

Compressed sensing with sparse, structured matrices

Maria Chiara Angelini

Dip. Fisica

Università La Sapienza

P.le Aldo Moro 5, 00185 Roma, Italy

Email: Maria.Chiara.Angelini@roma1.infn.it

Federico Ricci-Tersenghi

Dip. Fisica, CNR – IPCF, UOS Roma

INFN – Roma 1, Università La Sapienza

P.le Aldo Moro 5, 00185 Roma, Italy

Email: Federico.Ricci@roma1.infn.it

Yoshiyuki Kabashima

Dept. of Comput. Intell. & Syst. Sci.

Tokyo Institute of Technology

Yokohama 226-8502, Japan

Email: kaba@dis.titech.ac.jp

Abstract—In the context of the compressed sensing problem, we propose a new ensemble of sparse random matrices which allow one (i) to acquire and compress a ρ_0 -sparse signal of length N in a time linear in N and (ii) to perfectly recover the original signal, compressed at a rate α , by using a message passing algorithm (Expectation Maximization Belief Propagation) that runs in a time linear in N . In the large N limit, the scheme proposed here closely approaches the theoretical bound $\rho_0 = \alpha$, and so it is both optimal and efficient (linear time complexity). More generally, we show that several ensembles of dense random matrices can be converted into ensembles of sparse random matrices, having the same thresholds, but much lower computational complexity.

I. INTRODUCTION

Compressed sensing is a framework that enables an N -dimensional sparse signal $\mathbf{s} = (s_i)$ to be recovered from $M (< N)$ linear measurements of its elements, $\mathbf{y} = \mathbf{F}\mathbf{s}$, by exploiting the prior knowledge that \mathbf{s} contains many zero elements [1]. A simple consideration guarantees that ℓ_0 -recovery,

$$\hat{\mathbf{s}} = \underset{\mathbf{x}}{\operatorname{argmin}} \|\mathbf{x}\|_0 \quad \text{subj. to} \quad \mathbf{y} = \mathbf{F}\mathbf{x}, \quad (1)$$

where $\|\mathbf{x}\|_0$ denotes the number of non-zero elements in \mathbf{x} , is theoretically optimal in terms of minimizing the number of measurements M necessary for perfectly recovering any original signal \mathbf{s} . However, carrying out ℓ_0 -recovery for a general measurement matrix \mathbf{F} is NP-hard. To avoid such computational difficulties, an alternative approach, ℓ_1 -recovery

$$\hat{\mathbf{s}} = \underset{\mathbf{x}}{\operatorname{argmin}} \|\mathbf{x}\|_1 \quad \text{subj. to} \quad \mathbf{y} = \mathbf{F}\mathbf{x}, \quad (2)$$

where $\|\mathbf{x}\|_1 = \sum_{i=1}^N |x_i|$, is widely employed, as (2) is generally converted into a linear programming problem, and therefore, signal recovery is mathematically guaranteed in an $O(N^3)$ computational time through the use of the interior point method. Nevertheless, the $O(N^3)$ cost of computation can still be unacceptably high in many practical situations, and much effort is being put into finding more computationally feasible and accurate recovery schemes [2]–[7].

Among such efforts, the recovery scheme recently proposed by Krzakala et al. [7] is worth special attention. Their scheme basically follows the Bayesian approach. Namely, the

signal recovery problem is formulated as one of statistical inference from the posterior distribution,

$$P(\mathbf{x}|\mathbf{F}, \mathbf{y}) = \frac{\delta(\mathbf{F}\mathbf{x} - \mathbf{y})P(\mathbf{x})}{Z(\mathbf{F}, \mathbf{y})}, \quad (3)$$

where $Z(\mathbf{F}, \mathbf{y})$ is a normalization factor imposing the condition $\int d\mathbf{x}P(\mathbf{x}|\mathbf{F}, \mathbf{y}) = 1$ and a component-wise prior distribution $P(\mathbf{x}) = \prod_{i=1}^N [(1-\rho)\delta(x_i) + \rho\phi(x_i)]$ is assumed. ρ and $\phi(x)$ represent the density of the non-zero signal elements and a Gaussian distribution, respectively. Exactly inferring \mathbf{s} from (3) is NP-hard, similarly to (1). However, by employing the belief propagation (BP) in conjunction with the expectation-maximization (EM) algorithm for estimating ρ and the parameters of $\phi(x)$, they developed an approximation algorithm, termed EM-BP, which has better recovery performance than (2) with a computational cost of only $O(N^2)$. Furthermore, they showed that, by employing a peculiar type of “seeded” matrix \mathbf{F} , the threshold of the compression rate $\alpha = M/N$ of EM-BP, above which the original signal is typically recovered successfully, can approach very close to that of ℓ_0 -recovery, $\alpha_{s-EMBP} = \rho_0$, where ρ_0 is the actual signal density of \mathbf{s} and $s-EMBP$ stands for ‘seeded EM-BP’. The seeded matrix is composed of blocks along the diagonal densely filled with Gaussian random variables. It is important to remember that this result is achieved for the first time with an approach different from ℓ_0 -recovery, being the threshold for ℓ_1 -recovery that is much higher than the optimal one: $\alpha_{\ell_1} > \rho_0$. The optimality of the ℓ_0 -recovery is guaranteed for EM-BP, which means that this scheme can practically achieve the theoretically optimal threshold of signal recovery with an $O(N^2)$ computational cost. This remarkable property was recently proved in a mathematically rigorous manner in the case that the matrix entries satisfy certain conditions concerning their statistics [8]. However, it is still unclear whether their scheme is optimal in terms of the computational complexity; there might be a certain design of the measurement matrix \mathbf{F} that makes it possible to further reduce the computational cost while keeping the same signal recovery threshold.

The purpose of the present study is to explore such a possibility. For this, we focus on a class of matrices that are characterized by the following properties:

- **sparsity:** The matrix \mathbf{F} has only $O(1)$ non-zero elements per row and column. This implies that the

measurements can be performed in a time linear in the signal length. This situation is highly preferred for the sake of practicality, given that such an operation typically needs to be done in real time, during data acquisition.

- **integer values:** The matrix elements are not real valued, but take on small integer values. This means that an optimized code for the measurements can work with bitwise operations, thus achieving much better performance without any loss of precision.
- **no block structure:** The block structure used in [7] may not be necessary for reaching the optimal threshold. As an alternative possibility, we study a structure made of a square matrix in the upper left corner (the seed) plus a stripe along the diagonal. This structure is much more amenable for analytic computations, since it corresponds to a one-dimensional model homogeneous in space.

The use of sparse matrices for compressed sensing has already been suggested in several earlier studies [9]–[13]. Our approach is particularly similar to that of [13] in the sense that both sides are based on the Bayesian framework and use integer-valued sparse measurement matrices. Nevertheless, these two approaches differ considerably in the following two points. Firstly, we adopt EM-BP, which updates only a few variables per node for the signal recovery, while the recovery algorithm of [13] involves functional updates and needs significantly more computational time than ours. Secondly, we thoroughly explored a simple design of \mathbf{F} that achieves nearly optimal recovery performance. In contrast, the problem of the matrix design is not fully examined in [13]. By carrying out extensive numerical experiments in conjunction with an analysis based on density evolution [15], we show that a threshold close to the theoretical limit $\alpha = \rho_0$ can be achieved by using matrices with the above properties with an almost *linear* computational cost in the measurement and recovery stages.

This paper is organized as follows. In Section II, the EM-BP algorithm is briefly explained and the results for dense matrices are summarized. The algorithm is applied to homogeneously sparse matrices in Section III and to structured sparse block matrices in Section IV. A new type of “striped” sparse matrix without blocks is introduced in Section V. The last Section summarizes our work, focusing on its importance for practical use, and touches on future issues.

II. EXPECTATION MAXIMIZATION BELIEF PROPAGATION

The new algorithm based on BP in conjunction with the EM proposed in Ref. [7] starts from Eq. (3). A similar idea was also proposed in Ref. [14]. In order to solve it with BP, $O(MN)$ messages for the probability distributions of

the variables x_i are constructed in the following way:

$$m_{\mu \rightarrow i}(x_i) = \frac{1}{Z^{\mu \rightarrow i}} \int \prod_{j \neq i} dx_j m_{j \rightarrow \mu}(x_j) \delta\left(y_\mu - \sum_k F_{\mu k} x_k\right)$$

$$m_{i \rightarrow \mu}(x_i) = \frac{1}{Z^{i \rightarrow \mu}} \left[(1 - \rho) \delta(x_i) + \rho \phi(x_i) \right] \prod_{\gamma \neq \mu} m_{\gamma \rightarrow i}(x_i)$$

where $Z^{i \rightarrow \mu}$ and $Z^{\mu \rightarrow i}$ are normalization factors. This EM-BP equations are very complicated because the messages are distribution functions. In order to make them simpler, the messages can be approximated by assuming that they are Gaussian, thus obtaining the equations for the mean $a_{i \rightarrow \mu}$ and the variance $v_{i \rightarrow \mu}$ of $m_{i \rightarrow \mu}(x_i)$. This approximation was introduced for sparse matrices in Refs. [16]–[19], and it becomes asymptotically exact if \mathbf{F} is dense. In fact, it is derived from an expansion in small $F_{\mu i}$, and in the dense case $F_{\mu i} = O(1/\sqrt{N})$. Supposing that the elements of the original signal follow a Bernoulli-Gaussian distribution with parameters ρ_0 , \bar{x}_0 and σ_0 , the update rules for the messages are the following:

$$a_{i \rightarrow \mu} = f_a \left(\sum_{\gamma \neq \mu} A_{\gamma \rightarrow i}, \sum_{\gamma \neq \mu} B_{\gamma \rightarrow i} \right)$$

$$a_i = f_a \left(\sum_{\gamma} A_{\gamma \rightarrow i}, \sum_{\gamma} B_{\gamma \rightarrow i} \right)$$

$$v_{i \rightarrow \mu} = f_c \left(\sum_{\gamma \neq \mu} A_{\gamma \rightarrow i}, \sum_{\gamma \neq \mu} B_{\gamma \rightarrow i} \right)$$

$$v_i = f_c \left(\sum_{\gamma} A_{\gamma \rightarrow i}, \sum_{\gamma} B_{\gamma \rightarrow i} \right)$$

$$A_{\mu \rightarrow i} = \frac{F_{\mu i}^2}{\sum_{j \neq i} F_{\mu j}^2 v_{j \rightarrow \mu}}$$

$$B_{\mu \rightarrow i} = \frac{F_{\mu i} \left(y_\mu - \sum_{j \neq i} F_{\mu j} a_{j \rightarrow \mu} \right)}{\sum_{j \neq i} F_{\mu j}^2 v_{j \rightarrow \mu}} \quad (4)$$

where f_a and f_c are some analytical functions depending on the parameters ρ , \bar{x} and σ . For details, see Ref. [7].

In general, the original ρ_0 , \bar{x}_0 and σ_0 are not known, but one can use EM to derive the update rules for them, using the property that the partition function

$$Z(\rho, \bar{x}, \sigma) = \int d\mathbf{x} P(\mathbf{x}) \delta(\mathbf{y} - \mathbf{F}\mathbf{x})$$

is the likelihood of the parameters (ρ, \bar{x}, σ) and is maximized by the true parameters ρ_0 , \bar{x}_0 , and σ_0 . Thus, after the update of all the messages, the inferred parameters of the original distribution are updated following these rules:

$$\bar{\rho} \leftarrow \frac{1}{\rho N} \sum_i a_i, \quad \sigma^2 \leftarrow \frac{1}{\rho N} \sum_i (v_i + a_i^2) - \bar{x}^2,$$

$$\rho \leftarrow \frac{\sum_i \frac{1/\sigma^2 + U_i}{V_i + \bar{x}/\sigma^2} a_i}{\sum_i \left[1 - \rho + \frac{\rho}{\sigma(1/\sigma^2 + U_i)^{\frac{1}{2}}} e^{\frac{(V_i + \bar{x}/\sigma^2)^2 - \bar{x}^2}{2\sigma^2}} \right]^{-1}},$$

with $U_i = \sum_{\gamma} A_{\gamma \rightarrow i}$ and $V_i = \sum_{\gamma} B_{\gamma \rightarrow i}$. If the algorithm converges to the correct solution, $a_i = s_i$ and $v_i = 0$.

To reduce the number of messages from $O(NM)$ to $O(N)$, one can see that in the large N limit, the messages $a_{i \rightarrow \mu}$ and $v_{i \rightarrow \mu}$ are nearly independent of μ . Thus, we can derive the equations involving only a variable per each measurement node and a variable per signal node, if we are careful to keep the correcting Onsager reaction term as in the TAP equations of statistical physics [20]. This method was introduced in the context of compressed sensing in Ref. [6] and is called approximated message passing (AMP).

In general, the correct distribution of the original signal is unknown. However, in Ref. [7], it is demonstrated that if $\alpha > \rho_0$, the most probable configuration of \mathbf{x} with respect to $P(\mathbf{x}) = \prod_{i=1}^N [(1 - \rho)\delta(x_i) + \rho\phi(x_i)]$ with $\rho < 1$, restricted to the subspace $\mathbf{y} = \mathbf{F}\mathbf{x}$, is the original signal \mathbf{s} , even if the signal is not distributed according to $P(\mathbf{x})$. So our choice of a Gaussian distribution for $\phi(x)$ should be perfectly fine even if the original signal has a different distribution.

The free entropy $\Phi(D)$ at a fixed mean square error $D = (1/N)\sum_{i=1}^N (x_i - s_i)^2$ can be computed if a dense matrix is used. For $\alpha > \rho_0$, the global maximum of the function $\Phi(D)$ is at $D = 0$, that corresponds to the correct solution. However, below a certain threshold $\alpha < \alpha_{BP}$ that depends on the distribution $P(\mathbf{s})$, the free entropy develops a secondary, local maximum at $D \neq 0$. As a consequence, the EM-BP algorithm can not converge to the correct solution for $\rho_0 < \alpha < \alpha_{BP}$, because a dynamical transition occurs. Nonetheless, the threshold α_{BP} is lower than α_{ℓ_1} .

III. EM-BP WITH A SPARSE MATRIX

First of all, we want to verify if the use of a sparse matrix can reach the same results as the use of a dense one. For a sparse random matrix, the AMP equations can not be used; thus, we can use the update rules in Eq. (4) for inferring the original signal \mathbf{s} . In particular, we choose the matrix \mathbf{F} to have only $K = O(1)$ elements different from zero in each row and $H = \alpha K = O(1)$ elements in each column, extracting them from the distribution,

$$P(F_{\mu i}) = \frac{1}{2}\delta(F_{\mu i} - J) + \frac{1}{2}\delta(F_{\mu i} + J) \quad (5)$$

with $J = 1$. The use of the messages $a_{i \rightarrow \mu}$ and $v_{i \rightarrow \mu}$ instead of the AMP equations does not involve an extra cost in memory, because the number of the messages is $O(N)$ from the sparsity of the matrix. In principle, the messages $m_{i \rightarrow \mu}(x)$ are not Gaussian if the matrix is sparse, so the use of only the two parameters $a_{i \rightarrow \mu}$ and $v_{i \rightarrow \mu}$ is not exact. However, the convolution of K messages (with $K = 20$ in a typical matrix we use) is not far from a Gaussian, and indeed, we can verify a posteriori that this approximation is valid, because it gives good results.

In all our numerical simulations, we use a Bernoulli-Gaussian distributed signal and a compression rate $\alpha = 0.5$.

Figure 1 (top) shows the probability of perfect recovery as a function of the sparsity of the signal ρ_0 for different sizes, by applying the EM-BP algorithm using a sparse matrix

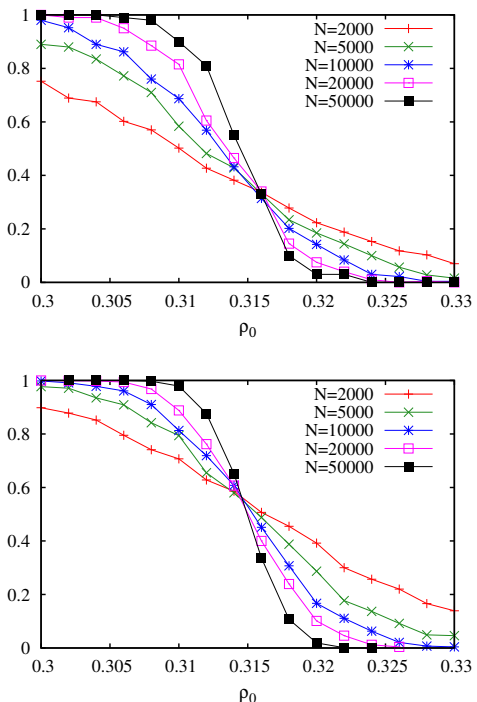


Fig. 1. Top: Probability of perfect recovery versus the signal sparsity ρ_0 using sparse matrices with $\alpha = 0.5$ and $K = 20$. The threshold is the same as with dense matrices. Bottom: Probability of perfect recovery computed with density evolution has the same threshold. Here, N is the population size.

with $K = 20$. The threshold for perfect recovery in the thermodynamic limit ($N \rightarrow \infty$) is $\rho_{BP} \simeq 0.315$, which is the same as the one obtained in Ref. [7] with a dense matrix. We can not analytically compute the free entropy $\Phi(D)$ as in [7], because we use sparse matrices and cannot use methods such as the saddle point one. However, we performed a numerical density evolution analysis, as shown in Fig. 1 (bottom), and found that the threshold is almost the same as the one computed with the matrices \mathbf{F} .

Next, we will verify that the correct solution is always the global maximum of $\Phi(D)$ and it is locally stable up to $\alpha \simeq \rho_0$ when using EM-BP with a sparse matrix. Since we can not analytically compute the free entropy, we must resort to a numerical method. We start EM-BP with an initial condition very close to the correct solution: $a_{i \rightarrow \mu}^0 = s_i + \delta_{i \rightarrow \mu}$, with $\delta_{i \rightarrow \mu}$ a random number uniformly distributed in $[-\Delta, \Delta]$. In this way, we have verified (see Fig. 2) that if Δ is sufficiently small, the correct solution can be found up to $\alpha \simeq \rho_0$, as in the case of a dense matrix.

For the algorithms based on the ℓ_1 minimization, it is known that the threshold with a sparse matrix is lower than that with a dense one. However, these algorithms are not optimal, because the correct solution disappears below the threshold α_{ℓ_1} . In this sense, the EM-BP algorithm is optimal, because the global maximum of the free entropy is always on the correct solution. Thus, one can expect that, if the rank of the sparse matrix is the same as that of the dense one, a




Fig. 2. Stability check of the solution determined by the EM-BP message passing algorithm. Starting the recovery process with a sparse matrix from an initial condition differing less than Δ from the correct solution, the latter is recovered as long as $\rho_0 < \alpha_{stab}(\Delta)$. In the limit $\Delta \rightarrow 0$, the stability limit $\alpha_{stab}(\Delta)$ tends to the theoretical bound α (which is 0.5).




Fig. 3. $\rho_c(N)$ with a sparse, structured matrix with blocks, for different sizes N and values of L (see text). The thresholds for the ℓ_1 -recovery and for EM-BP without any structure are also drawn. For comparison with the data in Fig. 5, we used the exponent $a \simeq 0.18$ best fitting those data.

similar threshold can be reached (as we have demonstrated numerically).

In summary, we can say that the EM-BP algorithm of Ref. [7] seems to reach the same threshold α_{BP} , either using a dense Gaussian matrix or a sparse binary one if the numbers of non-zero elements per row/column are $O(1)$ but sufficiently large. However, the use of a sparse matrix is computationally much faster than the use of a dense one. Moreover, the use of binary elements, instead of Gaussian real values, allows for better code optimization and eventually for hard-wired encoding of the compression process.

IV. BLOCK-STRUCTURED SPARSE MATRICES

To avoid the secondary maximum of the free entropy, in Ref. [7], the authors use a structured block matrix that helps to nucleate the correct solution. The idea is that the correct solution is found for the first variables, and then it propagates

to the whole signal. This idea is similar to the so-called spatial coupling that is very useful for solving many different problems [21]. With this trick, the authors of Ref. [7] reach perfect recovery for almost any $\alpha > \rho_0$ in the large N limit while Ref. [22] reports that the gain is quite small when different recovery algorithms are used. Here, we try to use a matrix with the same block structure, but sparsely filled. We divide the N variables into L groups of size N/L and M measurements into L groups of size $M_p = \alpha_p N/L$ in such a way that $M = \sum_{p=1}^L M_p = \alpha N$ and $1/L \sum_{p=1}^L \alpha_p = \alpha$.

In this way, the matrix F is divided into L^2 blocks, labeled with indices (p, q) . Each block is a sparse binary matrix with k elements different from zero for each row and $h_p = \alpha_p k$ elements for each column, distributed according to Eq. (5), with $J = J_{p,q}$. As in Ref. [7], we choose $J_{p,p-1} = J_1$, $J_{p,p} = 1$, $J_{p,p+1} = J_2$, and $J_{p,q} = 0$ otherwise. The important ingredient to nucleate the correct solution is that in the first block $\alpha_1 = (M_1/N)L > \alpha_{BP}$ holds. For simplicity, we can choose $\alpha_1 = 1$ and $\alpha_p = (L\alpha - 1)/(L - 1)$ for $p \neq 1$. The recovery strongly depends on the parameters J_1 and J_2 , and the best results for $\alpha = 0.5$ are obtained around $J_1 = 4$ and $J_2 = 1$. Moreover, we used these two values in the experiments described below because we wanted to work with matrices with elements having small integer values.

Similarly to the dense case, the use of a sparse structured matrix with blocks allows to overcome the dynamical transition at α_{BP} and to nucleate the correct solution until α is very close to ρ_0 . Figure 3 shows the mean critical threshold $\rho_c(N)$ for different signal lengths at a fixed compression rate $\alpha = 0.5$. The x axis uses the same scaling variable as in Fig. 5, and the best parameter a obtained from the fit of data in Fig. 5 also interpolates the data in Fig. 3 quite well. In the thermodynamic limit, ρ_c extrapolates to a value compatible with the optimal one, α , and it is certainly much higher than the thresholds for ℓ_1 -recovery and for EM-BP without any structure. We have also done a density evolution analysis that confirms this result.

For each value of N and L , the mean critical threshold ρ_c is computed as follows. We randomly generate a block structured matrix F with the given N and L . We start with a sufficiently sparse original signal s , which has been recovered by the algorithm; we then add non-zero entries to the signal and check whether the new signal can be recovered by the algorithm; we go on adding non-zero elements to the signal until a failure in a perfect recovery occurs. The previous to the last value for ρ_0 is the critical threshold for the matrix F . The mean critical threshold is obtained by averaging over many different random matrices and signals, with the same values of N and L . The number of such random extractions goes from 10^3 for the largest N value up to 10^4 for the smallest N value.

The values of (N, L, k) used for the simulations shown in Fig. 3 are the following: (2250, 10, 9), (19000, 20, 19), (31200, 40, 39), and (49000, 50, 49). We need to increase both N and L if we want to obtain good results in the thermodynamic limit. However, if we change L , we must change k too. Indeed, in order to have the same number of




Fig. 4. A nearly one-dimensional sparse matrix with a square block of size $L \times L$ at the top left and non-zero elements in the stripes around the diagonal can achieve compression and perfect recovery close to the theoretical bound in linear time.

elements per row and column for each of the L^2 blocks, we must satisfy the conditions: $(N/L)/M_p = k/h_p$ with k and h_p integer valued. Here, we have used the smaller possible value for k , that is $k = L - 1$. The fact that it is impossible to keep k constant while increasing L implies that these kinds of block-structured matrices always become dense in the thermodynamic limit. This is a limitation of the block structure that we want to eliminate with the matrix proposed in the following Section.

V. AND WITHOUT BLOCKS?

The matrix proposed in Ref. [7] is not the only one that allows the optimal threshold to be reached. Reference [23] analyzes the use of other good dense, block-structured matrices. However, the block-structure is not so simple to handle if one wants to do analytical calculations in the continuum limit. Moreover, in making these block-structured matrices sparse, one has to be careful to find the right values of $L, M, N, M_p, \alpha_p, k$. For these reasons, we want to know if the block structure is crucial, and, if not, we want to eliminate it.

We tried a different structured sparse matrix (see Fig. 4), that we called a striped matrix. It has one sparse square block of size L on the top left of the matrix with $K = O(1)$ elements for each row and column extracted from (5) with $J = 1$. This arrangement is fundamental for nucleating the correct solution. Apart from this first block, the residual compression rate is $\alpha' = \frac{M-L}{N-L}$. Then, we construct a one-dimensional structure around the diagonal of the remaining matrix. For each column $c > L$, we randomly place $2K\alpha'$ non-zero elements, again extracted from (5), in the interval of the width $2L\alpha'$ around the diagonal. One element with $J = 1$ is always placed on the diagonal (actually on the position closest to the diagonal). For the remaining elements, we use the following rules. If the element is at a distance $d \leq L\alpha'/3$ from the diagonal, we use $J = 1$. Otherwise, if its distance is $d > L\alpha'/3$, we use $J = J_1$ below the




Fig. 5. $\rho_c(N)$ with a sparse, striped matrix, as described in Section V for different sizes (from $N = 2000$ to $N = 40000$). The thresholds for ℓ_1 -recovery and for EM-BP without structure are drawn for comparison. The best fitting parameter is $a \simeq 0.18$, and it leads to an extrapolation of $\rho_c(N)$ in the thermodynamic limit compatible with α .

diagonal and $J = J_2$ above the diagonal. In this way, the number of elements per column is constant, while the number of elements per row is a truncated Poisson random variable with mean $2K$: indeed, there are no empty rows, thanks to the rule of placing the first element of each column closest to the diagonal. When constructing the matrix, we apply exactly the same rule to each column, but in the last L columns it may happen that a non-zero element has a row index larger than M : these elements are then moved below the first square matrix by changing the row and column indices as follows: $r \leftarrow r - (M - L)$ and $c \leftarrow c - (N - L)$.

In this way, we have some kind of continuous one-dimensional version of the block-structured matrix discussed in the previous Section. Within this striped matrix ensemble, the thermodynamic limit at a fixed matrix sparsity can be calculated without any problem, by sending $N, L \rightarrow \infty$ at a fixed L/N and fixed $K = O(1)$. In Fig. 5, we show the mean critical threshold reached by using striped matrices with a fixed ratio $L/N = 1/50$ (the same used in the plot of Fig. 4) and different signal lengths. Perfect decoding up to ρ_c is again achieved by using the EM-BP algorithm. We extrapolated the $\rho_c(N)$ data to the thermodynamic limit by assuming the following behavior in the large N limit:

$$\rho_c(N) = \rho_c(\infty) - bN^{-a} \quad (6)$$

The data in Fig. 5 are plotted with the best fitting parameter $a \simeq 0.18$, and the extrapolated value $\rho_c(\infty)$ is perfectly compatible with the theoretical bound α .

Hence, we can conclude that the important ingredient to reach optimality is not the block structure, but the nearly one-dimensional structure, associated with the initial block with $\alpha_1 > \alpha_{BP}$ to nucleate the correct solution.

It is worth noticing that the corresponding statistical mechanics model for these striped random matrices is a one-dimensional disordered model with an interaction range growing with the signal length, as in a Kac construction.

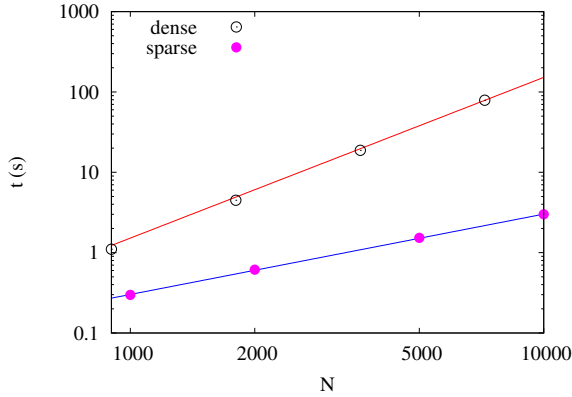


Fig. 6. Actual time (in seconds) for recovery of a signal with dense and sparse matrices for different data lengths N . The data are fitted respectively by a quadratic and a linear function.

Models of this kind are analytically solvable and have shown very interesting results [24].

The use of our striped sparse matrices allows for a great reduction in computational complexity. Indeed, the measurement and recovery times grow linearly with the size of the signal if sparse matrices are used, while they grow quadratically if dense matrices are used. Figure 6 shows the measurement and recovery times of a signal for different signal lengths N . For this test, we used dense block-structured matrices and sparse striped matrices. The number of EM-BP iterations to reach the solution is roughly constant for different N . A quadratic fit for the dense case and a linear fit for the sparse one perfectly interpolate the data.

VI. CONCLUSIONS AND FUTURE DEVELOPMENT

We introduced an ensemble of sparse random matrices F that, thanks to their particular structure (see Figure 4), allow us to perform the following operations in linear time:

- (i) measurement of a ρ_0 -sparse vector s of length N by using a linear transformation, $y = Fs$, to a vector y of length αN
- (ii) perfect recovery of the original vector s by using a message passing algorithm (Expectation Maximization Belief Propagation) for almost any parameter satisfying the theoretical bound $\rho_0 < \alpha$.

These striped sparse matrices that have such good performance because there is a ‘seeding’ sparse square matrix in the upper left corner that nucleates a seed for the right solution and the one-dimensional structure along the diagonal propagates the initial seed to the complete right solution. Both seeding and a one-dimensional structure have been used in the past [7], [21], but in our new ensemble, the matrices are sparse, and this permits us to perform all the operations in a time linear in the signal length.

We also checked that sparse matrices perform as well as dense ones in the case of block-structured matrices and for matrices with no structure at all.

Apart from the compressed sensing case, several other applications require a sparse matrix or equivalently linear time complexity [10]

In data streaming computing, one is typically interested in doing very quick measurements in constant time. For example, if the task is to measure the number of packets s_i with destination i passing through a network router, it is not possible to keep a vector s because it is generally too long. Instead, a much shorter sketch of it, $y = Fs$, is measured in such a way that a very sparse vector s can be recovered from y . The matrix F must be sparse in order to be able to update the sketch y in a constant time for each new packet passing through the router.

Another interesting application is the problem of group testing, where a very sparse vector $s \in \{0, 1\}^N$ is given and one is interested in performing the fewest linear measurements, $y = Fs$, that allow for detection of the defective elements ($s_i = 1$). In this case, the experimental constraints require a sparse matrix F : only if the tested compound $y_\mu = \sum_i F_{\mu i} s_i$ is made of a very few elements of s , the linear response holds and non-linear effects can be ignored.

However, in the more general case, one does not directly observe the sparse signal s but rather a linear transformation of it, $x = Ds$, made with a dictionary matrix D (which is typically a Fourier or wavelet transformation, and thus is a dense matrix). In this more difficult case, one would like to design a sparse measurement matrix A such that the measurement/compression operation, $y = Ax$, is fast, and the resulting observed data y is short, thanks to the sparseness of s . The conflicting requirement is to have a fast recovery scheme, because, now, to recover the original signal one should solve $\hat{s} = \arg\min \|s\|_0$ subject to $y = (AD)s$, where AD is typically dense (e.g. in case of Fourier and wavelet transformations). So a very interesting future development of the present approach is to extend it to this more complex case.

ACKNOWLEDGMENTS

FR-T is grateful for his useful discussions with F. Krzakala, M. Mézard and L. Zdeborova and financial support from the Italian Research Minister through the FIRB Project No. RBFR086NN1 on “Inference and Optimization in Complex Systems: From the Thermodynamics of Spin Glasses to Message Passing Algorithms”.

YK is supported by grants from the Japan Society for the Promotion of Science (KAKENHI No. 22300003) and the Mitsubishi foundation.

REFERENCES

- [1] E. J. Candès and M. B. Wakin, “An Introduction To Compressive Sampling,” *IEEE Signal Processing Magazine*, vol. 25, pp. 21–30, 2008.
- [2] M. A. T. Figueiredo, R. D. Nowak and S. J. Wright, “Gradient Projection for Sparse Reconstruction: Application to Compressed Sensing and Other Inverse Problems,” *IEEE Journal of Selected Topics in Signal Processing* vol. 1, pp. 586–597, 2007.
- [3] R. Cartrand and W. Yin, “Iteratively reweighted algorithms for compressive sensing,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2008, pp. 3869–3872, 2008.
- [4] W. Yin, S. Osher, D. Goldfarb and J. Darbon, “Bregman Iterative Algorithms for ℓ_1 -Minimization with Applications to Compressed Sensing,” *SIAM J. Imaging Sciences*, vol. 1, pp. 143–168, 2008.
- [5] T. Blumensath and M. E. Davies, “Iterative hard thresholding for compressed sensing,” *Applied and Computational Harmonic Analysis*, vol. 27, pp. 265–274, 2009.

- [6] D. L. Donoho, A. Maleki and A. Montanari, "Message-passing algorithms for compressed sensing," *PNAS*, vol. 106, 18914–18919, 2009.
- [7] F. Krzakala, M. Mézard, F. Sausset, Y. F. Sun and L. Zdeborova, "Statistical-physics-based reconstruction in compressed sensing," *Phys. Rev. X*, vol. 2, 021005 (18pages), 2012.
- [8] D. L. Donoho, A. Javanmard and A. Montanari, "Information-Theoretically Optimal Compressed Sensing via Spatial Coupling and Approximate Message Passing," arXiv:1112.0708, 2008.
- [9] R. Berinde and P. Indyk, "Sparse recovery using sparse random matrices," Preprint, 2008.
- [10] A. Gilbert and P. Indyk, "Sparse Recovery Using Sparse Matrices," *Proceedings of the IEEE*, vol. 98, pp. 937–947, 2010.
- [11] M. Akçakaya, J. Park and V. Tarokh, "A Coding Theory Approach to Noisy Compressive Sensing Using Low Density Frames," *IEEE Trans. on Signal Processing*, vol. 59, pp. 5369–5379, 2011.
- [12] Y. Kabashima and T. Wadayama, "A signal recovery algorithm for sparse matrix based compressed sensing," arXiv:1102.3220, 2011.
- [13] D. Baron, S. Sarvotham and R. G. Baranuik, "Bayesian Compressive Sensing via Belief Propagation," *IEEE Trans. on Signal Processing*, vol. 58, pp. 269–280, 2010.
- [14] J. P. Vila and P. Schniter, "Expectation-Maximization Bernoulli-Gaussian Approximate Message Passing," in Proc. Asilomar Conf. on Signals, Systems, and Computers (Pacific Grove, CA), Nov. 2011.
- [15] S.-Y. Chung, G. D. Forney Jr., T. J. Richardson and R. Urbanke, "On the Design of Low-Density Parity-Check Codes within 0.0045 dB of the Shannon Limit," *IEEE Comm. Lett.*, vol. 58 no.2, pp. 58-60, 2000.
- [16] Y. Kabashima, "A CDMA multiuser detection algorithm on the basis of belief propagation," *J. Phys. A*, vol. 36, pp. 11111–11121, 2003.
- [17] D. Guo and C. Wong, "Asymptotic Mean-Square Optimality of Belief Propagation for Sparse Linear Systems," in Proc. of Information Theory Workshop 2006, pp. 194–198, 2006.
- [18] A. Montanari and D. Tse, "Analysis of Belief Propagation for Non-Linear Problems: The Example of CDMA (or: How to Prove Tanaka's Formula)," arXiv:cs/0602028, 2006.
- [19] S. Rangan, "Estimation with random linear mixing, belief propagation and compressed sensing," *Proc. Conf. Inf. Sci. Syst.*, Princeton, NJ, pp. 1–6, 2010.
- [20] D. J. Thouless, P. W. Anderson and R. G. Palmer, "Solution of 'Solvable model of a spin glass'," *Phil. Mag.*, vol. 35, pp. 593–601, 1977.
- [21] S. H. Hassani, N. Macris and R. Urbanke, "Coupled graphical models and their thresholds," arXiv:1105.0785, 2011.
- [22] S. Kudekar and H. D. Pfister, "The Effect of Spatial Coupling on Compressive Sensing," in Proc. the 48th Annual Allerton Conference on Communication, Control, and Computing, 2010.
- [23] F. Krzakala, M. Mézard, F. Sausset, Y. Sun and L. Zdeborova, "Probabilistic Reconstruction in Compressed Sensing: Algorithms, Phase Diagrams, and Threshold Achieving Matrices," arXiv:1206.3953, 2012.
- [24] S. Franz and A. Montanari, "Analytic determination of dynamical and mosaic length scales in a Kac glass model," *J. Phys. A*, vol. 40, pp. F251–F257, 2007.