

MAXIMIN ANALYSIS OF MESSAGE PASSING ALGORITHMS FOR RECOVERING BLOCK SPARSE SIGNALS

ARMEEN TAEB*, ARIAN MALEKI†, CHRISTOPH STUDER‡, AND
RICHARD G. BARANIUK‡

Abstract. We consider the problem of recovering a block (or group) sparse signal from an underdetermined set of random linear measurements, which appear in compressed sensing applications such as radar and imaging. Recent results of Donoho, Johnstone, and Montanari have shown that approximate message passing (AMP) in combination with Stein’s shrinkage outperforms group LASSO for large block sizes. In this paper, we prove that, for a fixed block size and in the strong undersampling regime (i.e., having very few measurements compared to the ambient dimension), AMP cannot improve upon group LASSO, thereby complementing the results of Donoho *et al.*

Key words. Group sparsity; group LASSO; approximate message passing; phase transition

1. Introduction. The field of compressed sensing (CS) aims to recover a sparse signal from an undetermined systems of linear equations. Concretely, CS can be modeled as $\mathbf{y} = \mathbf{A}\mathbf{x}$, where \mathbf{y} is the n -dimensional measurement vector, \mathbf{A} is the (typically random) $n \times N$ measurement matrix, and \mathbf{x} is an N dimensional vector with at most k nonzero entries (often referred to as k -sparse vector). Our goal is to recover \mathbf{x} from this undetermined system.

A large class of signals of interest exhibit additional structure known as *block (or group) sparsity*, where the non-zero coefficients of the signal occur in clusters of size B [1, 2]. Such block-sparse signals naturally appear in genomics, radar, and communication applications. There has been a considerable amount of research on theory and algorithms for recovering such signals [1–11]. Perhaps the most popular recovery algorithm is *group LASSO* [1]—corresponding theoretical work has shown that under which conditions this algorithm recovers the exact group sparse solution [2–11]. While these results enable a qualitative characterization of the recovery performance of group LASSO, they do not provide an accurate performance analysis.

In order to arrive at a more accurate performance analysis of group LASSO, several authors have considered the asymptotic setting where $n, N \rightarrow \infty$, while their ratio $\delta = n/N$ is held constant [12–16]. Under this setting, the references [14–16] have shown that there exists a threshold on the normalized sparsity $\rho = k/n$, below which group LASSO recovers the correct signal vector \mathbf{x} with probability 1 and fails otherwise. Such a phase-transition (PT) analysis has led to the conclusion that group LASSO is sub-optimal, since there is a large gap between the PT of the information theoretic limit¹ and that of group LASSO (see Fig. 1.1).

There has been recent effort in using *approximate message passing* (AMP) to improve upon the performance of group LASSO. Schniter, for example, has experimentally shown in [18] that AMP combined with expectation maximization can outperform group LASSO. Kamilov *et al.* have taken the first step toward a theoretical understanding of such algorithms [19]. More recently, Donoho, Johnstone, and Monta-

*Dept. of Electrical, Computer, and Energy Engineering, University of Colorado at Boulder.

†Dept. of Statistics, Columbia University.

‡Dept. of Electrical and Computer Engineering, Rice University.

¹The information theoretic limit was only derived for regular sparse signals (with block size 1) in [17]. An extension of these results to block sparse signals with larger blocks is straightforward.




FIG. 1.1. Phase transition (PT) of group LASSO for various block sizes B . Evidently, there is a disparity between the phase transition of group LASSO and the information theoretic limit.

nari [16] have shown that AMP is able to outperform group LASSO if it employs Stein's shrinkage estimator. In fact, they demonstrate that for very large block sizes, i.e., for $B \rightarrow \infty$, the performance of this AMP variant is close to the information theoretic limit. However, in many applications, such as radar and communication systems, the group sizes are typically small and for fixed block sizes, Stein's estimator does not necessarily improve the performance of AMP. Consequently, the fundamental question remains whether AMP can outperform group LASSO for any block size. In this paper, we address this question in the high undersampling regime where $\delta \rightarrow 0$.

In particular, we show that, for $\delta \rightarrow 0$, there is no nonlinear shrinkage function that allows AMP to outperform group LASSO. We emphasize that this result does not contradict that in [16] as they considered a different limiting regime, i.e., where the block size B approaches infinity. A combination of these two results enables us to conclude that, for strong undersampling (i.e., small values of δ), AMP with Stein's estimator requires large block sizes in order to outperform group LASSO.

2. Background.

2.1. Notation. Lowercase boldface letters, such as \mathbf{v} , represent vectors and uppercase letters, such as V , represent matrices; lowercase letters, such as v represent scalars. We analyze the recovery of a block (or group) sparse signal $\mathbf{x} \in \mathbb{R}^N$ with at most k nonzero entries from the undersampled linear measurements $\mathbf{y} = \mathbf{A}\mathbf{x}$, where $\mathbf{A} \in \mathbb{R}^{n \times N}$ is i.i.d. zero-mean Gaussian with unit variance. We furthermore consider the asymptotic setting where $\delta = n/N$, $\rho = k/n$, and $N, n, k \rightarrow \infty$.

The notational conventions for block sparse signals are as follows. We assume that all the blocks have the same size, denoted by B . Extensions to signals with varying block sizes is straightforward. The signal \mathbf{x} is partitioned into M blocks where clearly, $N = MB$. In the remainder of the paper, we will denote \mathbf{x}_B as a particular block. Suppose that the elements of \mathbf{x}_B are drawn from a given distribution $F(\mathbf{x}_B) = (1 - \epsilon)\delta_0(\|\mathbf{x}_B\|_2) + \epsilon G(\mathbf{x}_B)$, where $\epsilon = \rho\delta$, and δ_0 is the Dirac delta function; G is a probability distribution that is typically unknown in practice.

The block soft-thresholding function used in this paper is defined as follows [16]:

$$\eta^{\text{soft}}(\mathbf{y}; \tau) = \frac{\mathbf{y}_B}{\|\mathbf{y}_B\|_2} (\|\mathbf{y}_B\|_2 - \tau)_+. \quad (2.1)$$

Here, $(z)_+ = \max(z, 0)$ and $\eta^{\text{soft}}(\mathbf{y}_B; \tau)$ sets its argument \mathbf{y}_B to zero if $\|\mathbf{y}_B\|_2 \leq \tau$, and shrinks the vector \mathbf{y}_B towards the origin by τ , otherwise.

2.2. Group LASSO and approximate message passing (AMP). A decade of research in sparse recovery has produced a plethora of algorithms for recovering block sparse signals from random linear measurements. Two popular algorithms are group LASSO and AMP. Group LASSO searches for a vector \mathbf{x} that minimizes the cost function, $\mathbf{x}^L \triangleq \arg \min_{\mathbf{x}} \{\sum_{B=1}^M \|\mathbf{x}_B\|_2 : \mathbf{y} = A\mathbf{x}\}$. AMP, on the other hand, is an iterative algorithm to recover the solution vector \mathbf{x} . Concretely, by initializing $\mathbf{x}^0 = \mathbf{0}$ and $\mathbf{z}^0 = \mathbf{0}$, AMP iteratively performs the following steps:

$$\mathbf{x}^{t+1} = \eta(\mathbf{x}^t + A^* \mathbf{z}^t) \quad \text{and} \quad \mathbf{z}^t = \mathbf{y} - A\mathbf{x}^t + \mathbf{c}^t. \quad (2.2)$$

Here, \mathbf{c}^t is a correction term that depends on the previous iterations, which significantly improves the convergence of AMP; \mathbf{x}^t is the (block) sparse estimate at iteration t , and η is a nonlinear function that imposes (block) sparsity. In particular, if $\eta(\cdot)$ is the block soft-thresholding function as in (2.1), then AMP is equivalent to group LASSO in the asymptotic setting [15, 16] (see [13] for the details).

One of the most appealing features of the AMP is that its operation can be viewed as a denoising problem at each iteration. That is, when $N \rightarrow \infty$, $\mathbf{x}^t + A^* \mathbf{z}^t$ can be modeled as the sparse signal \mathbf{x} plus zero-mean Gaussian noise, which is independent of the signal. This feature enables one to analytically predict the performance of AMP through a framework called *state evolution* (SE). Concretely, if the mean-square error (MSE) of AMP at iteration t is denoted by MSE^t , then

$$\text{MSE}^{t+1} = \frac{1}{\delta B} \mathbb{E} \left\{ \left\| \mathbf{x}_B - \eta^t(\mathbf{x}_B + \sqrt{\text{MSE}^t} \mathbf{z}_B) \right\|_2^2 \right\}, \quad (2.3)$$

where $\mathbf{z}_B \sim N(0, I_B)$ and the distribution of \mathbf{x}_B is the same as the empirical distribution of the blocks of the original vector \mathbf{x} . The expectation $\mathbb{E}\{\cdot\}$ is taken with respect to the vectors \mathbf{z}_B and \mathbf{x}_B .

2.3. Phase transition. The performance of CS recovery algorithms can be characterized accurately by their *phase transition* (PT) behavior. Specifically, we define a two-dimensional phase space $(\delta, \rho) \in [0, 1]$ that is partitioned into two regions: “success” and “failure”, with these regions separated by the PT curve $(\delta, \rho(\delta))$. For the same value of δ , algorithms with higher PT outperform algorithms with lower PT, i.e., guarantee the exact recovery for more nonzero entries k .

3. Main results. The thresholding function η^t determines the performance of AMP. Indeed, different choices of η^t may lead to fundamentally different performance. It has been shown in [15, 16] that if η^{soft} from (2.1) is used, then the performance of AMP is equivalent to that of group LASSO. Since η^{soft} is not necessarily the optimal thresholding function for group sparse signals, finding the optimal function is of significant practical interest.

In this paper, we characterize the optimal choice of the thresholding function η^t in the strong undersampling regime, i.e., for $\delta \rightarrow 0$. Before we proceed, let us define optimality. Suppose that each block \mathbf{x}_B is drawn independently from the distribution $F(\mathbf{x}_B)$ as defined in Section 2.1. We furthermore assume that η^t is applied to each block separately. As is evident from (2.3), each iteration of AMP is equivalent to a denoising problem, where the noise variance is equal to the MSE of the previous iteration. Therefore for a given initialization point, the effect of the thresholding function η^t on AMP can be characterized by a discrete set $\{\text{MSE}^t\}_{t=1}^{\infty}$. Thus, instead of considering a sequence of iteration-dependent thresholding functions $\{\eta^t\}_{t=1}^{\infty}$, we

can consider a sequence of thresholding functions $\bar{\eta}$ that depend on MSE^t where $\bar{\eta} : \mathbb{R}^B \times \mathbb{R} \rightarrow \mathbb{R}^B$ with the property $\eta^t(\mathbf{y}_B) = \bar{\eta}(\mathbf{y}_B, \text{MSE}^t)$. Since the MSE sequence is dependent on the initialization point of AMP (which can be chosen arbitrarily), we wish to optimize $\bar{\eta}$ with respect to *all possible* initializations. Hence at each iteration, the problem is simplified to finding the optimal $\bar{\eta}$ which is a function of \mathbf{y}_B and any MSE value greater than zero. Now, suppose that the PT of AMP with $\bar{\eta}$ is given by $\rho^{\bar{\eta}}(\delta, G)$. Then, we are interested in thresholding functions that achieve:

$$\rho^*(\delta) \triangleq \sup_{\bar{\eta}} \inf_G \rho^{\bar{\eta}}(\delta, G). \quad (3.1)$$

Such an $\bar{\eta}$ provides the best PT performance for the least favorable distribution—a reasonable assumption, since the distribution G is typically unknown in many practical applications. Our first contribution characterizes the behavior of $\rho^*(\delta)$ for strong undersampling, i.e., for small values of δ .

THEOREM 3.1. *The optimal PT $\rho^*(\delta)$ of AMP follows the behavior*

$$\rho^*(\delta) \sim \frac{B}{2 \log(1/\delta)} \quad \text{as } \delta \rightarrow 0.$$

As shown in Section 4.2, this behavior is determined when G is uniformly distributed on a sphere with infinite radius. We note that if the distribution G is unknown, this theorem does not provide any guidelines on how to choose η^t in (2.2). The next theorem shows that group LASSO follows exactly the same behavior, and hence, block soft thresholding is the optimal choice for η^t in the strong undersampling regime.

THEOREM 3.2. *The PT $\rho^L(\delta)$ of group LASSO follows the behavior*

$$\rho^L(\delta) \sim \frac{B}{2 \log(1/\delta)} \quad \text{as } \delta \rightarrow 0.$$

Combining Theorems 3.1 and 3.2 reveals that for a fixed B and in the strong undersampling regime, i.e., for $\delta \rightarrow 0$, the best achievable PT of AMP coincides with the phase transition of group LASSO. This result has two striking implications for strong undersampling and fixed block sizes: (i) Block soft thresholding is optimal and (ii) there is no thresholding function for AMP that outperforms group LASSO. Consequently, for decreasing δ , AMP equipped with a better thresholding operator (than block soft-thresholding) such as Stein’s shrinkage requires larger block sizes to outperform group LASSO in this regime. It is worth mentioning that these results do not contradict those of [16, Section 3.2], which show that AMP with Stein’s shrinkage outperforms group LASSO for large block sizes. In fact, combining our results with those in [16] provides a better picture of the potential benefits of Stein’s shrinkage within AMP.

4. Proofs of the main results. We next outline² the proofs of Theorems 3.2 and 3.1. Since the proof of Theorem 3.1 requires the result of Theorem 3.2, we begin by proving the latter.

4.1. Proof of Theorem 3.2. We start by deriving an implicit formula for the PT of group LASSO. We further use this formula and Laplace’s method to obtain the behavior of the PT in the strong undersampling regime. Since the performance of

²More details on these proofs can be found at sites.google.com/site/armeentaeb/spars/SPARS.pdf

group LASSO has been shown to be equivalent to AMP with block soft thresholding, we can use the state evolution formalism to obtain the PT for group LASSO. Due to the properties of SE (2.3), for any MSE and $\rho(\delta)$ below the PT, the following holds:

$$\frac{1}{\delta B} \mathbb{E} \left\{ \left\| \mathbf{x}_B - \eta^{\text{soft}}(\mathbf{x}_B + \sqrt{\text{MSE}} \mathbf{z}_B; \tau \sqrt{\text{MSE}}) \right\|_2^2 \right\} \leq \text{MSE}. \quad (4.1)$$

To ensure that (4.1) is satisfied while achieving the optimal PT with respect to τ , we use the minimax MSE M_B , which corresponds to [16]

$$M_B = \frac{1}{B} \inf_{\tau} \sup_G \mathbb{E} \left\{ \left\| \mathbf{x}_B - \eta^{\text{soft}}(\mathbf{x}_B + \mathbf{z}_B; \tau) \right\|_2^2 \right\} \quad (4.2)$$

in the asymptotic setting. Donoho *et al.* showed in [16] that:

$$M_B = \frac{1}{B} \inf_{\tau} \left\{ \epsilon(B + \tau^2) + (1 - \epsilon) \int_{\tau^2}^{\infty} (\sqrt{x} - \tau)^2 \frac{1}{2^{\frac{B}{2}} \Gamma(\frac{B}{2})} x^{\frac{B}{2}-1} e^{-\frac{x}{\tau^2}} dx \right\}.$$

One can rigorously show that the optimal PT obeys $M_B = \delta$ [13]. This enables us to obtain, $\rho^L(\delta)$ as follows.

LEMMA 4.1. *The PT of group LASSO is given by*

$$\rho^L(\delta) = \frac{B\delta - \int_{\tau^{*2}}^{\infty} (\sqrt{x} - \tau^*)^2 f(x) dx}{\delta(B + \tau^{*2}) - \int_{\tau^{*2}}^{\infty} (\sqrt{x} - \tau^*)^2 f(x) dx},$$

where $f(x)$ is the probability density function of $\Gamma(\frac{B}{2}, \frac{1}{2})$ and the optimal threshold parameter, τ^* , satisfies

$$\delta = \frac{-(B + \tau^{*2}) \int_{\tau^{*2}}^{\infty} (\tau^* - \sqrt{x}) f(x) dx + \tau^* \int_{\tau^{*2}}^{\infty} (\sqrt{x} - \tau^*)^2 f(x) dx}{B\tau^* - \int_{\tau^{*2}}^{\infty} (\tau^* - \sqrt{x}) f(x) dx}.$$

Note that $\rho^L(\delta)$ is independent of the distribution G (the proof is omitted as it is a mere extension of one provided in [15]). We are interested in observing the behavior of group LASSO for $\delta \rightarrow 0$. Intuitively, for such regime, a very sparse signal is recovered since $\epsilon \rightarrow 0$ ($\epsilon = \rho\delta$). From the definition of block soft thresholding (2.1), τ^* must be very large to promote sparse recovery. Using this knowledge, the integrals in Lemma 4.1 can be approximated via Laplace's method to obtain:

$$\delta \sim \frac{\tau^{*B-2} e^{-B\frac{\tau^{*2}}{2}}}{B2^{\frac{B}{2}-1} \Gamma(\frac{B}{2})}, \quad \rho^L(\delta) \sim \frac{B}{\tau^{*2}}.$$

Algebraic manipulations of these approximations lead to Theorem 3.2.

4.2. Proof of Theorem 3.1. We now derive an expression for $\rho^*(\delta)$ when $\delta \rightarrow 0$. By definition,

$$M_B \triangleq \frac{1}{B} \inf_{\eta} \sup_G \mathbb{E} \left\{ \left\| \mathbf{x}_B - \eta(\mathbf{x}_B + \mathbf{z}_B) \right\| \right\} = \frac{1}{B} \sup_G \inf_{\eta} \mathbb{E} \left\{ \left\| \mathbf{x}_B - \eta(\mathbf{x}_B + \mathbf{z}_B) \right\| \right\},$$

where \mathbf{x}_B is drawn from $F(\mathbf{x}_B)$ as defined in Section 2.1. Defining $\mathbf{y}_B = \mathbf{x}_B + \mathbf{z}_B$, the Bayes estimator, $\mathbb{E}[\mathbf{x}_B | \mathbf{y}_B]$ minimizes the risk for every G . Define

$$M_B(G) \triangleq \frac{1}{B} \mathbb{E} \left\{ \left\| \mathbf{x}_B - \mathbb{E}[\mathbf{x}_B | \mathbf{y}_B] \right\|_2^2 \right\},$$

and consider $F(\mathbf{x}_B) = (1-\epsilon)\delta_0(\|\mathbf{x}_B\|_2) + \epsilon G^*(\mathbf{x}_B)$, where G^* is a uniform distribution on a sphere with radius μ . Now, we characterize the risk associated with $F(\mathbf{x}_B)$.

LEMMA 4.2. *Suppose $\mathbf{x}_B \sim F(\mathbf{x}_B)$. Denote the associated risk function as $M_B(G^*)$. Let $\gamma > 0$, where γ can be made arbitrarily small. Set $(1-\gamma)\log\left(\frac{1-\epsilon}{\epsilon}\right) = \frac{1}{2}\mu^2$. Then,*

$$M_B(G^*) \sim \epsilon \frac{2(1-\gamma)}{B} \log\left(\frac{1-\epsilon}{\epsilon}\right) \quad \text{as } \epsilon \rightarrow 0.$$

The PT for the distribution $F(\mathbf{x}_B)$ is found by letting $M_B(G^*) = \delta$ to obtain $\rho(\delta) \sim B/(2(1-\gamma)\log(\frac{1-\epsilon}{\epsilon}))$. Since $\epsilon \rightarrow 0$ and γ can be made arbitrarily small, we have $\rho(\delta) \lesssim B/(2\log(\frac{1}{\delta}))$. From (3.1), we know that $\rho^*(\delta) \lesssim \rho(\delta) \lesssim B/(2\log(\frac{1}{\delta}))$. Moreover, recall from Theorem 3.2, $\rho^L(\delta) \sim B/(2\log(\frac{1}{\delta}))$ and is independent of the distribution G . By definition of optimality, $\rho^*(\delta) \gtrsim \rho^L(\delta) \sim B/(2\log(\frac{1}{\delta}))$. Comparing this with the upper bound for $\rho^*(\delta)$, we conclude that $\rho^*(\delta) \sim B/(2\log(\frac{1}{\delta}))$.

REFERENCES

- [1] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *J. Roy. Stat. Soc. Ser. B Stat. Methodol.*, 68(1):49–67, 2006.
- [2] Y. Eldar, P. Kuppinger, and H. Bölcskei. Block-sparse signals: Uncertainty relations and efficient recovery. *IEEE Sig. Proc.*, 58(6):3042–3054, June 2010.
- [3] C. Chesneau and M. Hebiri. Some theoretical results on the grouped variables LASSO. *Math. Methods Statist.*, 17(4):317–326, Dec. 2008.
- [4] Y. Nardi and A. Rinaldo. On the asymptotic properties of the group LASSO estimator for linear models. *Electron. J. Statist.*, 2:605–633, 2008.
- [5] F. R. Bach. Consistency of the group LASSO and multiple kernel learning. *J. Mach. Learn. Res.*, 9:1179–1225, 2008.
- [6] L. Meier, S. van de Geer, and P. Bühlmann. The group LASSO for logistic regression. *J. R. Statist. Soc. B*, 70(1):53–77, 2008.
- [7] K. Lounici, M. Pontil, A. B. Tsybakov, and S. van de Geer. Taking advantage of sparsity in multi-task learning. *arXiv:0903.1468v1*, 2009.
- [8] D. Malioutov, M. Cetin, and A. S. Willsky. A sparse signal reconstruction perspective for source localization with sensor arrays. *IEEE Sig. Proc.*, 53(8):3010–3022, Aug. 2005.
- [9] G. Obozinski, M. J. Wainwright, and M. I. Jordan. Support union recovery in high-dimensional multivariate regression. *Ann. Stat.*, 39:1–47, 2011.
- [10] M. F. Duarte, W. U. Bajwa, and R. Calderbank. The performance of group LASSO for linear regression of grouped variables. Technical report, TR-2010-10, Duke University, 2011.
- [11] X. Lv, G. Bi, and C. Wan. The group LASSO for stable recovery of block-sparse signal representations. *IEEE Sig. Proc.*, 59(4):1371–1382, 2011.
- [12] D. L. Donoho and J. Tanner. Counting the faces of randomly projected hypercubes and orthants with applications. *Discr. Comput. Geom.* to appear.
- [13] D. L. Donoho, A. Maleki, and A. Montanari. Message passing algorithms for compressed sensing. *PNAS*, 106:18914–18919, 2009.
- [14] M. Stojnic. Block-length dependent thresholds in block-sparse compressed sensing. *Preprint arXiv:0907.3679*, 2009.
- [15] A. Maleki, L. Anitori, A. Yang, and R. G. Baraniuk. Asymptotic analysis of complex LASSO via complex approximate message passing (CAMP). *Trans. Info. Theory*, to appear.
- [16] D. L. Donoho, I. M. Johnstone, and A. Montanari. Accurate prediction of phase transitions in compressed sensing via a connection to minimax denoising. *arXiv: 1108.0477*, 2011.
- [17] Y. Wu and S. Verdú. Rényi information dimension: Fundamental limits of almost lossless analog compression. *IEEE Trans. Info. Theory*, 56(8):3721–3748, Aug. 2010.
- [18] P. Schniter. Turbo reconstruction of structured sparse signals. In *Proc. IEEE Conf. Inform. Sci. Sys.*, pages 1–6, Aug. 2010.
- [19] U. Kamilov S. Rangan A. Fletcher, M. Unser. Approximate message passing with consistent parameter estimation and applications to sparse learning. *Submitted to IEEE Trans. Inf. Theory*, 2012.