \frac{1}{m} \sum_{i=1}^{m} [\min\{|x_i|, \xi\}]^2\ ,$$

as an estimator of the corresponding MSE.

4 Main Results

4.1 Standard Gaussian Design Model

We start by defining two estimators that are motivated by Proposition 3.2.

Definition 2. *Define*

$$\widehat{R}_\psi(x, \tau) \equiv -\tau^2 + 2\tau^2 \langle \psi'(x) \rangle + \langle\, (\psi(x) - x)^2 \rangle\,,$$

where $x \in \mathbb{R}^m$, $\tau \in \mathbb{R}_+$, and $\psi : \mathbb{R} \to \mathbb{R}$ is a suitable non-linear function. Also for $y \in \mathbb{R}^n$ and $X \in \mathbb{R}^{n\times p}$ denote by $\widehat{R}(y, X, \lambda, \tau)$, the estimator of the mean squared error of LASSO *where*

$$\widehat{R}(y, X, \lambda, \tau) \equiv \frac{\tau^2}{p}(2\|\widehat{\theta}\|_0 - p) + \frac{\|X^T(y - X\widehat{\theta})\|_2^2}{p(n - \|\widehat{\theta}\|_0)^2}\,.$$

Remark 2. *Note that $\widehat{R}(y, X, \lambda, \tau)$ is just a special case of $\widehat{R}_\psi(x, \tau)$ when $x = \widehat{\theta}^u$ and $\psi(\,\cdot\,) = \eta(\,\cdot\,;\xi\,)$ for $\xi = \lambda/(1 - \|\widehat{\theta}\|_0/p)$.*

We are now ready to state the following theorem on the asymptotic MSE of the AMP:

Theorem 4.1. *Let $\{\theta_0(n), X(n), \sigma^2(n)\}_{n\in\mathbb{N}}$ be a converging sequence of instances of the standard Gaussian design model. Denote the sequence of estimators of $\theta_0(n)$ by $\{\theta^t(n)\}_{t\ge 0}$, the pseudo-data by $\{y^t(n)\}_{t\ge 0}$, and residuals by $\{\epsilon^t(n)\}_{t\ge 0}$ produced by* AMP *algorithm using the sequence of Lipschitz continuous functions $\{\eta_t\}_{t\ge 0}$ as in Eq. 3.1.*

Then, as $n \to \infty$, the mean squared error of the AMP *algorithm at iteration $t+1$ has the same limit as $\widehat{R}_{\eta_t}(y^t, \widehat{\tau})$ where $\widehat{\tau}_t = \|\epsilon^t\|_2/n$. More precisely, with probability one,*

$$\lim_{n\to\infty} \|\theta^{t+1} - \theta_0\|_2^2/p(n) = \lim_{n\to\infty} \widehat{R}_{\eta_t}(y^t, \widehat{\tau}_t)\,. \tag{4.1}$$

In other words, $\widehat{R}_{\eta_t}(y^t, \widehat{\tau}_t)$ is a consistent estimator of the asymptotic mean squared error of the AMP *algorithm at iteration $t+1$.*

The above theorem allows us to accurately predict how far the AMP estimate is from the true signal at iteration $t+1$ and this can be utilized as a stopping rule for the AMP algorithm. Note that it was shown in [BM12b] that the left hand side of Eq. (4.1) is $\mathbb{E}[(\eta_t(\Theta_0+\tau_t Z)-\Theta_0)^2]$. Combining this with the above theorem, we easily obtain,

$$\lim_{n\to\infty} \widehat{R}_{\eta_t}(y^t,\widehat{\tau}_t) = \mathbb{E}[(\eta_t(\Theta_0+\tau_t Z)-\Theta_0)^2]\,.$$

We state the following version of Theorem 4.1 for the LASSO.

Theorem 4.2. *Let $\{\theta_0(n), X(n), \sigma^2(n)\}_{n\in\mathbb{N}}$ be a converging sequence of instances of the standard Gaussian design model. Denote the* LASSO *estimator of $\theta_0(n)$ by $\widehat{\theta}(n,\lambda)$. Then with probability one,*

$$\lim_{n\to\infty} \|\widehat{\theta}-\theta_0\|_2^2/p(n) = \lim_{n\to\infty} \widehat{R}(y,X,\lambda,\widehat{\tau})\,,$$

where $\widehat{\tau} = \|y-X\widehat{\theta}\|_2/[n-\|\widehat{\theta}\|_0]$. In other words, $\widehat{R}(y,X,\lambda,\widehat{\tau})$ is a consistent estimator of the asymptotic mean squared error of the LASSO.

Note that Theorem 4.2 enables us to assess the quality of the LASSO estimation without knowing the true signal itself or the noise (or their distribution). The following corollary can be shown using the above theorem and Eq. 3.3.

Corollary 4.3. *In the standard Gaussian design model, the variance of the noise can be accurately estimated by $\widehat{\sigma}^2/n \equiv \widehat{\tau}^2 - \widehat{R}(y,X,\lambda,\widehat{\tau})/\delta$ where $\delta = n/p$ and other variables are defined as in Theorem 4.2. In other words, we have*

$$\lim_{n\to\infty} \hat{\sigma}^2/n = \sigma_0^2\,, \tag{4.2}$$

almost surely, providing us a consistent estimator for the variance of the noise in the LASSO.

Remark 3. *Theorems 4.1 and 4.2 provide a rigorous method for selecting the regularization parameter optimally. Also, note that obtaining the expression in Theorem 4.2 only requires solving one solution path to LASSO problem versus k solution paths required by k-fold cross-validation methods. Additionally, using the exponential convergence of AMP algorithm for the standard gaussian design model, proved by [BM12b], one can use $O(\log(1/\epsilon))$ iterations of AMP algorithm and Theorem 4.1 to obtain the solution path with an additional error up to $O(\epsilon)$.*

4.2 General Gaussian Design Model

In Section 4.1, we devised our estimators based on the standard Gaussian design model. Motivated by Theorem 4.2, we state the following conjecture of [JM13].

Let $\{\Omega(n)\}_{n\in\mathbb{N}}$ be a sequence of inverse covariance matrices. Define the general Gaussian design model by the converging sequence of instances $\{\theta_0(n), X(n), \sigma^2(n)\}_{n\in\mathbb{N}}$ where for each n, rows of design matrix $X(n)$ are iid multivariate Gaussian, i.e. $\mathsf{N}_p(0,\Omega(n)^{-1})$.

Conjecture 4.4 ([JM13]). *Let $\{\theta_0(n), X(n), \sigma^2(n)\}_{n\in\mathbb{N}}$ be a converging sequence of instances under the general Gaussian design model with a sequence of proper inverse covariance matrices $\{\Omega(n)\}_{n\in\mathbb{N}}$. Assume that the empirical distribution of $\{(\theta_{0,i},\Omega_{ii}\}_{i=1}^p$ converges weakly to the distribution of a random vector (Θ_0,Υ). Denote the* LASSO *estimator of $\theta_0(n)$ by $\widehat{\theta}(n,\lambda)$ and the* LASSO *pseudo-data by $\widehat{\theta}^u(n,\lambda) \equiv \widehat{\theta} + \Omega X^T(y-X\widehat{\theta})/[n-\|\widehat{\theta}\|_0]$. Then, for some $\tau\in\mathbb{R}$, the empirical distribution of $\{\theta_{0,i},\widehat{\theta}_i^u,\Omega_{ii}\}$ converges weakly to the joint distribution of $(\Theta_0, \Theta_0+\tau\Upsilon^{1/2}Z,\Upsilon)$, where $Z\sim\mathsf{N}_1(0,1)$, and (Θ_0,Υ) are independent random variables. Further, the empirical distribution of $(y-X\widehat{\theta})/[n-\|\widehat{\theta}\|_0]$ converges weakly to $\mathsf{N}(0,\tau^2)$.*

A heuristic justification of this conjecture using the replica method from statistical physics is offered in [JM13]. Using the above conjecture, we define the following generalized estimator of the linearly transformed risk under the general Gaussian design model. The construction of the estimator is essentially the same as before i.e. apply SURE to unbiased pseudo-data.

Definition 3. *For an inverse covariance matrix Ω and a suitable matrix $V \in \mathbb{R}^{p\times p}$, let $W = V\Omega V^T$ and define an estimator of $\|V(\widehat{\theta}-\theta)\|_2^2/p$ as*

$$\widehat{\Gamma}_\Omega(y,X,\tau,\lambda,V) = \frac{\tau^2}{p}\left(\mathrm{Tr}\,(W_{SS}) - \mathrm{Tr}\,(W_{\tilde{S}\tilde{S}}) - 2\mathrm{Tr}\left(W_{\tilde{S}S}\Omega_{S\tilde{S}}\Omega^{-1}_{\tilde{S}\tilde{S}}\right)\right) + \frac{\|V\Omega X^T(y-X\widehat{\theta})\|_2^2}{p(n-\|\widehat{\theta}\|_0)^2}$$

where $y \in \mathbb{R}^n$ and $X \in \mathbb{R}^{n\times p}$ denote the linear observations and the design matrix, respectively. Further, $\widehat{\theta}(n,\lambda)$ is the LASSO *solution for penalty level λ and τ is a real number. $S \subset [p]$ is the support of $\widehat{\theta}$ and $\tilde{S}$ is $[p]\setminus S$. Finally, for a $p\times p$ matrix M and subsets D, E of $[p]$ the notation M_{DE} refers to the $|D|\times|E|$ sub-matrix of M obtained by intersection of rows with indices from D and columns with indices from E.*

Derivation of the above formula is rather complicated and we refer the reader to [BEM13] for a detailed argument. A notable case, when $V = I$, corresponds to the mean squared error of LASSO for the general Gaussian design and the estimator $\widehat{R}(y,X,\lambda,\tau)$ is just a special case of the estimator $\widehat{\Gamma}_\Omega(y,X,\tau,\lambda,V)$. That is, when $V = \Omega = I$, we have $\widehat{\Gamma}_I(y,X,\tau,\lambda,I) = \widehat{R}(y,X,\lambda,\tau)$.

Now, we state the following analog of Theorem 4.2.

Theorem 4.5. *Let $\{\theta_0(n), X(n), \sigma^2(n)\}_{n\in\mathbb{N}}$ be a converging sequence of instances of the general Gaussian design model with the inverse covariance matrices $\{\Omega(n)\}_{n\in\mathbb{N}}$. Denote the* LASSO *estimator of $\theta_0(n)$ by $\widehat{\theta}(n,\lambda)$. If Conjecture 4.4 holds, then, with probability one,*

$$\lim_{n\to\infty}\|\widehat{\theta}-\theta_0\|_2^2/p(n) = \lim_{n\to\infty}\widehat{\Gamma}_\Omega(y,X,\widehat{\tau},\lambda,I)$$

where $\widehat{\tau} = \|y-X\widehat{\theta}\|_2/[n-\|\widehat{\theta}\|_0]$. In other words, $\widehat{\Gamma}_\Omega(y,X,\widehat{\tau},\lambda,I)$ is a consistent estimator of the asymptotic MSE *of the* LASSO*.*

We will assume that a similar state evolution holds for the general design. In fact, for the general case, replica method suggests the relation

$$\lim_{n\to\infty}\|\Omega^{-\frac{1}{2}}(\widehat{\theta}-\theta)\|_2^2/p(n) = \delta(\tau^2-\sigma_0^2).$$

Hence motivated by the Corollary 4.3, we state the following result on the general Gaussian design model.

Corollary 4.6. *Assume that Conjecture 4.4 holds. In the general Gaussian design model, the variance of the noise can be accurately estimated by*

$$\hat{\sigma}^2(n,\Omega)/n \equiv \widehat{\tau}^2 - \widehat{\Gamma}_\Omega(y,X,\widehat{\tau},\lambda,\Omega^{-\frac{1}{2}})/\delta\,,$$

where $\delta = n/p$ and other variables are defined as in Theorem 4.5. Also, we have

$$\lim_{n\to\infty}\hat{\sigma}^2/n = \sigma_0^2\,,$$

almost surely, providing us a consistent estimator for the noise level in LASSO*.*

Corollary 4.6, extends the results stated in Corollary 4.3 to the general Gaussian design matrices. The derivation of formulas in Theorem 4.5 and Corollary 4.6 follows similar arguments as in the standard Gaussian design model. In particular, they are obtained by applying SURE to the distributional result of Conjecture 4.4 and using the stationary condition of the LASSO. Details of this derivation can be found in [BEM13].

References

[BC13] A. Belloni and V. Chernozhukov, *Least Squares after Model Selection in High-Dimensional Sparse Models*, Bernoulli (2013).

[BdG11] P. Bühlmann and S. Van de Geer, *Statistics for high-dimensional data*, Springer-Verlag Berlin Heidelberg, 2011.

[BEM13] M. Bayati, M. A. Erdogdu, and A. Montanari, *Estimating LASSO Risk and Noise Level*, long version (in preparation), 2013.

[BM12a] M. Bayati and A. Montanari, *The dynamics of message passing on dense graphs, with applications to compressed sensing*, IEEE Trans. on Inform. Theory **57** (2012), 764–785.

[BM12b] ———, *The LASSO risk for gaussian matrices*, IEEE Trans. on Inform. Theory **58** (2012).

[BRT09] P. Bickel, Y. Ritov, and A. Tsybakov, *Simultaneous Analysis of Lasso and Dantzig Selector*, The Annals of Statistics **37** (2009), 1705–1732.

[BS05] Z. Bai and J. Silverstein, *Spectral Analysis of Large Dimensional Random Matrices*, Springer, 2005.

[BT09] A. Beck and M. Teboulle, *A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems*, SIAM J. Imaging Sciences **2** (2009), 183–202.

[BY93] Z. D. Bai and Y. Q. Yin, *Limit of the Smallest Eigenvalue of a Large Dimensional Sample Covariance Matrix*, The Annals of Probability **21** (1993), 1275–1294.

[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. Càndes, 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. Càndes 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.

[DMM11] ———, *The noise-sensitivity phase transition in compressed sensing*, Information Theory, IEEE Transactions on **57** (2011), no. 10, 6920–6941.

[FGH12] J. Fan, S. Guo, and N. Hao, *Variance estimation using refitted cross-validation in ultrahigh dimensional regression*, Journal of the Royal Statistical Society: Series B (Statistical Methodology) **74** (2012), 1467–9868.

[JM13] A. Javanmard and A. Montanari, *Hypothesis testing in high-dimensional regression under the gaussian random design model: Asymptotic theory*, preprint available in arxiv:1301.4240, 2013.

[Joh12] I. Johnstone, *Gaussian estimation: Sequence and wavelet models*, Book draft, 2012.

[MB06] N. Meinshausen and P. Bühlmann, *High-dimensional graphs and variable selection with the lasso*, The Annals of Statistics **34** (2006), no. 3, 1436–1462.

[NSvdG10] P. Bühlmann N. Städler and S. van de Geer, *ℓ_1-penalization for Mixture Regression Models (with discussion)*, Test **19** (2010), 209–285.

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

[SJKG07] M. Lustig S. Boyd S. J. Kim, K. Koh and D. Gorinevsky, *An Interior-Point Method for Large-Scale l1-Regularized Least Squares*, IEEE Journal on Selected Topics in Signal Processing **4** (2007), 606–617.

[Ste81] C. Stein, *Estimation of the mean of a multivariate normal distribution*, The Annals of Statistics **9** (1981), 1135–1151.

[SZ12] T. Sun and C. H. Zhang, *Scaled sparse linear regression*, Biometrika (2012), 1–20.

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

[Wai09] M. J. Wainwright, *Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 constrained quadratic programming*, Information Theory, IEEE Transactions on **55** (2009), no. 5, 2183–2202.

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

Supplementary Material for

Estimating LASSO Risk and Noise Level

5 Proof of Main Results

The proof of main results will be build on the techniques developed in [BM12a] and [BM12b].

We start by proving Theorem 4.1. Then we will proceed to the main theorem on LASSO. Note that proof for the auxilary lemmas appear in Section 6.

Proof of Theorem 4.1. For any $t \geq 1$, $n \in \mathbb{N}$, we have

$$\begin{aligned}\left|\widehat{R}_{\eta_t}(y^t(n), \widehat{\tau}_t) - \frac{\|\eta_t(y^t) - \theta_0\|_2^2}{p}\right| =& \Big| - \widehat{\tau}_t^2 + 2\widehat{\tau}_t^2\langle \eta_t'(y^t)\rangle + \langle \eta_t(y^t) - y^t, \eta_t(y^t) - y^t\rangle \\ &- \langle \eta_t(y^t), \eta_t(y^t)\rangle + 2\langle \eta_t(y^t), \theta_0\rangle - \langle \theta_0, \theta_0\rangle\Big| \\ =& \Big| - \widehat{\tau}_t^2 + 2\widehat{\tau}_t^2\langle \eta_t'(y^t)\rangle - 2\langle \eta_t(y^t), y^t\rangle + \langle y^t, y^t\rangle \\ &+ 2\langle \eta_t(y^t), \theta_0\rangle - \langle \theta_0, \theta_0\rangle\Big|, \end{aligned} \tag{5.1}$$

with probability one. We will prove that the right hand side of Eq. 5.1 converges to 0 almost surely. We take a moment to state some useful results that are easily obtained by using Lemma 9.5. We have the following asymptotic results for the AMP outputs:

$$\lim_{n\to\infty} \langle \theta_0 - y^t, \eta_t(y^t)\rangle \overset{\text{a.s.}}{=} - \lim_{n\to\infty} \left(\langle \theta_0 - y^t, \theta_0 - y^t\rangle \langle \eta_t'(y^t)\rangle\right) \tag{5.2}$$

$$\lim_{n\to\infty} \langle \theta_0 - y^t, \theta_0 - y^t\rangle \overset{\text{a.s.}}{=} \tau_t^2 \overset{\text{a.s.}}{=} \lim_{n\to\infty} \widehat{\tau}_t^2 \tag{5.3}$$

$$\lim_{n\to\infty} \langle y^t, y^t\rangle \overset{\text{a.s.}}{=} \tau_t^2 + \mathbb{E}[\theta_0^2] \tag{5.4}$$

Eq. 5.2 can be obtained by applying Lemma Lemma 9.5d to the function $\varphi(a, b) = \eta_t(b-a)$ when $r = s = t$. Similarly, first equality in Eq. 5.3 and Eq. 5.4 can be obtained by applying Lemma 9.5a to the functions $\phi_h(a, b) = a^2$ and $\phi_h(a, b) = (b-a)^2$. Lastly, the second equality in Eq. 5.3 can be obtained by Lemma 9.3. Now we are ready to bound the right hand side of Eq. 5.1.

$$\begin{aligned}\left|\widehat{R}_{\eta_t}(y^t(n), \widehat{\tau}_t) - \frac{\|\eta_t(y^t) - \theta_0\|_2^2}{p}\right| \leq& \Big|\tau_t^2 - \widehat{\tau}_t^2\Big| + \Big|\langle \theta_0, \theta_0\rangle - \mathbb{E}[\theta_0^2]\Big| + \Big|\langle y^t, y^t\rangle - \tau_t^2 - \mathbb{E}[\theta_0^2]\Big| \\ &+ 2\Big|\langle \eta_t(y^t), \theta_0 - y^t\rangle + \langle \theta_0 - y^t, \theta_0 - y^t\rangle\langle \eta_t'(y^t)\rangle\Big| \\ &+ 2\left|\langle \eta_t'(y^t)\rangle\right| \, \Big|\langle \theta_0 - y^t, \theta_0 - y^t\rangle - \tau_t^2\Big| \end{aligned}$$

By using the definition of converging sequences and comparing the right-hand side of the above inequality with the Eqs. 5.2-5.4, we easily conclude that as $n \to \infty$, the right-hand side converges to 0 almost surely. ☐

Before we proceed to prove the main theorem, we will state two simple lemmas that are going to be used when we derive the main result. Proofs for the lemmas can be found in Section 6.

Lemma 5.1. *Let $\{\theta_0(n), w(n), A(n)\}_{n\in\mathbb{N}}$ be a converging sequence of instances of the standard Gaussian design model. Denote the sequence of estimators of θ_0 produced by AMP by $\{x^t(n)\}_{t\geq 1}$. Then with probability one,*

$$\lim_{n\to\infty} \frac{\|y - Ax^t\|_2^2}{n(1-\omega_t(n))^2} = \tau_t^2$$

where $\omega_t(n) \equiv \frac{1}{\delta}\langle \eta'(y^{t-1}; \theta_{t-1})\rangle$ and τ_t^2 is determined by the state evolution.

The following lemma shows that the mean squared errors of the AMP algorithm and the LASSO are asymptotically the same.

Lemma 5.2. *Let $\{\theta_0(n), w(n), A(n)\}_{n\in\mathbb{N}}$ be a converging sequence of instances of the standard Gaussian design model. Denote the sequence of estimators of θ_0 produced by AMP calibrated for λ by $\{x^t(n)\}_{t\geq 1}$. Also denote the* LASSO *estimator by $\hat{x}(n, \lambda)$. Then with probability one,*

$$\lim_{n\to\infty} \frac{\|y - A\hat{x}\|^2}{n} = \lim_{t\to\infty}\lim_{n\to\infty} \frac{\|y - Ax^t\|^2}{n}$$

Now we are ready to prove the main theorem.

Proof of Theorem 4.2. First note that $\widehat{\tau}$, $\widehat{\xi}$ and b_n are random variables and we have

$$b_\infty = \lim_{n\to\infty} b_n = \frac{1}{\delta}\mathbb{E}\big[\eta'(\theta_0 + \tau_* Z; \theta_*)\big] \tag{5.5}$$

where the convergence takes place almost surely. This follows from weak convergence of the empirical distribution of the LASSO solution and the fact that $\theta_0 + \tau_* Z$ has a density. Then we can approximate the discontinuous zero-"norm", with smooth pseudo-Lipschitz functions[3] and obtain Eq. 5.5. This result immediately implies

$$\lim_{n\to\infty} \widehat{\xi}(n) = \theta_* = \frac{\lambda}{1-b_\infty} = \frac{\lambda}{1-\frac{1}{\delta}\mathbb{E}\big[\eta'(\theta_0+\tau_* Z;\theta_*\big]}. \tag{5.6}$$

almost surely. It is also important to point out that as a simple application of dominated convergence theorem, we have $b_\infty = \omega_*^\infty = \lim_{t\to\infty}\lim_{n\to\infty}\omega_t(n)$ almost surely (See Eq. 7.4).

By using Lemmas 5.1 and 5.2, we obtain

$$\lim_{n\to\infty} \widehat{\tau}(n)^2 = \lim_{n\to\infty} \frac{\|y - A\hat{x}\|^2}{n(1-b_n)^2} = \tau_*^2 \frac{(1-\omega_*^\infty)^2}{(1-b_\infty)^2} = \tau_*^2 \tag{5.7}$$

almost surely. This proves the convergence of the first term.

For the second term, we define random variables Y_n and Y as the following: Denote the empirical distribution of $\{\widehat{\theta}_i^u\}_{i=1}^p$ with F_n. By Theorem 3.1 F_n converges weakly to F where F is the distribution function of the random variable $\theta_0 + \tau_* Z$. By the Skorohod's Theorem, there exists random variables on the same probability space, namely Y_n and Y so that Y_n follows distribution F_n and Y follows distribution F. Now we can apply Lemma 6.2 to $F_n(\widehat{\xi}(n))$ and with probability one, we obtain

$$\lim_{p\to\infty} \frac{1}{p}\sum_{i=1}^p \mathbb{1}_{\{|\hat{y}_i|\le\hat{\xi}\}} = F(\theta_*) - F(-\theta_*) = \mathbb{E}\big[\eta'(\theta_0+\tau_* Z;\theta_*)\big].$$

where we used the absolute continuity of the density of Y.

Combining with the previous result, the second term in the estimator $\widehat{R}_\eta(\widehat{\theta}^u(n,\lambda),\widehat{\tau},\widehat{\xi})$ converges almost surely to $2\tau_*^2\mathbb{E}\big[\eta'(\theta_0+\tau_* Z;\theta_*)\big]$.

For the last term, first note that $\widehat{\xi}(n)$ is a random variable that depends on n whereas θ_* is a deterministic constant. As $n\to\infty$, we have almost surely $\widehat{\xi}(n)^2 \to \theta_*^2$ (See Eq. 5.5 and Eq. 5.6). By using Theorem 3.1 with the Portmanteau theorem on the bounded function $(a,b)\to\min\{a^2,\theta_*^2\}$, we get

$$\lim_{p\to\infty} \frac{1}{p}\sum_{i=1}^p \Big[\min\{|\widehat{\theta}_i^u|,\theta_*^2\}\Big]^2 = \mathbb{E}[\min\{(\tau_* Z-\theta_0)^2,\theta_*^2\}]$$

almost surely. Now we continue by writing the following inequality:

$$\begin{aligned}\left|\frac{1}{p}\sum_{i=1}^p \min\{(\widehat{\theta}_i^u)^2,\widehat{\xi}^2\} - \frac{1}{p}\sum_{i=1}^p \min\{(\widehat{\theta}_i^u)^2,\theta_*^2\}\right| &= \left|\frac{1}{p}\sum_{i=1}^p \Big[\min\{(\widehat{\theta}_i^u)^2,\widehat{\xi}^2\} - \min\{(\widehat{\theta}_i^u)^2,\theta_*^2\}\Big]\right| \\ &\le \left|\widehat{\xi}^2 - \theta_*^2\right|\end{aligned}$$

For the inequality, we used the fact that when a and b are any Real numbers, we have $|\min\{a,b\} - \min\{a,c\}| \le |b-c|$.

By Eq. 5.6, we have $\lim_{n\to\infty}\widehat{\xi}(n) = \theta_*$ almost surely. Hence the right-hand side converges to 0, implying

$$\lim_{p\to\infty} \frac{1}{p}\sum_{i=1}^p \Big[\min\{|\widehat{\theta}_i^u|,\widehat{\xi}\}\Big]^2 = \mathbb{E}[\min\{(\tau_* Z-\theta_0)^2,\theta_*^2\}] \tag{5.8}$$

[3]A function $f:\mathbb{R}^m\to\mathbb{R}$ is called pseudo-Lipschitz if for all $x,y\in\mathbb{R}^m$ we have $|f(x)-f(y)| \le L(1+\|x\|_2+\|y\|_2)\|x-y\|_2$ for a universal positive constant L.

almost surely.

By combining our results, we get on the right-hand side,

$$\lim_{n\to\infty} \widehat{R}_\eta(\widehat{\theta}^u(n,\lambda),\widehat{\tau},\widehat{\xi}) = \tau_*^2 - 2\tau_*^2\mathbb{E}\big[\eta'(\theta_0+\tau_* Z;\theta_*)\big] + \mathbb{E}[\min\{|\tau_* Z-\theta_0|,\theta_*\}^2]$$

almost surely.

On the left-hand side, using Theorem 3.1 and the remark after it, we get

$$\lim_{p\to\infty} \frac{\|\widehat{\theta}-\theta_0\|_2^2}{p} = \mathbb{E}[(\eta(\theta_0+\tau_* Z;\theta_*)-\theta_0)^2]$$

as written explicitly in [BM12b]. Now by applying Lemma 6.1, we conclude the proof. □

6 Proof of Auxiliary Lemmas

6.1 Useful Probability Facts

The following elementary probability theory results will be useful.

Lemma 6.1. *For any random variable X with bounded second moment, $Z \sim N_1(0,1)$ independent of X, we have*

$$\mathbb{E}[(\eta(X+\tau Z;\theta)-X)^2] = \tau^2 - 2\tau^2\mathbb{E}\big[\eta'(X+\tau Z;\theta)\big] + \mathbb{E}[\min\{|\tau Z-X|,\theta\}^2],$$

where τ and θ are arbitrary positive constants.

Proof. This lemma is just an elemantary application of Proposition 3.2. If we start by conditioning on X, on the left-hand side, we get a random variable that is normally distributed around X with variance τ^2 (Note that X and Z are independent random variables). Given X, if we proceed by applying Stein's Proposition to one dimensional random variable $X+\tau Z \sim \mathsf{N}_1(X,\tau^2)$, we immediately get,

$$\mathbb{E}[(\eta(X+\tau Z;\theta)-X)^2|X] = \tau^2 - 2\tau^2\mathbb{E}\big[\eta'(X+\tau Z;\theta)|X\big] + \mathbb{E}[\min\{|\tau Z-X|,\theta\}^2|X].$$

The proposition is applicable since the soft thresholding function satisfies the constraints. Finally, the proof follows by taking expectation on both sides. □

Lemma 6.2. *Let μ_n and μ be probability measures on $(\mathbb{R}^1,\mathscr{R}^1)$ and $\mu_n \to \mu$ weakly. Let X_n be a random variable on a probability space $(\Omega,\mathscr{F},\mathbb{P})$ and $X_n \to c < \infty$ almost surely where c is a constant and a continuity point of $\mu(-\infty,x]$. Then, $\mu_n(-\infty,X_n] \to \mu(-\infty,c]$ almost surely.*

Proof. Define the subset of Ω,

$$A = \{\omega\in\Omega : X_n(\omega)\to c\},$$

where $\mathbb{P}(A)=1$ by construction. Since μ is a probability measure, the function $x\to\mu(-\infty,x]$ has at most countably many discontinuities. Hence for an $\epsilon>0$, there exists c_1 and c_2 continuity points of $\mu(-\infty,x]$ such that $c_1<c<c_2$ and

$$\mu(-\infty,c_2]-\mu(-\infty,c_1] < \epsilon/2.$$

Now for every $\omega\in A$, there exists $N_\omega\in\mathbb{N}$ such that $\forall n > N_\omega$, we have $c_1 < X_n(\omega) < c_2$, $|\mu_n(-\infty,c_2]-\mu(-\infty,c_2]| < \epsilon/2$ and $|\mu_n(-\infty,c_1]-\mu(-\infty,c_1]| < \epsilon/2$.

Now on the left-hand side we have,

$$\mu(-\infty,c]-\epsilon < \mu(-\infty,c_1]-\epsilon/2 < \mu_n(-\infty,c_1] \le \mu_n(-\infty,X_n(\omega)]$$

and on the right-hand side we have,

$$\mu(-\infty,c]+\epsilon > \mu(-\infty,c_2]+\epsilon/2 > \mu_n(-\infty,c_2] \ge \mu_n(-\infty,X_n(\omega)]$$

which implies $|\mu(-\infty,c]-\mu_n(-\infty,X_n(\omega)]| < \epsilon$. Hence we have $\forall\omega\in A$, we have $\mu_n(-\infty,X_n(\omega)] \to \mu(-\infty,c]$ which concludes the proof. □

6.2 Proof of Lemmas 5.1 and 5.2

Proof of Lemma 5.1. For any $t \geq 0$ and $n \in \mathbb{N}$, we have

$$\begin{aligned}\frac{\|y - Ax^t\|_2^2}{n(1-\omega_t)^2} &= \frac{\|y - Ax^t + \omega_t z^{t-1} - \omega_t z^{t-1}\|_2^2}{n(1-\omega_t)^2} \\ &= \frac{\|z^t - \omega_t z^{t-1}\|_2^2}{n(1-\omega_t)^2} \\ &= \frac{1}{(1-\omega_t)^2}\left(\frac{1}{n}\|z^t\|_2^2 + \frac{1}{n}\omega_t^2\|z^{t-1}\|_2^2 - 2\omega_t\langle z^t, z^{t-1}\rangle\right)\end{aligned}$$

Then , as $n \to \infty$, by Lemmas 9.3 and 9.4, the terms $\|z^t\|_2^2/n$ and $\langle z^t, z^{t-1}\rangle$ on the right-hand side, converges to τ_t^2. Hence the proof is completed. ☐

Proof of Lemma 5.2. The proof simply follows from Theorem 9.2. For any $t \geq 0$ and $n \in \mathbb{N}$,

$$\begin{aligned}\frac{1}{n}\left|\|y - A\hat{x}\|_2^2 - \|y - Ax^t\|_2^2\right| &= \frac{1}{n}\left|(x^t - \hat{x})^T[2A^Ty + A^TA(x^t - \widehat{\theta}) - 2A^TAx^t]\right| \\ &\leq \frac{1}{n}\left[2\|x^t - \widehat{\theta}\|_2\|A^Ty\|_2 + \|A(x^t - \widehat{\theta})\|_2^2 + 2\|x^t - \widehat{\theta}\|_2\|A^TAx^t\|_2\right] \\ &\leq \frac{2}{n}\sigma_{max}(A)\|x^t - \widehat{\theta}\|_2\|y\|_2 + \frac{1}{n}\sigma_{max}^2(A)\|x^t - \widehat{\theta}\|_2^2 + \frac{2}{n}\sigma_{max}^2(A)\|x^t - \widehat{\theta}\|_2\|x^t\|_2\end{aligned}$$

First inequality follows from Cauchy-Schwartz and the second one follows from Proposition 10.1. Note that as t and n goes to ∞, by Theorem 9.2, $\|x^t - \hat{x}\|^2/n \to 0$. By using standard asymptotic estimate on the singular values of random matrices, together with the fact that $\{\theta_0(n), w(n), A(n)\}_{n\in\mathbb{N}}$ is a converging sequence, all the other terms are bounded. Hence the right hand side converges to 0 almost surely. ☐

7 Proof of Normality for the Pseudo-data

In this section, we will prove the distributional result for the LASSO pseudo-data. For the greater convenience of the reader, we start by stating the following theorem which was first established in [BM12a].

Theorem 7.1. *[BM12a] Let $\{\theta_0(n), w(n), A(n)\}_{n\in\mathbb{N}}$ be a converging sequence of instances of the standard Gaussian design model. Denote by z^t the residual and by y^t the pseudo-data at iteration step t produced by the* AMP *algorithm, given as in Eq. 3.1.*

Then for a fixed t, as $n \to \infty$, the empirical distribution of $\{x_{0,i}, y_i^{t+1}\}_{i=1}^p$ weakly converges to the joint distribution of $(X_0, X_0 + \tau_t Z)$ where $\theta_0 \sim p_{\theta_0}$, $Z \sim \mathsf{N}_1(0,1)$ and θ_0 and Z are independent random variables in the same probability space. τ_t is determined by the state evolution given in Eq. 3.2. Also, the empirical distribution of $\{z_i^t\}_{i=1}^n$ weakly converges to $\mathsf{N}_1(0, \tau_t^2)$.

Note that the above theorem is quite intuitive about its LASSO connection. We now state the following theorem.

Theorem 7.2. *Let $\{y^t\}_{t\geq 1}$ be the sequence of pseudo-data produced by AMP calibrated for λ and $\widehat{\theta}^u(n,\lambda) \equiv \widehat{\theta} + A^T(y - A\widehat{\theta})/(1 - b_n)$ where $\widehat{\theta}(n,\lambda)$ is the* LASSO *solution and $b_n = \|\hat{x}\|_0/n$. Then,*

$$\lim_{t\to\infty}\lim_{n\to\infty}\frac{1}{n}\|y^t(n) - \widehat{\theta}^u(n,\lambda)\|_2^2 = 0$$

almost surely.

Proof. For any $t \geq 0$, $n \in \mathbb{N}$,

$$\frac{1}{n}\|y^t - \widehat{\theta}^u\|_2^2 = \frac{1}{n}\|x^t + A^Tz^t - \widehat{\theta} - A^T(y - A\widehat{\theta})/(1-b_n)\|_2^2 \tag{7.1}$$

$$\leq \frac{2}{n}\left(\|x^t - \widehat{\theta}\|_2^2 + \left\|A^T\left(z^t - \frac{y - A\widehat{\theta}}{1-b_n}\right)\right\|_2^2\right) \tag{7.2}$$

By Theorem 9.2, first term on the right hand site converges to 0 as $t, n \to \infty$. If the second term also converges to 0, the proof will be completed. But obviously,

$$
\begin{aligned}
\frac{1}{n}\left(\left\|A^T\left(z^t-\frac{y-A\widehat{\theta}}{1-b_n}\right)\right\|_2^2\right) &\leq \frac{1}{n}\sigma_{max}^2(A)\left\|z^t-\frac{y-A\widehat{\theta}}{1-b_n}\right\|_2^2 \\
&= \frac{\sigma_{max}^2(A)}{(1-b_n)^2}\frac{1}{n}\left\|z^t(1-b_n)-y+A\widehat{\theta}\right\|_2^2 \\
&= \frac{\sigma_{max}^2(A)}{(1-b_n)^2}\frac{1}{n}\left\|z^t-\omega_t z^{t-1}+\omega_t z^{t-1}-b_n z^t-y+A\widehat{\theta}\right\|_2^2 \\
&= \frac{\sigma_{max}^2(A)}{(1-b_n)^2}\frac{1}{n}\left\|\omega_t z^{t-1}-b_n z^t+A\widehat{\theta}-Ax^t\right\|_2^2 \\
&\leq \frac{\sigma_{max}^2(A)}{(1-b_n)^2}\left(\frac{2}{n}b_n^2\left\|\frac{\omega_t}{b_n}z^{t-1}-z^t\right\|_2^2+\frac{2}{n}\left\|A(x^t-\widehat{\theta})\right\|_2^2\right) \qquad (7.3)
\end{aligned}
$$

First, note that by Lemma 9.5, we have

$$
\lim_{n\to\infty}\omega_t(n)=\omega_t^\infty \equiv \frac{1}{\delta}\mathbb{E}\big[\eta'(\theta_0+\tau_{t-1}Z;\theta_{t-1})\big] \qquad (7.4)
$$

Notice that the function $\eta'(\,\cdot\,;\theta_t)$ is discontinuous and therefore Theorem 7.1 does not apply immediately. On the other hand, Lemma 9.5 implies that the empirical distribution of $\{(A^* z_i^{t-1}+x_i^{t-1}, x_{0,i})\}_{1\leq i\leq p}$ converges weakly to the distribution of $(\theta_0+\tau_{t-1}Z,\theta_0)$. The claim follows from the fact that $\theta_0+\tau_{t-1}Z$ has a density, together with the standard properties of weak convergence.

Similar to Eq. 7.4, we state the following equation to show right hand side of Eq. 7.8 converges to 0. Note that this equation appeared before when we were proving the main theorem.

Under the conditions of Theorem 7.2, we have

$$
\lim_{n\to\infty} b_n = \frac{1}{\delta}\mathbb{E}\big[\eta'(\theta_0+\tau_* Z;\theta_*)\big] \qquad (7.5)
$$

almost surely. The proof of is a simple exercise of convergence in distribution. It appears immediately when one approximates $\eta'(\,\cdot\,;\theta_*\,)$ with nice pseudo-Lipschitz functions. The above equation proves that $\lim_{n\to\infty} b_n = \lim_{t\to\infty}\lim_{n\to\infty}\omega_t(n)$ where the limit simply follows from dominated convergence theorem. Since the soft thresholding denoiser will produce a point mass at 0, right-hand side of Eq. 7.5 will be greater than 0 almost surely. Now on the right-hand side of Eq. 7.3, as $t\to\infty$, $n\to\infty$, the first term goes to 0 by Lemmas 9.3 and 7.5.

For the second term, we have

$$
\frac{1}{n}\left\|A(x^t-\widehat{\theta})\right\|_2^2 \leq \sigma_{max}^2(A)\frac{1}{n}\|x^t-\widehat{\theta}\|_2^2 \qquad (7.6)
$$

where $\sigma_{max}^2(A)$ is bounded and the other term converges to 0 by Theorem 9.2. Hence the proof is completed. □

Now the proof for Theorem 3.1 will follow immediately from Theorem 7.2.

Proof of Theorem 3.1. By Lemma 9.5b, we have the following result. For any $t\geq 0$ and any pseudo-Lipschitz function $\psi:\mathbb{R}^2\to\mathbb{R}$ of order 2, we have

$$
\lim_{p\to\infty}\frac{1}{p}\sum_{i=1}^{p}\psi\big(y_i^t, x_{0,i}\big)=\mathbb{E}\Big[\psi(\theta_0+\tau_t Z,\theta_0)\Big], \qquad (7.7)
$$

almost surely. This result follows by considering the iterations 9.4 and applying Lemma 9.5b to the function $(h_i^{t+1}, x_{0,i})\to\psi(x_{0,i}-h_i^{t+1}, x_{0,i})$.

Now for any $\epsilon > 0$ and $t \geq 0$, for some $L > 0$ we have,

$$\begin{aligned}
\left| \frac{1}{p} \sum_{i=1}^{p} \psi\big(y_i^t, x_{0,i}\big) - \frac{1}{p} \sum_{i=1}^{p} \psi\big(\widehat{\theta}_i^u, x_{0,i}\big) \right| &\leq \frac{L}{p} \sum_{i=1}^{p} |y_i^t - \widehat{\theta}_i^u| \big(1 + 2|x_{0,i}| + |y_i^t| + |\widehat{\theta}_i^u|\big) \\
&\leq \frac{L}{p} \|y^t - \widehat{\theta}^u\|_2 \sqrt{\sum_{i=1}^{p} \big(1 + 2|x_{0,i}| + |y_i^t| + |\widehat{\theta}_i^u|\big)^2} \\
&\leq L \frac{\|y^t - \widehat{\theta}^u\|_2}{\sqrt{p}} \sqrt{4 + \frac{8\|\theta_0\|_2^2}{p} + \frac{4\|y^t\|_2^2}{p} + \frac{4\|\widehat{\theta}^u\|_2^2}{p}}\,, \qquad (7.8)
\end{aligned}$$

where the first inequality follows from the pseudo-Lipschitz property of ψ, and the second one follows from Cauchy-Schwarz inequality. As $t \to \infty$, $n \to \infty$, the first term on the right-hand side goes to 0 by Theorem 7.2. We need the following lemma to conclude the proof.

Lemma 7.3. *Under the conditions of 7.2, there is a constant $B < \infty$, such that*

$$\begin{aligned}
\lim_{t\to\infty} \lim_{p\to\infty} \frac{1}{p} \|y^t\|_2^2 <& B, \\
\lim_{p\to\infty} \frac{1}{p} \|\widehat{\theta}^u\|_2^2 <& B,
\end{aligned}$$

with probability one.

Proof of this lemma will appear later in this section.
For the second term in Eq. 7.8, we have $\|\theta_0\|_2^2/p$ bounded by definition, $|\widehat{\theta}^u\|_2^2/p$ and $\|y^t\|_2^2/p$ are bounded by the lemma above.

Finally, the distributional result follows from Portmanteau Theorem and the fact that any Lipschitz continuous function is also pseudo-Lipschitz continuous. □

Proof of Lemma 7.3. We will use the definition of pseudo-data and state evolution to obtain a bound for pseudo-data. The first one has already appeared in the proof of main theorem. By equation 5.4, we have,

$$\lim_{n\to\infty} \langle y^t, y^t \rangle = \tau_t^2 + \mathbb{E}[\theta_0^2]$$

almost surely. $\lim_{t\to\infty} \lim_{n\to\infty} \langle y^t, y^t \rangle$ is bounded since the right hand side is bounded by definition.

For the pseudo-data generated by LASSO, we have,

$$\begin{aligned}
\frac{1}{p} \|\widehat{\theta}^u\|_2^2 &= \frac{1}{p} \left\| \widehat{\theta} + \frac{A^T(y - A\widehat{\theta})}{1 - b_n} \right\|_2^2 \\
&\leq \frac{2}{p} \|\widehat{\theta}\|_2^2 + \frac{2}{p} \frac{\|A^T(y - A\widehat{\theta})\|_2^2}{(1 - b_n)^2}.
\end{aligned}$$

By Theorem 3.1, $\|\widehat{\theta}\|_2^2/p$ converges to $\mathbb{E}\left[\eta(\theta_0 + \tau_* Z, \theta_*)^2\right]$ almost surely, as $p \to \infty$. But,

$$\mathbb{E}\left[\eta(\theta_0 + \tau_* Z, \theta_*)^2\right] \leq \mathbb{E}\left[(\theta_0 + \tau_* Z)^2\right] = \tau_*^2 + \mathbb{E}[\theta_0^2].$$

Since the right hand side is bounded by definition, we have $\lim_{p\to\infty} \|\widehat{\theta}\|_2^2/p$ bounded almost surely. For the second term, we have

$$\begin{aligned}
\frac{1}{2p} \|y - A\widehat{\theta}\|_2^2 \leq \frac{1}{p} \mathcal{C}(\widehat{\theta}) \leq \frac{1}{p} \mathcal{C}(0) &= \frac{1}{2p} \|y\|_2^2 \\
&= \frac{1}{2p} \|A\theta_0 + w\|_2^2 \\
&\leq \frac{\|w\|_2^2}{p} + \frac{\sigma_{\max}(A)^2 \|\theta_0\|_2^2}{p}.
\end{aligned}$$

$\|\theta_0\|_2^2/p$ and $\|w\|_2^2/p$ are bounded by definition. $\sigma_{\max}(A)^2$ is bounded by Theorem 10.1. We finalize the proof by taking B the maximum of two bounds obtained for two pseudo-data. □

8 Calibrating AMP for the LASSO

In order to establish the LASSO connection for the AMP algorithm, we need an appropriate calibration to the state evolution.

Denote by $\eta : \mathbb{R} \times \mathbb{R}_+ \to \mathbb{R}$ the soft thresholding denoiser

$$\eta(x;\theta) = \begin{cases} x-\theta & \text{if } x > \theta, \\ 0 & \text{if } -\theta \le x \le \theta, \\ x+\theta & \text{if } x < -\theta, \end{cases} \tag{8.1}$$

and denote by $\eta'(\,\cdot\,;\,\cdot\,)$, the derivative of the soft thresholding function with respect to its first argument. We will use the AMP algorithm with the soft-thresholding denoiser $\eta_t(\,\cdot\,) = \eta(\,\cdot\,;\theta_t\,)$ with a suitable sequence of thresholds $\{\theta_t\}_{t\ge 0}$ in order to obtain a connection to the LASSO problem.

This modifies the state evolution formula as

$$\tau_{t+1}^2 = \mathsf{F}(\tau_t^2, \theta_t)\,, \tag{8.2}$$

$$\mathsf{F}(\tau^2,\theta) \equiv \sigma^2 + \frac{1}{\delta}\,\mathbb{E}\{\,[\eta(\theta_0+\tau Z;\theta) - \theta_0]^2\}\,, \tag{8.3}$$

where the dependence of F_t to t in Eq. 3.2 is undertaken by θ_t. Now, at every iteration t in AMP, we apply the threshold $\theta_t = \alpha\tau_t$ to the pseudo-data. We have the following proposition from [DMM09].

Proposition 8.1. *[DMM09] Let $\phi(x)$ and $\Phi(x)$ be the standard Gaussian density and distribution functions, respectively. Let $\alpha_{\min} = \alpha_{\min}(\delta)$ be the unique non-negative solution of the equation*

$$(1+\alpha^2)\Phi(-\alpha) - \alpha\phi(\alpha) = \frac{\delta}{2}\,. \tag{8.4}$$

Then 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 where F is as in Eq. 8.2. Denote the fixed point by $\tau_* = \tau_*(\alpha)$. Then we have $\lim_{t\to\infty}\tau_t = \tau_*(\alpha)$. Further the convergence takes place for any initial condition and is monotone. Finally $\left|\frac{\mathrm{d}\mathsf{F}}{\mathrm{d}\tau^2}(\tau^2,\alpha\tau)\right| < 1$ at $\tau = \tau_*$.

The above proposition relates τ_* to α. Next, define the function $\alpha \mapsto \lambda(\alpha)$ on $(\alpha_{\min}(\delta), \infty)$, by

$$\lambda(\alpha) \equiv \alpha\tau_*\left(1 - \frac{1}{\delta}\mathbb{E}\big[\eta'(\theta_0 + \tau_* Z; \alpha\tau_*)\big]\right)\,. \tag{8.5}$$

This equation defines a calibration between the threshold $\theta_* \equiv \alpha\tau_*$ and the regularization parameter λ. Now, we will invert this function in order to obtain a mapping from λ to α. Define $\alpha : (0,\infty) \to (\alpha_{\min}, \infty)$ such that

$$\alpha(\lambda) \in \big\{\, a \in (\alpha_{\min},\infty) \,:\, \lambda(a) = \lambda \big\}\,. \tag{8.6}$$

The following proposition from [**?**] states that the above mapping $\lambda \mapsto \alpha(\lambda)$ is well defined.

Proposition 8.2. *[**?**] The function $\alpha \mapsto \lambda(\alpha)$ is continuous on the interval $(\alpha_{\min}, \infty)$ with $\lambda(\alpha_{\min}+) = -\infty$ and $\lim_{\alpha\to\infty}\lambda(\alpha) = \infty$. Hence the function $\lambda \mapsto \alpha(\lambda)$ satisfying Eq. (8.6) exists.*

Note that the definition of $\alpha(\lambda)$ does not imply uniqueness. But this property will simply follow from Theorem 3.1 which was stated in [BM12b]. Hence we get the following result:

Proposition 8.3. *[BM12b] 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 as in Eq. (8.5).*

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

The above statements rigorously define the relation between the fixed point of state evolution τ_* and the regularization parameter λ.

9 Useful Results from [BM12a] and [BM12b]

Our proof uses the results of [BM12a] and [BM12b]. We state copy here the crucial technical lemmas in those papers.

Theorem 9.1. *[BM12a] Let $\{\theta_0(n), w(n), A(n)\}_{n\in\mathbb{N}}$ be a converging sequence of instances of order k with the entries of $A(n)$ iid normal with mean 0 and variance $1/n$. Let $\{\eta_t\}_{t\ge 0}$ be a sequence of Lipschitz continuous functions and $\psi : \mathbb{R}\times\mathbb{R} \to \mathbb{R}$ be any pseudo-Lipschitz function of order k. Then, almost surely*

$$\lim_{p\to\infty}\frac{1}{p}\sum_{i=1}^{p}\psi\big(x_i^{t+1}, x_{0,i}\big) = \mathbb{E}\Big\{\psi\big(\eta_t(\theta_0 + \tau_t Z), \theta_0\big)\Big\}\,, \tag{9.1}$$

where $Z \sim \mathsf{N}(0,1)$ is independent of $\theta_0 \sim p_{\theta_0}$.

Theorem 9.2. *[BM12b] Let $\{\theta_0(n), w(n), A(n)\}_{n\in\mathbb{N}}$ be a converging sequence of instances of the standard Gaussian design model. Denote the sequence of estimators of θ_0 produced by AMP by $\{x^t(n)\}_{t\geq 1}$. Also denote the* LASSO *estimator by $\hat{x}(n,\lambda)$. Then with probability one,*

$$\lim_{t\to\infty}\lim_{n\to\infty}\frac{\|x^t-\hat{x}\|^2}{n}=0$$

where $y^t = x^t + A^* z^t$. θ_t and τ_t are determined by state evolution.

Lemma 9.3. *[BM12b] Under the condition of Theorem 9.1, if $\{z^t\}_{t\geq 0}$ are the* AMP *residuals, then*

$$\lim_{n\to\infty}\frac{1}{n}\,\|z^t\|_2^2=\tau_t^2\,. \tag{9.2}$$

Lemma 9.4. *[BM12b] Under the condition of Theorem 9.1, the estimates $\{x^t\}_{t\geq 0}$ and residuals $\{z^t\}_{t\geq 0}$ of* AMP *almost surely satisfy*

$$\lim_{t\to\infty}\lim_{p\to\infty}\frac{1}{p}\|x^t-x^{t-1}\|^2=0\,,\qquad \lim_{t\to\infty}\lim_{p\to\infty}\frac{1}{p}\|z^t-z^{t-1}\|^2=0\,. \tag{9.3}$$

AMP, cf. Eq. (3.1) is a special case of the general iterative procedure given by Eq. (3.1) of [BM12a]. The general case takes the general form

$$\begin{aligned} h^{t+1} &= A^* m^t - \xi_t\, q^t\,, \qquad & m^t = g_t(b^t,w)\,,\\ b^t &= A\,q^t - \lambda_t m^{t-1}\,, \qquad & q^t = f_t(h^t,\theta_0)\,, \end{aligned} \tag{9.4}$$

where $\xi_t = \langle g'(b^t,w)\rangle$, $\lambda_t = \frac{1}{\delta}\langle f_t'(h^t,x^0)\rangle$ (both derivatives are with respect to the first argument).

The general state evolution can be written for the quantities $\{\tau_t^2\}_{t\geq 0}$ and $\{\sigma_t^2\}_{t\geq 0}$ via

$$\tau_t^2 = \mathbb{E}\big\{g_t(\sigma_t Z, W)^2\big\}\,,\qquad \sigma_t^2=\frac{1}{\delta}\,\mathbb{E}\big\{f_t(\tau_{t-1}Z,\theta_0)^2\big\}\,, \tag{9.5}$$

where $W\sim p_W$ and $\theta_0\sim p_{\theta_0}$ are independent of $Z\sim \mathsf{N}(0,1)$.

The connection to the AMP can be seen by defining

$$h^{t+1} = \theta_0 - (A^* z^t + x^t)\,, \tag{9.6}$$

$$q^t = x^t - \theta_0\,, \tag{9.7}$$

$$b^t = w - z^t\,, \tag{9.8}$$

$$m^t = -z^t\,, \tag{9.9}$$

where

$$f_t(s,\theta_0) = \eta_{t-1}(\theta_0 - s) - \theta_0\,,\qquad\qquad g_t(s,w) = s-w\,, \tag{9.10}$$

and the initial condition is $q^0 = -\theta_0$.

Regarding h^t, b^t as column vectors, the equations for $b^0,\dots,b^{t-1}$ and $h^1,\dots,h^t$ can be written in matrix form as:

$$\underbrace{\big[h^1+\xi_0 q^0|h^2+\xi_1 q^1|\cdots|h^t+\xi_{t-1}q^{t-1}\big]}_{X_t} = A^*\underbrace{\big[m^0|\dots|m^{t-1}\big]}_{M_t}, \tag{9.11}$$

$$\underbrace{\big[b^0|b^1+\lambda_1 m^0|\cdots|b^{t-1}+\lambda_{t-1}m^{t-2}\big]}_{Y_t} = A\underbrace{\big[q^0|\dots|q^{t-1}\big]}_{Q_t}\,. \tag{9.12}$$

or in short $Y_t = AQ_t$ and $X_t = A^* M_t$.

Following [BM12a], we define $\mathfrak{S}_t$ as the σ-algebra generated by $b^0,\dots,b^{t-1}$, $m^0,\dots,m^{t-1}$, $h^1,\dots,h^t$, and $q^0,\dots,q^t$. The conditional distribution of the random matrix A given the σ-algebra $\mathfrak{S}_t$, is given by

$$A|_{\mathfrak{S}_t} \stackrel{\mathrm{d}}{=} E_t + \mathcal{P}_t(\tilde{A}). \tag{9.13}$$

Here $\tilde{A}\stackrel{\mathrm{d}}{=} A$ is a random matrix independent of $\mathfrak{S}_t$, and $E_t = \mathbb{E}(A|\mathfrak{S}_t)$ is given by

$$E_t = Y_t(Q_t^*Q_t)^{-1}Q_t^* + M_t(M_t^*M_t)^{-1}X_t^* - M_t(M_t^*M_t)^{-1}M_t^*Y_t(Q_t^*Q_t)^{-1}Q_t^*\,. \tag{9.14}$$

Further, $\mathcal{P}_t$ is the orthogonal projector onto subspace $\mathrm{V}_t = \{A|AQ_t = 0, A^*M_t = 0\}$, defined by

$$\mathcal{P}_t(\tilde{A}) = P_{M_t}^{\perp}\tilde{A}P_{Q_t}^{\perp}.$$

Here $P_{M_t}^{\perp} = I - P_{M_t}$, $P_{Q_t}^{\perp} = I-P_{Q_t}$, and P_{Q_t}, P_{M_t} are orthogonal projector onto column spaces of Q_t and M_t respectively.

Lemma 9.5. *Let $\{q_0(p)\}_{p\geq 0}$ and $\{A(p)\}_{p\geq 0}$ be, respectively, a sequence of initial conditions and a sequence of matrices $A \in \mathbb{R}^{n\times p}$ indexed by p with iid entries $A_{ij} \sim \mathsf{N}(0, 1/n)$. Assume $n/p \to \delta \in (0,\infty)$. Consider sequences of vectors $\{\theta_0(n), w(n)\}_{p\geq 0}$, whose empirical distributions converge weakly to probability measures p_{θ_0} and p_W on $\mathbb{R}$ with bounded $(2k-2)^{th}$ moment, and assume:*

(i) $\lim_{p\to\infty} \mathbb{E}_{\hat{p}_{\theta_0(p)}}(\theta_0^{2k-2}) = \mathbb{E}_{p_{\theta_0}}(\theta_0^{2k-2}) < \infty$.

(ii) $\lim_{p\to\infty} \mathbb{E}_{\hat{p}_{w(p)}}(W^{2k-2}) = \mathbb{E}_{p_W}(W^{2k-2}) < \infty$.

(iii) $\lim_{p\to\infty} \mathbb{E}_{\hat{p}_{q_0(p)}}(X^{2k-2}) < \infty$.

Let $\{\sigma_t, \tau_t\}_{t\geq 0}$ be defined uniquely by the recursion (9.5) *with initialization $\sigma_0^2 = \delta^{-1}\lim_{n\to\infty}\langle q^0, q^0\rangle$. Then the following hold for all $t \in \mathbb{N}\cup\{0\}$*

(a)

$$h^{t+1}|_{\mathfrak{S}_{t+1,t}} \stackrel{\mathrm{d}}{=} \sum_{i=0}^{t-1} \alpha_i h^{i+1} + \tilde{A}^* m_\perp^t + \tilde{Q}_{t+1}\vec{o}_{t+1}(1)\,, \tag{9.15}$$

$$b^t|_{\mathfrak{S}_{t,t}} \stackrel{\mathrm{d}}{=} \sum_{i=0}^{t-1} \beta_i b^i + \tilde{A} q_\perp^t + \tilde{M}_t\vec{o}_t(1)\,, \tag{9.16}$$

where $\tilde{A}$ is an independent copy of A and the matrix $\tilde{Q}_t$ ($\tilde{M}_t$) is such that its columns form an orthogonal basis for the column space of Q_t (M_t) and $\tilde{Q}_t^*\tilde{Q}_t = N\,\mathbf{I}_{t\times t}$ ($\tilde{M}_t^*\tilde{M}_t = n\,\mathbf{I}_{t\times t}$).

(b) *For all pseudo-Lipschitz functions $\phi_h, \phi_b : \mathbb{R}^{t+2} \to \mathbb{R}$ of order k*

$$\lim_{p\to\infty} \frac{1}{p}\sum_{i=1}^{p} \phi_h(h_i^1, \ldots, h_i^{t+1}, x_{0,i}) \stackrel{\text{a.s.}}{=} \mathbb{E}\big\{\phi_h(\tau_0 Z_0, \ldots, \tau_t Z_t, \theta_0)\big\}\,, \tag{9.17}$$

$$\lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n} \phi_b(b_i^0, \ldots, b_i^t, w_i) \stackrel{\text{a.s.}}{=} \mathbb{E}\big\{\phi_b(\sigma_0 \hat{Z}_0, \ldots, \sigma_t \hat{Z}_t, W)\big\}\,, \tag{9.18}$$

where $(Z_0, \ldots, Z_t)$ and $(\hat{Z}_0, \ldots, \hat{Z}_t)$ are two zero-mean gaussian vectors independent of θ_0, W, with $Z_i, \hat{Z}_i \sim \mathsf{N}(0,1)$.

(c) *For all $0 \leq r, s \leq t$ the following equations hold and all limits exist, are bounded and have degenerate distribution (i.e. they are constant random variables):*

$$\lim_{p\to\infty} \langle h^{r+1}, h^{s+1}\rangle \stackrel{\text{a.s.}}{=} \lim_{n\to\infty} \langle m^r, m^s\rangle\,, \tag{9.19}$$

$$\lim_{n\to\infty} \langle b^r, b^s\rangle \stackrel{\text{a.s.}}{=} \frac{1}{\delta}\lim_{N\to\infty} \langle q^r, q^s\rangle\,. \tag{9.20}$$

(d) *For all $0 \leq r, s \leq t$, and for any Lipschitz function $\varphi : \mathbb{R}^2 \to \mathbb{R}$, the following equations hold and all limits exist, are bounded and have degenerate distribution (i.e. they are constant random variables):*

$$\lim_{p\to\infty} \langle h^{r+1}, \varphi(h^{s+1}, \theta_0)\rangle \stackrel{\text{a.s.}}{=} \lim_{p\to\infty} \langle h^{r+1}, h^{s+1}\rangle\langle \varphi'(h^{s+1}, \theta_0)\rangle, \tag{9.21}$$

$$\lim_{n\to\infty} \langle b^r, \varphi(b^s, w)\rangle \stackrel{\text{a.s.}}{=} \lim_{n\to\infty} \langle b^r, b^s\rangle\langle \varphi'(b^s, w)\rangle\,. \tag{9.22}$$

Here φ' denotes derivative with respect to the first coordinate of φ.

(e) *For $\ell = k-1$, the following hold almost surely*

$$\limsup_{p\to\infty} \frac{1}{p}\sum_{i=1}^{p} (h_i^{t+1})^{2\ell} < \infty\,, \tag{9.23}$$

$$\limsup_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n} (b_i^t)^{2\ell} < \infty. \tag{9.24}$$

(f) *For all $0 \leq r \leq t$:*

$$\lim_{p\to\infty} \frac{1}{p}\langle h^{r+1}, q^0\rangle \stackrel{\text{a.s.}}{=} 0\,. \tag{9.25}$$

(g) *For all $0 \le r \le t$ and $0 \le s \le t-1$ the following limits exist, and there exist strictly positive constants ρ_r and ς_s (independent of p, n) such that almost surely*

$$\lim_{N\to\infty} \langle q_\perp^r, q_\perp^r \rangle > \rho_r\,, \tag{9.26}$$

$$\lim_{n\to\infty} \langle m_\perp^s, m_\perp^s \rangle > \varsigma_s\,. \tag{9.27}$$

10 Singular values of random matrices

We have used the limit behavior of extreme singular values of Gaussian matrices. The following more general result from [BY93] can be used to justify our statements. (see also [BS05]).

Theorem 10.1 ([BY93])**.** *Let $A \in \mathbb{R}^{n\times N}$ be a matrix with iid entries such that $\mathbb{E}\{A_{ij}\} = 0$, $\mathbb{E}\{A_{ij}^2\} = 1/n$, and $n = N\delta$. Let $\sigma_{\max}(A)$ be the largest singular value of A, and $\hat{\sigma}_{\min}(A)$ be its smallest non-zero singular value. Then*

$$\lim_{N\to\infty} \sigma_{\max}(A) \overset{\text{a.s.}}{=} \frac{1}{\sqrt{\delta}} + 1\,, \tag{10.1}$$

$$\lim_{N\to\infty} \hat{\sigma}_{\min}(A) \overset{\text{a.s.}}{=} \frac{1}{\sqrt{\delta}} - 1\,. \tag{10.2}$$

We have also used the following simple fact that follows from the standard singular value decomposition

$$\min\big\{\|Ax\|_2 \;:\; x \in \ker(A)^\perp,\; \|x\| = 1\big\} = \sigma_{\min}(A)\,. \tag{10.3}$$