Generalized Approximate Survey Propagation
for High-Dimensional Estimation

Luca Saglietti [1 2]   Yue M. Lu [3]   Carlo Lucibello [4 1]

Abstract

In Generalized Linear Estimation (GLE) problems, we seek to estimate a signal that is observed through a linear transform followed by a component-wise, possibly nonlinear and noisy, channel. In the Bayesian optimal setting, Generalized Approximate Message Passing (GAMP) is known to achieve optimal performance for GLE. However, its performance can significantly degrade whenever there is a mismatch between the assumed and the true generative model, a situation frequently encountered in practice. In this paper, we propose a new algorithm, named Generalized Approximate Survey Propagation (GASP), for solving GLE in the presence of prior or model mis-specifications. As a prototypical example, we consider the phase retrieval problem, where we show that GASP outperforms the corresponding GAMP, reducing the reconstruction threshold and, for certain choices of its parameters, approaching Bayesian optimal performance. Furthermore, we present a set of State Evolution equations that exactly characterize the dynamics of GASP in the high-dimensional limit.

1. Introduction

Approximate message passing (AMP) algorithms have become a well established tool in the study of inference problems (Donoho et al., 2009; Donoho & Montanari, 2016; Advani & Ganguli, 2016) that can be represented by dense graphical models. An important feature of AMP is that its dynamical behavior in the large system limit can be exactly

[1]Microsoft Research New England, Cambridge, MA 02142, USA [2]Italian Institute for Genomic Medicine, Turin, Italy [3]John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA [4]Bocconi Institute for Data Science and Analytics, Bocconi University, Milan, Italy. Correspondence to: Carlo Lucibello <carlo.lucibello@unibocconi.it>, Luca Saglietti <luca.saglietti@gmail.com>.

predicted through a dynamical system involving only scalar quantities called *State Evolution* (SE) (Bayati & Montanari, 2011). This relationship paved the way for a series of rigorous results (Rangan & Fletcher, 2012; Deshpande & Montanari, 2014; Deshpande et al., 2016). It also helps clarify the connection to several fascinating predictions obtained through the *replica analysis* in statistical physics (Mézard et al., 1987). In the optimal Bayesian setting, where one has perfect information on the process underlying data generation, AMP has been empirically shown to achieve optimal performances among polynomial algorithms for many different problems. However, in the more realistic case of mismatch between the assumed and the true generative model, i.e. when AMP is not derived on the true posterior distribution, it may become sub-optimal. A possible source of problems for the AMP class of algorithms is the outbreak of *Replica Symmetry Breaking* (Mézard et al., 1987), a scenario where an exponential number of fixed point and algorithmic barriers dominate the free energy landscape explored by AMP. This phenomena can be accentuated in case of model mismatch: a notable example is maximum likelihood estimation (as opposed to estimation by the posterior mean, which corresponds to the low temperature limit of a statistical physics model.

These considerations are well known within the physics community of disordered systems (Krzakala et al., 2016), where the problem of signal estimation is informally referred to as "crystal hunting". Estimation problems in high dimensions are characterized by a complex energy-entropy competition where the true signal is hidden in a vast and potentially rough landscape. In a wide class of problems, one observes the presence of a algorithmically "*hard*" phase for some range of values for the parameters defining the problem (e.g. signal-to-noise ration). In this regime, all known polynomial complexity algorithms fail to saturate the information theoretic bound (Ricci-Tersenghi et al., 2019). While reconstruction is possible in principle, algorithms are trapped in a region of the configuration space with low overlap with the signal and many local minima (Antenucci et al., 2019a; Ros et al., 2019).

In a recent work (Antenucci et al., 2019b), a novel message-passing algorithm, *Approximate Survey Propagation* (ASP),

was introduced in the context of low-rank. The algorithm is based on the *1-step Replica Symmetry Breaking* (1RSB) ansatz from spin glass theory (Mézard et al., 1987), which was specifically developed to deal with landscapes populated by exponentially many local minima. It was shown that ASP on the mismatched model could reach the performance of (but not improve on) matched AMP and do far better than mismatched AMP (Antenucci et al., 2019a;b). In the present paper, we build upon these previous works and derive the ASP algorithm for *Generalized Linear Estimation* (GLE) models. Since the extension of AMP to GLE problems is commonly known as GAMP, we call *Generalized Approximate Survey Propagation* (GASP), our extension of ASP. We will show that also in this case, in presence of model mismatch, (G)ASP improves over the corresponding (G)AMP.

2. Model specification

An instance of the general class of models to which GASP can be applied is defined, for some integer M and N, by an *observed signal* $\boldsymbol{y} \in \mathbb{R}^M$ and an $M \times N$ *observation matrix* $\boldsymbol{F}$. Clearly, this scenario encompasses also GLE. We denote with $\boldsymbol{F}^\mu$, $\mu \in [M]$, the rows of $\boldsymbol{F}$ and refer to the ratio $\alpha = M/N$ as the *sampling ratio* of $\boldsymbol{F}$. We consider a probability density distribution $p(\boldsymbol{x})$ on a (possibly discrete) space χ^N, $\chi \subseteq \mathbb{R}$, defined as:

$$p(\boldsymbol{x}) = \frac{1}{Z} \, e^{-\beta \mathcal{H}_{\boldsymbol{y},\boldsymbol{F}}(\boldsymbol{x})}, \tag{1}$$

where, following statistical physics jargon, β plays the role of an inverse temperature, Z is a normalization factor called partition function (both Z and p implicitly depend on $\beta, \boldsymbol{y}$ and $\boldsymbol{F}$), and $\mathcal{H}_{\boldsymbol{y},\boldsymbol{F}}$ is the Hamiltonian of the model, that in our setting takes the form:

$$\mathcal{H}_{\boldsymbol{y},\boldsymbol{F}}(\boldsymbol{x}) = \sum_{\mu=1}^{M} \ell\big(y_\mu, \langle \boldsymbol{F}^\mu, \boldsymbol{x}\rangle\big) + \sum_{i=1}^{N} r(x_i). \tag{2}$$

Here $\langle \bullet, \bullet \rangle$ denotes the scalar product and we call ℓ and r the loss function and the regularizer of the problem respectively.

In this quite general context, the purpose of GASP is to approximately compute the marginal distribution $p(x_i)$, along with some expected quantities such as e.g. $\hat{\boldsymbol{x}} = \mathbb{E}_p \boldsymbol{x}$. The approximation entailed in GASP turns out to be exact under some assumptions in the large N limit, as we shall later see. A crucial assumption in the derivation of the GASP algorithm (and of GAMP as well), is that the entries of $\boldsymbol{F}$ are independently generated according to some zero mean and finite variance distribution.

Although the general formulation of GASP, presented in Sec. 2 of the SM, is able to deal with any model of the form (1), we will here restrict the setting to discuss Generalized Linear Estimation (Rangan, 2011).

In GLE problems, $p(x)$ is sensibly chosen in order to infer a *true signal* $\boldsymbol{x}_0 \in \mathbb{R}^N$, whose components are assumed to be independently extracted from some prior P_0, $x_{0,i} \sim P_0 \; \forall i \in [N]$. The observations are independently produced by a (probabilistic) scalar channel P^{out}: $y_\mu \sim P^{\text{out}}(\bullet \,|\, \langle \boldsymbol{F}^\mu, \boldsymbol{x}_0\rangle)$.

It is then reasonable to choose $\ell(y, z) = -\log P^{\text{out}}(y|z)$, $r(x) = -\log P_0(x)$ and $\beta = 1$, so that the probability density $p(\boldsymbol{x})$ corresponds to the true posterior $P(\boldsymbol{x}|\boldsymbol{F}, \boldsymbol{y}) \propto P^{\text{out}}(\boldsymbol{y}|\boldsymbol{x}, \boldsymbol{F})P_0(\boldsymbol{x})$, where $\propto$ denotes equality up to a normalization factor. We refer to this setting as to the Bayesian-optimal or matched setting (Barbier et al., 2018). Notice that in the $\beta \uparrow \infty$ limit $p(x)$ concentrates around the maximum-a-posteriori (MAP) estimate. If $\beta \neq 1$ or if the Hamiltonian doesn't correspond to the minus log posterior (e.g, when P_0 and P^{out} used in the Hamiltonian do not correspond to true ones) we talk about model mismatch.

As a testing ground for GASP, and the corresponding State Evolution, we here consider the phase retrieval problem, which has undergone intense investigation in recent years (Candes et al., 2015; Dhifallah & Lu, 2017; Chen et al., 2018; Goldstein & Studer, 2018; Mondelli & Montanari, 2018; Sun et al., 2018; Mukherjee & Seelamantula, 2018). We examine its noiseless and real-valued formulation, where observations are generated according to the process

$$\boldsymbol{x}_0 \sim \mathcal{N}(0, I_N), \tag{3}$$
$$F_i^\mu \sim \mathcal{N}(0, 1/N) \quad \forall \mu \in [M], \forall i \in [N], \tag{4}$$
$$y_\mu \sim |\langle \boldsymbol{F}^\mu, \boldsymbol{x}_0\rangle|. \tag{5}$$

for some M and N, such that $\alpha = M/N > 1$. For such generative model, we will focus on the problem of recovering $\boldsymbol{x}_0$ by minimizing the energy function $\mathcal{H}_{\boldsymbol{y},\boldsymbol{F}}(\boldsymbol{x})$ of Eq. (2), in the case

$$\ell(y, z) = (y - |z|)^2, \tag{6}$$
$$r(x) = \frac{1}{2}\lambda \, x^2. \tag{7}$$

Since the setting assumed for inference corresponds to MAP estimation in presence of a noisy channel, we are dealing with a case of model mismatch. The effect of the parameter λ on the estimation shall be explored in Sec. 7, but we assume $\lambda = 0$ until then. The optimization procedure will be performed using the zero-temperature (i.e. $\beta \uparrow \infty$) version of the GASP algorithm.

3. Previous work on Approximate Message Passing for Phase Retrieval

Generalized approximate message passing (GAMP) was developed and rigorously analyzed in Refs. (Rangan, 2011) and (Javanmard & Montanari, 2013). It was then applied for the first time to the (complex-valued) phase retrieval

problem in Ref. (Schniter & Rangan, 2015). In Ref. (Barbier et al., 2018) the authors report an algorithmic threshold for the perfect recovery of $\alpha_{\text{alg}} \approx 1.13$, when using matched AMP on the real-valued version of the problem. This is to be compared to the information theoretic bound $\alpha_{\text{IT}} = 1$.

The performance of GAMP in the MAP estimation setting, instead, was investigated in Ref. (Ma et al., 2018; 2019). A "vanilla" implementation of the zero temperature GAMP equations for the absolute value channel was reported to achieve perfect recovery for real-valued signals above $\alpha_{\text{alg}} \approx 2.48$. The authors were able to show that the algorithmic threshold of GAMP in the mismatched case can however be drastically lowered by introducing regularization a regularization term ultimately continued to zero. The AMP.A algorithm proposed in (Ma et al., 2018; 2019) uses an adaptive L_2 regularization that improves the estimation threshold and also makes the algorithm more numerically robust compensating a problematic divergence that appears in the message-passing equations (see Sec. 1.3 in the SM for further details).

Another important ingredient for AMP.A's performance is initialization: in order to achieve perfect recovery one has to start from a configuration that falls within the basin of attraction of the true signal, which rapidly shrinks as the sampling ratio α decreases. A well-studied method for obtaining a configuration correlated with the signal is *spectral initialization*, introduced and studied in Refs. (Jain et al., 2013; Candes et al., 2015; Chen & Candes, 2015): in this case the starting condition is given by the principal eigenvector of a matrix obtained from the data matrix $\boldsymbol{F}$ and the labels $\boldsymbol{y}$ passed through a nonlinear processing function. The asymptotic performance of this method was analyzed in (Lu & Li, 2017), while the form of the optimal processing function was described in (Mondelli & Montanari, 2018; Luo et al., 2019). However, since the SE description is based on the assumption of the initial condition being uncorrelated with the data, in AMP.A the authors revisited the method, proposing a modification that guarantees "enough independency" while still providing high overlap between the starting point and the signal.

With the combination of these two heuristics, AMP.A is able to reconstruct the signal down $\alpha_{\text{alg}} \approx 1.5$. In the present paper we will show that, with a basic continuation scheme, the 1RSB version of the zero temperature GAMP can reach the Bayes-optimal threshold $\alpha_{\text{alg}} \approx 1.13$ also in the mismatched case, without the need of spectral initialization.

3.1. GAMP equations at zero temperature

Here we provide a brief summary of the AMP equations for the general graphical model of Eq. (1), in the $\beta \uparrow \infty$ limit. This is both to allow an easy comparison with our novel GASP algorithm and to introduce some notation that

will be useful in the following discussion. There is some degree of model dependence in the scaling of the messages when taking the zero-temperature limit: here we adopt the one appropriate for over-constrained models in continuous space. Details of the derivation can be found in Sec. 1 of the SM, along with the specialization of the equations for phase retrieval.

First, we introduce two free entropy functions associated to the input and output channels (Rangan, 2011):

$$\varphi^{\text{in}}(B, A) = \max_x \; -r(x) - \frac{1}{2}Ax^2 + Bx \qquad (8)$$

$$\varphi^{\text{out}}(\omega, V, y) = \max_u \; -\frac{(u - \omega)^2}{2V} - \ell(y, u). \qquad (9)$$

We define for convenience $\varphi_i^{\text{in},t} = \varphi^{\text{in}}(B_i^t, A^t)$ and $\varphi_\mu^{\text{out},t} = \varphi^{\text{out}}(\omega_\mu^t, V^{t-1}, y_\mu)$. In our notation the GAMP message passing equations read:

$$\omega_\mu^t = \sum_i F_i^\mu \hat{x}_i^{t-1} - g_\mu^{t-1} V^{t-1} \qquad (10)$$

$$g_\mu^t = \partial_\omega \varphi_\mu^{\text{out},t} \qquad (11)$$

$$\Gamma_\mu^t = -\partial_\omega^2 \varphi_\mu^{\text{out},t} \qquad (12)$$

$$A^t = c_{\boldsymbol{F}} \sum_\mu \Gamma_\mu^t \qquad (13)$$

$$B_i^t = \sum_\mu F_i^\mu g_\mu^t + \hat{x}_i^{t-1} A^t \qquad (14)$$

$$\hat{x}_i^t = \partial_B \varphi_i^{\text{in},t} \qquad (15)$$

$$\Delta_i^t = \partial_B^2 \varphi_i^{\text{in},t} \qquad (16)$$

$$V^t = c_{\boldsymbol{F}} \sum_i \Delta_i^t \qquad (17)$$

where $c_{\boldsymbol{F}} = \frac{1}{MN} \sum_{\mu,i} (F_i^\mu)^2$. It is clear from the equations that the two free entropy functions are supposed to be twice differentiable. This is not the case for phase retrieval, where GAMP encounters some non-trivial numerical stability issues: during the message-passing iterations one would have to approximately evaluate an empirical average of $\partial_\omega^2 \varphi_\mu^{\text{out},t}$, containing Dirac's δ-function. This is the problem encountered in AMP.A of Ref. (Ma et al., 2018). We will see that this problem is not present in GASP thanks to a Gaussian smoothing of the denoising function.

4. Generalized Approximate Survey Propagation

The (G)ASP algorithm builds on decades of progress within the statistical physics community in understanding and dealing with rough high-dimensional landscapes. The starting point for the derivation of the algorithm is the partition

function of m replicas (or clones) of the system $\{\boldsymbol{x}^a\}_{a=1}^m$:

$$Z_{\boldsymbol{y},\boldsymbol{F}}^m = \int \prod_{a=1}^m \prod_{i=1}^N dx_i^a \; e^{-\beta \sum_{a=1}^m \mathcal{H}_{\boldsymbol{y},\boldsymbol{F}}(\boldsymbol{x}^a)}. \qquad (18)$$

Note that, while this probability measure factorizes trivially, setting $m \neq 1$ can introduce many important differences with respect to the standard case, both from the algorithmic and from the physics standpoints (Monasson, 1995; Antenucci et al., 2019b).

We write down the Belief Propagation (BP) equations associated to the replicated factor graph, where messages are probability distributions associated to each edge over the single-site replicated variables $\bar{\boldsymbol{x}}_i = (x_i^1, \ldots, x_i^m)$. We make the assumption that the messages are symmetric under the group of replica indexes permutations. This allows for a parametrization of the message passing that can be continued analytically to any real value of m. The resulting algorithm goes under the name of 1RSB Cavity Method or, more loosely speaking, of Survey Propagation (with reference in particular to a zero temperature version of the 1RSB cavity method in discrete constraint satisfaction problems), and led to many algorithmic breakthroughs in combinatorial optimization on sparse graphical models (Mézard et al., 2002; Braunstein et al., 2005; Krzakała et al., 2007). One possible derivation of the (G)ASP algorithm is as the dense graph limit of the Survey Propagation equations, in the same way as AMP is obtained starting from BP. The derivation requires two steps. First, BP messages are projected by moment-matching onto (replica-symmetric) multivariate Gaussian distributions on the replicated variables $\bar{\boldsymbol{x}} \in \mathbb{R}^m$, which we express in the form

$$\nu(\bar{\boldsymbol{x}}) \propto \int d\, h \; e^{-\frac{1}{2\Delta_0}(h-\hat{x})^2} \prod_{a=1}^m e^{-\frac{1}{2\Delta_1}(x^a-h)^2}; \qquad (19)$$

Then, messages on the edges are conveniently expressed in term of single site quantities. We note that, some statistical independence assumptions on the entries of the measurement matrix are crucial for the derivation, as goes for AMP as well. While the starting point of the derivation assumed integer m, the resulting message passing can be analytically continued to any real m. Applying this procedure to the GLE graphical model of Eq. (1) we obtain the GASP equations. Here we consider the $\beta \uparrow \infty$ limit to deal with the MAP estimation problem. Details of the GASP derivation and the finite β GASP equations are given in Sec. 2 of the SM. Particular care has to be taken in the limit procedure, as a proper rescaling with β is needed for each parameter. For instance, as the range of sensible choices for m shrinks towards zero for increasing β, we rescale m through the substitution $m \leftarrow m/\beta$.

Relying on the definitions given Eqs. (8) and (9), we intro-

Algorithm 1 GASP(m) for MAP
___
  initialize $g_\mu = 0 \; \forall \mu$
  initialize $V_0, V_1, \hat{x}_i \; \forall i$ to some values
  for $t = 1$ **to** $t_{\max}$ **do**
    compute $\omega_\mu, g_\mu, \Gamma_\mu^0, \Gamma_\mu^1 \; \forall \mu$ using (22,23, 24, 25)
    compute A_0, A_1 using (26,27)
    compute $B_i, \hat{x}_i, \Delta_{0,i}, \Delta_{1,i} \; \forall i$ using (28, 29, 30, 31)
    compute V^0, V^1 using (32, 33)
  end for
___

duce the two 1RSB free entropies:

$$\phi^{\text{in}}(B, A_0, A_1, m) = \frac{1}{m} \log \int Dz \; e^{m\varphi^{\text{in}}(B+\sqrt{A_0}z, A_1)} \qquad (20)$$

$$\phi^{\text{out}}(\omega, V_0, V_1, y, m) = \frac{1}{m} \log \int Dz \; e^{m\varphi^{\text{out}}(\omega+\sqrt{V_0}z, V_1, y)}. \qquad (21)$$

Here $\int Dz$ denotes the standard Gaussian integration $\int dz \exp(-z^2/2)/\sqrt{2\pi}$. Using the shorthand notations $\phi_i^{\text{in},t} = \phi^{\text{in}}(B_i^t, A_0^t, A_1^t, m)$ and $\phi_\mu^{\text{out},t} = \phi^{\text{out}}(\omega_\mu^t, V_0^{t-1}, V_1^{t-1}, y_\mu, m)$ (notice the shift in the time indexes), and using again the definition $c_{\boldsymbol{F}} = \frac{1}{MN} \sum_{\mu,i} (F_i^\mu)^2$ (hence $\mathbb{E} c_{\boldsymbol{F}} = 1/N$ in our setting), the GASP equations read:

$$\omega_\mu^t = \sum_i F_i^\mu \hat{x}_i^{t-1} - g_\mu^{t-1} (mV_0^{t-1} + V_1^{t-1}) \qquad (22)$$

$$g_\mu^t = \partial_\omega \phi_\mu^{\text{out},t} \qquad (23)$$

$$\Gamma_{0,\mu}^t = 2\partial_{V_1} \phi_\mu^{\text{out},t} - (g_\mu^t)^2 \qquad (24)$$

$$\Gamma_{1,\mu}^t = -\partial_\omega^2 \phi_\mu^{\text{out},t} + m\Gamma_{0,\mu}^t \qquad (25)$$

$$A_0^t = c_{\boldsymbol{F}} \sum_\mu \Gamma_{0,\mu}^t \qquad (26)$$

$$A_1^t = c_{\boldsymbol{F}} \sum_\mu \Gamma_{1,\mu}^t \qquad (27)$$

$$B_i^t = \sum_\mu F_i^\mu g_\mu^t - \hat{x}_i^{t-1} (mA_0^t - A_1^t) \qquad (28)$$

$$\hat{x}_i^t = \partial_B \phi_i^{\text{in},t} \qquad (29)$$

$$\Delta_{0,i}^t = -2\partial_{A_1} \phi_i^{\text{in},t} - (\hat{x}_i^t)^2 \qquad (30)$$

$$\Delta_{1,i}^t = \partial_B^2 \phi_i^{\text{in},t} - m\Delta_{0,i}^t. \qquad (31)$$

$$V_0^t = c_{\boldsymbol{F}} \sum_i \Delta_{0,i}^t \qquad (32)$$

$$V_1^t = c_{\boldsymbol{F}} \sum_i \Delta_{1,i}^t \qquad (33)$$

The computational time and memory complexity per iteration of the algorithm is the same of GAMP and is determined

by the linear operations in Eqs. (22) and (28). With respect to GAMP, we have the additional (but sub-leading) complexity due to the integrals in the input and output channels. In some special cases, the integrals in Eqs. (20) and (21) can be carried out analytically (e.g. in the phase retrieval problem).

Notice that GASP iteration reduces to standard GAMP iterations if V_0 and A_0 are initialized (or shrink) to zero, but can produce non-trivial fixed points depending on the initialization condition and on the value of m.

We remark the importance of setting the time-indices correctly in order to allow convergence (Caltagirone et al., 2014). The full algorithm is detailed in Alg. 1.

The expressions for the special case of the absolute value channel (6) and L_2 regularization (7) can be found in Sec. 2.4 of the SM. An important comment is that the divergence issue arising in AMP.A, in the same setting, does not affect GASP: the discontinuity in the expression for the minimizer of Eq. (9) is smoothed out in the 1RSB version by the Gaussian integral in Eq. (20). We also note that, in phase retrieval, a problematic initialization can be obtained by choosing configurations that are exactly orthogonal to the signal, since the message-passing will always be trapped in the uninformative fixed-point (due to the Z_2 symmetry of the problem). However, for finite size instances, a random Gaussian initial condition will have an overlap $\rho \equiv \langle \hat{x}, x_0 \rangle$ of order $\mathcal{O}\left(1/\sqrt{N}\right)$ with the signal, which allows to escape the uninformative fixed point whenever it is unstable (i.e. for high α).

In Fig. 1 (Top and Middle), we show the probability of a perfect recovery and convergence times of GASP for the real-valued phase retrieval problem, for different sampling ratios α and values of the symmetry-breaking parameter m, with $\lambda = 0$. The initial condition is given by $V_0^{t=0} = V_1^{t=0} = 1$ and $\hat{x}^{t=0} \sim \mathcal{N}(0, I_N)$. Notice that standard Gaussian initialization is able to break the symmetry of the channel and, at large t, GASP matches the fixed points predicted by SE (see next Section) with a small initial overlap with the true signal ($\rho^{t=0} = 10^{-3}$). In order to achieve signal recovery at low α, the symmetry-breaking parameter has to be increased. In correspondence of values $m \approx 100$, we report an algorithmic threshold around $\alpha_{\text{alg}}^{\lambda=0} \approx 1.5$. This threshold is comparable to the one of AMP.A, without exploiting adaptive regularization and spectral initialization as AMP.A (and which could be employed also for GASP).

We report that, at fixed m, when α is increased above a certain value the message-passing will stop converging. The oscillating/diverging behavior of the messages can however be exploited for hand-tuning m, in the absence of a replica analysis to support the selection of its most appropriate value. More details can be found in Sec. 3 of the SM.
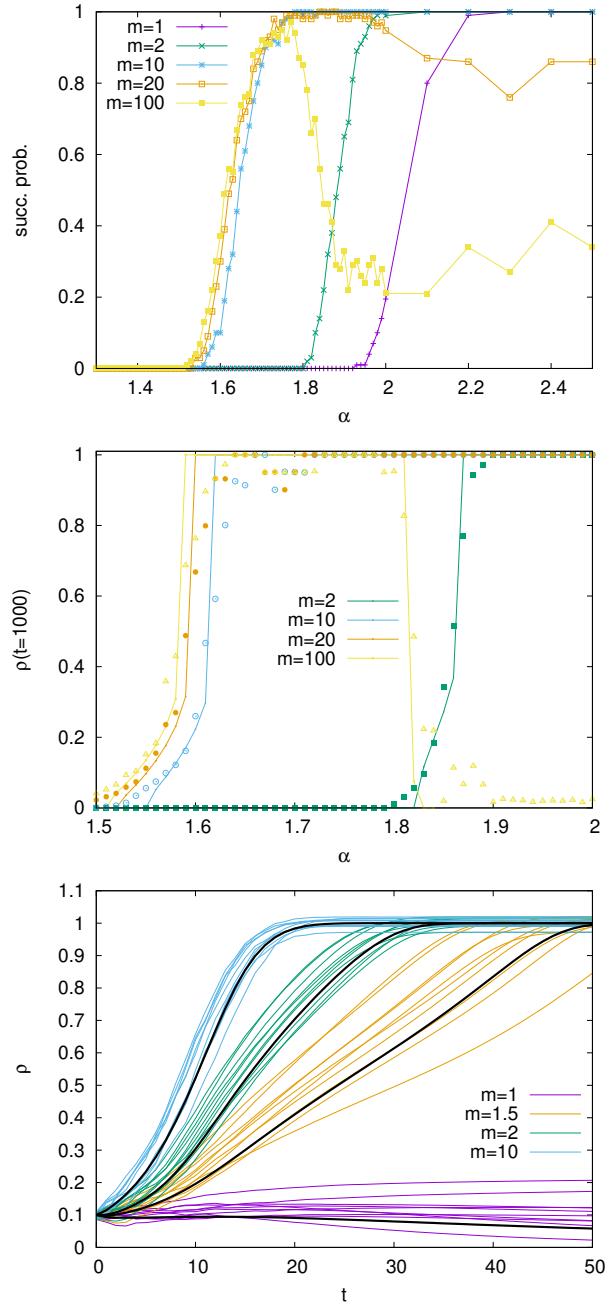

Figure 1. (Top) Probability of perfect recovery of the true signal using GASP (Alg. 1), as a function of the sampling ratio $\alpha = M/N$. (Middle) GASP and SE result after $t = 10^3$ iterations. Start at $t = 0$ with $\rho = 10^{-3}$ for SE and $\hat{x} \sim \mathcal{N}(0, I_N)$ for GASP ($N = 10^3$, averaged over 100 samples). (Bottom) Overlap ρ^t with the true signal predicted by SE dynamics at $\alpha = 2$ and initial overlap $\rho = 0.1$ (black lines) compared to 10 GASP trajectories for each value of m. Here $N = 10^4$.

We presented here the zero-temperature limit of the GASP message-passing to solve the MAP problem. Refer to Sec. 2 of the SM for a more general formulation dealing with the class of graphical models in the form of Eq. 1.

5. State Evolution for GASP

State Evolution (SE) is a set of iterative equations involving a few scalar quantities, that were rigorously proved to track the (G)AMP dynamics, in the sense of almost sure convergence of empirical averages (Javanmard & Montanari, 2013) in the large N limit and with fixed sampling ratio $\alpha = M/N$. Following the analysis of Ref. (Rangan, 2011) for GAMP, in order to present the SE equations for GASP we assume that the observation model $y_\mu \sim P^{\text{out}}(\bullet \,|\, \langle \boldsymbol{F}^\mu, \boldsymbol{x}_0 \rangle)$ is such that can be expressed in the form $y^\mu \sim h(\langle \boldsymbol{F}^\mu, \boldsymbol{x}_0 \rangle, \xi^\mu)$ for some function $h(z,\xi)$, with ξ^μ a scalar- or vector-valued random variable modeling the noise and sampled according to some distribution P_ξ. We also set $F_i^\mu \sim \mathcal{N}(0, 1/N)$ i.i.d.. The recursion is a closed set of equations over the variables $\hat{\rho}^t, \hat{q}_0^t, A_0^t, A_1^t, \rho^t, q_0^t, V_0^t$ and V_1^t Initializing at time $t = 0$ the variables ρ, q_0, V_0 and V_1, the SE equations for $t \geq 1$:

$$\hat{\rho}^t = \alpha \mathbb{E}\big[\partial_z \partial_\omega \phi^{\text{out}}(\omega^t, V_0^{t-1}, V_1^{t-1}, h(z,\xi), m)\big] \quad (34)$$

$$\hat{q}_0^t = \alpha \mathbb{E}\big[\big(\partial_\omega \phi^{\text{out}}(\omega^t, V_0^{t-1}, V_1^{t-1}, y, m)\big)^2\big] \quad (35)$$

$$A_0^t = \alpha \mathbb{E}\big[2\partial_{V_1} \phi^{\text{out}}(\omega^t, V_0^{t-1}, V_1^{t-1}, y, m)\big] - \hat{q}_0^t \quad (36)$$

$$A_1^t = \alpha \mathbb{E}\big[-\partial_\omega^2 \phi^{\text{out}}(\omega^t, V_0^{t-1}, V_1^{t-1}, y, m)\big] + m A_0^t, \quad (37)$$

where the expectation is over the process $\omega^t \sim \mathcal{N}(0, q_0^{t-1})$, $z \sim \mathcal{N}(\rho^{t-1}/q_0^{t-1} \omega^t, \mathbb{E}[x_0^2] - (\rho^{t-1})^2/q_0^{t-1}), \xi \sim P_\xi$ and $y \sim h(z,\xi)$. Also, we have a second set of equations that read:

$$\rho^t = \mathbb{E}\big[x_0\, \partial_B \phi^{\text{in}}(B^t, A_0^t, A_1^t)\big] \quad (38)$$

$$q_0^t = \mathbb{E}\big[\big(\partial_B \phi^{\text{in}}(B^t, A_0^t, A_1^t, m)\big)^2\big] \quad (39)$$

$$V_0^t = \mathbb{E}\big[-2\partial_{A_1} \phi^{\text{in}}(B^t, A_0^t, A_1^t, m)\big] - q_0^t \quad (40)$$

$$V_1^t = \mathbb{E}\big[\partial_B^2 \phi^{\text{in}}(B^t, A_0^t, A_1^t)\big] - m V_0^t \quad (41)$$

where the expectation is over the Markov chain $x_0 \sim P_0$, $B^t \sim \mathcal{N}(\hat{\rho}^t x_0, \hat{q}_0^t)$.

The trajectories of V_0^t, V_1^t, A_0^t and A_1^t in GASP concentrate for large N on their expected value given by the SE dynamics. In order to frame the GASP State Evolution in the rigorous setting of Ref.(Javanmard & Montanari, 2013), we define a slightly different message-passing by replacing their GASP values for a given realization of the problem with the correspondent sample-independent SE values. Also, we replace $c_{\boldsymbol{F}}$ with the expected value $1/N$. Let us define the denoising functions:

$$\eta^{\text{out}}(\omega, y, t) = \partial_\omega \phi^{\text{out}}(\omega, V_0^{t-1}, V_1^{t-1}, y, m) \quad (42)$$

$$\eta^{\text{in}}(B, t) = \partial_B \phi^{\text{in}}(B, A_0^t, A_1^t, m) \quad (43)$$

and their vectorized extensions $\boldsymbol{\eta}^{\text{out}}(\boldsymbol{\omega}, \boldsymbol{y}, t) = (\eta^{\text{out}}(\omega_1, y_1, t), \ldots, \eta^{\text{out}}(\omega_M, y_M, t))$ and $\boldsymbol{\eta}^{\text{in}}(\boldsymbol{\omega}, \boldsymbol{y}, t) = (\eta^{\text{in}}(B_1, t), \ldots, \eta^{\text{in}}(B_N, t))$. The modified GASP message-passing then reads

$$\boldsymbol{\omega}^t = \boldsymbol{F}\,\boldsymbol{\eta}^{\text{in}}(\boldsymbol{B}^{t-1}, t-1) - d_{t-1}^{\text{in}} \boldsymbol{\eta}^{\text{out}}(\boldsymbol{\omega}^{t-1}, \boldsymbol{y}, t-1) \quad (44)$$

$$\boldsymbol{B}^t = \boldsymbol{F}^T\,\boldsymbol{\eta}^{\text{out}}(\boldsymbol{\omega}^t, \boldsymbol{y}, t) - d_t^{\text{out}} \boldsymbol{\eta}^{\text{in}}(\boldsymbol{B}^{t-1}, t-1) \quad (45)$$

where the divergence terms are given by

$$d_t^{\text{in}} = \frac{1}{N} \sum_{i=1}^{N} \partial_B \eta^{\text{in}}(B_i^t, t)$$
$$d_t^{\text{out}} = \frac{1}{N} \sum_{\mu=1}^{M} \partial_\omega \eta^{\text{out}}(\omega_\mu^t, y_\mu, t) \quad (46)$$

Message-passing (44, 45) falls within the class of AMP algorithms analyzed in Ref. (Javanmard & Montanari, 2013) (under some further technical assumptions, see Proposition 5 there). Therefore, it can be rigorously tracked by the SE Eqs. (34,41) in the sense specified in that work. In particular, denoting here $\hat{\boldsymbol{x}}^t = \boldsymbol{\eta}^{\text{in}}(\boldsymbol{B}^t, t)$, we have have almost sure converge in the large system limit of the overlap with the true signal and of the norm of $\hat{\boldsymbol{x}}^t$ to their SE estimates:

$$\lim_{N \to \infty} \frac{1}{N} \langle \hat{\boldsymbol{x}}^t, \boldsymbol{x}_0 \rangle = \rho^t \quad (47)$$

$$\lim_{N \to \infty} \frac{1}{N} \langle \hat{\boldsymbol{x}}^t, \hat{\boldsymbol{x}}^t \rangle = q_0^t \quad (48)$$

In Fig. 1(Bottom), we compare the SE dynamics to the original GASP one (Alg. 1). We compare SE prediction for the evolution of the overlap ρ to that observed in 10 sample trajectories of GASP at $N = 1000$, for a sampling ratio of $\alpha = 2$ and different values of m. The initial estimate $\hat{\boldsymbol{x}}^{t=0}$ in GASP was set to be a mixture $\hat{\boldsymbol{x}}^{t=0} \sim \mathcal{N}(0, I_N) + 0.1 \boldsymbol{x}_0$. Therefore we initialize SE with $\rho^{t=0} = 0.1$, and $q^{t=0} = 1 + (\rho^{t=0})^2$. Moreover, we set $V_0^{t=0} = V_1^{t=0} = 1$ for both. As expected, we observe a good agreement between the two dynamics.

6. Effective Landscape and Message-Passing Algorithms

The posterior distribution of statistical models in the hard phase is known to be riddled with glassy states (Antenucci et al., 2019a) preventing the retrieval of the true signal, a situation which is exacerbated in the low temperature limit corresponding to MAP estimation.

Within the replica formalism, the 1RSB free energy provides a description of this complex landscape. The Parisi parameter m allows to select the contributions of different families of states. More specifically m acts as an inverse temperature

coupled to the internal free energy of the states: increasing m selects families of states with lower complexity (i.e., states that are less numerous) and lower free energy.

The fixed points of the State Evolution of GASP are in one-to-one correspondence to the stationary points of the 1RSB free energy, and while the role of m in the dynamics of SE is difficult to analyze, some insights can be gained from the static description given by the free energy.

For phase retrieval in the MAP setting without regularization, a stable fixed-point of GAMP can be found in the space orthogonal to the signal (i.e. at overlap $\rho = 0$) for values of the sampling ratio below $\alpha^{\text{GAMP}} \approx 2.48$ (Ma et al., 2018), which is the algorithmic threshold for GAMP. For GASP instead, it is possible to see that the uninformative fixed-point is stable only below $\alpha^{\text{GASP}} \approx 1.5$, a noticeable improvement of the threshold with respect to GAMP. This is obtained by choosing the m corresponding to lowest complexity states according to the 1RSB free energy (see Sec. 3 of the SM for further details). As we will see in the following, both these thresholds can be lowered by employing a continuation strategy for the regularizer.

A thorough description of the results of the replica analysis and of the landscape properties for GLE models will be presented in a more technical future work.

7. MAP estimation with an L_2 regularizer

The objective function introduced in Eq. (2) contains a L_2 regularization term weighted by an intensity parameter λ.

Regularization plays and important role in reducing the variance of the inferred estimator, and can be crucial when the observations are noise-affected, since it lowers the sensitivity of the learned model to deviations in the training set. However, as observed in (Ma et al., 2018; 2019; Balan, 2016), regularization is also useful for its smoothing effect, and can be exploited in non-convex optimization problems even in the noiseless setting. When the regularization term is turned up, the optimization landscape gradually simplifies and it becomes easier to reach a global optimizer. However, the problem of getting stuck in bad local minima is avoided at the cost of introducing a bias. The *continuation* strategy is based on the fact that such biased estimator might be closer than the random initial configuration to the global optimizer of the unregularized objective : in a multi-stage approach, regularization is decreased (down to zero) after each warm restart.

Among the many possible continuation schedules for λ (a little decrease after each minimization, or, as in AMP.A, at the end of each iteration) in this paper we choose a simple two-stage approach: first we run GASP till convergence with a given value of $\lambda > 0$, then we set $\lambda = 0$ in the
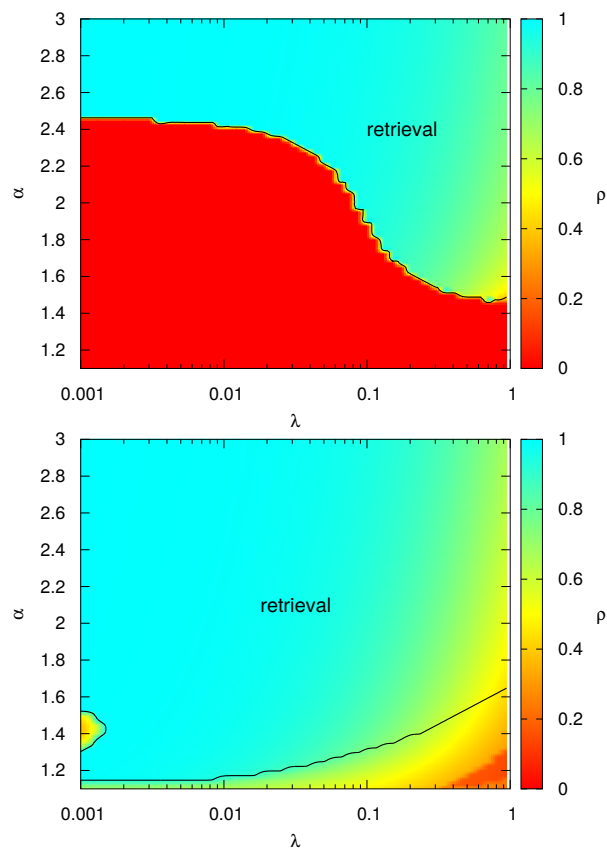
successive iterations.

Figure 2. Phase diagrams corresponding to the SE asymptotic analysis of GAMP (top) and GASP (bottom). The color maps indicate the overlap ρ reached at convergence in the presence of an L_2 regularizer of intensity λ.

In Fig.2, we compare the asymptotic performance (tracked by SE) of GAMP and GASP for the phase retrieval problem with an L_2 regularization. The color map indicates the overlap with the signal reached at the end of the first stage of our continuation strategy (with $\lambda \neq 0$), while the black curves delimit the perfect retrieval regions, where the overlap reached at the end of stage two (with $\lambda = 0$) is $\rho = 1$.

In both cases we set the initial variances Δ to 1, and consider an initial condition with a small positive overlap with the signal, $\rho = 0.1$. An assumption of this kind is indeed needed to ensure that we avoid lingering on the fixed-point at $\rho = 0$; however, the specific value of ρ can be chosen arbitrarily (e.g., it could be taken much smaller without affecting the phase diagrams). Even in real-world applications, it is often the case that the non-orthogonality requirement is easily met, for example in many imaging applications the signal is known to be real non-negative. As explained in the previous section, we also set $q_0 = 1 + \rho^2$ in the initialization of the self-overlap parameter.

In the GASP phase diagram, for each α and λ, the value of m was set to the thermodynamic optimum value $m^\star$ (obtained at $\rho = 0$), and was kept fixed throughout the two stages of our continuation strategy. This $m^\star$ can be obtained by optimizing the 1RSB free energy over the symmetry-breaking parameter; the numerical values of m, corresponding to the points in the plot, can be found in Sec. 3 of the SM, in Fig. 1. It is not strictly necessary to fix m to this specific value, as any value in a broad range of around $m^\star$ will still be effective (see for example Fig. 2 in the SM). As expected from the numerical experiments at $\lambda = 0$, we can see from Fig. 2 that when the regularizer becomes too small an uninformative fixed-point (in $\rho = 0$) becomes attractive for the dynamics of GASP and signal recovery becomes impossible below $\alpha_{alg} \sim 1.5$ (we expect also the recovery region with $\alpha \in [1.13, 1.3]$ at $\lambda = 0.001$ to shrink and close when the regularizer is further decreased).

It is clear that the introduction of an L_2-norm is crucial for reducing the algorithmic gap of both GAMP and GASP (the information theoretic threshold is $\alpha_{IT} = 1$), as previously observed in (Ma et al., 2018; 2019). In this work we find that also in GLE problems, when the mismatched setting is considered (and inference happens off the Nishimori line (Nishimori, 2001; Antenucci et al., 2019b)), the more fitting geometrical picture provided by the 1RSB ansatz can be exploited algorithmically: with a simple continuation strategy it is possible to lower the algorithmic threshold of GASP down to the Bayes-optimal value $\alpha = 1.13$.

8. Discussion

We presented Generalized Approximate Survey Propagation, a novel algorithm designed to improve over AMP in the context of GLE inference problems, when faced with a mismatch between assumed and true generative model. The algorithm, parametrized by the symmetry-breaking parameter m, allows one to go beyond some symmetry assumptions at the heart of the previous algorithms, and proves to be more suited for the MAP estimation task considered in this work.

In the prototypical case of real-valued phase retrieval, we have shown that with little tuning of m it is possible to modify the effective landscape explored during the message-passing dynamics and avoid getting stuck in otherwise attractive uninformative fixed points. Furthermore, we have seen that, even in the noiseless case, a simple continuation strategy, based on the introduction of an L_2 regularizer, can guide GASP close enough to the signal and allow its recovery, extending the region of parameters where GASP is more effective than GAMP. In some cases we observed that GASP can achieve perfect retrieval until the Bayes-optimal threshold, at the sampling ratio $\alpha \sim 1.13$. We also derived the 1RSB State Evolution equations, and showed that they can be used as a simple tool for tracking the asymptotic

behaviour of GASP.

We delay a comprehensive analysis of the landscape associated to GLE models to a more technical publication, where we will also deal with the case of noisy observation channels. A straightforward follow-up of the present work could focus on the search for an adaptation scheme for the L_2 regularizer, possibly extending the work of Refs. (Ma et al., 2018; 2019), and more importantly, for a criterion to identify the best setting for the symmetry-breaking parameter. Another possible future line of work could go in the direction of relaxing some of the assumptions made in deriving the GASP algorithm over the observation matrix. This could motivate the derivation of a 1RSB version of the Vector Approximate Message Passing equations (Schniter et al., 2016). Also, the extension of GASP to deep non-linear inference model, along the lines of Ref. (Manoel et al., 2017; Fletcher et al., 2018) seems to be promising and technically feasible.

References

Advani, M. and Ganguli, S. Statistical mechanics of optimal convex inference in high dimensions. *Physical Review X*, 6(3):031034, 2016.

Antenucci, F., Franz, S., Urbani, P., and Zdeborová, L. Glassy nature of the hard phase in inference problems. *Physical Review X*, 9(1):011020, 2019a.

Antenucci, F., Krzakala, F., Urbani, P., and Zdeborová, L. Approximate survey propagation for statistical inference. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(2):023401, 2019b.

Balan, R. Reconstruction of signals from magnitudes of redundant representations: The comple case. *Foundations of Computational Mathematics*, 16(3):677–721, 2016.

Barbier, J., Krzakala, F., Macris, N., Miolane, L., and Zdeborová, L. Optimal errors and phase transitions in high-dimensional generalized linear models. In *Conference On Learning Theory*, pp. 728–731, 2018.

Bayati, M. and Montanari, A. The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Transactions on Information Theory*, 57 (2):764–785, 2011.

Braunstein, A., Mézard, M., and Zecchina, R. Survey propagation: An algorithm for satisfiability. *Random Structures & Algorithms*, 27(2):201–226, 2005.

Caltagirone, F., Zdeborová, L., and Krzakala, F. On convergence of approximate message passing. In *Information*

Theory (ISIT), 2014 IEEE International Symposium on, pp. 1812–1816. IEEE, 2014.

Candes, E. J., Li, X., and Soltanolkotabi, M. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015.

Chen, Y. and Candes, E. Solving random quadratic systems of equations is nearly as easy as solving linear systems. In *Advances in Neural Information Processing Systems*, pp. 739–747, 2015.

Chen, Y., Chi, Y., Fan, J., and Ma, C. Gradient descent with random initialization: Fast global convergence for nonconvex phase retrieval. *Mathematical Programming*, pp. 1–33, 2018.

Deshpande, Y. and Montanari, A. Information-theoretically optimal sparse pca. In *2014 IEEE International Symposium on Information Theory*, pp. 2197–2201. IEEE, 2014.

Deshpande, Y., Abbe, E., and Montanari, A. Asymptotic mutual information for the binary stochastic block model. In *Information Theory (ISIT), 2016 IEEE International Symposium on*, pp. 185–189. IEEE, 2016.

Dhifallah, O. and Lu, Y. M. Fundamental limits of phasemax for phase retrieval: A replica analysis. In *Computational Advances in Multi-Sensor Adaptive Processing (CAM-SAP), 2017 IEEE 7th International Workshop on*, pp. 1–5. IEEE, 2017.

Donoho, D. and Montanari, A. High dimensional robust m-estimation: Asymptotic variance via approximate message passing. *Probability Theory and Related Fields*, 166 (3-4):935–969, 2016.

Donoho, D. L., Maleki, A., and Montanari, A. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.

Fletcher, A. K., Rangan, S., and Schniter, P. Inference in deep networks in high dimensions. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pp. 1884–1888. IEEE, 2018.

Goldstein, T. and Studer, C. Phasemax: Convex phase retrieval via basis pursuit. *IEEE Transactions on Information Theory*, 2018.

Jain, P., Netrapalli, P., and Sanghavi, S. Low-rank matrix completion using alternating minimization. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pp. 665–674. ACM, 2013.

Javanmard, A. and Montanari, A. State evolution for general approximate message passing algorithms, with applications to spatial coupling. *Information and Inference: A Journal of the IMA*, 2(2):115–144, 2013.

Krzakała, F., Montanari, A., Ricci-Tersenghi, F., Semerjian, G., and Zdeborová, L. Gibbs states and the set of solutions of random constraint satisfaction problems. *Proceedings of the National Academy of Sciences*, 104(25):10318–10323, 2007.

Krzakala, F., Ricci-Tersenghi, F., Zdeborova, L., Zecchina, R., Tramel, E. W., and Cugliandolo, L. F. *Statistical Physics, Optimization, Inference, and Message-Passing Algorithms: Lecture Notes of the Les Houches School of Physics-Special Issue, October 2013*. Oxford University Press, 2016.

Lu, Y. M. and Li, G. Phase transitions of spectral initialization for high-dimensional nonconvex estimation. *arXiv preprint arXiv:1702.06435*, 2017.

Luo, W., Alghamdi, W., and Lu, Y. M. Optimal spectral initialization for signal recovery with applications to phase retrieval. *IEEE Transactions on Signal Processing*, 2019.

Ma, J., Xu, J., and Maleki, A. Approximate message passing for amplitude based optimization. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3365–3374, Stockholmsmssan, Stockholm Sweden, 10–15 Jul 2018. PMLR. URL http://proceedings.mlr.press/v80/ma18e.html.

Ma, J., Xu, J., and Maleki, A. Optimization-based amp for phase retrieval: The impact of initialization and l2-regularization. *IEEE Transactions on Information Theory*, 2019.

Manoel, A., Krzakala, F., Mézard, M., and Zdeborová, L. Multi-layer generalized linear estimation. In *Information Theory (ISIT), 2017 IEEE International Symposium on*, pp. 2098–2102. IEEE, 2017.

Mézard, M., Parisi, G., and Virasoro, M. *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*, volume 9. World Scientific Publishing Company, 1987.

Mézard, M., Parisi, G., and Zecchina, R. Analytic and algorithmic solution of random satisfiability problems. *Science*, 297(5582):812–815, 2002.

Monasson, R. Structural glass transition and the entropy of the metastable states. *Physical review letters*, 75(15): 2847, 1995.

Mondelli, M. and Montanari, A. Fundamental limits of weak recovery with applications to phase retrieval. *Foundations of Computational Mathematics*, pp. 1–71, 2018.

Mukherjee, S. and Seelamantula, C. S. Phase retrieval from binary measurements. In *IEEE Signal Processing Letters*, volume 25, pp. 348–352. IEEE, 2018.

Nishimori, H. *Statistical physics of spin glasses and information processing: an introduction*, volume 111. Clarendon Press, 2001.

Rangan, S. Generalized approximate message passing for estimation with random linear mixing. In *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*, pp. 2168–2172. IEEE, 2011.

Rangan, S. and Fletcher, A. K. Iterative estimation of constrained rank-one matrices in noise. In *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*, pp. 1246–1250. IEEE, 2012.

Ricci-Tersenghi, F., Semerjian, G., and Zdeborová, L. Typology of phase transitions in bayesian inference problems. *Physical Review E*, 99(4):042109, 2019.

Ros, V., Arous, G. B., Biroli, G., and Cammarota, C. Complex energy landscapes in spiked-tensor and simple glassy models: Ruggedness, arrangements of local minima, and phase transitions. *Physical Review X*, 9(1):011003, 2019.

Schniter, P. and Rangan, S. Compressive phase retrieval via generalized approximate message passing. *IEEE Transactions on Signal Processing*, 63(4):1043–1055, 2015.

Schniter, P., Rangan, S., and Fletcher, A. K. Vector approximate message passing for the generalized linear model. In *Signals, Systems and Computers, 2016 50th Asilomar Conference on*, pp. 1525–1529. IEEE, 2016.

Sun, J., Qu, Q., and Wright, J. A geometric analysis of phase retrieval. *Foundations of Computational Mathematics*, 18 (5):1131–1198, 2018.