Chapter 18

The cavity method: from exact solutions to algorithms

Alfredo Braunstein[1], Guilhem Semerjian[2]

[1] Politecnico di Torino, Corso Duca degli Abruzzi, 24, I-10129, Torino, Italy, Italian Institute for Genomic Medicine, IRCCS Candiolo, SP-142, I-10060, Candiolo (TO), Italy and INFN, Sezione di Torino, Italy

[2] Laboratoire de Physique de l'École Normale Supérieure, ENS, Université PSL, CNRS, Sorbonne Université, Université Paris Cité, F-75005 Paris, France

18.1   Introduction

The quest of an analytic solution for the simplest mean-field spin-glass model (the Sherrington-Kirkpatrick (SK) one [Sherrington and Kirkpatrick (1975)]) led Giorgio Parisi to the invention of the replica method [Parisi (1980)]. This method is able to describe and handle the complicated structure of the configuration space of the SK model, with a hierarchical division of the configurations into nested pure states, through the analytical parametrization of matrices of size $n \times n$, in the limit where $n \to 0$, which is, to say the least, a questionable mathematical construction (its predictions have been nevertheless confirmed rigorously later on [Guerra and Toninelli (2002); Talagrand (2006); Panchenko (2013)]). In the physics literature an alternative method to solve the SK model was proposed in [Mézard *et al.* (1986)], and subsequently dubbed the cavity method. In a nutshell the idea is to consider the effect of the addition of one spin in a large SK model, or equivalently to create a "cavity" by isolating one spin and modeling the influence that the rest of the system has on it in a self-consistent way. The replica and the cavity methods yield the same predictions for the SK model,

with complementary insights on its structure, the cavity method bypassing the "analytic continuation" from integer values of n to 0.

Even if the replica and cavity methods have had an impact inside physics, in particular in the context of structural glasses, they have also been very fruitful in fields which at first sight could seem unrelated, and in particular in computer science, information theory and discrete mathematics. Roughly speaking, the reason for their versatility lies in the rather universal character of the structure of the configuration space evoked above, that appears not only in the SK model but in many other problems with a non-physical origin, notably some random constraint satisfaction problems and error correcting codes. It turns out indeed that these problems can be viewed as mean-field spin-glasses, but slightly different from the SK one: the degrees of freedom in these problems interact strongly with a finite number of neighbors, whereas in the SK all degrees of freedom interact with each other weakly, in a "fully-connected" manner. The mean-field character of these sparse, or diluted, models arise from the choice of the neighbors, which is done uniformly at random, without the geometrical constraints of an Euclidean space. In physics terms such a network of interaction is called a Bethe lattice, in mathematics a random graph. This type of model appeared in the physics literature relatively shortly after the fully-connected ones [Viana and Bray (1985)], but it became quickly clear that they were much more challenging to solve, some simplifications of the diverging connectivity (of a central limit theorem flavor) being absent in this case. A line of research extended the replica method to this sparse setting, see in particular [Monasson (1998); Biroli *et al.* (2000)] and references therein, at the price of a rather complicated order parameter. It turned out that the cavity method is a more convenient framework than the replica one for these problems, the complex configuration space encoded by the replica symmetry breaking being formulated in a more transparent manner through the cavity approach, as first discussed in [Mézard and Parisi (2001)]; in addition the formalism of the cavity method can be used to develop algorithms that provide informations on a single sample of mean-field spin-glasses, not only on average thermodynamic quantities.

The goal of this chapter is to review the main ideas that underlie the cavity method for models defined on random graphs, as well as present some of its outcomes, focusing on the random constraint satisfaction problems for which it provided both a better understanding of the phase transitions they undergo, and suggestions for the development of algorithms to solve them. It is organized as follows; section 18.2 focuses on the analytic aspects

of the method. It contains an introduction to models defined on random graphs (in Sec. 18.2.1), then the equations of the cavity method at the so-called replica symmetric (RS) level and one step of replica symmetry breaking (1RSB) are presented in Sec. 18.2.2 and 18.2.3, before reviewing in Sec. 18.2.4 their outcomes concerning the phase diagram of random constraint satisfaction problems. Algorithmic consequences of this approach are detailed in Sec. 18.3.

18.2   The cavity method for sparse mean-field models

18.2.1   *Models on random graphs*

We shall consider systems made of N elementary degrees of freedom (spins) σ_i, which take values in some finite alphabet χ, and whose global configuration will be denoted $\underline{\sigma} = (\sigma_1, \ldots, \sigma_N) \in \chi^N$. They interact through an energy function (also called Hamiltonian, or cost function), that we decompose as

$$E(\underline{\sigma}) = \sum_{a=1}^{M} \varepsilon_a(\underline{\sigma}_{\partial a}) \ , \tag{18.1}$$

where the sum runs over the M basic interactions terms ε_a. We denote $\partial a \subset \{1, \ldots, N\}$ the set of variables involved in the a'th constraint, and for a subset S of the variables $\underline{\sigma}_S$ means $\{\sigma_i | i \in S\}$. In what follows we assume that all interactions involves a subset of k variables, for a given $k \geq 2$. This framework encompasses usual Ising spin-glass models, with $\chi = \{-1, 1\}$, $k = 2$ and $\varepsilon_a(\underline{\sigma}_{\partial a}) = -J_a \sigma_{i_a} \sigma_{j_a}$, J_a being the coupling constant between the spins i_a and j_a. It also allows to deal with Potts spins when $\chi = \{1, \ldots, q\}$ for a number $q \geq 2$ of spin states, also interpreted as colors; in this case a relevant energy function corresponds to pairwise interactions ($k = 2$), with $\varepsilon_a(\underline{\sigma}_{\partial a}) = \delta_{\sigma_{i_a}, \sigma_{j_a}}$. This yields the Hamiltonian of the Potts antiferromagnetic model, corresponding in the perspective of computer science to the q-coloring problem, the cost function counting the number of monochromatic edges among the interacting ones. More generically a constraint satisfaction problem (CSP) corresponds to a cost function of the form (18.1) with ε_a taking values 0 or 1, and being interpreted as the indicator function of the event "the a-th constraint is not satisfied by the configuration of the variables in $\underline{\sigma}_{\partial a}$". In particular the k-SAT and k-XORSAT problems can be described in this way with Ising spins and k-wise interactions. One calls solution of a CSP a configuration $\underline{\sigma}$ satisfying

simultaneously all the constraints, i.e. a zero-energy groundstate, and one says that the CSP is satisfiable if and only if it admits at least one solution.

The Gibbs-Boltzmann probability measure associated to this Hamiltonian for an inverse temperature β reads

$$\mu(\underline{\sigma}) = \frac{1}{Z} \prod_{a=1}^{M} w_a(\underline{\sigma}_{\partial a}) \ , \quad Z = \sum_{\underline{\sigma} \in \mathcal{X}^N} \prod_{a=1}^{M} w_a(\underline{\sigma}_{\partial a}) \ , \quad \Phi = \frac{1}{N} \ln Z \ . \quad (18.2)$$

where the partition function Z ensures the normalization of the probability law, and $w_a(\underline{\sigma}_{\partial a}) = e^{-\beta \varepsilon_a(\underline{\sigma}_{\partial a})}$. We introduced the thermodynamic potential Φ which we shall call a free-entropy, as we did not include the constant $-1/\beta$ that would make it a free-energy. This choice allows to handle the uniform measure over the solutions of a CSP (assumed to be satisfiable), that corresponds to $w_a(\underline{\sigma}_{\partial a}) = (1 - \varepsilon_a(\underline{\sigma}_{\partial a}))$, in which case Z counts the number of solutions and Φ is the associated entropy rate. It amounts to set formally $\beta = \infty$ in the Gibbs-Boltzmann definition, in other words to work directly at zero temperature.

A convenient representation of a probability measure μ of the form (18.2) is provided by a factor graph [Kschischang *et al.* (2001)], see Fig. 18.1 for an example, which is a bipartite graph where each of the N variables σ_i is represented by a circle vertex, while the M weight functions w_a are associated to square vertices. An edge is drawn between a variable i and an interaction a if and only if w_a actually depends on σ_i, i.e. $i \in \partial a$. In a similar way we shall denote ∂i the set of interactions in which σ_i appears, i.e. the graphical neighborhood of i in the factor graph, and call $|\partial i|$ the degree of the i-th variable. One has a natural notion of graph distance between two variable nodes i and j, defined as the minimal number of interaction nodes on a path linking i and j.

Our interest lies in disordered systems, in which the probability measure μ is itself a random object. Suppose indeed that the weight functions w_a are built by drawing, independently for each a, the k-uplet of variables ∂a uniformly at random among the $\binom{N}{k}$ possible choices (and also the coupling constants defining the interaction if necessary). We will denote $\mathbb{E}[\bullet]$ the average with respect to this quenched randomness (let us emphasize that there are two distinct level of probabilities in these systems: the spins $\underline{\sigma}$ are random variables with the probability law μ, and μ is random because of the stochastic choices in the construction of the factor graph). For $k = 2$ the resulting factor graph is drawn from nothing but the celebrated Erdős-Rényi $G(N, M)$ random graph ensemble, the case $k > 2$ corresponding to its natural hypergraph generalization. The large size (thermodynamic)
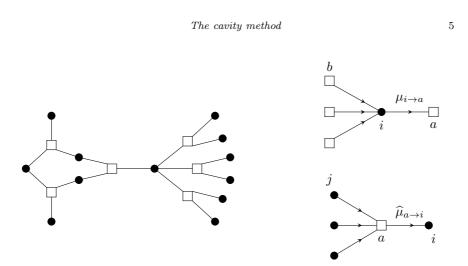
Fig. 18.1   Left: an example of a factor graph. Right: illustration of Eqs. (18.4,18.5).

limit we shall consider corresponds to $N, M \to \infty$, with $\alpha = M/N$ a fixed parameter. Let us recall some elementary properties of these random factor graphs in this limit:

- the probability that a randomly chosen variable i has degree $|\partial i| = d$ is $q_d = e^{-\alpha k}(\alpha k)^d/d!$, the Poisson law of mean αk.
- if one chooses randomly an interaction a, then a variable $i \in \partial a$, the probability that i appears in d interactions *besides* a, i.e. that $|\partial i \setminus a| = d$, is $\widetilde{q}_d = e^{-\alpha k}(\alpha k)^d/d!$.
- the random factor graphs are locally tree-like: choosing at random a vertex i, the subgraph made of all nodes at graph distance from i smaller than some threshold t is, with a probability going to 1 in the thermodynamic limit with t fixed, a tree.

More general ensembles of random factor graphs can be constructed, by fixing a degree distribution q_d and drawing at random from the set of all graphs of size N with Nq_0 isolated vertices, Nq_1 vertices of degree 1, and so on and so forth. Then the two distributions q_d and $\widetilde{q}_d$ are different in general, and related through $\widetilde{q}_d = (d+1)q_{d+1}/\sum_{d'} d' q_{d'}$. An important example in this class corresponds to random regular graphs, where q_d is supported by a single integer.

18.2.2   *The replica symmetric (RS) cavity method*

The goal of the cavity method is to describe the properties of the random measure μ constructed above, for typical samples of the random graph ensemble. The free-entropy Φ is self-averaging in the thermodynamic limit, its typical value concentrates around its average, the quenched free-entropy ϕ defined as

$$\phi = \lim_{N \to \infty} \mathbb{E}[\Phi] = \lim_{N \to \infty} \frac{1}{N} \mathbb{E}[\ln Z] \ . \tag{18.3}$$

The computation of this quantity is thus the objective of the cavity method, along with a local description of the measure μ, in terms of its marginal distributions on a finite number of spins.

The cavity method relies crucially on the local convergence of random factor graph models to random trees explained at the end of Sec. 18.2.1. Let us assume momentarily that the factor graph representing the model under study is a finite tree. Then the problem of characterizing the measure (18.2) and computing the associated partition function Z can be solved exactly in a simple, recursive way: one can break the tree into independent subtrees, solve the problems on these substructures, and combine them together to get the solution on the larger problem. This is nothing but a generalization of the transfer matrix method used in physics to solve unidimensional problems, a form of what is known as dynamic programming in computer science. More precisely, for each edge between a variable i and an adjacent interaction a one introduces two directed "messages", $\mu_{i \to a}$ and $\widehat{\mu}_{a \to i}$, which are probability measures on the alphabet χ, that would be the marginal probability of σ_i if, respectively, the interaction a were removed from the graph, or if all interactions around i except a were removed. A moment of thought reveals that these messages obey the following recursive (so-called Belief Propagation (BP)) equations (see the right part of Fig. 18.1 for an illustration),

$$\mu_{i \to a}(\sigma_i) = \frac{1}{z_{i \to a}} \prod_{b \in \partial i \setminus a} \widehat{\mu}_{b \to i}(\sigma_i) \ , \tag{18.4}$$

$$\widehat{\mu}_{a \to i}(\sigma_i) = \frac{1}{\widehat{z}_{a \to i}} \sum_{\underline{\sigma}_{\partial a \setminus i}} w_a(\underline{\sigma}_{\partial a}) \prod_{j \in \partial a \setminus i} \mu_{j \to a}(\sigma_j) \ , \tag{18.5}$$

with $z_{i \to a}$ and $\widehat{z}_{a \to i}$ ensuring the normalization of the laws. On a tree factor graph there exists a single solution of these equations, which is easily determined starting from the leaves of the graph (for which the empty product above is conventionally equal to 1) and sweeping towards the inside of the

graph. Once the messages have been determined all local averages with respect to μ can be computed, as well as the partition function, in terms of the solutions of these BP equations. The Belief Propagation algorithm consists in looking for a fixed-point solution of (18.4,18.5), iteratively, even if the factor graph is not a tree; in this case the formula giving Φ in terms of the messages is only an approximation, known as the Bethe formula for the free-entropy (see for instance [Yedidia *et al.* (2003)] for more details on the connections between the stationary points of the Bethe free-entropy and the solutions of the Belief Propagation equations). These equations were discovered independently in Statistical Physics as the Bethe-Peierls approximation, in artificial intelligence as the Belief Propagation algorithm, and in Information Theory as the Sum-Product algorithm [Mézard and Montanari (2009)].

Of course random graphs are only locally tree-like, they do possess loops, even if their lengths typically diverge in the thermodynamic limit. The cavity method amounts thus to a series of prescriptions to handle these long loops and to describe the boundary condition they impose on the local tree neighborhoods inside a large random graph. The simplest prescription, that goes under the name of replica symmetric (RS) and that is valid for weakly interacting models (i.e. small α and/or large temperature), assumes some spatial correlation decay properties of the probability measure μ. When one removes an interaction a from a factor graph the variables around it becomes strictly independent if one starts from a tree, and asymptotically independent provided only long enough loops join them in absence of a, and provided the correlation decays fast enough along these loops. To compute the average thermodynamic potential (18.3) it is enough in this case to study the statistics with respect to the quenched disorder of the messages $\mu_{i\to a}$, $\widehat{\mu}_{a\to i}$ on the edges of the random factor graph. In other words the order parameter of the RS cavity method is the law of the random variables η, $\widehat{\eta}$, which are equal to the random messages one obtains by drawing at random a sample, solving the BP equations on it, choosing at random an edge $a-i$, and observing the value of $\mu_{i\to a}$ and $\widehat{\mu}_{a\to i}$. With the assumption of independence underlying the RS cavity method the equations (18.4,18.5) translate into Recursive Distributional Equations (RDE) of the form:

$$\eta \stackrel{\mathrm{d}}{=} f(\widehat{\eta}_1,\ldots,\widehat{\eta}_d) \ , \qquad \widehat{\eta} \stackrel{\mathrm{d}}{=} \widehat{f}(\eta_1,\ldots,\eta_{k-1}) \ . \qquad (18.6)$$

In this equation all the η_i's and $\widehat{\eta}_i$'s are independent copies of the random variables η and $\widehat{\eta}$, $\stackrel{\mathrm{d}}{=}$ denotes the equality in distribution between random variables, d is drawn according to the law $\widetilde{q}_d$, and the functions f and $\widehat{f}$ are

defined by the right hand sides of equations (18.4,18.5) (with possibly an additional random draw of the weight w). The RS prediction for ϕ can then be expressed as the average over random copies of η and $\widehat{\eta}$ of the local free-entropy contributions obtained from the exact computation of the partition function of a finite tree. Note that the equation (18.6), if it has in general no analytic solution, lends itself to a very natural numerical resolution where the law of η is approximately represented as an empirical distribution over a set of representatives η (a population representation) [Abou-Chacra *et al.* (1973); Mézard and Parisi (2001)].

The exactness of the predictions of the RS cavity method has been proven rigorously for some models which are not too frustrated (e.g. ferromagnetic systems, or matching models), see for instance [Dembo and Montanari (2010); Bordenave and Lelarge (2010); Bordenave *et al.* (2012)]. But in general the correlation decay assumption fails, in this case one has to turn to a more sophisticated version of the cavity method, which will be introduced in the next section.

18.2.3   *Handling the replica symmetry breaking (RSB) with the cavity method*

As a matter of fact for low enough temperature, and high enough density of interactions α, the configuration space of frustrated random models gets fractured in a large number of pure states (or clusters), and the correlation decay hypothesis only holds for the Gibbs measure restricted to one pure state, not for the complete Gibbs measure. In the replica method this phenomenon shows up as a breaking of the equivalence between different replicas, we will now explain how the cavity method is able to handle this structure of the configuration space. It amounts to make further self-consistent hypotheses on the correlated boundary conditions this induces on the tree-like portions of the factor graph. Inside each pure state the RS computation is assumed to hold true, and the RSB computation is then a study of the statistics of the pure states. Let us explain how this is done in practice at the first level of RSB (1RSB cavity method). The partition function is written as a sum over the pure states γ, that form a partition of the configuration space, $Z = \sum_\gamma Z_\gamma$, where Z_γ is the partition function restricted to the pure state γ. It can be written in the thermodynamic limit as $Z_\gamma = e^{N f_\gamma}$, with f_γ the internal free-entropy density of a given pure state. One further assumes that the number of pure states with a given value of f is, at the leading exponential order, $e^{N\Sigma(f)}$, with the so-called

configuration entropy, or complexity, Σ a concave function of f, positive on the interval $[f_{\min}, f_{\max}]$. In order to compute Σ one introduces a parameter m (called Parisi breaking parameter) conjugated to the internal thermodynamic potential, and the generating function of the Z_γ as $\mathcal{Z}(m) = \sum_\gamma Z_\gamma^m$. In the thermodynamic limit its dominant behavior is captured by the 1RSB potential $\phi_{1\mathrm{RSB}}(m)$,

$$\phi_{1\mathrm{RSB}}(m) = \lim_{N\to\infty} \frac{1}{N} \log \mathcal{Z}(m) = \sup_f \left[\Sigma(f) + mf\right] , \qquad (18.7)$$

where the last expression is obtained by a saddle-point evaluation of the sum over γ. The complexity function is then accessible via the inverse Legendre transform of $\phi_{1\mathrm{RSB}}(m)$ [Monasson (1995)], or in a parametric form

$$f(m) = \phi'_{1\mathrm{RSB}}(m) , \qquad \Sigma(f(m)) = \phi_{1\mathrm{RSB}}(m) - m\phi'_{1\mathrm{RSB}}(m) , \quad (18.8)$$

where $f(m)$ denotes the point where the supremum is reached in Eq. (18.7). One has $\Sigma'(f(m)) = -m$, i.e. the introduction of the parameter m allows to explore the complexity curve by tuning the tangent slope of the selected point.

The actual computation of $\phi_{1\mathrm{RSB}}(m)$ is done as follows [Mézard and Parisi (2001)]. One introduces on each edge of the factor graph two distributions $P_{i\to a}$ and $\widehat{P}_{i\to a}$ of messages, which are the probability over the different pure states γ, weighted proportionally to Z_γ^m, to observe a given value of $\mu_{i\to a}^\gamma$ and $\widehat{\mu}_{a\to i}^\gamma$ respectively, where $\mu_{i\to a}^\gamma$ and $\widehat{\mu}_{a\to i}^\gamma$ are the messages that appear in Eq. (18.4,18.5), for the measure restricted to the pure state γ. Because $P_{i\to a}$ and $\widehat{P}_{a\to i}$ are themselves random objects with respect to the choices in the generation of the instance of the factor graph, the order parameter of the 1RSB cavity method becomes the distributions of $P_{i\to a}$ and $\widehat{P}_{a\to i}$ with respect to the disorder. The latter is solution of a self-consistent functional equation written as

$$P \stackrel{\mathrm{d}}{=} F(\widehat{P}_1, \ldots, \widehat{P}_d) , \qquad \widehat{P} \stackrel{\mathrm{d}}{=} \widehat{F}(P_1, \ldots, P_{k-1}) , \qquad (18.9)$$

that parallels the equation (18.6) of the RS cavity method, with again independent copies of the distributions P_i and $\widehat{P}_i$. The right hand sides of these distributional equalities stand for:

$$P(\eta) = \frac{1}{Z} \int \prod_{i=1}^d \mathrm{d}\widehat{P}_i(\widehat{\eta}_i) \; \delta(\eta - f(\{\widehat{\eta}_i\})) \; z(\{\widehat{\eta}_i\})^m , \qquad (18.10)$$

$$\widehat{P}(\eta) = \frac{1}{\widehat{Z}} \int \prod_{i=1}^{k-1} \mathrm{d}P_i(\eta_i) \; \delta(\widehat{\eta} - \widehat{f}(\{\eta_i\})) \; \widehat{z}(\{\eta_i\})^m , \qquad (18.11)$$

with the functions f and $\widehat{f}$ corresponding to the recursion functions at the RS level, see Eq. (18.4,18.5), and z and $\widehat{z}$ the associated normalization factors. From the solution of this equation (that again can be found numerically with the population dynamics method [Mézard and Parisi (2001)]) one computes the 1RSB potential $\phi_{1\text{RSB}}(m)$ via an expression similar to the one giving the expression of ϕ at the RS level, with now averages over random distributions P and $\widehat{P}$.

There are different justifications for the appearance of the "reweighting factors" z^m and $\widehat{z}^m$ in Eqs. (18.10,18.11). The argument in [Mézard and Parisi (2001)] is based on the exponential distribution of the free-entropies Nf_γ of the pure states with respect to some reference value, and on consistency requirements on the evolution of the pure states when the cavity factor graph is modified. One can also study the statistics of the many fixed point solutions of the BP equations (18.4,18.5) and devise a dual factor graph for the counting of these fixed points [Mézard and Montanari (2009)], the reweighting factor allowing to select the fixed points associated to some internal free-entropy. Another interpretation was proposed in [Krzakala *et al.* (2007)], associating the pure states of a large but finite factor graph model to boundary conditions on trees. This interpretation is particularly relevant in the case $m = 1$, for which these boundary conditions are actually drawn from the Gibbs measure itself, and reveals a deep connection between the 1RSB cavity method and the reconstruction on tree problem, as first unveiled in [Mézard and Montanari (2006)], and with the point-to-set correlations of the Gibbs measure [Montanari and Semerjian (2006)].

This construction can be generalized to higher levels of replica symmetry breaking [Parisi (1980)], with a hierarchical partition of the configuration space into nested pure spaces; the resulting equations for models on sparse random graphs involve a recursive tower of probability distributions over probability distributions, whose numerical resolution becomes extremely challenging beyond 1RSB.

18.2.4   *Some analytic outcomes of the cavity method*

As presented above the cavity method is quite versatile, in the sense that it can address a variety of models defined on random graphs, and it has indeed been applied to several different problems. As an illustration of some of its outcomes we shall now present some results it has provided on the phase diagram of random constraint satisfaction problems (see also chapter 31),

and sketch the connections between this qualitative understanding and the quantitative formalism we have introduced before.

In the case of a constraint satisfaction problem the cost function defined in Eq. (18.1) is made of a sum of indicator functions of events that the a-th constraint is unsatisfied, for instance the number of monochromatic edges in the q-coloring problem. The natural questions in this context are: does an instance of the problem admit at least one solution? if yes, how are the solutions organized in the configuration space? It turns out that the answers to these questions have drastically different answers depending on the value of the density of constraints α, in other words there exist, in the thermodynamic limit, sharp phase transitions for some threshold values of this parameter.
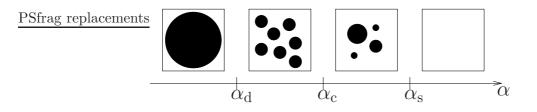
PSfrag replacements

Fig. 18.2   Schematic representation of the phase transitions in a random CSP ensemble.

The main transitions that occur for generic ensembles of random CSPs are represented in a schematic way on Fig. 18.2. The squares represent the full configuration space, for four different values of α (obviously the representation of this N-dimensional hypercube on a two-dimensional drawing is only a cartoon), while the black area stands for the solutions. For $\alpha > \alpha_{\mathrm{s}}$, the satisfiability transition, the square is empty, which translates the absence of solution in typical instances for these density of constraints. The satisfiable regime $\alpha < \alpha_{\mathrm{s}}$ is further divided in three regions, separated by structural phase transitions at which the organization of the set of solutions changes qualitatively. For $\alpha < \alpha_{\mathrm{d}}$, the so-called clustering, or dynamic transition, all solutions are somehow close to each other, while in the rest of the satisfiable regime they are broken in clusters of nearby solutions, each cluster being separated from the other ones. The number and size of the relevant clusters further change at the condensation threshold α_{c}: for $\alpha_{\mathrm{d}} < \alpha < \alpha_{\mathrm{c}}$ most solutions are contained in an exponential number of clusters which have all roughly the same size, while in the regime $\alpha_{\mathrm{c}} < \alpha < \alpha_{\mathrm{s}}$ most solutions are found in a sub-exponential number of

clusters with strongly fluctuating sizes.

These qualitative predictions, along with quantitative numerical values for some definite random CSPs families, have been obtained by the analysis of the solutions of the 1RSB cavity equations, according to the following criteria [Mézard and Zecchina (2002a); Krzakala *et al.* (2007)]:

- α_{d} is the smallest value of α such that the 1RSB equations at $m = 1$ admit a non-trivial solution.
- in the regime $[\alpha_{\mathrm{d}}, \alpha_{\mathrm{c}}]$ the configurational entropy, or complexity, associated to the $m = 1$ solution, is positive, whereas it becomes negative for $\alpha > \alpha_{\mathrm{c}}$.
- the satisfiability transition is marked by the vanishing of the complexity computed at $m = 0$, in the so-called energetic version of the 1RSB cavity method [Mézard and Parisi (2003)], that counts all clusters irrespectively of their sizes.

18.3   Some algorithmic outcomes of the cavity method

18.3.1   *Algorithmic applications of the cavity method*

As mentioned above the equations (18.4)-(18.5) can be used on a single instance to compute (approximately) several properties of the distribution (18.2), including single-site marginals, joint marginals of variables in a common factor, the free energy and Shannon's entropy. This approach has been applied to Bayesian networks, in the decoding phase of communication codes (syndrome-based decoding, Turbo Codes [Benedetto *et al.* (1996)]) and in stereo image reconstruction. More recently, it has found applications in a large variety of fields that we shall now review.

BP applications in notable models   In [Kabashima (2003)], a Belief Propagation algorithm for CDMA decoding has been presented. Interestingly, it shows how BP can be efficiently applied to dense models (i.e. in which constraints involve an extensive number of variables) through an application of the Central Limit Theorem (the basis of a BP derivative called AMP, see Chapter 19), and it is also shown that solutions are also fixed points of the famous Thouless-Anderson-Palmer (TAP) equations [Thouless *et al.* (1977)] while showing superior iterative convergence properties. A similar approach has been employed in [Braunstein and Zecchina (2006)] for the binary discrete perceptron learning problem.

In [Frey and Dueck (2005)], the Affinity Propagation (AP) algorithm was presented. AP is a BP algorithm for variables with an extensive number of states. The AP algorithm solves approximately a clustering problem which is similar in spirit to K-means, but with the important difference of only relying on a distance matrix instead of the original, possibly high-dimensional, data representation. Auxiliary variables with a large number of states can be employed to locally enforce global constraints such as connectivity, by representing in the variables state the discrete time of an underlying dynamics. BP has been applied to the resulting extended model [Bayati *et al.* (2008)].

The dynamic cavity method [Neri and Bollé (2009)] is an application of BP to study a certain class of out-of equilibrium dynamical models. The method can be understood as an application of BP to an auxiliary model in which a variable consists in a couple of time-dependent quantities: one is a single spin trajectory, the other a local field. Subsequent works showed that a slightly simpler but equivalent representation can be obtained with a pair of spin trajectories. On certain models such as discrete, microscopically irreversible ones (i.e. ones in which a variable can never go back to a visited state, including the Bootstrap percolation model [Altarelli *et al.* (2013)], SI or SIR epidemic models [Altarelli *et al.* (2014)]), single trajectories can be efficiently represented by the transition times. In other cases, some approximations must be employed [Aurell and Mahmoudi (2012)]. A somehow related variant of the cavity method deals with quantum models, the basic degrees of freedom becoming imaginary-time spin trajectories [Bapst *et al.* (2012)].

Exactness of BP on single instances   Some rigorous results have been proven regarding the exactness of BP algorithms. For certain models and sufficiently large temperature, the BP update equation becomes a contractive mapping, guaranteeing the existence and uniqueness of its fixed point and the convergence towards it under iterations thanks to the Banach theorem. Moreover, this condition guarantees exactness in the thermodynamical limit on graphs with large girth [Bayati and Nair (2006)].

On the other side of the spectrum, some exactness results exist in the small temperature limit as well. Equations to analize models explicitly at zero temperature can be devised by taking the $T \to 0$ limit of the BP equations under an an opportune change of variables, resulting in equations for energy-shifts instead of probabilities. These had been known in coding theory as Max-Sum algorithms. Existing proofs of exactness (on some

models) rely on a local optimality condition for BP fixed points. [Bayati *et al.* (2005); Weiss and Freeman (2001a); Gamarnik *et al.* (2012)].

Gaussian BP (GaBP) [Weiss and Freeman (2001b)] is an application of BP for a continuous model with positive definite quadratic potential, i.e. a Multivariate Gaussian. It is shown under certain conditions on the precision matrix that the GaBP equations converge and give the correct estimation of the means (but wrong estimation of the variances in general), effectively solving a linear system iteratively, with convergence properties that make the method competitive. Note that due to the fact that the mode is equal to the mean in a Gaussian distribution, this result can be again thought of as the exactness of the computation of the maximum.

Survey Propagations and the RSB Phase   Survey propagation (SP) is the algorithmic counterpart of the 1RSB cavity method. It has seen its first applications to study the k-SAT [Mézard and Zecchina (2002b); Braunstein *et al.* (2005)] and $q-$coloring [Krzakala *et al.* (2004)] problems in the replica symmetry broken phase. SP can be thought as BP for the combinatorial problem of solutions of a lower order message passing system (typically Max-Sum or some coarsened version of it). Such a hierarchical approach can also be employed to analyze problems that possess explicitely such a nested structure, such as the ones coming from (stochastic) control problems (e.g. the Stochastic Matching problem [Altarelli *et al.* (2011)]).

It should also be noted that BP can be used in the RSB phase of constraint satisfaction problems. In [Braunstein and Zecchina (2006)] BP has been applied successfully to the perceptron learning problem with binary synapses, even in the regime in which it shows a RSB phase. The solution to this conundrum has been clarified in [Baldassi *et al.* (2016)], where it was shown that BP describes an exponentially small portion of the solution space that is still exponentally large and has a non-clustered geometry akin to the dominant region of the solution space in the RS phase.

Decimation and reinforcement.   An algorithm estimating marginal distributions such as BP can be employed for sampling, and in particular to find solutions to a constraint satisfaction problem. The main idea is ancestral sampling, i.e. given an arbitrary permutation of variable indices π, one can estimate the marginal distribution $p\left(x_{\pi_1}\right)$ and sample $x^*_{\pi_1}$ from it, then restrict the solution space to solutions with $x_{\pi_1} = x^*_{\pi_1}$ and reiterate, effectively sampling $x^*_{\pi_i} \sim p\left(x_{\pi_i}|x^*_{\pi_1},\ldots,x^*_{\pi_{i-1}}\right)$ for $i = 1,\ldots,n$.

As $p\left(\underline{x}\right) = \prod_{i=1}^{n} p\left(x_{\pi_i}|x_{\pi_1}, \ldots, x_{\pi_{i-1}}\right)$, this solution provides a fair sample $\underline{x}^*$ if the estimation of the marginals is exact. The analysis of ancestral sampling with BP has been performed in [Montanari *et al.* (2007); Ricci-Tersenghi and Semerjian (2009); Coja-Oghlan (2011); Coja-Oghlan and Pachon-Pinzon (2012)]. When one is merely interested in finding *any* solution to a contraint satisfaction problem, and remembering that marginal estimations are only approximate, it is convenient to iteratively fix the variable that reduces the solution space the *less,* which corresponds to fixing the most polarized variable in the direction of the largest probability of its marginal. This process is called *decimation.* In practice, decimation corresponds to iteratively selecting the variable with the largest local field and applying an infinite external field to it with the same sign (and then making the equations converge again and reiterating). A soft version of decimation, called reinforcement, can also be conceived, in which a field is applied iteratively to all variables with the same sign of their local field and an intensity that is either a constant [Chavas *et al.* (2005)] or proportional to its magnitude [Braunstein and Zecchina (2006); Bayati *et al.* (2008)]. This dynamics slowly drives the system to one with sufficiently large external fields that becomes trivially polarized on one solution. As an additional twist, a backtracking procedure can be implemented on top of decimation, in which variables are occasional freed from their external field when that choice enlarges the solution space sufficiently. This has been implemented for SP, with excellent results [Marino *et al.* (2016)].

18.4   Conclusions

The Cavity method is a powerful and versatile approach to the description of disordered systems, that has been shown so far to provide the exact asymptotic solution for many models. For given (finite) system instances, its algorithmic counterpart has many practical applications, ranging from a statistical description of the Boltzman-Gibbs distribution to the individuation of single solutions of a CSP. Moreover, at variance with more traditional methods for inference such as MCMC sampling, it can provide an analytical description, given implicitly by the solution(s) of the cavity equations. This fact enables many possibilities, such as its recursive application (SP), and a functional expression of statistical features as a function of the disorder parameters (see for instance chapter 21 for a discussion of inverse problems).

Bibliography

Abou-Chacra, R., Thouless, D., and Anderson, P. (1973). A selfconsistent theory of localization, *J. Phys. C* **6**, p. 1734.

Altarelli, F., Braunstein, A., Dall'Asta, L., Lage-Castellanos, A., and Zecchina, R. (2014). Bayesian inference of epidemics on networks via belief propagation, *Physical Review Letters* **112**, 11.

Altarelli, F., Braunstein, A., Dall'Asta, L., and Zecchina, R. (2013). Optimizing spread dynamics on graphs by message passing, *Journal of Statistical Mechanics: Theory and Experiment* **2013**, 9.

Altarelli, F., Braunstein, A., Ramezanpour, A., and Zecchina, R. (2011). Stochastic matching problem, *Physical review letters* **106**, 19, p. 190601.

Aurell, E. and Mahmoudi, H. (2012). Dynamic mean-field and cavity methods for diluted Ising systems, *Phys. Rev. E* **85**, 3, p. 031119, doi:10.1103/PhysRevE.85.031119, `https://link.aps.org/doi/10.1103/PhysRevE.85.031119`.

Baldassi, C., Borgs, C., Chayes, J. T., Ingrosso, A., Lucibello, C., Saglietti, L., and Zecchina, R. (2016). Unreasonable effectiveness of learning neural networks: From accessible states and robust ensembles to basic algorithmic schemes, *Proceedings of the National Academy of Sciences* **113**, 48, pp. E7655–E7662, doi:10.1073/pnas.1608103113, `https://www.pnas.org/doi/abs/10.1073/pnas.1608103113`.

Bapst, V., Foini, L., Krzakala, F., Semerjian, G., and Zamponi, F. (2012). The quantum adiabatic algorithm applied to random optimization problems: The quantum spin glass perspective, *Physics Reports* , 0, pp. –, doi: 10.1016/j.physrep.2012.10.002, `http://www.sciencedirect.com/science/article/pii/S037015731200347X?v=s5`.

Bayati, M., Borgs, C., Braunstein, A., Chayes, J., Ramezanpour, A., and Zecchina, R. (2008). Statistical mechanics of steiner trees, *Physical Review Letters* **101**, 3.

Bayati, M. and Nair, C. (2006). A rigorous proof of the cavity method for counting matchings, *arXiv:cond-mat/0607290* `http://arxiv.org/abs/cond-mat/0607290`, arXiv: cond-mat/0607290.

Bayati, M., Shah, D., and Sharma, M. (2005). Maximum weight matching via

max-product belief propagation, in *Proceedings. International Symposium on Information Theory, 2005. ISIT 2005.*, pp. 1763–1767, doi:10.1109/ISIT.2005.1523648, iSSN: 2157-8117.

Benedetto, S., Montorsi, G., Divsalar, D., and Pollara, F. (1996). Soft-output decoding algorithms in iterative decoding of turbo codes, *The Telecommunications and Data Acquisition Report* .

Biroli, G., Monasson, R., and Weigt, M. (2000). A variational description of the ground state structure in random satisfiability problems, *Eur. Phys. J. B* **14**, p. 551.

Bordenave, C. and Lelarge, M. (2010). Resolvent of large random graphs, *Random Structures & Algorithms* **37**, 3, pp. 332–352.

Bordenave, C., Lelarge, M., and Salez, J. (2012). Matchings on infinite graphs, *Probability Theory and Related Fields* , pp. 1–26doi:10.1007/s00440-012-0453-0.

Braunstein, A., Mézard, M., and Zecchina, R. (2005). Survey propagation: An algorithm for satisfiability, *Random Structures & Algorithms* **27**, 2, pp. 201–226, doi:10.1002/rsa.20057, `https://onlinelibrary.wiley.com/doi/abs/10.1002/rsa.20057`.

Braunstein, A. and Zecchina, R. (2006). Learning by message passing in networks of discrete synapses, *Physical Review Letters* **96**, 3.

Chavas, J., Furtlehner, C., Mézard, M., and Zecchina, R. (2005). Survey-propagation decimation through distributed local computations, *J. Stat. Mech.* **2005**, 11, pp. P11016–P11016, doi:10.1088/1742-5468/2005/11/P11016, `https://doi.org/10.1088/1742-5468/2005/11/p11016`.

Coja-Oghlan, A. (2011). On belief propagation guided decimation for random k-sat, *Proc. 22nd SODA* , p. 957.

Coja-Oghlan, A. and Pachon-Pinzon, A. Y. (2012). The decimation process in random *k*-sat, *SIAM Journal on Discrete Mathematics* **26**, pp. 1471–1509.

Dembo, A. and Montanari, A. (2010). Ising models on locally tree-like graphs, *Ann. Appl. Probab.* **20**, pp. 565–592.

Frey, B. J. and Dueck, D. (2005). Mixture Modeling by Affinity Propagation, in *Advances in Neural Information Processing Systems*, Vol. 18 (MIT Press), `https://proceedings.neurips.cc/paper/2005/hash/327708dd10d68b1361ad3addbaca01f2-Abstract.html`.

Gamarnik, D., Shah, D., and Wei, Y. (2012). Belief Propagation for Min-Cost Network Flow: Convergence and Correctness, *Operations Research* **60**, 2, pp. 410–428, doi:10.1287/opre.1110.1025, `https://pubsonline.informs.org/doi/abs/10.1287/opre.1110.1025`.

Guerra, F. and Toninelli, F. L. (2002). The thermodynamic limit in mean field spin glass models, *Comm. Math. Phys.* **230**, pp. 71–79.

Kabashima, Y. (2003). A CDMA multiuser detection algorithm on the basis of belief propagation, *J. Phys. A: Math. Gen.* **36**, 43, pp. 11111–11121, doi:10.1088/0305-4470/36/43/030, `https://doi.org/10.1088/0305-4470/36/43/030`.

Krzakala, F., Montanari, A., Ricci-Tersenghi, F., Semerjian, G., and Zdeborová, L. (2007). Gibbs states and the set of solutions of random constraint satis-

faction problems, *Proc. National Academy of Sciences* **104**, 25, pp. 10318–10323.

Krzakala, F., Pagnani, A., and Weigt, M. (2004). Threshold values, stability analysis, and high-q asymptotics for the coloring problem on random graphs, *Phys. Rev. E* **70**, 4, p. 046705.

Kschischang, F., Frey, B., and Loeliger, H. (2001). Factor graphs and the sum-product algorithm, *IEEE Transactions on Information Theory* **47**, 2, p. 498.

Marino, R., Parisi, G., and Ricci-Tersenghi, F. (2016). The backtracking survey propagation algorithm for solving random K-SAT problems, *Nat Commun* **7**, 1, p. 12996, doi:10.1038/ncomms12996, `https://www.nature.com/articles/ncomms12996`.

Mézard, M. and Montanari, A. (2006). Reconstruction on trees and spin glass transition, *J. Stat. Phys.* **124**, 6, pp. 1317–1350.

Mézard, M. and Montanari, A. (2009). *Information, Physics and Computation* (Oxford University Press).

Mézard, M. and Parisi, G. (2001). The Bethe lattice spin glass revisited, *Eur. Phys. J. B* **20**, p. 217.

Mézard, M. and Parisi, G. (2003). The cavity method at zero temperature. *J. Stat. Phys.* **111**, 1-2, pp. 1–34.

Mézard, M., Parisi, G., and Virasoro, M. A. (1986). SK model: The replica solution without replicas, *Europhysics Letters (EPL)* **1**, 2, pp. 77–82.

Mézard, M. and Zecchina, R. (2002a). Random k-satisfiability problem: From an analytic solution to an efficient algorithm, *Phys. Rev. E* **66**, 5, p. 056126.

Mézard, M. and Zecchina, R. (2002b). Random \$K\$-satisfiability problem: From an analytic solution to an efficient algorithm, *Phys. Rev. E* **66**, 5, p. 056126, doi:10.1103/PhysRevE.66.056126, `https://link.aps.org/doi/10.1103/PhysRevE.66.056126`.

Monasson, R. (1995). Structural glass transition and the entropy of the metastable states, *Phys. Rev. Lett.* **75**, pp. 2847–2850.

Monasson, R. (1998). Optimization problems and replica symmetry breaking in finite connectivity spin glasses, *J. Phys. A* **31**, 2, pp. 513–529.

Montanari, A., Ricci-Tersenghi, F., and Semerjian, G. (2007). Solving constraint satisfaction problems through belief propagation-guided decimation, `arXiv:0709.1667`, Proceedings of the 45th Allerton Conference.

Montanari, A. and Semerjian, G. (2006). Rigorous inequalities between length and time scales in glassy systems. *J. Stat. Phys.* **125**, 1, pp. 23–54.

Neri, I. and Bollé, D. (2009). The cavity approach to parallel dynamics of Ising spins on a graph, *J. Stat. Mech.* **2009**, 08, p. P08009, doi:10.1088/1742-5468/2009/08/P08009, `https://doi.org/10.1088/1742-5468/2009/08/p08009`.

Panchenko, D. (2013). *The Sherrington-Kirkpatrick Model* (SpringerNew York, NY).

Parisi, G. (1980). A sequence of approximated solutions to the S-K model for spin glasses, *J. of Phys. A* **13**, 4, p. L115.

Ricci-Tersenghi, F. and Semerjian, G. (2009). On the cavity method for deci-

mated random constraint satisfaction problems and the analysis of belief propagation guided decimation algorithms, *Journal of Statistical Mechanics: Theory and Experiment* , p. P09001.

Sherrington, D. and Kirkpatrick, S. (1975). Solvable Model of a Spin-Glass, *Phys. Rev. Lett.* **35**, 26, pp. 1792–1796.

Talagrand, M. (2006). The Parisi formula, *Annals of Mathematics* **163**, p. 221.

Thouless, D., Anderson, P., and Palmer, R. (1977). Solution of'Solvable model of a spin glass', *Philosophical Magazine* **35**, 3, pp. 593–601.

Viana, L. and Bray, A. (1985). Phase diagrams for dilute spin glasses, *J. Phys. C* **18**, 15, pp. 3037–3051.

Weiss, Y. and Freeman, W. (2001a). On the optimality of solutions of the max-product belief-propagation algorithm in arbitrary graphs, *IEEE Transactions on Information Theory* **47**, 2, pp. 736–744, doi:10.1109/18.910585.

Weiss, Y. and Freeman, W. T. (2001b). Correctness of Belief Propagation in Gaussian Graphical Models of Arbitrary Topology, *Neural Computation* **13**, 10, pp. 2173–2200, doi:10.1162/089976601750541769, `https://doi.org/10.1162/089976601750541769`.

Yedidia, J. S., Freeman, W. T., and Weiss, Y. (2003). Understanding belief propagation and its generalisations, in *Exploring Artificial Intelligence in the New Millennium*, p. 239.