P H Y S I C A L   R E V I E W   L E T T E R S

Landscape of Solutions in Constraint Satisfaction Problems

Marc Mézard,[1] Matteo Palassini,[1,2] and Olivier Rivoire[1]

[1]*Laboratoire de Physique Théorique et Modèles Statistiques, Université Paris-Sud, F-91405 Orsay, France*
[2]*Departament de Física Fonamental, Universitat de Barcelona, Diagonal 647, E-08028 Barcelona, Spain*

We present a theoretical framework for characterizing the geometrical properties of the space of solutions in constraint satisfaction problems, together with practical algorithms for studying this structure on particular instances. We apply our method to the coloring problem, for which we obtain the total number of solutions and analyze in detail the distribution of distances between solutions.

Constraint satisfaction problems (CSPs) offer a unified language describing many complex systems. Originally investigated by computer scientists in relation with algorithmic complexity [1], CSPs have recently attracted much interest within the physics community, following the discovery of their close ties with spin-glass theory [2,3]. They are currently used to tackle systems as diverse as, among others, error-correcting codes [4], rigidity models [5], and regulatory genetic networks [6]. The ubiquity of CSPs stems from their general nature: Given a set of N discrete variables subject to M constraints, a CSP consists in deciding whether there are assignments of the variables satisfying all the constraints. Of special interest is the class of NP-complete problems [1], for which no algorithm is known that guarantees to decide the satisfiability of a problem instance in a time polynomial in N. A well-studied example is the q-coloring problem (q-COL): Given a graph with N nodes and M edges connecting certain pairs of nodes, and given q colors, can we color the nodes so that no two connected nodes share a common color?

Much insight into CSPs is gained by focusing on typical instances drawn from an ensemble with a fixed density of constraints $\alpha = M/N$. As α is varied, a threshold phenomenon is generically observed. Below a critical value α_c, instances are typically satisfiable (SAT phase): At least one satisfying assignment (or *solution*) exists with probability one when $N \rightarrow \infty$; above α_c, they are typically unsatisfiable (UNSAT phase). Rigorous bounds on α_c have been derived [7]. The running time of algorithms often increases greatly near α_c [8].

CSPs enter the standard framework of statistical physics by associating to each assignment of the N variables $\sigma \equiv \{\sigma_i\}_{i=1}^{N}$ an energy $E[\sigma]$ defined as the number of constraints violated by σ. The satisfiability problem reduces to the determination of the ground-state energy $E_0 = \min_\sigma E[\sigma]$: If $E_0 > 0$, the instance is UNSAT; if $E_0 = 0$, it is SAT. In recent years, several methods borrowed from statistical physics [3,9,10] have pointed to the existence of a second threshold $\alpha_d < \alpha_c$, associated with clustering of the space S of all solutions. For $\alpha < \alpha_d$ (easy-SAT phase), S is typically connected: Any two solutions are joined by a path of moves involving a finite number of variables. For

$\alpha_d < \alpha < \alpha_c$ (hard-SAT phase), S is typically disconnected: Solutions gather into clusters far apart from each other [as illustrated in Fig. 1(a)], which can be joined only by moves involving a finite fraction of the variables. This scenario, which has been confirmed rigorously in some cases [11], suggests that computational hardness may be caused by the trapping of local algorithms in metastable clusters, which are exponentially more numerous than clusters of solutions.

In this Letter, we introduce methods to analyze in detail the structure of the solution space of CSPs in the hard-SAT phase. The first aspect we analyze is the entropic structure. A cluster λ typically contains an exponential number of solutions $\mathcal{M}_\lambda \asymp \exp(Ns_\lambda)$, where Ns_λ is the internal entropy of λ (we write $a_N \asymp b_N$ when $\ln a_N / \ln b_N \rightarrow 1$ as $N \rightarrow \infty$). We introduce the *entropic complexity* $\Sigma_s(s)$ that counts the number $\mathcal{N}_N(s) \asymp \exp[N\Sigma_s(s)]$ of clusters with internal entropy Ns and a method for computing $\Sigma_s(s)$ and the total entropy density s_{tot}, yielding the total number of solutions $|S| \asymp \exp(Ns_{\text{tot}})$ for individual instances of any CSP. The problem of counting the number of solutions of a CSP is, in general, #P-complete [1,12], a class of problems even harder than NP-complete. Estimating $|S|$ is important in applications such as graph reliability [12] and computing partition functions.

A second, related aspect of the structure of S is its geometry. We introduce a method to compute the *geomet-*
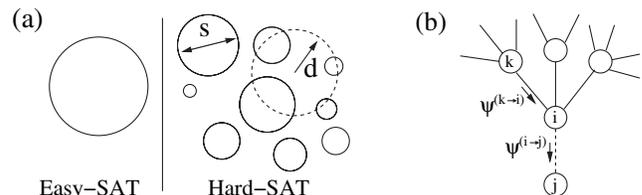

FIG. 1. (a) Illustration of the clustering phenomenon. In the easy-SAT phase $\alpha < \alpha_d$, all the solutions are connected. In the hard-SAT phase $\alpha_d < \alpha < \alpha_c$, solutions separate into distinct clusters. (b) Notations used in the cavity approach: The message (cavity field) $\psi_\sigma^{(i \rightarrow j)}$ gives the probability that node i has color σ in the absence of node j.

ric complexity $\Sigma_d(d)$, which counts the number of clusters at a given distance Nd from a reference assignment [see Fig. 1(a)], and the related weight enumerator function, of direct interest in coding theory [13]. Finally, we indicate several generalizations of these methods.

Our methods are based on extensions of the "energetic" cavity method (CM) of Ref. [14]. We illustrate them for q-COL and show numerical results for $q = 3$, but we emphasize that any CSP can be studied along the same lines. The energy function associated to q-COL is that of the antiferromagnetic Potts model, $E[\sigma] = \sum_{(i,j)} \delta_{\sigma_i,\sigma_j}$, where $\sigma_i \in \{1, \ldots, q\}$ and the sum is over the M graph edges. We study Erdős-Rényi random graphs [15], constructed by connecting any pair of nodes with probability $2\alpha/N$. For large N, this gives $M = \alpha N$ and a Poisson-distributed connectivity with mean 2α.

In the unclustered phase ("replica symmetric" phase in the language of spin-glass theory), the zero-temperature energetic CM [14] computes the ground-state energy recursively by adding one node at a time. For large enough α, the recursion no longer admits a unique solution and is generalized, via the one-step replica symmetry breaking ansatz (1RSB), to a distributional recursion which can be solved self-consistently, yielding the energetic complexity $\Sigma_\epsilon(\epsilon)$, which counts the number $\mathcal{N}_N(\epsilon) \asymp \exp[N\Sigma_\epsilon(\epsilon)]$ of clusters of local minima with energy $E = N\epsilon$. In particular, $\Sigma_\epsilon(0)$ is found positive in an interval $\alpha \in [\alpha_d^{(f)}, \alpha_c]$. The method was applied to q-COL in Refs. [16,17], which report $\alpha_d^{(f)} \simeq 2.21$, $\alpha_c \simeq 2.34$ for $q = 3$ (see also Fig. 3). The validity of the 1RSB ansatz in an interval $[\alpha_m, \alpha_{SP}]$ containing α_c was established for q-COL in Ref. [18], using the stability analysis of Ref. [19], with $\alpha_m \simeq 2.25$, $\alpha_{SP} \simeq 2.50$ for $q = 3$.

Counting solutions.—The energetic CM has the virtue of being simple enough, and it thus allows a precise determination of α_c and the development of a powerful new class of algorithms (survey propagation [10]). This simplicity is obtained because one focuses only on clusters in which some of the variables are frozen, i.e., constrained to adopt a unique color. Computing the entropy requires a more detailed information and, thus, a different formalism, as first identified in Ref. [9] within the replica framework. Our approach to computing entropies is illustrated in Fig. 1(b). The basic quantity we consider is the number $Z_{\sigma_i}^{(i \to j)}$ of solutions for the "cavity" graph obtained from the original graph by removing node j, when the color of node i is fixed to σ_i. In the unclustered phase, due to the locally treelike structure of large random graphs, the quantities $Z_{\sigma_k}^{(k \to i)}$, with k denoting any of the nodes connected to i except j (in symbols, $k \in i - j$), are independent of each other for large N. Hence, a recursion relation holds, $Z_{\sigma_i}^{(i \to j)} = \prod_{k \in i - j} \sum_{\sigma_k \neq \sigma_i} Z_{\sigma_k}^{(k \to i)}$. By defining a cavity field as the probability of having color σ on node i in the absence of j, $\psi_\sigma^{(i \to j)} \equiv Z_\sigma^{(i \to j)} / \sum_\tau Z_\tau^{(i \to j)}$, the recursion relation translates to

$$\psi_\sigma^{(i \to j)} = \hat{\psi}_\sigma^{(i \to j)}(\{\psi^{(k \to i)}\}) \equiv \mathcal{Z}^{-1} \prod_{k \in i - j} (1 - \psi_\sigma^{(k \to i)}), \quad (1)$$

with $\mathcal{Z}$ fixed by normalization. The ensemble of these equations on all oriented links, known as belief propagation equations [20], has a unique solution for $\alpha < \alpha_d$. In general, $\alpha_d \leq \alpha_d^{(f)}$ holds, since $\alpha_d^{(f)}$ refers to the onset of clusters with frozen variables, while at α_d clusters without frozen variables may also appear. It is not difficult to show that the total entropy of the whole graph is given by $N s_{tot} = \sum_i \Delta S^{(i)} - \sum_{(i,j)} \Delta S^{(i,j)}$ where, similarly to the energetic CM [3], we need to substract the link contributions $\Delta S^{(i,j)} = \ln(1 - \sum_\tau \psi_\tau^{(i \to j)} \psi_\tau^{(j \to i)})$ from the node contributions $\Delta S^{(i)} = \ln \sum_\tau \prod_{k \in i} (1 - \psi_\tau^{(k \to i)})$ to avoid double counting. Above α_d, following the 1RSB ansatz [21], we assume the existence of many clusters. We then compute a potential $\phi(x)$ related to the entropic complexity $\Sigma_s(s)$ through

$$e^{N\phi(x)} = \int_{s_{min}}^{s_{max}} e^{N[\Sigma_s(s) + xs]} ds, \quad (2)$$

where x is a Lagrange multiplier which fixes the internal entropy and s_{min}, s_{max} are the points at which $\Sigma_s(s)$ vanishes. Assuming the independence of the quantities $Z_{\sigma_k}^{(k \to i)}$ within each cluster, we introduce probability distributions of the cavity fields $P^{(i \to j)}(\psi^{(i \to j)})$ with respect to the clusters and generalize the cavity recursion to

$$P^{(i \to j)}(\psi^{(i \to j)}) \propto \int \prod_{k \in i - j} dP^{(k \to i)}(\psi^{(k \to i)}) Z(\{\psi^{(k \to i)}\})^x$$
$$\times \delta(\psi^{(i \to j)} - \hat{\psi}^{(i \to j)}(\{\psi^{(k \to i)}\})). \quad (3)$$

After solving Eq. (3), the potential is computed as

$$N\phi(x) = \sum_i \ln \int \prod_{j \in i} dP^{(k \to i)}(\psi^{(k \to i)}) e^{x\Delta S^{(i)}(\{\psi^{(k \to i)}\})}$$
$$- \sum_{(i,j)} \ln \int \prod_{a = (i \to j),(j \to i)} dP^{(a)}(\psi^{(a)}) e^{x\Delta S^{(i,j)}(\psi^{(i \to j)}, \psi^{(j \to i)})}, \quad (4)$$

where $\Delta S^{(i)}$, $\Delta S^{(i,j)}$ are given above. A saddle point evaluation of Eq. (2) gives $x = -\partial_s \Sigma_s(s)$. Hence, from $\phi(x)$ we obtain $\Sigma_s(s)$ via the Legendre transform $s(x) = \partial_x \phi(x)$, $\Sigma_s(x) = \phi(x) - xs(x)$. We solve numerically Eq. (3) on individual graphs by representing the distributions $P^{(i \to j)}$ with a population of N_P cavity fields on each oriented link. The resulting message passing algorithm is an entropic generalization of survey propagation [10]. Our entropic CM provides greater information at the price of greater computational difficulty, due to the continuous nature of the cavity fields.

Figure 2 displays some of our results for 3-COL, for three individual graphs with $N = 5000$. In particular, the total entropy of solutions is given by $s_{tot} = \phi(1) = s_{max}$, where the last equality holds because, according to our numerical results in Fig. 2, $\Sigma_s(x)$ vanishes at $x = x^* < 1$, with $s_{max} = \partial_x \phi(x = x^*)$. Therefore, for 3-COL the total

entropy is dominated by a subexponential number of giant clusters: A randomly chosen solution falls almost surely in one of such rare clusters. We also find that the fraction of frozen variables is finite in the interval $[s_{\min}, s_{\max}]$.

We also implemented a version of Eq. (3) averaged over Erdős-Rényi graphs, by considering a population of links with Poisson connectivity and a population of cavity fields on each link. Figure 3 shows the graph averages obtained in this way for s_{tot}, the typical internal entropy $s_{\text{typ}} = \text{argmax}_s \Sigma_s(s)$, and the typical complexity $\Sigma_{\text{typ}} = \Sigma_s(s_{\text{typ}})[= \Sigma_\epsilon(0)]$, as a function of α. The graph-averaged complexity curves $\Sigma_s(s)$ (not shown) resemble those in Fig. 2. (Graph-to-graph fluctuations for $N = 5000$ are significant: For $\alpha = 2.3$, the standard deviation of Σ_{typ} is about 27% of the mean.) The negative complexities in Fig. 2 have no direct interpretation on individual graphs, but for the graph-averaged case they are related to rare atypical graphs [22].

The above formalism can be generalized to yield $\Sigma_{\epsilon,s}(\epsilon, s)$, the complexity associated with metastable clusters of energy $N\epsilon > 0$ and entropy Ns, with $\Sigma_{\epsilon,s}(0, s) = \Sigma_s(s)$, by adding a second multiplier y [23]. Equivalent information is contained in the finite temperature complexity $\Sigma_f(f; \beta)$ [21], where f is the free energy and β the inverse temperature, based on the identity

$$\int e^{N[\Sigma_{\epsilon,s}(\epsilon,s)-y\epsilon+xs]}d\epsilon\, ds = \int e^{N[\Sigma_f(f;\beta)-x\beta f]}df, \quad (5)$$

with $f = \epsilon - s/\beta$ and $y = \beta x$. The energetic CM is recovered for $\beta \to \infty$ and $x \to 0$ with $y = \beta x$ fixed, which amounts to ignoring all entropic effects [24].

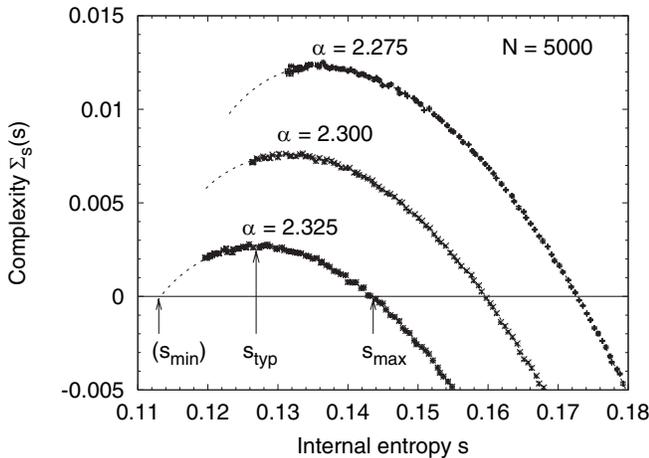*Counting clusters at a given distance.*—We now turn to the geometric structure and show how the CM can be used to investigate intercluster distances. We illustrate this by addressing the problem of counting the number of clusters as a function of their distance from a fixed reference configuration ꜱ, which we rephrase as a new CSP, named dCSP, whose thermodynamics reflect the geometry of the solution space of the initial CSP. The valid assignments of dCSP are the solutions $\sigma \in S$ of the initial CSP: These are configurations of zero energy, and in this sense dCSP concentrates on the zero-temperature case of the original problem. But we introduce in dCSP a new energy function which is the Hamming distance from ꜱ, $E_D[\sigma] \equiv \sum_{i=1}^N (1 - \delta_{\varsigma_i, \sigma_i})$. Therefore, the clusters (assignments) of dCSP with energy E_D are the zero-energy clusters (solutions) at distance E_D from ꜱ in the initial CSP.

The optimization problem for dCSP consists in finding the maximal (or the minimal) distance between ꜱ and a solution of the original problem. By applying the energetic CM to this problem [25], we obtain a geometric complexity $\Sigma_d(d)$ giving the number of clusters at distance Nd of ꜱ, $\mathcal{N}_N(d) \asymp \exp[N\Sigma_d(d)]$. Figure 4 shows results for 3-COL on individual graphs. Two features are worth noticing: (i) $\Sigma_d(d)$ becomes positive only above a threshold $d_{\min}$, reflecting the fact that clusters are well separated; (ii) a plateau appears between d_1 and d_2, reflecting the finite diameter of clusters. We have verified that the size of this plateau coincides with the typical diameter computed within the entropic CM [23].

Generalizations.—The above method can be extended to count the number of solutions at distance Nd from ꜱ, known as the weight enumerator function $A_N(d)$ in coding theory [13]. This can be deduced from the complexity $\Sigma_{d,s}(d, s)$, which gives the number of clusters with internal entropy Ns at distance Nd from the reference configuration ꜱ. Such a complexity can be obtained by studying the dCSP

FIG. 2. Entropic complexity for three individual Erdős-Rényi graphs with $N = 5000$ and different values of $\alpha = M/N$. Data obtained with a population size $N_P = 512$ on each oriented link (we verified that using $N_P = 4096$ gives a change smaller than the error bars). The dotted lines are obtained by a polynomial fit of the potential $\phi(x)$, the symbols by direct computation of the derivative $\partial_x \phi(x)$ in Eq. (4) [27].

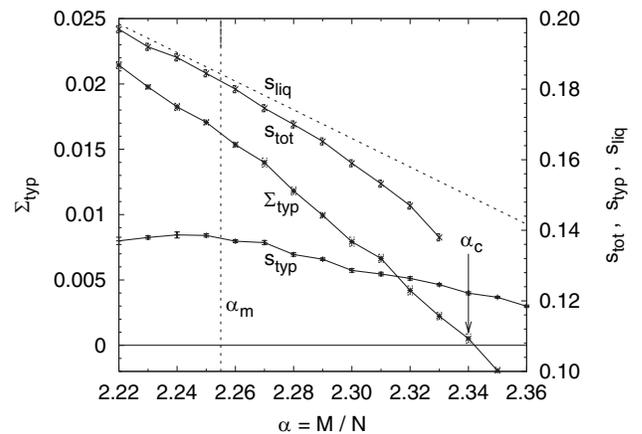FIG. 3. Graph-averaged Σ_{typ}, s_{typ}, and s_{tot}. Notice the different vertical scales. Data obtained with a population of 16 000 links and $N_P = 512$ fields on each link. The vertical line shows the threshold $\alpha_m \simeq 2.255$ below which the 1RSB ansatz is unstable [18]. The straight line is the "liquid" or infinite-temperature solution $s_{\text{liq}} = (1 - \alpha)\ln(q) + \alpha \ln(q - 1)$.
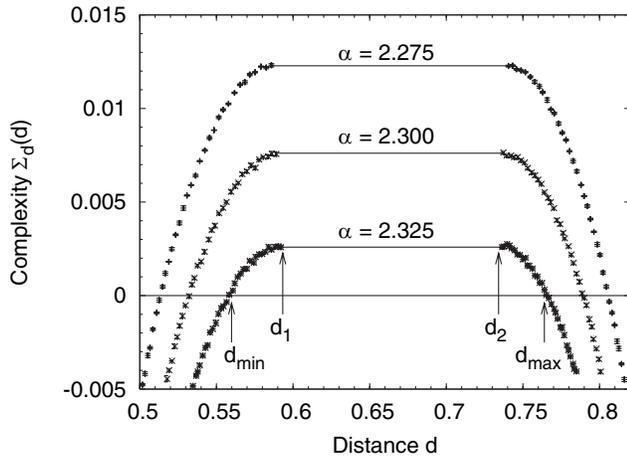
FIG. 4. Geometric complexity for the same individual graphs as in Fig. 2. Within the error bars, $\Sigma_d(d_1) = \Sigma_d(d_2) = \Sigma_s(s_{\text{typ}})$. The horizontal lines are a guide to the eye.

with a finite value of a new inverse temperature β_D, which is conjugate to the energy E_D (keeping the original temperature β^{-1} to zero) [23]. Once $\Sigma_{d,s}(d, s)$ has been found, one obtains the leading behavior of the weight enumerator as $A_N(d) \asymp \exp[N\max_s(\Sigma_{d,s}(d, s) + s)]$. In the same spirit, our analysis can be extended to metastable configurations: In order to compute the complexity $\Sigma_{\epsilon,d,s}(\epsilon, d, s)$ counting clusters with energy $N\epsilon$, entropy Ns, at distance Nd from ς, one needs to introduce three Lagrange multipliers x, y, z. All the previous complexities are particular limits of this more general framework [23].

Conclusions.—We have presented methods to analyze the entropic and geometric structure of the clustered phase in q-COL, which give access to quantities such as internal cluster entropies not accessible to previous methods. Our results for 3-COL show the existence of giant, atypical clusters which contain the majority of solutions. Generalization to other CSPs such as k-SAT, where a similar picture may hold, is straightforward.

Notice that the present results were obtained within a 1RSB ansatz, and the stability of our solution should thus be checked (extending the method of Ref. [19]) to assess whether the solution is exact or only an approximation to a more complicated one involving higher order RSB.

The new information extracted with our entropic CM could be exploited to design new algorithms for finding solutions to individual instances, improving on present survey propagation algorithms which use only energetic information [10]. We also envision applications to inference problems such as Bayesian belief networks [26].

Note added in proof.—The recent paper [28] addresses similar questions in the simpler unclustered case.

[1] C. H. Papadimitriou, *Computational Complexity* (Addison-Wesley, Reading, MA, 1994).
[2] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky, Nature (London) **400**, 133 (1999).
[3] M. Mézard, G. Parisi, and R. Zecchina, Science **297**, 812 (2002).
[4] D. Battaglia, A. Braunstein, J. Chavas, and R. Zecchina, Phys. Rev. E **72**, 015103 (2005).
[5] J. Barré, A. R. Bishop, T. Lookman, and A. Saxena, J. Stat. Phys. **118**, 1057 (2005).
[6] L. Correale, M. Leone, A. Pagnani, M. Weigt, and R. Zecchina, cond-mat/0412443.
[7] D. Achlioptas, A. Naor, and Y. Peres, Nature (London) **435**, 759 (2005), and references therein.
[8] P. Cheeseman, B. Kanefsky, and W. M. Taylor, in *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91)* (Morgan Kaufmann, San Mateo, CA, 1991), pp. 331–337.
[9] G. Biroli, R. Monasson, and M. Weigt, Eur. Phys. J. B **14**, 551 (2000).
[10] M. Mézard and R. Zecchina, Phys. Rev. E **66**, 056126 (2002).
[11] M. Mézard, T. Mora, and R. Zecchina, Phys. Rev. Lett. **94**, 197205 (2005).
[12] L. G. Valiant, SIAM J. Comput. **8**, 410 (1979).
[13] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press, Cambridge, England, 2003).
[14] M. Mézard and G. Parisi, J. Stat. Phys. **111**, 1 (2003).
[15] B. Bollobás, *Random Graphs* (Cambridge University Press, Cambridge, England, 2001), 2nd ed.
[16] R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina, Phys. Rev. Lett. **89**, 268701 (2002).
[17] A. Braunstein, R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina, Phys. Rev. E **68**, 036702 (2003).
[18] F. Krzakala, A. Pagnani, and M. Weigt, Phys. Rev. E **70**, 046705 (2004).
[19] A. Montanari and F. Ricci-Tersenghi, Eur. Phys. J. B **33**, 339 (2003).
[20] F. R. Kschischang, B. Frey, and H.-A. Loeliger, IEEE Trans. Inf. Theory **47**, 498 (2001).
[21] M. Mézard and G. Parisi, Eur. Phys. J. B **20**, 217 (2001).
[22] O. Rivoire, J. Stat. Mech. (2005) P07004.
[23] M. Mézard, M. Palassini, and O. Rivoire (to be published).
[24] When $\beta \to \infty$ at fixed x, the cavity field $\psi(\beta)$ used in computing $\sum_f(f; \beta)$ in Eq. (5) yields either the purely entropic field $\psi = \lim_{\beta\to\infty}\psi(\beta)$ or the energetic field $h = \lim_{\beta\to\infty}\{\ln[\psi(\beta)]/\beta\}$ (the value of x is irrelevant in the latter case). The supplement of entropic information conveyed by ψ defines the "evanescent field" [9] $h' = \lim_{\beta\to\infty}\{\ln[\psi(\beta)]/ -\beta h\}$.
[25] If $\varsigma \in S$, the minimal distance is zero, but by focusing on the metastable (glassy) states, the cavity method allows one to ignore this trivial (crystalline) ground state.
[26] D. Roth, Artif. Intell. **82**, 273 (1996).
[27] The numerical solution of Eq. (3) becomes difficult for small s as the distributions become broad. We display data only conservatively down to the lowest s, at which $\partial_x\phi(x)$ computed by direct derivation of Eq. (4) coincides with the derivative of a polynomial fit of $\phi(x)$.
[28] A. Bandyopadhyay and D. Gamarnik, math.PR/0510471.