

High dimensional robust M-estimation: asymptotic variance via approximate message passing

David Donoho¹ · Andrea Montanari^{1,2}

Received: 14 November 2013 / Revised: 28 May 2015 / Published online: 7 November 2015
© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract In a recent article, El Karoui et al. (Proc Natl Acad Sci 110(36):14557–14562, 2013) study the distribution of robust regression estimators in the regime in which the number of parameters p is of the same order as the number of samples n . Using numerical simulations and ‘highly plausible’ heuristic arguments, they unveil a striking new phenomenon. Namely, the regression coefficients contain an extra Gaussian noise component that is not explained by classical concepts such as the Fisher information matrix. We show here that that this phenomenon can be characterized rigorously using techniques that were developed by the authors for analyzing the Lasso estimator under high-dimensional asymptotics. We introduce an *approximate message passing* (AMP) algorithm to compute M-estimators and deploy *state evolution* to evaluate the operating characteristics of AMP and so also M-estimates. Our analysis clarifies that the ‘extra Gaussian noise’ encountered in this problem is fundamentally similar to phenomena already studied for regularized least squares in the setting $n < p$.

Mathematics Subject Classification 62F10 · 62F12 · 60F99

1 M-estimation under high dimensional asymptotics

Consider the traditional linear regression model

$$Y = \mathbf{X}\theta_0 + W, \quad (1)$$

✉ Andrea Montanari
montanari@stanford.edu

¹ Department of Statistics, Stanford University, Stanford, USA

² Department of Electrical Engineering, Stanford University, Stanford, USA

with $Y = (Y_1, \dots, Y_n)^\top \in \mathbb{R}^n$ a vector of responses, $\mathbf{X} \in \mathbb{R}^{n \times p}$ a known design matrix, $\theta_0 \in \mathbb{R}^p$ a vector of parameters, and $W \in \mathbb{R}^n$ random noise having zero-mean components $W = (W_1, \dots, W_n)^\top$ i.i.d. with distribution $F = F_W$ having finite second moment.¹

We are interested in estimating θ_0 from observed data² (Y, \mathbf{X}) using a traditional M-estimator, defined by a non-negative convex function $\rho : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$:

$$\hat{\theta}(Y; \mathbf{X}) \equiv \arg \min_{\theta \in \mathbb{R}^p} \mathcal{L}(\theta; Y, \mathbf{X}), \quad \mathcal{L}(\theta; Y, \mathbf{X}) \equiv \sum_{i=1}^n \rho(Y_i - \langle X_i, \theta \rangle), \quad (2)$$

where $\langle u, v \rangle = \sum_{i=1}^m u_i v_i$ is the standard scalar product in \mathbb{R}^m , and $\hat{\theta}$ is chosen arbitrarily if there is multiple minimizers.

Although this is a completely traditional problem, we consider it under *high-dimensional asymptotics* where the number of parameters p and the number of observations n are both tending to infinity, at the same rate. This is becoming a popular asymptotic model owing to the modern awareness of ‘big data’ and ‘data deluge’; but also because it leads to entirely new phenomena.

1.1 Extra Gaussian noise due to high-dimensional asymptotics

Classical statistical theory considered the situation where the number of regression parameters p is fixed and the number of samples n is tending to infinity. The asymptotic distribution was found by Huber [2, 18] to be normal $\mathbf{N}(0, \mathbf{V})$ where the asymptotic variance matrix \mathbf{V} is given by

$$\mathbf{V} = V(\psi, F_W)(\mathbf{X}^\top \mathbf{X})^{-1} \quad (3)$$

here $\psi = \rho'$ is the score function of the M-estimator and $V(\psi, F) = (\int \psi^2 dF) / (\int \psi' dF)^2$ the asymptotic variance functional of [17], and $(\mathbf{X}^\top \mathbf{X})$ the usual Gram matrix associated with the least-squares problem. Importantly, it was found that for efficient estimation—i.e. the smallest possible asymptotic variance—the optimal M-estimator depended on the probability distribution F_W of the errors W . Choosing $\psi(x) = (\log f_W(x))'$ (with f_W the density of W), the asymptotic variance functional yields $V(\psi, F_W) = 1/I(F_W)$, with $I(F)$ denoting the Fisher information. This achieves the fundamental limit on the accuracy of M-estimators [18].

In modern statistical practice there is increasing interest in applications where the number of explanatory variables p is very large, and comparable to n . Examples of

¹ With a slight abuse of notation, we shall use W to denote a random variable with the same distribution F_W .

² We denote by X_1, \dots, X_n the rows of \mathbf{X} . We often omit the arguments Y, \mathbf{X} as this dependency will hold throughout. Without loss of generality, we assume that the columns of \mathbf{X} are normalized so that $\|\mathbf{X} e_i\|_2 \approx 1$. (A more precise assumption will be formulated below.)

this new regime can be given, spanning bioinformatics, machine learning, imaging, and signal processing (a few research areas in the last domains include [8, 23, 28, 29]).

This paper considers the properties of M-estimators in the high-dimensional asymptotic $n \rightarrow \infty$, $n/p(n) \rightarrow \delta \in (1, \infty)$. In this regime, the asymptotic distribution of M-estimators no longer needs to obey the classical formula (3) in widespread use. We make a random-design assumption on the \mathbf{X} 's detailed below. We show that the asymptotic covariance matrix of the parameters is now of the form

$$\mathbf{V} = V(\tilde{\Psi}, \tilde{F}_W) \left(\mathbb{E}\{\mathbf{X}^T \mathbf{X}\} \right)^{-1}, \quad (4)$$

where V is still Huber's asymptotic variance functional, but $\tilde{\Psi}$ is the *effective score function*, which is different from ψ under high-dimensional asymptotics and \tilde{F}_W is the *effective error distribution*, which is different from F_W under high-dimensional asymptotics. In the limit $\delta \rightarrow \infty$, the effective score and the effective error distribution both tend to their classical counterparts, and one recovers $V(\psi, F_W)$.

The effective error distribution \tilde{F}_W is a convolution of the noise distribution with an extra Gaussian noise component, not seen in the classical setting (here \star denotes convolution):

$$\tilde{F}_W \equiv F_W \star \mathbf{N}(0, \tau_*^2(\psi, F_W, \delta)). \quad (5)$$

The extra Gaussian noise depends in a complex way on ψ , F_W , δ , which we characterize fully below in Theorem 4.2.

Several important insights follow immediately:

1. Existing formulas are inadequate for confidence statements about M-estimates under high dimensional asymptotics, and will need to be systematically broadened.
2. Classical maximum likelihood estimates are inefficient under high-dimensional asymptotics. The idea dominating theoretical statistics since R.A. Fisher to use $\psi = (-\log f_W)'$ as a scoring rule, does not yield the efficient estimator.
3. The usual Fisher Information bound is not necessarily attainable in the high-dimensional asymptotic, as $I(\tilde{F}_W) < I(F_W)$.

M-estimation in this high-dimensional asymptotic setting was considered in a recent article by El Karoui et al. [15], who studied the distribution of $\hat{\theta}$ for Gaussian design matrices \mathbf{X} . In short they observed empirically the basic phenomenon of extra Gaussian noise appearing in high-dimensional asymptotics and rendering classical inference incorrect. The dependence of the additional variance τ_*^2 on δ , ψ and F was characterized by [15] through a non-rigorous heuristics.³ that the authors describe as 'highly plausible and buttressed by simulations.'⁴

³ To the reader familiar with the mathematical theory of spin glasses, the argument of [15] appears analogous to the cavity method from statistical physics [24, 26, 33]. (We refer to Sect. 5 for further discussion of related work.)

⁴ After the first version of our manuscript was posted on ArXiv, Nouredine El Karoui announced an independent proof of related results, using a completely different approach.

1.2 Asymptotic variance: a formal statement

In order to provide a preview of our rigorous results, we state here formally our asymptotic variance result. As explained below, this follows as a corollary of our technical results, and—for the reader's convenience—will be restated and proved as Theorem 4.2. We will work with random Gaussian designs, defined as follows. (We refer to Sect. 5.)

Definition 1.1 We say that a sequence of random design matrices $\{\mathbf{X}(n)\}_n$, with $n \rightarrow \infty$ is a *Gaussian design* if each $\mathbf{X} = \mathbf{X}(n)$ has dimensions $n \times p$, and entries $(X_{ij})_{i \in [n], j \in [p]}$ that are i.i.d. $\mathcal{N}(0, 1/n)$. Further, $p = p(n)$ is such that $\lim_{n \rightarrow \infty} n/p(n) = \delta \in (0, \infty)$.

Associated to ρ , we introduce the family ρ_b of regularizations of ρ :

$$\rho_b(z) \equiv \min_{x \in \mathbb{R}} \left\{ b\rho(x) + \frac{1}{2}(x - z)^2 \right\}, \quad (6)$$

in words, this is the min-convolution of the original loss with a square loss. Each ρ_b has a corresponding score function

$$\Psi(z; b) = \rho'_b(z). \quad (7)$$

We are now in position to state formally our asymptotic variance result.

Theorem 1.2 Assume that the loss function ρ is strongly convex and smooth, that the sequence of matrices $\{\mathbf{X}(n)\}_n$ is a standard Gaussian design with $\delta > 1$. Further assume that F_W has finite second moment.

Then, the asymptotic variance of $\widehat{\theta}$ obeys

$$\lim_{n, p \rightarrow \infty} \text{Ave}_{i \in [p]} \text{Var}(\widehat{\theta}_i) =_{\text{a.s.}} V(\tilde{\Psi}, \tilde{F}), \quad (8)$$

where $\text{Ave}_{i \in [p]}$ denotes the average across indices i , $V(\psi, F)$ denotes the usual Huber asymptotic variance formula for M -estimates— $V(\psi, F) = (\int \psi^2 dF) / (\int \psi' dF)^2$ —and the effective score $\tilde{\Psi}$ is

$$\tilde{\Psi}(\cdot) = \Psi(\cdot; b_*),$$

while the effective noise distribution \tilde{F} is

$$\tilde{F} = F_W \star \mathcal{N}(0, \tau_*^2).$$

Here (τ_*, b_*) is defined to be the solution of the two equations below (which is proved to exist and be unique):

$$\tau^2 = \delta \mathbb{E} \left\{ \Psi(W + \tau Z; b)^2 \right\}, \quad (9)$$

$$\frac{1}{\delta} = \mathbb{E} \left\{ \Psi'(W + \tau Z; b) \right\}. \quad (10)$$

As mentioned above, this is proved as Theorem 4.2 below.

1.3 Proof strategy: approximate message passing

In the present paper, we show that this important statistical phenomenon can be *characterized rigorously*, in a way that we think fully explains the main new concepts of extra Gaussian noise, effective noise and the effective score. Our proof strategy has three steps

- Introduce an *Approximate Message Passing* (AMP) algorithm for M-estimation; an iterative procedure with the M-estimator as a fixed point, and having the effective score function $\tilde{\Psi}$ as its score function at algorithm convergence.
- Introduce *state evolution* for calculating properties of the AMP algorithm iteration by iteration. We show that these calculations are exact at each iteration in the large- n limit where we freeze the iteration number and let $n \rightarrow \infty$.

At the center of the state evolution calculation is precisely an extra Gaussian noise term that is tracked from iteration to iteration, and which is shown to converge to a nonzero noise level. In this way, state evolution makes very explicit that AMP faces at each iteration and even in the limit, an effective noise that differs from the noise W by addition of an appreciable extra independent Gaussian noise.

- Show that the AMP algorithm converges to the solution of the M-estimation problem in mean square, from which it follows that the asymptotic variance of the M-estimator is identical to the asymptotic variance of the AMP algorithm. More specifically, the asymptotic variance of the M-estimator is given by a formula involving the effective score function and the effective noise.

As it turns out, our formula for the asymptotic variance coincides with the one derived heuristically in [15, Corollary 1] although our technique is remarkably different, and our proof provides a very clear understanding of the operational significance of the terms appearing in the asymptotic variance. It also allows explicit calculation of many other operating characteristics of the M-estimator, for example when used as an outlier detector.⁵

1.4 Underlying tools

At the heart of our analysis, we are simply applying an approach developed [4, 5] for rigorous analysis of solutions to convex optimization problems under high-dimensional asymptotics.

⁵ The slightly more general [15, Result 1] covers heteroscedastic noise is not covered by the analysis of this paper, but should be provable by adapting our argument.

That approach grew out of a series of earlier papers studying the compressed sensing problem [5, 9, 11, 13]. From the perspective of this paper, those papers considered the same regression model (1) as here; however, they emphasized the challenging asymptotic regime where there are fewer observations than predictors, (i.e. $n/p(n) \rightarrow \delta \in (0, 1)$) so that even in the noiseless case, the equations $Y = \mathbf{X}\theta$ would be underdetermined. In the $p > n$ setting, it became popular to use ℓ_1 -penalized least squares (Lasso [7, 34]). That series of papers considered the Lasso convex optimization problem in the case of \mathbf{X} with iid $\mathcal{N}(0, 1/n)$ entries (just as here) and followed the same 3-step strategy we use here; namely, (1) Introducing an AMP algorithm; (2) Obtaining the asymptotic distribution of AMP by state evolution; and (3) Showing that AMP agrees with the Lasso solution in the large- n limit. This procedure proved that the Lasso solution has the asymptotic distribution

$$\widehat{\theta}^u \sim \mathcal{N}(\theta_0, (\sigma^2 + \tau_{\text{Lasso}}^2) \mathbf{I}_{p \times p}) \quad (11)$$

where σ^2 is the variance of the noise in the measurements, and τ_{Lasso}^2 is the variance of an extra Gaussian noise, not appearing in the classical setting where $p(n)/n \rightarrow 0$. The variance of this extra Gaussian noise was obtained by state evolution and shown to depend on the distribution of the coefficients being recovered, and on the noise level in a seemingly complicated way that can be characterized by a fixed-point relation, see [5, 13]. At the center of the rigorous analysis stand the papers [4, 5] which analyze recurrences of the type used by AMP and establish the validity of State Evolution in considerable generality. Those same papers stand at the center of our analysis in this paper.

Apart from allowing a simple treatment, this provides a unified understanding of the phenomenon of high-dimensional extra Gaussian noise.

1.5 The role of AMP

This paper introduces a new first-order algorithm for computing the M-estimator $\widehat{\theta}$ which is uniquely appropriate for the random-design case. This algorithm fits within the class of approximate message passing (AMP) algorithms introduced in [4, 11] (see also [27] for extensions). This algorithm is of independent interest because of its low computational complexity.

AMP has a deceptive simplicity. As an iterative procedure for convex optimization, it looks almost the same as the ‘standard’ application of simple fixed-stepsize gradient descent. However, it is intended for use in the random-design setting, and it has an extra memory term (aka reaction term) that modifies the iteration in a profound and beneficial way. In the Lasso setting, AMP algorithms have been shown to have remarkable fast convergence properties [11], far outperforming more complex-looking iterations like Nesterov and FISTA.

In the present paper, AMP has an second important wrinkle—it solves a convex optimization problem associated to minimizing ρ with iterations based on gradient descent with an objective ρ_{b_t} which varies from one iteration to the next, as b_t changes, but which does not tend to ρ in the limit.

In the present paper, AMP is mainly used as a proof device, one component of the three-part strategy outlined earlier. However, a key benefit produced by the curious features of AMP is strong heuristic insight, which would not be available for a ‘standard’ gradient-descent algorithm.

The AMP proof strategy makes visible the extra Gaussian noise appearing in the M-estimator $\hat{\theta}$. Elementary considerations show that such extra noise is present at iteration zero of AMP. State evolution faithfully tracks the dynamics of this extra noise across iterations. State evolution proves that the extra noise level does not go to zero asymptotically with increasing iterations, but instead that the extra noise level tends to a fixed nonzero value. Because AMP is solving the M-estimation problem, the M-estimator must be infected by this extra noise.

The AMP algorithm and its state evolution analysis shows that the extra noise in parameter $\hat{\theta}_i^t$ at iteration t is due to cross-parameter estimation noise leakage, where errors in the estimation of all other parameters at the previous iteration $(t - 1)$ cause extra noise to appear in $\hat{\theta}_i^t$. In the classical setting no such effect is visible. One could say that the central fact about the high-dimensional setting revealed here as well as in our earlier work [5, 9, 11, 13], is that when there are so many parameters to estimate, one cannot really insulate the estimation of any one parameter from the errors in estimation of all the other parameters.

2 Approximate message passing (AMP)

This section introduces the Approximate Message Passing algorithm relevant to the general M-estimation problem (2). Its analysis through state evolution is discussed in the next section.

The interested reader can find relevant background on AMP [4, 12], and in the tutorial papers [25, 35]. In particular, the techniques used here have been already applied to a large number of problems. An list of simple examples includes:

1. Bolthausen [6] considers the Sherrington–Kirkpatrick model of mean-field spin glasses. It develops an iterative scheme to construct a solution of the celebrated TAP (Thouless–Anderson–Palmer) equations. This scheme and its analysis can be regarded as special cases of AMP and State Evolution.
2. Our early work [4] develops the general state evolution techniques on a rigorous basis, building on ideas from [6]. It illustrates it in a few examples: (1) Linear estimation (Sect. 2.1); (2) Compressed sensing via soft-thresholding (Sect. 2.2); (3) Multi-user detection (Sect. 2.3).
A different proof technique, and a universality result are proved in [3].
3. The tutorial chapter [35] contains an expository presentation of a simple AMP algorithm for the hidden clique problem, alongside its state evolution analysis (Sect. 6). More details on the same problem can be found in [10].

2.1 A family of score functions

For the rest of the paper, we make the following smoothness assumption on ρ :

Definition 2.1 We call the loss function $\rho : \mathbb{R} \rightarrow \mathbb{R}$ *smooth* if it is continuously differentiable, with absolutely continuous derivative $\psi = \rho'$ having an a.e. derivative ψ' that is bounded: $\sup_{u \in \mathbb{R}} \psi'(u) < \infty$.

Our assumption excludes some interesting cases, such as $\rho(u) = |u|$, but includes for instance the Huber loss⁶

$$\rho_H(z; \lambda) = \begin{cases} z^2/2 & \text{if } |z| \leq \lambda, \\ \lambda|z| - \lambda^2/2 & \text{otherwise.} \end{cases} \quad (12)$$

Recall the definition of regularized loss ρ_b , and corresponding score function $\Psi(z; b)$, given in Eqs. (6) and (7). The effective score of the M-estimator belongs to this family, for a particular choice of b , explained below.

In the classical M-estimation literature [16], monotonicity and differentiability of the score function ψ is frequently useful; our assumptions on ρ guarantee these properties for the nominal score function ψ . The score family $\Psi(\cdot; b)$ has such properties as well: for any b , $\Psi(\cdot; b)$ is a strictly monotone increasing function; second, for any $b > 0$, $\Psi(\cdot; b)$ is a contraction. With Ψ' denoting differentiation with respect to the first variable, we have $\Psi'(z; b) \in (0, 1)$. For proof and further discussion, see Appendix 1.

Before proceeding, we give an example. Consider the Huber loss $\rho_H(z; \lambda)$, with score function $\psi(z; \lambda) = \min(\max(-\lambda, z), \lambda)$. We have

$$\Psi(z; b) = b\psi\left(\frac{z}{1+b}; \lambda\right).$$

In particular the shape of each Ψ is similar to ψ , but the slope of the central part is now $\|\Psi'(\cdot; b)\|_\infty = \frac{b}{1+b} < 1$.

2.2 AMP algorithm

Our proposed approximate message passing (AMP) algorithm for the optimization problem (2) is iterative, starting at iteration 0 with an initial estimate $\hat{\theta}^0 \in \mathbb{R}^p$. At iteration $t = 0, 1, 2, \dots$ it applies a simple procedure to update its estimate $\hat{\theta}^t \in \mathbb{R}^p$, producing $\hat{\theta}^{t+1}$. The procedure involves three steps.

Adjusted residuals Using the current estimate $\hat{\theta}^t$, we compute the vector of *adjusted residuals* $R^t \in \mathbb{R}^n$,

$$R^t = Y - \mathbf{X}\hat{\theta}^t + \Psi\left(R^{t-1}; b_{t-1}\right); \quad (13)$$

⁶ We expect that the proof technique developed in this paper should be generalizable to a broader class of functions ρ , at the cost of additional technical complications.

where to the ordinary residuals $Y - \mathbf{X}\hat{\theta}^t$ we here add the extra term⁷ $\Psi(R^{t-1}; b_{t-1})$.

Effective score We choose a scalar $b_t > 0$, so that the effective score $\Psi(\cdot; b_t)$ has empirical average slope $p/n \in (0, 1)$. Setting $\delta = \delta(n) = n/p > 1$, we take any solution⁸ (for instance the smallest solution) to⁹:

$$\frac{1}{\delta} = \frac{1}{n} \sum_{i=1}^n \Psi'(R_i^t; b). \quad (14)$$

Scoring We apply the effective score function $\Psi(R^t; b_t)$:

$$\hat{\theta}^{t+1} = \hat{\theta}^t + \delta \mathbf{X}^\top \Psi(R^t; b_t). \quad (15)$$

The Scoring step of the AMP iteration (15) is similar to traditional iterative methods for M-estimation, compare [2]. Indeed, using the traditional residual $z^t = Y - \mathbf{X}\theta^t$, the traditional method of scoring at iteration t would read

$$\hat{\theta}^{t+1} = \hat{\theta}^t + \frac{1}{\frac{1}{n} \sum_{i=1}^n \psi'(z_i^t)} (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \psi(z^t), \quad (16)$$

and one can see correspondences of individual terms to the method of scoring used in AMP. Of course the traditional term $[\sum_{i=1}^n \psi'(z_i^t)/n]^{-1}$ corresponds to AMP's $[\sum_{i=1}^n \Psi'(R_i^t; b_t)/n]^{-1} \equiv \delta$ (because of step 14), while the traditional term $(\mathbf{X}^\top \mathbf{X})^{-1}$ corresponds to AMP's implicit $I_{p \times p}$ —which is appropriate in the present context because our random-design assumption below makes $\mathbf{X}^\top \mathbf{X}$ behave approximately like the identity matrix.

2.3 Relation to M-estimation

The next lemma explains the reason for using the effective score $\Psi(\cdot; b_t)$ in the AMP algorithm: this is what connects the AMP iteration to M-estimation (2).

Lemma 2.2 *Let $(\hat{\theta}_*, R_*, b_*)$ be a fixed point of the AMP iteration (13), (14), (15) having $b_* > 0$. Then $\hat{\theta}_*$ is a minimizer of the problem (2). Viceversa, any minimizer $\hat{\theta}_*$ of the problem (2) corresponds to one (or more) AMP fixed points of the form $(\hat{\theta}_*, R_*, b_*)$.*

⁷ Here and below, given $f: \mathbb{R} \rightarrow \mathbb{R}$ and $v = (v_1, \dots, v_m)^\top \in \mathbb{R}^m$, we define $f(v) \in \mathbb{R}^m$ by applying f coordinate-wise to v , i.e. $f(v) \equiv (f(v_1), \dots, f(v_m))^\top$.

⁸ This equation always admits at least one solution since $b \mapsto \Psi'(r; b)$ is continuous in $b \geq 0$, with $\Psi'(r; 0) = 0$ and (for ρ strictly convex) $\Psi'(r; \infty) = 1$, cf. Proposition 6.3.

⁹ Under this prescription, the sequence b_t depends on the instance (Y, \mathbf{X}) . As explained in the next section, for the proof of our main result we will use a slightly different prescription, that is independent of the problem instance.

Proof By differentiating Eq. (2), and omitting the arguments Y, \mathbf{X} for simplicity from $\mathcal{L}(\theta; Y, \mathbf{X})$, we get

$$\nabla_{\theta} \mathcal{L}(\theta) = - \sum_{i=1}^n \rho'(Y_i - \langle X_i, \theta \rangle) X_i = -\mathbf{X}^T \rho'(Y - \mathbf{X}\theta), \quad (17)$$

where as usual ρ' is applied component-wise to vector arguments. The minimizers of $\mathcal{L}(\theta)$ are all the vectors θ for which the right hand side vanishes.

Consider then a fixed point $(\hat{\theta}_*, R_*, b_*)$, of the AMP iteration (13), (15). This satisfies the equations

$$R_* = Y - \mathbf{X}\hat{\theta}_* + \Psi(R_*; b_*), \quad (18)$$

$$0 = \delta \mathbf{X}^T \Psi(R_*; b_*). \quad (19)$$

The first equation can be written as

$$Y - \mathbf{X}\theta_* = R_* - \Psi(R_*; b_*), \quad (20)$$

Using Proposition 6.4 below, (20) implies that $\Psi(R_*; b_*) = b_* \rho'(Y - \mathbf{X}\hat{\theta}_*)$. Hence the second equation reads

$$0 = \delta b_* \mathbf{X}^T \rho'(Y - \mathbf{X}\hat{\theta}_*), \quad (21)$$

which coincides with the stationarity condition (17) for $b_* > 0$. This concludes the proof. \square

2.4 Example

To make the AMP algorithm concrete, we consider an example with $n = 1000$, $p = 200$, so $\delta = 5$. For design matrix we let $X_{i,j} \sim \mathcal{N}(0, \frac{1}{n})$, and we draw θ_0 a random vector of norm $\|\theta_0\|_2 = 6\sqrt{p}$. For the distribution $F = F_W$ of errors, we use Huber's contaminated normal distribution $\mathbf{CN}(0.05, 10)$, so that $F = 0.95\Phi + 0.05H_{10}$, where H_x denotes a unit atom at x . For the loss function, we use the Huber's $\rho_H(z; \lambda)$ with $\lambda = 3$. Starting the AMP algorithm with $\hat{\theta}^0 = 0$, we run 20 iterations.

Separately, we solved the M-estimation problem using CVX, obtaining $\hat{\theta}$.

Figure 1 (left panel) shows the progress of the AMP algorithm across iterations, presenting

$$\text{RMSE}(\hat{\theta}^t; \theta_0) \equiv \frac{1}{\sqrt{p}} \|\hat{\theta}^t - \theta_0\|_2,$$

while Fig. 1 (right panel) shows the progress of AMP in approaching the M-estimate $\hat{\theta}$, as measured by




Fig. 1 Left panel RMSE of AMP versus iteration (black curve), and its convergence to RMSE of M-estimation (constant green curve). Right panel discrepancy of AMP from M-estimate, versus iteration

$$\text{RMSE}(\hat{\theta}^t; \hat{\theta}) \equiv \frac{1}{\sqrt{p}} \|\hat{\theta}^t - \hat{\theta}\|_2.$$

As is evident, the iterations converge rapidly, and they converge to the M-estimator, both in the sense of convergence of risks—measured here by $\text{RMSE}(\hat{\theta}^t; \theta_0) \rightarrow \text{RMSE}(\hat{\theta}; \theta_0) \approx 1.6182$ —and, more directly, in convergence of the estimates themselves: $\text{RMSE}(\hat{\theta}^t; \hat{\theta}) \rightarrow 0$.

Figure 2 (left panel) shows the process by which the effective score parameter \hat{b}_t is obtained at iteration $t = 3$, while the right panel shows how \hat{b}_t behaves across iterations. In fact it converges quickly towards a limit $b_\infty \approx 0.2710$.

2.5 Contrast to iterative M-estimation

Earlier we pointed to resemblances between AMP (15) and the traditional method of scoring for obtaining M-estimators (16). In reality the two approaches are very different:

- The precise form of various terms in (13), (14) (15) is dictated by the statistical assumptions that we are making on the design \mathbf{X} . In particular the memory terms are crucial for the state evolution analysis to hold. Several papers document this point [22, 25, 27, 30, 31].
- Under classical asymptotics, where p is fixed and $n \rightarrow \infty$, it is sufficient to run a single step of such an algorithm [2], in the high-dimensional setting it is necessary




Fig. 2 Left panel determining the regularization parameter at iteration 3. Blue curve average slope (vertical) versus regularization parameter b (horizontal). The blue curve intersects desired level $0.2 = 1/\delta$ near 0.3. Right panel regularization parameter b_t versus iteration; it converges rapidly to roughly 0.2710

to iterate numerous times. The resulting analysis is considerably more complex because of correlations arising as the algorithm evolves.

3 State evolution description of AMP

State evolution is a method for computing the operating characteristics of the AMP iterates $\hat{\theta}^t$ and R^t for arbitrary fixed t , under the high-dimensional asymptotic limit $n, p \rightarrow \infty, n/p \rightarrow \delta$.

In this section we initially describe a purely formal procedure which *assumes* that the AMP adjusted residuals $R^t = Y - \mathbf{X}\hat{\theta}^t + \Psi(R^t; b_t)$ really behave as $W + \tau_t Z$, with W the error distribution and Z an independent standard normal, for $t = 0, 1, 2, \dots$. The variable τ_t^2 thus quantifies the extra Gaussian noise supposedly present in the adjusted residuals of AMP; we show how this ansatz allows one to calculate τ_t^2 for each $t = 0, 1, 2, 3, \dots$, and to calculate the limit of τ_t as $t \rightarrow \infty$. Later in the section we present a rigorous result validating the method under the random Gaussian design assumption.

3.1 Initialization of the extra variance

Under the Gaussian design assumption, suppose that \mathbf{u} is a vector in \mathbb{R}^p with norm $\|\mathbf{u}\|_2$. Then $\{\mathbb{E}\|\mathbf{X}\mathbf{u}\|_2^2\} = \|\mathbf{u}\|_2^2$. Moreover, $\mathbf{X}\mathbf{u}$ is a Gaussian random vector with entries iid $N(0, \|\mathbf{u}\|_2^2/n)$.

It will be convenient to introduce for any estimator $\tilde{\theta}$ the notation

$$\text{MSE}(\tilde{\theta}, \theta_0) = \frac{1}{p} m \|\tilde{\theta} - \theta_0\|_2^2. \quad (22)$$

So initialize AMP with a deterministic estimate $\hat{\theta}^0$, and take $R^{-1} = 0$. Then the initial residual is $R^1 = Y - \mathbf{X}\hat{\theta}^0 = W + \mathbf{X}(\theta_0 - \hat{\theta}^0)$. The terms W and $\mathbf{X}(\theta_0 - \hat{\theta}^0)$ are independent, and $\mathbf{X}(\theta_0 - \hat{\theta}^0)$ is Gaussian with variance $\tau_0^2 = \|\hat{\theta}^0 - \theta_0\|_2^2/n = \text{MSE}(\hat{\theta}^0, \theta_0)/\delta$. Consider some fixed coordinate $R^1(i)$ of R^1 . Then

$$\text{Var}(R_i^1) = \text{Var}(W) + \text{Var}(\mathbf{X}(\theta_0 - \hat{\theta}^0)) = \text{Var}(W) + \text{MSE}(\theta^0, \theta_0)/\delta.$$

Hence, when AMP is started this way, we see that the adjusted residuals initially contain an extra Gaussian noise of variance $\tau_0^2 = \text{MSE}(\hat{\theta}^0, \theta_0)/\delta$.

3.2 Evolution of the extra Gaussian variance to its ultimate limit

Assuming the adjusted residuals continue, at later iterations, to behave as $W + \tau_t Z$ with Z an independent standard normal, we now calculate τ_t^2 for each $t = 1, 2, 3, \dots$, and eventually identify the limit of τ_t as $t \rightarrow \infty$.

For a given $\tau > 0$, $\delta = n/p$ and noise distribution F_W , define the *variance map*

$$\mathcal{V}(\tau^2, b; \delta, F_W) = \delta \mathbb{E} \left\{ \Psi(W + \tau Z; b)^2 \right\},$$

where $W \sim F_W$, and, independently, $Z \sim \mathcal{N}(0, 1)$. In this display, the reader can see that extra Gaussian noise of variance τ^2 is being added to the underlying noise W , and \mathcal{V} measures the δ -scaled variance of the resulting output. Evidently for $b > 0$, $0 \leq \mathcal{V}(\tau^2, b) \cdot \delta \leq (\text{Var}(W) + \tau^2) \cdot \delta$.

Under our assumptions for Ψ , for each given specification $(\tau; \delta, F_W)$ of the ingredients besides b that go into \mathcal{V} , there is (as clarified by Lemma 6.5) a well-defined value $b = b(\tau; \delta, F_W)$ giving the smallest solution $b \geq 0$ to

$$\frac{1}{\delta} = \mathbb{E} \left\{ \Psi'(W + \tau \cdot Z; b) \right\}. \quad (23)$$

Definition 3.1 *State evolution* is an iterative process for computing the scalars $\{\tau_t^2\}_{t \geq 0}$, starting from an initial condition $\tau_0^2 \in \mathbb{R}_{\geq 0}$ following

$$\tau_{t+1}^2 = \mathcal{V}(\tau_t^2, b(\tau_t)) = \mathcal{V}(\tau_t, b(\tau_t; \delta, F_W); \delta, F_W). \quad (24)$$

Defining $\tilde{\mathcal{V}}(\tau^2) = \mathcal{V}(\tau^2, b(\tau))$, we see that the evolution of τ_t follows the iterations of the map $\tilde{\mathcal{V}}$. In particular, we make these observations:

- $\tilde{\mathcal{V}}(0) > 0$,
- $\tilde{\mathcal{V}}(\tau^2)$ is a continuous, nondecreasing function of τ .
- $\tilde{\mathcal{V}}(\tau^2) < \tau^2$ as $\tau \rightarrow \infty$.

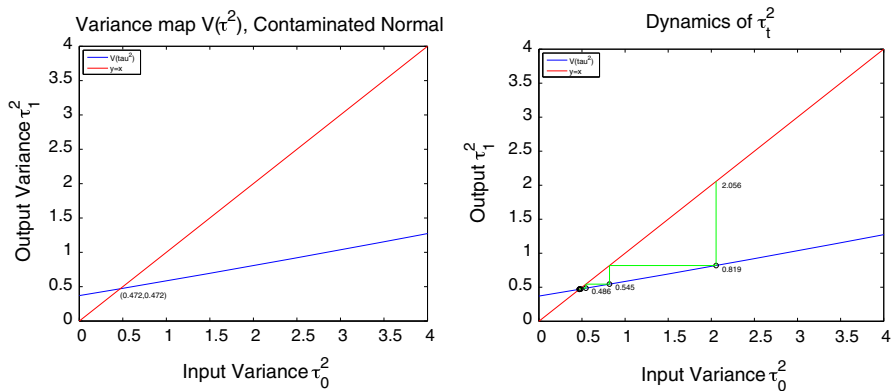


Fig. 3 The state evolution variance mapping. *Left panel* Blue curve \tilde{V} versus τ^2 . *Red curve*: diagonal; unique fixed point at about 0.472. *Right panel* the iteration history of state evolution, starting from $\tau_0^2 = 2.0556$

Figure 3, left panel, considers the case where W again follows the Huber's contaminated normal distribution $\text{CN}(0.05, 10)$ and ψ is the standard Huber estimator with parameter $\lambda = 3$. The ratio $n/p = \delta = 2$, and the parameter vector has $\|\theta_0\|_2^2/p = 6^2$. It displays the function $\tilde{V}(\tau^2)$ as a function of τ .

Evidently, there is a stable fixed point $\tau_* = \tau_*(\delta, F_W)$, i.e. a point obeying $\tilde{V}(\tau_*^2) = \tau_*^2$, such that $\tau^2 \mapsto \tilde{V}(\tau^2)$ has a derivative less than 1 at τ_*^2 . We conclude that τ_t evolves under state evolution to a nonzero limit. Figure 3, right panel, shows how τ_t^2 evolves to the fixed point near 0.472 starting from $\tau_0^2 = 2.056$.

3.3 Predicting operating characteristics from state evolution

State Evolution offers a formal¹⁰ procedure for predicting operating characteristics of the AMP iteration at any fixed iteration t or in the limit $t \rightarrow \infty$. Later in this section, we will provide rigorous validation of these predictions.

Call the tuple $S = (\tau; b, \delta, F)$ a *state*; in running the AMP algorithm we assume that the algorithm is initialized with $\hat{\theta}^0$ so that $\tau_0^2 = \text{MSE}(\hat{\theta}^0, \theta_0)/\delta$, so that AMP starts in state $S = (\tau_0; b^0, \delta, F)$, and visits $S_1 = (\tau_1; b^1, \delta, F)$, $S_2 = (\tau_2; b^2, \delta, F)$, ...; eventually AMP visits states arbitrarily close to the equilibrium state $S_* = (\tau_*; b^*, \delta, F)$.

SE predictions of operating characteristics are provided by two rules assigning predictions to certain classes of observables, based on the state that AMP is in.

Definition 3.2 The *state evolution formalism* assigns predictions \mathcal{E} to two types of observables under specific states.^x

Observables involving $\hat{\theta} - \theta_0$ Given a univariate test function $\xi : \mathbb{R} \mapsto \mathbb{R}$, assign the predicted value for $p^{-1} \sum_{i \in p} \xi(\hat{\theta}_i - \theta_{0,i})$ under state S by the rule

¹⁰ By *formal*, we mean a rule-based procedure which we can follow to get a prediction, without any guarantees that the prediction is correct.

$$\mathcal{E}(\xi(\hat{\theta} - \vartheta)|S) \equiv \mathbb{E}\left\{\xi(\sqrt{\delta} \tau Z)\right\},$$

where expectation on the right hand side is with respect to $Z \sim \mathbf{N}(0, 1)$.

Observables involving residual, error Let R denote some coordinate of the adjusted residual for AMP in state S and W the same coordinate of the underlying error. Given a bivariate test function $\xi_2 : \mathbb{R}^2 \mapsto \mathbb{R}$, assign the prediction of $n^{-1} \sum_{i=1}^n \xi_2(R_i, W_i)$ in state S by

$$\mathcal{E}(\xi_2(R, W)|S) \equiv \mathbb{E}\xi_2(W + \tau Z, W)$$

where $Z \sim \mathbf{N}(0, 1)$ and $W \sim F_W$ is independent of Z .

The two most important predictions of operating characteristics are undoubtedly:

- *MSE at iteration t .* We let $S_t = (\tau_t, b(\tau_t), \delta, F_W)$ denote the state of AMP at iteration t , and predict

$$\text{MSE}(\hat{\theta}^t, \theta_0) \approx \mathcal{E}((\hat{\vartheta} - \vartheta)^2|S_t) = \mathbb{E}\left\{(\sqrt{\delta} \tau_t Z)^2\right\} = \delta \tau_t^2.$$

- *MSE at convergence.* With $\tau_* > 0$ the limit of τ_t , let $S_* = (\tau_*, b(\tau_*), \delta, F_W)$ denote the state of AMP at convergence. and predict

$$\text{MSE}(\hat{\theta}_*, \theta_0) \approx \mathcal{E}((\hat{\vartheta} - \vartheta)^2|S_*) = \mathbb{E}\left\{(\sqrt{\delta} \tau_* Z)^2\right\} = \delta \tau_*^2.$$

Other predictions might also be of interest. Thus, concerning the mean absolute error $MAE(\hat{\theta}^t, \theta_0) = \|\hat{\theta}^t - \theta_0\|_1/p$, state evolution predicts $MAE \approx \sqrt{2\delta\tau_t^2/\pi}$. Concerning functions of (R, W) , consider the ordinary residuals $Y - X\hat{\theta}^*$ at AMP convergence. These residuals will of course in general not have the distribution of the errors W . Setting $\eta(z; b) = z - \Psi(z; b)$, we have $Y - X\hat{\theta}^* = \eta(R; b_*)$. State evolution predicts that the ordinary residuals will have the same distribution as $\eta(W + \tau_* Z; b_*)$.

3.4 Example of state evolution predictions

Continuing with our running example, we again consider the case of contaminated normal data $W \sim \mathbf{CN}(0.05, 10)$ and Huber ρ with $\lambda = 3$. If we start AMP with the all-zero estimate $\hat{\theta}^0 = 0$, then since $\|\theta_0\|_2 = 6\sqrt{p}$ we start SE with $\tau_0 = 2.056$. Figure 4 presents predictions by state evolution for the MSE (left panel) and for the mean absolute error MAE.

Again in our running example, these predictions can be tested empirically. For illustration, we conducted a very small experiment, generating 10 independent realizations of the running model at $n = 1000$ and $p = 200$, and comparing the actual evolutions of observables during AMP iterations with the predicted evolutions. Figure 5 shows that the predictions from SE are very close to the averages across realizations.




Fig. 4 State evolution predictions for $\text{CN}(0.05, 10)$, with Huber ψ , $\lambda = 3$. Predicted evolutions of two observables of $\hat{\theta}^t - \theta_0$: *Left* MSE, mean squared error. *Right* MAE mean absolute error

3.5 Correctness of state evolution predictions

The predictions of state evolution can be validated in the large-system limit $n, p \rightarrow \infty$, under the random Gaussian design assumption of Definition 1.1. We impose regularity conditions on the observables whose behavior we attempt to predict:

Definition 3.3 A function $\xi : \mathbb{R}^k \rightarrow \mathbb{R}$ is *pseudo-Lipschitz* if there exists $L < \infty$ such that, for all $x, y \in \mathbb{R}^k$, $|\xi(x) - \xi(y)| \leq L(1 + \|x\|_2 + \|y\|_2) \|x - y\|_2$.

In particular, $\xi(x) = x^2$ is pseudo-Lipschitz.

Recall also the definition of MSE in Eq. (22). For a sequence of estimators $\tilde{\theta}$, define the per-coordinate asymptotic mean squared error (AMSE) as the following large-system limit:

$$\text{AMSE}(\tilde{\theta}; \theta_0) = \text{a.s.} \lim_{n, p_n \rightarrow \infty} \text{MSE}(\tilde{\theta}; \theta_0), \quad (25)$$

when the indicated limit exists.

The following result validates the predictions of state evolution for pseudo-Lipschitz observables. Our proof is deferred to Appendix 2.

Theorem 3.4 Assume that the loss function ρ is convex and smooth, that the sequence of matrices $\{\mathbf{X}(n)\}_n$ is a standard Gaussian design, and that $\theta_0, \hat{\theta}^0$ are deterministic sequences such that $\text{AMSE}(\theta_0, \hat{\theta}^0) = \delta \tau_0^2$. Further assume that F_W has finite second moment and let $\{\tau_t^2\}_{t \geq 0}$ be the state evolution sequence with initial condition τ_0^2 . Let $\{\hat{\theta}^t, R^t\}_{t \geq 0}$ be the AMP trajectory with parameters b_t as per Eq. (23).

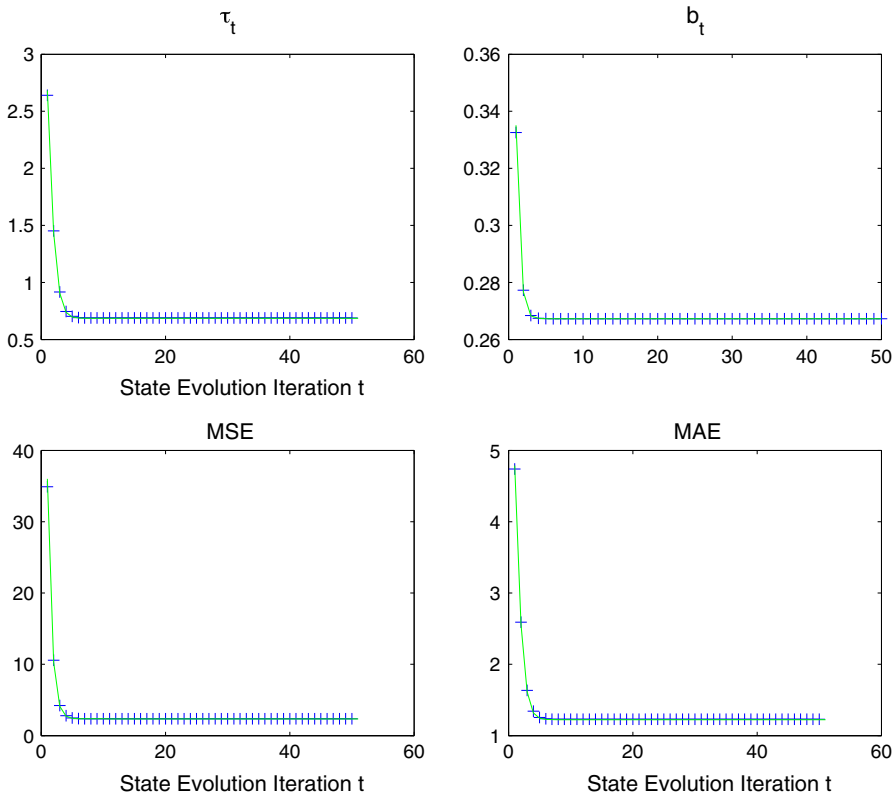


Fig. 5 Experimental means from 10 simulations compared with state evolution predictions under $CN(0.05, 10)$, with Huber ψ , $\lambda = 3$. Upper left $\hat{\tau}_t = \|\hat{\theta}^t - \theta_0\|_2 / \sqrt{n}$. Upper right \hat{b}_t . Lower left MSE mean squared error. Lower right MAE mean absolute error. Blue '+' symbols empirical means of AMP observables. Green curve theoretical predictions by SE

Let $\xi : \mathbb{R} \rightarrow \mathbb{R}$, $\xi_2 : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be pseudo-Lipschitz functions. Then, for any $t > 0$, we have, for $Z \sim N(0, 1)$ independent of $W \sim F_W$

$$\lim_{n \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p \xi(\hat{\theta}_i^t - \theta_{0,i}) =_{a.s.} \mathbb{E} \left\{ \xi(\sqrt{\delta} \tau_t Z) \right\}, \quad (26)$$

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \xi_2(R_i^t, W_i) =_{a.s.} \mathbb{E} \left\{ \xi_2(W + \tau_t Z, W) \right\}. \quad (27)$$

In particular, we may take $\xi(x) = x^2$ and obtain for the AMP iteration

$$AMSE(\hat{\theta}^t, \theta_0) = \delta \tau_t^2,$$

in full agreement with the predictions of state evolution in Definition 3.2.

4 Convergence and characterization of M-estimators

The key step for characterizing the distribution of the M-estimator $\hat{\theta}$, cf. Eq. (2), is to prove that the AMP iterates $\hat{\theta}^t$ converge to $\hat{\theta}$. We will prove that this is indeed the case, at least in the limit $n, p \rightarrow \infty$, and for suitable initial conditions.¹¹

Throughout this section, we shall assume that ρ is *strongly convex*, i.e. that $\inf_{x \in \mathbb{R}} \rho''(x) > 0$. This corresponds to assuming $\inf_{x \in \mathbb{R}} \psi'(x) > 0$, which is rather natural from the point of view of robust statistics since it ensures uniqueness of the M estimator.¹²

The key step is to establish the following high-dimensional convergence result.

Theorem 4.1 (Convergence of AMP to the M-Estimator) *Assume the same setting as in Theorem 3.4, and further assume that ρ is strongly convex and that $\delta > 1$.*

Let (τ_, b_*) be a solution of the two equations*

$$\tau^2 = \delta \mathbb{E} \left\{ \Psi(W + \tau Z; b)^2 \right\}, \quad (28)$$

$$\frac{1}{\delta} = \mathbb{E} \left\{ \Psi'(W + \tau Z; b) \right\}. \quad (29)$$

and assume that $\text{AMSE}(\hat{\theta}^0, \theta_0) = \delta \tau_^2$. Then*

$$\lim_{t \rightarrow \infty} \text{AMSE}(\hat{\theta}^t, \hat{\theta}) = 0. \quad (30)$$

From this and Theorem 3.4, the desired characterization of $\hat{\theta}$ immediately follows.

To tie back to the introduction, we prove formula (4), that we restate here for the reader's convenience.

Theorem 4.2 (Asymptotic Variance Formula under High-Dimensional Asymptotics) *Assume the setting of Theorem 3.4, and further assume that ρ is strongly convex and $\delta > 1$. The asymptotic variance of $\hat{\theta}$ obeys*

$$\lim_{n, p \rightarrow \infty} \text{Ave}_{i \in [p]} \text{Var}(\hat{\theta}_i) =_{\text{a.s.}} V(\tilde{\Psi}, \tilde{F}), \quad (31)$$

where $\text{Ave}_{i \in [p]}$ denotes the average across indices i , $V(\psi, F)$ denotes the usual Huber asymptotic variance formula for M-estimates – $V(\psi, F) = (\int \psi^2 dF) / (\int \psi' dF)^2$ – and the effective score $\tilde{\Psi}$ is

$$\tilde{\Psi}(\cdot) = \Psi(\cdot; b_*),$$

¹¹ We expect convergence for arbitrary initial conditions (as long as they are independent of (W, \mathbf{X})), but proving this claim is not needed for our main goal, and we leave it for future study. Proving this claim would require showing convergence of the state evolution recursion (24).

¹² The Huber estimator is not covered by the result of this section; although we expect our approach to apply in such generality. We focus here on the strongly convex case to avoid un-necessary complications.

while the effective noise distribution \tilde{F} is

$$\tilde{F} = F_W \star N(0, \tau_*^2).$$

Here (τ_*, b_*) are the unique solutions of the Eqs. (28), (29).

Proof By symmetry, $\text{Ave}_{i \in [p]} \text{Var}(\hat{\theta}_i) = \mathbb{E} \text{MSE}(\hat{\theta}, \theta_0)$. Theorem 4.1 and state evolution show that $\text{AMSE}(\hat{\theta}, \theta_0) = \delta \tau_*^2$. By (28) and (29)

$$V(\tilde{\Psi}, \tilde{F}) = \frac{\mathbb{E} \Psi^2(W + \tau_* Z; b_*)}{[\mathbb{E} \Psi'(W + \tau_* Z; b_*)]^2} = \frac{\tau_*^2 / \delta}{\delta^{-2}} = \delta \tau_*^2.$$

□

Corollary 4.3 Assume the setting of Theorem 3.4, and further assume that ρ is strongly convex and $\delta > 1$. Then for any pseudo-Lipschitz function $\xi : \mathbb{R} \rightarrow \mathbb{R}$, we have, for $Z \sim N(0, 1)$

$$\lim_{n \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p \xi(\hat{\theta}_i^t - \theta_{0,i}) =_{a.s.} \mathbb{E} \left\{ \xi(\sqrt{\delta} \tau_* Z) \right\}. \quad (32)$$

In particular, the solution of Eqs. (28) and (29) is necessarily unique.

Among other applications, this result can be used to bound the suboptimality of AMP after a fixed number of iterations. Combining Theorems 3.4 and 4.1 gives:

Corollary 4.4 Assume the same setting as in Theorem 3.4, and further assume that ρ is strongly convex and $\delta > 1$. Then the almost sure limits $\text{AMSE}(\hat{\theta}^t; \theta_0)$ and $\text{AMSE}(\hat{\theta}; \theta_0)$ exist, and obey

$$\text{AMSE}(\hat{\theta}^t; \theta_0) - \text{AMSE}(\hat{\theta}; \theta_0) = \delta(\tau_t^2 - \tau_*^2). \quad (33)$$

Theorem 4.1 extends to cover general Gaussian matrices \mathbf{X} with i.i.d. rows.

Definition 4.5 We say that a sequence of random design matrices $\{\mathbf{X}(n)\}_n$, with $n \rightarrow \infty$, is a *general Gaussian design* if each $\mathbf{X} = \mathbf{X}(n)$ has dimensions $n \times p$, and rows $(X_i)_{i \in [n]}$ that are i.i.d. $N(0, \Sigma/n)$, where $\Sigma = \Sigma(n) \in \mathbb{R}^{p \times p}$ is a strictly positive definite matrix. Further, $p = p(n)$ is such that $\lim_{n \rightarrow \infty} n/p(n) = \delta \in (0, \infty)$.

Notice that, if \mathbf{X} is a general Gaussian design, then $\mathbf{X}\Sigma^{-1/2}$ is a standard Gaussian design. The following then follows from Corollary 4.6 together with a simple change of variables argument, cf. [15, Lemma 1].

Corollary 4.6 Assume the same setting as in Theorem 3.4, but with $\{\mathbf{X}(n)\}_{n \geq 0}$ being a general Gaussian design with covariance Σ , and further assume that ρ is strongly convex and $\delta > 1$. There is a scalar random variable T_n so that

$$\hat{\theta} = \theta_0 + \sqrt{\delta} T_n \Sigma^{-1/2} \mathbf{Z}, \quad (34)$$

where $\mathbf{Z} \sim \mathbf{N}(0, \mathbf{I}_{p \times p})$ and we have the almost-sure limit $\lim_{n \rightarrow \infty} T_n =_{a.s.} \tau_*$, where τ_* solves Eqs. (28), (29).

This result coincides with *Corollary 1* in [15] apart from a factor \sqrt{n} in the random part of Eq. (34) that arises because of a difference in the normalization of \mathbf{X} .

5 Discussion

The asymptotic variance formula proven in Theorem 1.2 establishes rigorously the main conjecture of [15]: the two formulae for asymptotic variance coincide. As mentioned above, [15] derived this formula exclusively on the basis of heuristic arguments.¹³

Apart from providing a rigorous proof of the asymptotic variance formula, we believe that our treatment is of independent interest. In particular, we develop an efficient algorithm (AMP) that converges exponentially fast to the M-estimator, at least for random designs. This can be regarded as a iterative version of classical ‘one-step’ estimators [2] and its risk can be analyzed sharply after any number of iterations. In contrast, [15] resorted to generic convex optimization algorithms.

Several generalizations of the present proof technique should be possible, and would be of interest. We list a few in order of increasing difficulty:

1. Generalize the i.i.d. Gaussian rows model for \mathbf{X} by allowing different rows to be randomly scaled copies of a common $X \sim \mathbf{N}(0, \Sigma/n)$. This is the setting of [15, Result 1].
2. Remove the smoothness and strong convexity assumptions on ρ .
3. Add a regularization term to the objective function $\mathcal{L}(\theta)$ cf. Eq. (2), of the form $\sum_{i=1}^p J(\theta_i)$, with $J: \mathbb{R} \rightarrow \mathbb{R}$ a convex penalty. For ℓ_1 penalty and ℓ_2 loss, this reduces to the Lasso, studied [5].
4. Generalize the present results to non-Gaussian designs. We expect—for instance—that they should hold universally across matrices \mathbf{X} with i.i.d. entries (under suitable moment conditions). A similar universality result was established in [3] for compressed sensing.

A possible approach would be to use the Lindeberg principle following the route traced in [21].

Let us mention that alternative proof techniques would be worth exploring as well. In particular, Shcherbina and Tirozzi [32] define a statistical mechanics model with energy function that is analogous to the loss $\mathcal{L}(\theta)$, cf. Eq. (2), and Talagrand [33, Chapter 3] proves further results on the same model. While this treatment focuses on estimating a certain partition function, in the case of strongly convex ρ it should be possible to extract properties of the minimizer from a ‘zero-temperature’ limit.

¹³ However, as we pointed out at the end of Sect. 4, [15] proves that the case of general Gaussian designs $X_i \sim \mathbf{N}(0, \Sigma/n)$ can be reduced to standard Gaussian designs $X_i \sim \mathbf{N}(0, \mathbf{I}_{p \times p}/n)$ by a change of variables $\theta^{\text{new}} - \theta_0 = \Sigma^{1/2}(\theta - \theta_0)$.

Rangan [27] considers a similar regression model to the one studied here using approximate message passing algorithms, albeit from a Bayesian point of view.

Shortly after the present paper appeared as preprint and was submitted for publication, El Karoui [20] independently proved a result analogous to our Theorem 1.2. The proof technique is very different from ours, and of independent interest.

6 Duality between robust regression and regularized least squares

The reader might have noticed many analogies between the analysis in the last pages and earlier work on estimation in the underdetermined regime $n < p$ using the Lasso [5, 9, 11, 13]. Most specifically, the central tool in our proof of the correctness of State Evolution is a set of lemmas and theorems about analysis of recursive systems that were developed to understand the Lasso. That the same machinery directly gives results in robust regression—see for example our proof of correctness of state evolution in Appendix 2 below—might seem particularly unexpected. In this section we briefly point out that the two problems are so closely linked that phenomena which appear in one situation are bound to appear in the other.

6.1 Duality of optimization problems

In a very strong sense, solving an M-estimation problem with $p < n$ is the very same thing as solving a related penalized regression problem in $\tilde{p} > \tilde{n}$. Given a convex function $J : \mathbb{R} \rightarrow \mathbb{R}$, define the ρ function

$$\rho_J(z) \equiv \min_{x \in \mathbb{R}} \left\{ \frac{1}{2} (z - x)^2 + J(x) \right\} \quad (35)$$

We then have the M-Estimation problem

$$(M_J) \quad \min_{\theta \in \mathbb{R}^p} \sum_{i=1}^n \rho_J(Y_i - \langle X_i, \theta \rangle) \quad (36)$$

This problem has $p < n$ and is generically a determined problem. We now construct a corresponding underdetermined problem with the ‘same’ solution. Set $\tilde{n} = n - p$, $\tilde{p} = n$. We soon will construct a vector/matrix pair $(\tilde{Y} \in \mathbb{R}^{\tilde{n}}, \tilde{\mathbf{X}} \in \mathbb{R}^{\tilde{n} \times \tilde{p}})$ obeying $\tilde{n} < \tilde{p}$, where \tilde{Y} and $\tilde{\mathbf{X}}$ are related to Y and \mathbf{X} in a specific way. With this pair we pose the J -penalized least squares problem

$$(L_J) \quad \min_{\beta \in \mathbb{R}^{\tilde{p}}} \frac{1}{2} \|\tilde{Y} - \tilde{\mathbf{X}}\beta\|_2^2 + \sum_{i=1}^{\tilde{p}} J(\beta_i). \quad (37)$$

with solution $\hat{\beta}(\tilde{Y}; \tilde{\mathbf{X}})$, say.

Here is the specific pair that links (M_J) with (L_J) . We let $\tilde{\mathbf{X}}$ be a matrix with orthonormal rows such that $\tilde{\mathbf{X}}\mathbf{X} = 0$, i.e.

$$\text{null}(\tilde{\mathbf{X}}) = \text{image}(\mathbf{X}), \quad (38)$$

finally, we set $\tilde{Y} = \tilde{\mathbf{X}}Y$.

6.1.1 The Lasso-Huber connection

Of special interest is the case $J(x) = \lambda |x|$ in which case (L_J) of (37) defines the Lasso estimator. Then $\rho_J(x) = \rho_H(x; \lambda)$ is the Huber loss and (M_J) of (36) defines the Huber M-estimate. Indeed, in that case (L_J) is more classically presented as

$$(\text{Lasso}_\lambda) \quad \min_{\beta \in \mathbb{R}^{\tilde{p}}} \frac{1}{2} \|\tilde{Y} - \tilde{\mathbf{X}}\beta\|_2^2 + \lambda \sum_{i=1}^{\tilde{p}} |\beta_i|, \quad (39)$$

while (M_J) is more classically presented as

$$(\text{Huber}_\lambda) \quad \min_{\beta \in \mathbb{R}^p} \sum_{i=1}^n \rho_H(Y_i - \langle X_i, \beta \rangle; \lambda) \quad (40)$$

In this special case, our general result from the next section implies the following:

Proposition 6.1 *With problem instances (Y, X) and $(\tilde{Y}, \tilde{\mathbf{X}})$ related as above, the optimal values of the Lasso problem (Lasso_λ) and the Huber problem (Huber_λ) are identical. The solutions of the two problems are in one-one-relation. In particular, we have*

$$\hat{\theta} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top (Y - \hat{\beta}). \quad (41)$$

In a sense the Lasso problem solution $\hat{\beta}$ is finding the outliers in Y ; once the solution is known, the solution of the M-estimation problem is simply a least squares regression on adjusted data $Y_{\text{adj}} \equiv (Y - \hat{\beta})$ with outliers removed.

6.1.2 General duality result

We will now show that the problem (36) is dual to (37) under or special choice of $(\tilde{Y}, \tilde{\mathbf{X}})$, via (38).

Notation. For $x \in \mathbb{R}^n$, we denote by $\partial\rho(x)$ the subgradient of the convex function $\sum_{i=1}^n \rho(x_i)$, at x . Analogously, for $z \in \mathbb{R}^{\tilde{p}}$, we denote by $\partial J(z)$ the subgradient of the convex function $\sum_{i=1}^{\tilde{p}} J(z_i)$, at z .

Proposition 6.2 *Assume that $\rho(\cdot) = \rho_J(\cdot)$, that $\tilde{\mathbf{X}}$ has orthonormal rows with $\text{null}(\tilde{\mathbf{X}}) = \text{image}(\mathbf{X})$, and finally that $\tilde{Y} = \tilde{\mathbf{X}}Y$. Then the solutions of the regularized least squares problem (37) are in one-to-one correspondence with the solutions of the*

robust regression problem (2), via the mappings

$$\hat{\beta} = Y - \mathbf{X}\hat{\theta} - u, \quad u \in \text{null}(\mathbf{X}^T) \cap \partial\rho(Y - \mathbf{X}\hat{\theta}), \quad (42)$$

$$\hat{\theta} = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T(Y - \hat{\beta}). \quad (43)$$

Proof ‘Differentiating’ Eq. (35) it is easy to see that

$$u \in \partial\rho(x) \quad \text{if and only if} \quad u \in \partial J(x - u). \quad (44)$$

First assume $\hat{\theta}$ is a minimizer of problem (36). This happens if and only if there exists $u \in \mathbb{R}^n$ such that

$$\mathbf{X}^T u = 0, \quad u \in \partial\rho(Y - \mathbf{X}\hat{\theta}). \quad (45)$$

We then claim that $\hat{\beta} \equiv Y - \mathbf{X}\hat{\theta} - u$ is a minimizer of Eq. (37). Indeed

$$\tilde{\mathbf{X}}^T(\tilde{Y} - \tilde{\mathbf{X}}\hat{\beta}) = \tilde{\mathbf{X}}^T\tilde{\mathbf{X}}(Y - \hat{\beta}) \quad (46)$$

$$= \tilde{\mathbf{X}}^T\tilde{\mathbf{X}}(\mathbf{X}\hat{\theta} + u) = u, \quad (47)$$

where the last identity follows since, by Eq. (38), $\text{null}(\mathbf{X}^T) = \text{image}(\tilde{\mathbf{X}}^T)$, and hence $u \in \text{image}(\tilde{\mathbf{X}}^T)$ by Eq. (45). Using again Eqs. (45) and (44), we deduce that $u \in \partial J(\hat{\beta})$, i.e.

$$\tilde{\mathbf{X}}^T(\tilde{Y} - \tilde{\mathbf{X}}\hat{\beta}) \in \partial J(\hat{\beta}), \quad (48)$$

which is the stationarity condition for the problem (37).

Viceversa a similar argument shows that, given $\hat{\beta}$ that minimizes Eq. (37), and $\hat{\theta} \equiv (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T(Y - \hat{\beta})$ is a minimizer of the robust regression problem (36). \square

6.2 Comparison to AMP in the $p > n$ case

The last section raises the possibility that the phenomena found in this paper for M-estimation in the $p < n$ case are actually isomorphic to those found in our previous work on penalized regression in the $p > n$ case; [5, 9, 11, 13]. Here we merely content ourselves with sketching a few similarities.

To be definite, consider robust regression using the Huber loss [16, 17] $\rho(x) = x^2/2$ for $|x| \leq \lambda$ and $\rho(x) = \lambda|x| - \lambda^2/2$ otherwise. In this case it is easy to see that

$$\Psi(z; b) = \begin{cases} \lambda b & \text{if } z > \lambda(1 + b), \\ b z / (1 + b) & \text{if } |z| \leq \lambda(1 + b), \\ -\lambda b & \text{if } z < -\lambda(1 + b). \end{cases} \quad (49)$$

In order to make contact with the Lasso, recall the definition of soft thresholding operator $\eta(x; \alpha) = \text{sign}(x) (|x| - \alpha)_+$. We have the relationship

$$\Psi(z; b) = \frac{bz}{1+b} - \eta\left(\frac{bz}{1+b}; \lambda b\right). \quad (50)$$

Letting $c_t \equiv b_t/(1 + b_t)$, the state evolution Eq. (24), then reads

$$\tau_{t+1}^2 = \delta c_t^2 \mathbb{E} \left\{ \left[\eta\left(W + \tau_t Z; \lambda(1 + b_t)\right) - W - \tau_t Z \right]^2 \right\}. \quad (51)$$

This is very close to the state evolution equation in compressed sensing for reconstructing a sparse signal whose entries have distribution F_W , from an underdetermined number of linear measurements; indeed in that setting we have the state evolution recursion

$$\tau_{t+1}^2 = \delta \mathbb{E} \left\{ \left[\eta\left(W + \tau_t Z; \lambda \tau_t\right) - W \right]^2 \right\}; \quad (52)$$

[5, 9, 11, 13]. The connection is quite suggestive: while in compressed sensing we look for the few non-zero coefficients in the signal, in robust regression we try to identify the few outliers contaminating the linear relation. A similar duality was already pointed out in [14], although in a specific setting.

Acknowledgments This work was partially supported by the NSF CAREER award CCF-0743978, the NSF Grant DMS-0806211, and the Grants AFOSR/DARPA FA9550-12-1-0411 and FA9550-13-1-0036.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix 1: Properties of the functions Prox , Ψ

Throughout this section $\rho : \mathbb{R} \rightarrow \mathbb{R}$ is convex bounded below and smooth (i.e. with bounded second derivative). Recall the definition of $\text{Prox} : \mathbb{R} \times \mathbb{R}_{>0} \rightarrow \mathbb{R}$ and $\Psi : \mathbb{R} \times \mathbb{R}_{>0} \rightarrow \mathbb{R}$, given by

$$\text{Prox}(z; b) \equiv \arg \min_{x \in \mathbb{R}} \left\{ \rho(x) + \frac{1}{2b}(x - z)^2 \right\}, \quad (53)$$

$$\Psi(z; b) \equiv b \rho'(\text{Prox}(z; b)). \quad (54)$$

Proposition 6.3 *The function $\text{Prox} : \mathbb{R} \times \mathbb{R}_{>0} \rightarrow \mathbb{R}$ is differentiable in its domain, with partial derivatives*

$$\frac{\partial \text{Prox}}{\partial z}(z; b) = \frac{1}{1 + b\rho''(x)} \Big|_{x=\text{Prox}(z; b)}, \quad \frac{\partial \text{Prox}}{\partial b}(z; b) = -\frac{\rho'(x)}{1 + b\rho''(x)} \Big|_{x=\text{Prox}(z; b)}. \quad (55)$$

In particular, letting $\|\rho''\|_\infty \equiv \sup_{x \in \mathbb{R}} \rho''(x)$, and for any fixed b , $z \mapsto \text{Prox}(z; b)$ is strictly increasing and Lipschitz continuous, with

$$\frac{1}{1 + b\|\rho''\|_\infty} \leq \frac{\partial \text{Prox}}{\partial z}(z; b) \leq 1 \quad (56)$$

Proof Since, for $b > 0$, $x \mapsto \rho(x) + (x - z)^2/(2b)$ is differentiable and strongly convex, $x = \text{Prox}(z; b)$ is uniquely determined by setting to zero the first derivative:

$$x + b\rho'(x) - z = 0. \quad (57)$$

The claim then follows from the implicit function theorem. \square

Proposition 6.4 For $(z, b) \in \mathbb{R} \times \mathbb{R}_+$, we have

$$\Psi(z; b) = z - \text{Prox}(z, b), \quad (58)$$

and hence Ψ is differentiable, with partial derivatives

$$\frac{\partial \Psi}{\partial z}(z; b) = \frac{b\rho''(x)}{1 + b\rho''(x)} \Big|_{x=\text{Prox}(z; b)}, \quad \frac{\partial \Psi}{\partial b}(z; b) = \frac{\rho'(x)}{1 + b\rho''(x)} \Big|_{x=\text{Prox}(z; b)}. \quad (59)$$

In particular, for any fixed b , $z \mapsto \Psi(z; b)$ is strictly increasing and Lipschitz continuous, with

$$\frac{b \inf_{x \in \mathbb{R}} \rho''(x)}{1 + b \inf_{x \in \mathbb{R}} \rho''(x)} \leq \frac{\partial \Psi}{\partial z}(z; b) \leq \frac{b\|\rho''\|_\infty}{1 + b\|\rho''\|_\infty}. \quad (60)$$

Proof Using again the stationarity condition (57) that holds for $x = \text{Prox}(z; b)$, we have

$$\text{Prox}(z; b) + b\rho'(\text{Prox}(z; b)) - z = 0, \quad (61)$$

which is our first claim. The other claims immediately follow by calculus. \square

Finally, we prove that Eq. (23) that defines b_t as a function of τ_t always has at least one solution.

Lemma 6.5 For $\tau > 0$ fixed, let $G : \mathbb{R}_{>0} \rightarrow \mathbb{R}$ be defined by

$$G(b) \equiv \mathbb{E} \left\{ \Psi'(W + \tau Z; b) \right\}. \quad (62)$$

Then for any $a \in (0, 1)$, the set of solutions

$$\mathcal{S}_a \equiv \{b \in \mathbb{R}_{>0} : G(b) = a\}, \quad (63)$$

is closed and non-empty.

Proof It follows immediately from the continuity properties of Ψ that $b \mapsto G(b)$ is continuous. The claim follows by proving that $\lim_{b \rightarrow 0} G(b) = 0$ and $\lim_{b \rightarrow \infty} G(b) = 1$.

By Proposition 6.4 Eq. (60) $0 \leq \Psi'(z; b) \leq 1$. The limit $b \rightarrow 0$ follows from dominated convergence since, by the upper bound in (60) $\lim_{b \rightarrow 0} \Psi'(z; b) = 0$ for each z .

In order to obtain the limit as $b \rightarrow \infty$, note that by Stein Lemma:

$$G(b) = \frac{1}{\tau} \mathbb{E} \left\{ Z \Psi(W + \tau Z; b) \right\}. \quad (64)$$

Since $0 \leq \Psi'(z, b) \leq 1$, the integrand is bounded in modulus by an integrable quantity. We can therefore use again dominated convergence. Now $\lim_{b \rightarrow \infty} \text{Prox}(z; b) = \arg \min_{x \in \mathbb{R}} \rho(x) \equiv c_0$ and hence $\lim_{b \rightarrow \infty} \Psi(z; b) = z - c_0$. By dominated convergence we obtain

$$\lim_{b \rightarrow \infty} G(b) = \frac{1}{\tau} \mathbb{E} \left\{ Z (W + \tau Z - c_0) \right\} = 1. \quad (65)$$

□

Appendix 2: Proof of correctness of state evolution (Theorem 3.4)

We will show correctness of state evolution for the AMP algorithm using analytically defined b_t . Namely, we suppose that with b_t defined recursively as the smallest positive solution of the second equation in this system:

$$\tau_{t+1}^2 = \delta \mathbb{E} \left\{ \Psi(W + \tau_t Z; b_t)^2 \right\}, \quad (66)$$

$$\frac{1}{\delta} = \mathbb{E} \left\{ \Psi'(W + \tau_t Z; b_t) \right\}. \quad (67)$$

For analysis purposes, we consider a recursion equivalent to the AMP recursion, in which the data are recentered and the recursion is recast around recentered variables. We change the initial condition of the AMP iteration by letting $\hat{\theta}^{\text{cen},0} = \hat{\theta}^0 - \theta_0$, and change data by letting $Y^{\text{cen}} = Y - \mathbf{X}\theta_0 \equiv W$. Applying the AMP recursion in these new coordinates gives the new trajectory $\hat{\theta}^{\text{cen},t} = \hat{\theta}^t - \theta_0$ for all t , and $R^{\text{cen},t} = R^t$ for all t .

The new trajectory follows the recursion

$$R^{\text{cen},t} = W - \mathbf{X}\hat{\theta}^{\text{cen},t} + \Psi(R^{\text{cen},t-1}; b_{t-1}), \quad (68)$$

$$\hat{\theta}^{\text{cen},t+1} = \hat{\theta}^{\text{cen},t} + \delta \mathbf{X}^\top \Psi(R^{\text{cen},t}; b_t), \quad (69)$$

In this form, the recursion can be reduced to a recursion studied [4], for which state evolution has been proven correct. The reduction is to introduce a recursion generating

iterates $\{\vartheta^t, S^t\}$ that approximates closely the iterates $\{\widehat{\theta}^{\text{cen},t}, R^{\text{cen},t}\}$ defined by (68), (69). The new sequence is defined by letting $\vartheta^0 = \widehat{\theta}^0 - \theta_0$ and, for all $t \geq 0$

$$S^t = -\mathbf{X}\vartheta^t + \Psi(W + S^{t-1}; b_{t-1}), \quad (70)$$

$$\vartheta^{t+1} = \delta \mathbf{X}^T \Psi(W + S^t; b_t) + q_t \vartheta^t, \quad (71)$$

where

$$q_t = \delta \left\{ \frac{1}{n} \sum_{i=1}^n \Psi'(W_i + S_i^t; b_t) \right\}. \quad (72)$$

The only difference between this recursion and the previous one cf. Eqs. (68) and (69), lies in the new coefficient q_t , which was identically equal to 1 in the previous recursion. The benefit of this specific recursion is that we already know that state evolution is correct.

Lemma 6.6 *Under the assumptions of Theorem 3.4, we have, for any fixed $t \geq 0$,*

$$\lim_{n \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p \xi(\vartheta_i^t) =_{\text{a.s.}} \mathbb{E}\{\xi(\sqrt{\delta} \tau_t Z)\} \quad (73)$$

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \xi_2(S_i^t, W_i) =_{\text{a.s.}} \mathbb{E}\{\xi_2(\tau_t Z, W)\}. \quad (74)$$

Proof This is an immediate application of Theorem 2 in [4]. \square

Theorem 3.4 now follows from the equivalence of the last two recursions.

Lemma 6.7 *Under the assumptions of Theorem 3.4, we have, for any fixed $t \geq 0$,*

$$\lim_{n \rightarrow \infty} \frac{1}{p} \|\widehat{\theta}^{\text{cen},t} - \vartheta^t\|_2^2 =_{\text{a.s.}} 0, \quad \lim_{n \rightarrow \infty} \frac{1}{n} \|R^{\text{cen},t} - S^t - W\|_2^2 =_{\text{a.s.}} 0. \quad (75)$$

Appendix 2.1: Proof of Lemma 6.7 (Equivalence of recursions)

Throughout this proof, we will drop the superscript ‘cen’ from $R^{\text{cen},t}$ and $\widehat{\theta}^{\text{cen},t}$. Define $S_+^t \equiv W + S^t$, whence

$$S_+^t = W - \mathbf{X}\vartheta^t + \Psi(S_+^{t-1}; b_{t-1}), \quad (76)$$

$$\vartheta^{t+1} = \delta \mathbf{X}^T \Psi(S_+^t; b_t) + q_t \vartheta^t, \quad (77)$$

Comparing the first of these equations with Eq. (68), and using triangular inequality, we get

$$\|R^t - S_+^t\|_2 \leq \|\mathbf{X}\|_2 \|\widehat{\theta}^t - \vartheta^t\|_2 + \|\Psi(R^{t-1}; b_{t-1}) - \Psi(S_+^{t-1}; b_{t-1})\|_2 \quad (78)$$

$$\leq \|\mathbf{X}\|_2 \|\widehat{\theta}^t - \vartheta^t\|_2 + \|R^{t-1} - S_+^{t-1}\|_2, \quad (79)$$

where the last inequality follows since $\Psi(\cdot; b) : \mathbb{R} \rightarrow \mathbb{R}$ is Lipschitz continuous with Lipschitz constant at most 1, cf Proposition 6.4.

Comparing analogously Eqs. (69) and (77), we obtain

$$\|\widehat{\theta}^{t+1} - \vartheta^{t+1}\|_2 \leq \delta \|\mathbf{X}\|_2 \|\Psi(R^t; b_t) - \Psi(S_+^t; b_t)\|_2 + \|\widehat{\theta}^t - \vartheta^t\|_2 + |q_t - 1| \|\vartheta^t\|_2 \quad (80)$$

$$\leq \delta \|\mathbf{X}\|_2 \|R^t - S_+^t\|_2 + \|\widehat{\theta}^t - \vartheta^t\|_2 + |q_t - 1| \|\vartheta^t\|_2. \quad (81)$$

Iterating the upper bounds (79) and (81), and using the fact that $\vartheta^0 = \widehat{\theta}^0$, we conclude that there exists a constant $A = A(\delta) < \infty$ such that

$$\|\widehat{\theta}^t - \vartheta^t\|_2 \leq (A \|\mathbf{X}\|_2)^{2t} \sum_{\ell=0}^{t-1} |q_\ell - 1| \|\vartheta^\ell\|_2 \quad (82)$$

By Lemma 6.6, we have, almost surely

$$\lim_{n \rightarrow \infty} \frac{1}{\sqrt{p}} \|\vartheta^\ell\|_2 = \tau_\ell < \infty, \quad (83)$$

and

$$\lim_{n \rightarrow \infty} q_t = \delta \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \Psi'(W_i + S_i^t; b_t) \quad (84)$$

$$= \delta \mathbb{E}\{\Psi'(W + \tau_t Z; b_t)\} = 1, \quad (85)$$

where the second identity follows from Lemma 6.6 and, in the third, we used the definition of b_t . (Note that we are applying here Lemma 6.6 to $\xi(\cdot) = \Psi'(\cdot; b_t)$ which is bounded and non-negative but not necessarily continuous. However, since $W + \tau_t Z$ has a density for every $\tau_t > 0$, the limit holds by a standard weak convergence argument, approximating ξ by simple functions. Namely, we construct a sequence of simple functions ξ_ℓ such that $\xi_\ell(t) \leq \xi(t) \leq \xi_\ell(t) + (1/\ell)$ for all t , and apply Lemma 6.6—which implies weak convergence of the empirical distribution of $\{W_i + S_i^t\}$ to $\xi_\ell(\cdot)$.)

Finally, it is a standard result in random matrix theory [1] that $\lim_{n \rightarrow \infty} \|\mathbf{X}\|_2 = C(\delta) < \infty$. Hence, by taking the limit of Eq. (82) we get, almost surely,

$$\lim_{n \rightarrow \infty} \frac{1}{\sqrt{p}} \|\widehat{\theta}^t - \vartheta^t\|_2 = 0. \quad (86)$$

The norm $\|R^t - S_+^t\|_2$ is then controlled using Eq. (79).

Appendix 3: Proof that AMP converges to the M-estimator (Theorem 4.1)

Notice first of all that, by construction, $\tau_t^2 = \tau_*^2$, $b_t = b_*$ for all t .

Given δ, ρ as in the statement of the theorem and τ_*, b_* a solution of the fixed point Eqs. (28) and (29), we define the doubly infinite matrix $\Gamma = (\Gamma_{t,s})_{t,s \geq 0}$ by letting, recursively for $t, s \geq 0$

$$\Gamma_{t+1,s+1} = \delta \mathbb{E}\{\Psi(W + Z_t; b_*)\Psi(W + Z_s; b_*)\}, \quad (87)$$

where the expectation is with respect to (Z_t, Z_s) jointly Gaussian, with zero means and covariance $\mathbb{E}\{Z_t^2\} = \Gamma_{t,t}$, $\mathbb{E}\{Z_s^2\} = \Gamma_{s,s}$, $\mathbb{E}\{Z_t Z_s\} = \Gamma_{t,s}$, independent of $W \sim F_W$. This is supplemented with the boundary condition $\Gamma_{0,0} = \tau_*^2$ and $\Gamma_{0,t} = \Gamma_{t,0} = 0$ for $t > 0$.

Notice that, in particular, $\Gamma_{s,t} = \Gamma_{t,s}$ for all $s, t \geq 0$ and $\Gamma_{t,t} = \tau_*^2$ for all t .

The significance of these quantities is clarified by the following result.

Lemma 6.8 *Under the hypotheses of Theorem 3.4, further assume that τ_*^2 and Γ are defined as above. Then, for any $t, s \geq 0$,*

$$\lim_{n \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p \xi_2(\widehat{\theta}_i^t - \theta_{0,i}, \widehat{\theta}_i^s - \theta_{0,i}) =_{a.s.} \mathbb{E} \xi_2(\sqrt{\delta} Z_t, \sqrt{\delta} Z_s), \quad (88)$$

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \xi_2(R_i^t - W_i, R_i^s - W_i) =_{a.s.} \mathbb{E} \xi_2(Z_t, Z_s), \quad (89)$$

where the expectation is with respect to (Z_t, Z_s) jointly Gaussian, with zero means and covariance $\mathbb{E}\{Z_t^2\} = \Gamma_{t,t}$, $\mathbb{E}\{Z_s^2\} = \Gamma_{s,s}$, $\mathbb{E}\{Z_t Z_s\} = \Gamma_{t,s}$, independent of $W \sim F_W$.

The proof is deferred to Sect. 1.

As a special case of the latter result, we have

$$\lim_{n \rightarrow \infty} \frac{1}{p} \|\widehat{\theta}^t - \widehat{\theta}^s\|_2^2 =_{a.s.} 2\delta (\tau_*^2 - \Gamma_{t,s}), \quad (90)$$

$$\lim_{n \rightarrow \infty} \frac{1}{n} \|R^t - R^s\|_2^2 =_{a.s.} 2 (\tau_*^2 - \Gamma_{t,s}). \quad (91)$$

The following lemma provides information about the asymptotic behavior of $\Gamma_{t,s}$. Its proof is deferred to Sect. 1.

Lemma 6.9 *Let τ_*, Γ be defined as above for $\delta > 1$. Then*

$$\lim_{t \rightarrow \infty} \Gamma_{t,t+1} = \tau_*^2 \quad (92)$$

Applying this result to Eqs. (90) and (91) we get, for any fixed $h \in \mathbb{N}$,

$$\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{p} \|\hat{\theta}^{t+h} - \hat{\theta}^t\|_2^2 =_{a.s.} 0, \quad (93)$$

$$\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{p} \|R^{t+h} - R^t\|_2^2 =_{a.s.} 0. \quad (94)$$

(The case $h > 1$ follows from $h = 1$ by the triangle inequality.)

We are now ready to prove Theorem 4.1. Recall that $\mathcal{L}(\theta) = \mathcal{L}(\theta; Y, \mathbf{X})$ denotes the loss function defined in Eq. (2), and that its gradient and Hessian are given by

$$\nabla_{\theta} \mathcal{L}(\theta) = - \sum_{i=1}^n \rho'(Y_i - \langle X_i, \theta \rangle) X_i, \quad (95)$$

$$\nabla_{\theta}^2 \mathcal{L}(\theta) = \sum_{i=1}^n \rho''(Y_i - \langle X_i, \theta \rangle) X_i X_i^{\top}. \quad (96)$$

In particular, letting $\sigma_{\min}(\mathbf{X})$ denote the minimum non-zero singular value of \mathbf{X} , we have

$$\lambda_{\min}(\nabla_{\theta}^2 \mathcal{L}(\theta)) \geq \inf_{x \in \mathbb{R}} \rho''(x) \cdot \sigma_{\min}(\mathbf{X})^2. \quad (97)$$

Using the hypothesis of strong convexity and standard concentration of measure for the singular values of Wishart matrices [36], there exists constants $c_0, c_1, n_0 > 0$ for $\delta > 1$ such that for any $n \geq n_0$,

$$\mathbb{P}(\nabla_{\theta}^2 \mathcal{L}(\theta) \succeq c_0 \mathbf{I} \quad \forall \theta \in \mathbb{R}^p) \geq 1 - e^{-c_1 n}. \quad (98)$$

As a consequence, with probability at least $1 - e^{-c_1 n}$, we have

$$\mathcal{L}(\hat{\theta}^t) \geq \mathcal{L}(\hat{\theta}) \geq \mathcal{L}(\hat{\theta}^t) + \langle \nabla_{\theta} \mathcal{L}(\hat{\theta}^t), \hat{\theta} - \hat{\theta}^t \rangle + \frac{1}{2} c_0 \|\hat{\theta} - \hat{\theta}^t\|_2^2. \quad (99)$$

Hence using Cauchy–Schwartz

$$\|\hat{\theta} - \hat{\theta}^t\|_2 \leq \frac{2}{c_0} \|\nabla_{\theta} \mathcal{L}(\hat{\theta}^t)\|_2. \quad (100)$$

The last step of the proof consists in showing that, almost surely

$$\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{p} \|\nabla_{\theta} \mathcal{L}(\hat{\theta}^t)\|_2^2 = 0. \quad (101)$$

In order to prove this claim, reconsider Eq. (13), for time $t + 1$, with $b_t = b_*$. Using the fact that $\Psi(z; b_*) = z - \text{Prox}(z; b_*)$, this can be rewritten as

$$\text{Prox}(R^t; b_*) = Y - \mathbf{X}\hat{\theta}^{t+1} + R^t - R^{t+1}. \quad (102)$$

By Eq. (15), and recalling that $\Psi(z; b) = b \rho'(\text{Prox}(z; b))$, we have

$$\frac{1}{b_*\delta} (\hat{\theta}^{t+1} - \hat{\theta}^t) = \mathbf{X}^T \rho'(\text{Prox}(R^t; b_*)) \quad (103)$$

$$= \mathbf{X}^T \rho'(Y - \mathbf{X}\hat{\theta}^{t+1} + R^t - R^{t+1}), \quad (104)$$

where the last identity followed by Eq. (102). Using the triangle inequality and noting that, by the smoothness assumption $C \equiv \sup_{z \in \mathbb{R}} \rho''(z) < \infty$, we get

$$\|\mathbf{X}^T \rho'(Y - \mathbf{X}\hat{\theta}^{t+1})\|_2 \leq \frac{1}{b_*\delta} \|\hat{\theta}^{t+1} - \hat{\theta}^t\|_2 + C \|\mathbf{X}\|_2 \|R^t - R^{t+1}\|_2. \quad (105)$$

Hence, using Eqs. (93) and (93), and recalling that $\lim_{n \rightarrow \infty} \|\mathbf{X}\|_2 < \infty$ almost surely [1], we get

$$\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{p} \|\mathbf{X}^T \rho'(Y - \mathbf{X}\hat{\theta}^{t+1})\|_2^2 = 0. \quad (106)$$

This is equivalent to the claim (101) since $\nabla_{\theta} \mathcal{L}(\theta) = -\mathbf{X} \rho'(Y - \mathbf{X}\theta)$.

Appendix 3.1: Proof of Lemma 6.8

First of all note that, due to Lemma 6.7, it is sufficient to prove that

$$\lim_{n \rightarrow \infty} \frac{1}{p} \sum_{i=1}^p \xi_2(\vartheta_i^t, \vartheta_i^s) =_{a.s.} \mathbb{E} \xi_2(\sqrt{\delta} Z_t, \sqrt{\delta} Z_s), \quad (107)$$

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \xi_2(s_i^t, s_i^s) =_{a.s.} \mathbb{E} \xi_2(Z_t, Z_s). \quad (108)$$

Note that a similar statement is proved in [5, Theorem 4.2] for characterizing the Lasso estimator. While the same argument can be followed here, we outline an alternative argument that is based on a reduction to the setting of [19].

We fix an even number $q \in \mathbb{N}$, and will prove the claim for all $t, s \leq T \equiv (q/2) - 1$. Let $N \equiv n + p$. For $t \in \{0, \dots, T\}$, we introduce a vector $z^t \in (\mathbb{R}^q)^N$, which we think of as a vector with entries in \mathbb{R}^q : $z^t = (\mathbf{z}_1^t, \dots, \mathbf{z}_N^t)$ $\mathbf{z}_i^t \in \mathbb{R}^q$. Its entries are

defined as follows:

$$\mathbf{z}_i^t = (S_i^0, 0, S_i^1, 0, S_i^2, 0, \dots, S_i^t, 0, 0, 0, \dots, 0) \quad \text{if } 1 \leq i \leq n, \quad (109)$$

$$\begin{aligned} \mathbf{z}_i^{t+1} &= (0, \vartheta_j^1, 0, \vartheta_j^2, 0, \vartheta_j^3, \dots, 0, \vartheta_i^{t+1}, 0, 0, 0, \dots, 0) \\ &\quad \text{if } n+1 \leq i = j+n \leq n+p. \end{aligned} \quad (110)$$

Further, we let $A \in \mathbb{R}^{N \times N}$ be a symmetric matrix with $A_{ii} = 0$, $A_{ij} = \sqrt{n/N} X_{i,j-n}$ for $1 \leq i \leq n$ and $n+1 \leq j \leq n+p$, and all the other entries A_{ij} $i < j$ i.i.d. $\mathcal{N}(0, 1/N)$. It is then easy to see that the iteration in Eqs. (70) and (71) is equivalent to the following

$$\mathbf{z}^{t+1} = A f(\mathbf{z}^t; t) - \mathbf{B}_t f(\mathbf{z}^{t-1}; t-1). \quad (111)$$

Here, for each t , $f(\cdot; t) : (\mathbb{R}^q)^N \rightarrow (\mathbb{R}^q)^N$ is separable in the following sense

$$f(\mathbf{z}; t) = (f_1(\mathbf{z}_1; t); f_2(\mathbf{z}_2; t); \dots; f_N(\mathbf{z}_N; t)), \quad (112)$$

with $f_i(\cdot; t) : \mathbb{R}^q \rightarrow \mathbb{R}^q$. These are defined as follows (letting $\Psi_{t,i}(x) = \Psi(W_i + x; b_t)$ and $h = \sqrt{(1+\delta)/\delta}$)

$$\begin{aligned} f_i(\mathbf{z}_i; t) &= (0, \delta h \Psi_{0,i}(\mathbf{z}_{i,1}), 0, \delta h \Psi_{1,i}(\mathbf{z}_{i,3}), 0, \dots, 0, \delta h \Psi_{t,i}(\mathbf{z}_{i,2t-1}), 0, 0, 0, \dots, 0) \\ &\quad \text{if } 1 \leq i \leq n, \end{aligned} \quad (113)$$

$$\begin{aligned} f_i(\mathbf{z}_i; t) &= (0, 0, -h \mathbf{z}_{i,2}, 0, -h \mathbf{z}_{i,4}, 0, \dots, -h \mathbf{z}_{i,2t}, 0, 0, 0, 0, \dots, 0) \\ &\quad \text{if } n+1 \leq i \leq n+p. \end{aligned} \quad (114)$$

The matrix multiplication in Eq. (111) operates in the natural way over $(\mathbb{R}^q)^N$, namely we identified A with the Kronecker product $A \otimes \mathbf{I}_{q \times q}$. Explicitly, Eq. (111) reads

$$\mathbf{z}_i^{t+1} = \sum_{j \in [N]} A_{ij} f_j(\mathbf{z}_j^t; t) - \mathbf{B}_t f_i(\mathbf{z}_i^{t-1}; t-1). \quad (115)$$

Finally $\mathbf{B}_t \in \mathbb{R}^{q \times q}$ is given by

$$\mathbf{B}_t = \frac{1}{N} \sum_{i=1}^N \frac{\partial f_i}{\partial \mathbf{z}}(\mathbf{z}_i^t; t). \quad (116)$$

The recursion (111) is characterized in [19, Theorem 1], which establishes—for instance—that, for $\xi : \mathbb{R}^q \rightarrow \mathbb{R}$ pseudo-Lipschitz, we have, almost surely,

$$\lim_{N \rightarrow \infty} \frac{1}{p} \sum_{i=n+1}^{n+p} \xi(\mathbf{z}_i^t) = \mathbb{E}\{\xi(\mathbf{Z}_t)\}. \quad (117)$$

Here \mathbf{Z}_t is a Gaussian random vector whose covariance is fully specified in [19]. The proof of the lemma is finished by comparing the expressions in [19] for the covariance with the ones in the statement of the lemma.

Appendix 3.2: Proof of Lemma 6.9

First of all we introduce the notation $q_t \equiv \Gamma_{t,t+1}/\tau_*^2$. We then have the recursion

$$q_{t+1} = H(q_t), \quad (118)$$

$$H(q) = \frac{\delta}{\tau_*^2} \mathbb{E}_q \{ \Psi(W + \tau_* Z_1; b_*) \Psi(W + \tau_* Z_2; b_*) \}, \quad (119)$$

where expectation \mathbb{E}_q is with respect to the centered Gaussian vector (Z_1, Z_2) with $\mathbb{E}_q\{Z_1^2\} = \mathbb{E}_q\{Z_2^2\} = 1$ and $\mathbb{E}_q\{Z_1 Z_2\} = q$, independent of $W \sim F_W$. We claim that:

- (i) $H(1) = 1$;
- (ii) $H(q)$ is increasing for $q \in [0, 1]$;
- (iii) $H(q)$ is strictly convex for $q \in [0, 1]$.

In order to prove (i), note that, for $q = 1$, $Z_1 = Z_2 \equiv Z \sim N(0, 1)$ and hence

$$H(1) = \frac{\delta}{\tau_*^2} \mathbb{E}_q \{ \Psi(W + \tau_* Z; b_*)^2 \}, \quad (120)$$

which is equal to 1 since b_*, τ_* satisfy Eq. (28).

In order to prove (ii), (iii), define

$$h_W(z) \equiv \Psi(W + \tau_* z; b_*), \quad (121)$$

$$\mathcal{H}(q) \equiv \mathbb{E}_q \{ h_W(Z_1) h_W(Z_2) | W \}, \quad (122)$$

We will prove that \mathcal{H} is strictly increasing and convex for any W , whence claims (ii) and (iii) follow by linearity. The argument is the same as in [5, Lemma C.1] Let $\{X_t\}_{t \geq 0}$ be the stationary Ornstein–Uhlenbeck process with covariance $\mathbb{E}(X_0 X_t) = e^{-t}$, and denote by \mathbb{E} expectation with respect to X . Then

$$\mathcal{H}(q) = \mathbb{E} \{ h_W(X_0) h_W(X_t) \} \Big|_{t=\log(1/q)}, \quad (123)$$

Then we have the spectral representation (for $t = \log(1/q)$)

$$\mathcal{H}(q) = \sum_{\ell=0}^{\infty} c_{\ell}^2 e^{-\ell t} = \sum_{\ell=0}^{\infty} c_{\ell}^2 q^{\ell}, \quad (124)$$

whence the claim follows since $c_{\ell} \neq 0$ for some $\ell \geq 2$ as long as $h_W(x)$ is non-linear.

Because of the remarks (i)–(iii) just proven, it follows that $\lim_{t \rightarrow \infty} q_t = 1$ (and hence $\lim_{t \rightarrow \infty} \Gamma_{t,t+1} = \tau_*^2$) if and only if $H'(1) \leq 1$. A simple calculation yields

$$H'(1) = \delta \mathbb{E} \{ \Psi'(W + \tau_* Z; b_*)^2 \}, \quad (125)$$

where $Z \sim N(0, 1)$. Recalling that $\Psi'(z; b) \in (0, 1)$, we have $(\Psi')^2 \leq \Psi'$ and so

$$H'(1) \leq \delta \mathbb{E} \{ \Psi'(W + \tau_* Z; b_*) \} = 1, \quad (126)$$

where the last identity follows because (τ_*, b_*) solve Eq. (29). This finishes the proof.

References

1. Anderson, G.W., Guionnet, A., Zeitouni, O.: An introduction to random matrices. Cambridge University Press, (2009)
2. Bickel, Peter J.: One-step Huber estimates in the linear model. *Journal of the American Statistical Association* **70**(350), 428–434 (1975)
3. Bayati, Mohsen, Lelarge, Marc, Montanari, Andrea: Universality in polytope phase transitions and message passing algorithms. *Annals of Applied Probability* **25**, 753–822 (2015)
4. Bayati, M., Montanari, A.: The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Trans. on Inform. Theory* **57**, 764–785 (2011)
5. Bayati, M., Montanari, A.: The LASSO risk for gaussian matrices. *IEEE Trans. on Inform. Theory* **58**, 1997–2017 (2012)
6. Bolthausen, Erwin: An iterative construction of solutions of the tap equations for the sherrington-kirkpatrick model. *Communications in Mathematical Physics* **325**(1), 333–366 (2014)
7. Chen, S.S., Donoho, D.L.: Examples of basis pursuit. *Proceedings of Wavelet Applications in Signal and Image Processing III* (San Diego, CA), 1995
8. Chang, C.-I.: *Hyperspectral Imaging: Techniques for Spectral Detection and Classification*. Springer, (2003)
9. Donoho, David, Johnstone, Iain, Maleki, Arian, Montanari, Andrea: Compressed sensing over ℓ_p -balls: Minimax mean square error. *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*, IEEE, 129–133 (2011)
10. Deshpande, Yash, Montanari, Andrea: Finding Hidden Cliques of Size \sqrt{N}/e in Nearly Linear Time. *Foundations of Computational Mathematics*, 1–60 (2013)
11. Donoho, D.L., Maleki, A., Montanari, A.: Message Passing Algorithms for Compressed Sensing. *Proceedings of the National Academy of Sciences* **106**, 18914–18919 (2009)
12. Donoho, D.L., Maleki, A., Montanari, A.: Message Passing Algorithms for Compressed Sensing: I. Motivation and Construction. *Proceedings of IEEE Inform. Theory Workshop (Cairo)*, (2010)
13. Donoho, D.L., Maleki, A., Montanari, A.: The Noise Sensitivity Phase Transition in Compressed Sensing. *IEEE Trans. on Inform. Theory* **57**, 6920–6941 (2011)
14. Donoho, David, Tanner, Jared: Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **367**(1906), 4273–4293 (2009)
15. El Karoui, Nouredine El, Bean, Derek, Bickel, Peter J., Lim, Chingwayand, Yu, Bin: On robust regression with high-dimensional predictors. *Proceedings of the National Academy of Sciences* **110**(36), 14557–14562 (2013)
16. Huber, P.J., Ronchetti, E.: *Robust statistics*, 2nd edn. J. Wiley and Sons, (2009)
17. Huber, P.J.: Robust estimation of a location parameter. *The Annals of Mathematical Statistics* **35**(1), 73–101 (1964)
18. Huber, Peter J.: Robust regression: asymptotics, conjectures and monte carlo. *The Annals of Statistics* **1**(5), 799–821 (1973)
19. Javanmard, A., Montanari, A.: State Evolution for General Approximate Message Passing Algorithms, with Applications to Spatial Coupling. [arXiv:1211.5164v1](https://arxiv.org/abs/1211.5164v1) (2012)
20. Karoui, Nouredine El: Asymptotic behavior of unregularized and ridge-regularized high-dimensional robust regression estimators: rigorous results, [arXiv:1311.2445](https://arxiv.org/abs/1311.2445) (2013)
21. Korada, S., Montanari, A.: Applications of Lindeberg Principle in Communications and Statistical Learning. *IEEE Trans. on Inform. Theory* **57**, 2440–2450 (2011)
22. Krzakala, Florent, Mézard, Marc, Zdeborová, Lenka: Phase diagram and approximate message passing for blind calibration and dictionary learning. *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*, IEEE, pp. 659–663 (2013)
23. Lustig, M., Donoho, D.L., Santos, J.M., Pauly, J.M.: Compressed sensing mri. *IEEE Signal Processing Magazine* **25**, 72–82 (2008)
24. Mézard, M., Montanari, A.: *Information. Physics and Computation*, Oxford (2009)

25. Montanari, A.: Graphical Models Concepts in Compressed Sensing. In: Eldar, Y.C., Kutyniok, G. (eds.) *Compressed Sensing: Theory and Applications*. Cambridge University Press, (2012)
26. Mézard, M., Parisi, G., Virasoro, M.A.: *Spin glass theory and beyond*. World Scientific, (1987)
27. Rangan, S.: Generalized Approximate Message Passing for Estimation with Random Linear Mixing. *IEEE Intl. Symp. on Inform. Theory (St. Petersburg)*, August 2011
28. Richards, M.A.: *Fundamentals of Radar Signal Processing*. McGraw-Hill, (2005)
29. Scales, J.A.: *Theory of Seismic Imaging*. Samizdat Press, (1997)
30. Schniter, P.: Turbo Reconstruction of Structured Sparse Signals. *Proceedings of the Conference on Information Sciences and Systems (Princeton)*, (2010)
31. Potter, L.C., Som, S., Schniter, P.: On Approximate Message Passing for Reconstruction of Non-Uniformly Sparse Signals. *Proceedings of the National Aerospace and Electronics Conference (Dayton, OH)*, (2010)
32. Shcherbina, Mariya, Tirozzi, Brunello: Rigorous solution of the gardner problem. *Communications in mathematical physics* **234**(3), 383–422 (2003)
33. Talagrand, M.: *Mean field models for spin glasses: Volume i*. Springer-Verlag, Berlin (2010)
34. Tibshirani, R.: Regression shrinkage and selection with the Lasso. *J. Royal. Statist. Soc B* **58**, 267–288 (1996)
35. Tramel, Eric W., Kumar, Santhosh, Giurgiu, Andrei, Montanari, Andrea: Statistical estimation: From denoising to sparse regression and hidden cliques. In: Krzakala, F., Ricci-Tersenghi, F., Zdeborova, L., Zecchina, R., Tramel, E.W., Cugliandolo, L.F. (eds.) *Statistical Physics, Optimization, Inference, and Message-Passing Algorithms*. Oxford University Press, (2015)
36. Vershynin, R.: Introduction to the non-asymptotic analysis of random matrices. In: Eldar, Y.C., Kutyniok, G. (eds.) *Compressed Sensing: Theory and Applications*, pp. 210–268. Cambridge University Press, (2012)