

Random Subcubes as a Toy Model for Constraint Satisfaction Problems

Thierry Mora · Lenka Zdeborová

Received: 23 October 2007 / Accepted: 4 April 2008 / Published online: 19 April 2008
© Springer Science+Business Media, LLC 2008

Abstract We present an exactly solvable random-subcube model inspired by the structure of hard constraint satisfaction and optimization problems. Our model reproduces the structure of the solution space of the random k -satisfiability and k -coloring problems, and undergoes the same phase transitions as these problems. The comparison becomes quantitative in the large- k limit. Distance properties, as well the x -satisfiability threshold, are studied. The model is also generalized to define a continuous energy landscape useful for studying several aspects of glassy dynamics.

Keywords Constraint satisfaction problems · Clustering of solutions · Exactly solvable models

1 Introduction

Combinatorial optimization and Constraint Satisfaction Problems (CSPs) arise in a wide array of scientific branches, including statistical physics, information theory, inference and machine learning. These problems, which involve a large number of variables interacting through a large number of constraints or cost terms, are in general very hard to solve, and in most cases no algorithm seems to be able to find a solution within a reasonable time, as formalized by the $P \neq NP$ conjecture [1].

In order to circumvent this intrinsic difficulty and possibly to identify (or avoid) “hard” instances, random ensembles of optimization problems were introduced and used as test beds for theories and algorithms. This line of research has considerably benefited from

T. Mora (✉)
Lewis Sigler Institute for Integrative Genomics, Princeton University, Princeton, NJ 08544, USA
e-mail: tmora@princeton.edu

L. Zdeborová
Université Paris-Sud, LPTMS, UMR8626, Bât. 100, Université Paris-Sud, 91405 Orsay cedex, France

L. Zdeborová
CNRS, LPTMS, UMR8626, Bât. 100, Université Paris-Sud, 91405 Orsay cedex, France

the methods and concepts of statistical mechanics [2–5]. In particular, a spectacular breakthrough was made by the development of the survey propagation algorithm [5, 6] which is able to solve large random instances of CSPs in the so-called “hard-SAT” region. The key to the success of the physics approach lies in the understanding of the rugged energy landscape (reminiscent of glassy phases) exhibited by these problems, which survey propagation exploits and integrates into a sophisticated message-passing [7] scheme.

The structure and organisation of solutions has been analyzed in detail for several CSPs, including the satisfiability problem (k -SAT) [8], the colorability of random graphs (k -COL) [8–10], and systems of linear Boolean equations (k -XORSAT) [11–13]. It was shown that as the density of constraints is increased the space of solutions undergoes several phase transitions. At low density of constraints the solution space is concentrated in one big ergodic component, called “cluster” or “state”. For higher densities the systems undergoes a clustering transition, whereby the solution space breaks into an exponential number of well-separated clusters (this separation can be energetic or entropic). For even higher densities a second transition occurs, which reduces these clusters to a finite number. Finally, all clusters disappear at the SAT-UNSAT threshold.

Since the clustering phenomenon is one of the main building blocks underlying the statistical physics approach,¹ substantial efforts have been made to give a rigorous base to it [15, 16]. Some mathematical results were also obtained in the ergodic phase [17], and in the simple case of linear Boolean equations [18]. Remarkably, pure states play an central role in spin-glass theory, and they have been extensively studied in that context. However, the geometrical organization of glassy phases is not yet fully understood, and the classical picture of complex energy landscapes with many “valleys” still lacks an appropriate representation.

In this paper we introduce an exactly solvable random-subcube model² (RSM), in the spirit of Derrida’s Random Energy Model (REM) [19]. This model is inspired by the structure of hard CSPs and optimization problems, and reproduces most of their phenomenology. It can be also thought of as an attempt to construct a minimal setting that is able to reproduce the structure of solutions in hard CSPs. Its purpose is mainly pedagogical, and it offers an excellent testing playground for ideas and methods in combinatorial optimization and glass physics, while being fully tractable.

Despite its simplicity, the RSM undergoes the same phase transitions as those observed in random CSPs such as k -SAT [6, 8] or k -COL [8, 10]. The connection even becomes quantitative in the large- k limit of these problems. So far only the zeroth order of this limit was intuitively known and related to Shannon’s random code model [20–22], in which clusters are uniformly distributed singletons. One of our most notable results is that the RSM provides the first-order approximation of this large- k limit, reproducing the cluster size distribution and freezing properties of the original models.

We also generalize the RSM to deal with continuous energy landscapes resembling those observed in glassy systems and hard optimization problems, and show how static and dynamical properties can be related explicitly. This energetic RSM displays temperature chaos, undergoes a dynamical transition, and has a Kauzmann temperature.

The paper is organized as follows: In Sect. 2 we define the model and describe its basic properties, as well as the connection with the random k -SAT and k -COL problems. We also analyze the behaviour of physics-guided decimation schemes. In Sect. 3 we discuss

¹Technically, clustering is closely related to the one-step replica-symmetry breaking ansatz used in the replica/cavity method [6, 14].

²Independently of our work A. Montanari inspired by an idea of D. Achlioptas also introduced this model and worked out some parts of our Sect. 2.

the relation between its dynamical and geometrical properties. In Sect. 4 we extend the definition to energetic landscapes, compute static and dynamical properties, and comment on some ideas from the physics of glassy systems. Finally we present a general discussion of our results in Sect. 5.

2 The Random-Subcube Model

2.1 Definition

Most constraint satisfaction problems are defined by a set of constraints on N variables $\sigma = (\sigma_1, \dots, \sigma_N)$ with a finite alphabet, e.g. $\{0, 1\}$. In contrast, the random-subcube model is defined directly by its solution space $S \subset \{0, 1\}^N$; we define S as the union of $\lfloor 2^{(1-\alpha)N} \rfloor$ random clusters (where $\lfloor x \rfloor$ denotes the integer value of x). A random cluster A being defined as:

$$A = \{\sigma \mid \forall i \in \{1, \dots, N\}, \sigma_i \in \pi_i^A\}, \tag{1}$$

where π^A is a random mapping:

$$\pi^A : \{1, \dots, N\} \longrightarrow \{\{0\}, \{1\}, \{0, 1\}\} \tag{2}$$

$$i \longmapsto \pi_i^A \tag{3}$$

such that for each variable i , $\pi_i^A = \{0\}$ with probability $p/2$, $\{1\}$ with probability $p/2$, and $\{0, 1\}$ with probability $1 - p$. A cluster is thus a random subcube of $\{0, 1\}^N$. If $\pi_i^A = \{0\}$ or $\{1\}$, variable i is said “frozen” in A ; otherwise it is said “free” in A . One given configuration σ might belong to zero, one or several clusters. A solution belongs to at least one cluster.

The parameter α is analogous to the density of constraints in CSPs; clearly the SAT-UNSAT transition occurs at $\alpha_s := 1$, where clusters cease to exist. The parameter p gives the probability that a variable is frozen, and plays a role similar to the clause size k in k -SAT, or to the number of colors k in k -coloring, as we will see later. Note that in the special case $p = 1$, the RSM is equivalent to the random code model with rate $R = 1 - \alpha$ [21].

Frozen variables and the structure of the solution space have been introduced to mimic the situation observed in random CSPs. However there are important differences between the RSM and models like k -SAT. First, in real CSPs the clusters are not necessarily subcubes of $\{0, 1\}^N$. We stress here that when speaking about clusters in the RSM we have in mind the above definition, whereas in the context of the CSPs the notion of cluster is more general [5, 8]. Further, in real CSPs the sets of frozen variables associated with clusters are correlated by the underlying graph, instead of being distributed uniformly. Moreover, free variables in CSPs do not enjoy the same freedom as in the RSM, as clusters usually do not fill up the whole subcube allowed by the frozen variables. In fact, free variables can be correlated within each cluster in a highly nontrivial way, and these correlations may even be so strong that they create clusters without the help of frozen variables. Clusters without frozen variables are indeed very important, as discussed recently in [10, 23].

2.2 The Basic Structural Phase Transitions

We now describe the static properties of the RSM in the thermodynamic limit $N \rightarrow \infty$ (the two parameters α and p being fixed and independent of N). The internal entropy s of a

cluster A is defined as $\frac{1}{N} \log_2 |A|$, i.e. the fraction of free variables in A . The probability $\mathcal{P}(s)$ that a cluster has internal entropy s follows the binomial distribution

$$\mathcal{P}(s) = \binom{N}{sN} (1-p)^{sN} p^{(1-s)N}. \tag{4}$$

Let $\mathcal{N}(s)$ be number of clusters of entropy s . This number follows a binomial law of parameter $\mathcal{P}(s)$ with $2^{N(1-\alpha)}$ terms. Then the mean and the variance of $\mathcal{N}(s)$ read:

$$\mathbb{E}\mathcal{N}(s) = 2^{N(1-\alpha)}\mathcal{P}(s), \quad \text{Var}\mathcal{N}(s) = 2^{N(1-\alpha)}\mathcal{P}(s)[1 - \mathcal{P}(s)]. \tag{5}$$

By Markov’s inequality:

$$\mathbb{P}[\mathcal{N}(s) \geq 1] \leq \mathbb{E}[\mathcal{N}(s)], \tag{6}$$

and by Chebyshev’s inequality:

$$\mathbb{P}\left\{\left|\frac{\mathcal{N}(s)}{\mathbb{E}\mathcal{N}(s)} - 1\right| > \varepsilon\right\} \leq \frac{\text{Var}\mathcal{N}(s)}{[\mathbb{E}\mathcal{N}(s)]^2\varepsilon^2} \leq \frac{1}{2^{N(1-\alpha)}\varepsilon^2\mathcal{P}(s)} \quad \text{for all } \varepsilon > 0, \tag{7}$$

we get, with high probability (w.h.p.: with probability going to 1 as $N \rightarrow \infty$):

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log_2 \mathcal{N}(s) = \begin{cases} \Sigma(s) := 1 - \alpha - D(s \parallel 1 - p) & \text{if } \Sigma(s) \geq 0, \\ -\infty & \text{otherwise,} \end{cases} \tag{8}$$

where $D(x \parallel y) := x \log_2 \frac{x}{y} + (1-x) \log_2 \frac{1-x}{1-y}$ is the binary Kullback-Leibler divergence. Throughout the paper, the same Markov/Chebyshev argument will apply every time we will have to deal with a number of clusters with a specific property.

We now compute the total entropy $s_{\text{tot}} = \frac{1}{N} \log_2 |S|$. First note a random configuration belongs on average to $2^{N(1-\alpha)}(1 - \frac{p}{2})^N$ clusters. Therefore, if

$$\alpha < \alpha_d := \log_2(2 - p), \tag{9}$$

then with high probability the total entropy is $s_{\text{tot}} = 1$.

Now assume $\alpha > \alpha_d$. The total entropy is given by a saddle-point estimation:

$$\sum_A 2^{s(A)N} = [1 + o(1)]N \int_{\Sigma(s) \geq 0} ds 2^{N[\Sigma(s)+s]}, \tag{10}$$

$$\text{whence } s_{\text{tot}} = \max_s [\Sigma(s) + s \mid \Sigma(s) \geq 0]. \tag{11}$$

We denote by $s^* = \text{argmax}[\dots]$ the fraction of free variables in the clusters that dominate the sum. Note that in this sum solutions belonging to several clusters have been counted too many times. This does not affect the validity of our estimation, since in every cluster the fraction of solutions belonging to more than one cluster is exponentially small as long as $\alpha > \alpha_d$.

Define $\tilde{s} := 2(1 - p)/(2 - p)$ such that $\partial_s \Sigma(\tilde{s}) = -1$. The complexity of clusters with entropy \tilde{s} reads:

$$\Sigma(\tilde{s}) = \frac{p}{2 - p} + \log_2(2 - p) - \alpha. \tag{12}$$




Fig. 1 (Color online) *Top*: graphical construction of the maximum of $\Sigma(s) + s$ by a Legendre transformation. In the *top left* figure ($\alpha < \alpha_c$), the line of slope -1 tangent to the complexity function gives the saddle-point s^* , as well as the total entropy s_{tot} by the intercept on the s -axis. In the *top right* figure ($\alpha > \alpha_c$), the supporting line of slope -1 gets “stuck” at s_M , where the derivative is $-m > -1$. *Bottom*: represented as a function of α : total entropy s_{tot} , total complexity $\Sigma_{tot} = 1 - \alpha$, typical entropy s^* , complexity of dominating clusters $\Sigma^* = \Sigma(s^*)$, and $m = -\partial_s \Sigma(s^*)$. The condensation point α_c marks the separation between the two regimes illustrated above

\tilde{s} maximizes (11) as long as $\Sigma(\tilde{s}) \geq 0$, that is if

$$\alpha \leq \alpha_c := \frac{p}{(2 - p)} + \log_2(2 - p). \tag{13}$$

Then the total entropy reads

$$s_{tot} = 1 - \alpha + \log_2(2 - p) \quad \text{for } \alpha \leq \alpha_c. \tag{14}$$

For $\alpha > \alpha_c$, the maximum in (11) is realized by the largest possible cluster entropy s_M , which is given by the largest root of $\Sigma(s)$. Then $s_{tot} = s^* = s_M$. In this phase the dominating clusters³ have size $2^{Ns^* + \Delta}$, where $\Delta = O(1)$ is asymptotically distributed according to a Poisson point process of rate $2^{-m\Delta}$, i.e., for $d\Delta \ll \Delta$ the probability that there is at least one state of size between $2^{Ns^* + \Delta}$ and $2^{Ns^* + \Delta + d\Delta}$ is $2^{-m\Delta} d\Delta$, where $m = -\partial_s \Sigma(s^*)$. Extreme value study of this process leads to the Poisson-Dirichlet [24–26] distribution of weights of clusters.⁴ In particular it follows that an arbitrary large fraction of the solutions can be covered by a finite number of clusters. Such a phase is called condensed.

³The “dominating clusters” are the minimal set of clusters covering almost all solutions.

⁴Strictly speaking in the RSM $Ns^* + \Delta$ can only be an integer, then Δ is a discrete random variable. Some detailed properties of the extreme values might thus differ from the properties of the well known Poisson-Dirichlet process.

In summary, for a fixed value of the parameter p , and for increasing values of α , four different phases can be distinguished:

- Liquid phase, $\alpha < \alpha_d$: almost all configurations are solutions.
- Clustered phase with many states, $\alpha_d < \alpha < \alpha_c$: an exponential number of clusters is needed to cover almost all solutions.
- Condensed clustered phase, $\alpha_c < \alpha < \alpha_s = 1$: a finite number of the biggest clusters cover almost all solutions.
- Unsatisfiable phase, $\alpha > \alpha_s$: no cluster, hence no solution, exists.

The very same series of phase transitions is observed in random k -coloring and random k -satisfiability, where α is the density of constraints [8, 10]. The condensation transition at α_c corresponds to the Kauzmann temperature [27] in the theory of glasses; at α_c the total entropy $s_{\text{tot}}(\alpha)$ has a discontinuity in its second derivative with respect to α , in analogy with the discontinuity of the specific heat at the Kauzmann temperature.

Keeping in mind the similarity between the RSM and real CSPs, it can be useful to mention some of the properties that are commonly discussed in the statistical physics analysis of these problems (for brevity we omit the proofs of these statements). Among them is the probability distribution of mutual overlaps $P(q)$, where $q(\sigma, \sigma') = \sum_{i=1}^N [2\delta(\sigma_i, \sigma'_i) - 1]/N$. Below the condensation transition, $\alpha < \alpha_c$, we have $P(q) = \delta(q)$, reflecting the fact that random pairs of solutions are uncorrelated. For $\alpha > \alpha_c$, the overlap function consists of intra-cluster and inter-cluster overlaps: $P(q) = w\delta[q - (1 - s_{\text{tot}})] + (1 - w)\delta(q)$, where w is the sum of squares of weights of all the clusters, it is a non-self-averaging random variable, the distribution of which can be computed from the Poisson-Dirichlet process [26, 28].

An equivalent way of characterizing the condensed phase is to consider the k -point correlation function, with $k \geq 2$:

$$\sum_{x_1, \dots, x_k} |\mathbb{P}(\sigma_1 = x_1, \sigma_2 = x_2, \dots, \sigma_k = x_k) - \mathbb{P}(\sigma_1 = x_1) \dots \mathbb{P}(\sigma_k = x_k)|. \quad (15)$$

This quantity decays to zero as N goes to infinity in the non-condensed phase, whereas it remains bounded away from zero in the condensed phase.

2.3 The Large k -Limit of Random k -SAT and k -COL

One of the most interesting properties of the random-energy model [19] is its equivalence with the large- p limit of the p -spin glass [28]. In the same spirit, although the justification is slightly different, the random-subcube model is found to be equivalent to random k -SAT, random k -COL, and presumably other constraint satisfaction problems in the limit $k \rightarrow \infty$, for connectivities close to the satisfiability threshold. Let us detail this statement.

It was already known that at zeroth order (when $k \rightarrow \infty$) random k -SAT and k -COL behave as a random-code model (random-subcube model with $p = 1$), in which clusters are uniformly distributed singletons. Recent large- k calculations of the cluster size distribution $\Sigma(s)$ in k -SAT and k -COL [8, 10] allow a direct comparison with the RSM.

The control parameters of the RSM are rescaled as:

$$p = 1 - \varepsilon, \quad \alpha = 1 + \varepsilon \frac{1 + \gamma}{\ln 2}, \quad (16)$$

with $\varepsilon \ll 1$ and $\gamma = \Theta(1)$. The cluster size distribution in the RSM is then, at leading order:

$$\Sigma(s) \ln(2) = s \left[1 - \ln \frac{s}{\varepsilon} \right] - \varepsilon(2 + \gamma) + o(\varepsilon). \quad (17)$$

while the condensation and satisfiability thresholds read, in terms of the rescaled variable γ :

$$\gamma_c = -2 \ln 2, \quad \gamma_s = -1. \tag{18}$$

In k -SAT and k -COL, identifying ε and γ in the following way:

$$\text{SAT: } \varepsilon = \frac{1}{2^{k+1}}, \quad \frac{M}{N} = 2^k \ln 2 - \frac{\ln 2}{2} + \frac{\gamma}{2}, \tag{19}$$

$$\text{COL: } \varepsilon = \frac{1}{2k}, \quad \frac{M}{N} = k \ln k - \frac{\ln k}{2} + \frac{\gamma}{2}, \tag{20}$$

where M is the number of constraints (edges in coloring), and N the number of variables, gives a perfect match⁵ for the complexity function (17), as computed in [8, 10]. This illustrates the analogy between α and the density of constraints M/N , as well as between p and k .

This equivalence goes further than simply having identical cluster size distributions. The cavity analysis of k -SAT shows that the fraction of free variables in a cluster scales exactly as its internal entropy s . The entropy of clusters is thus maximal, from which we infer that clusters fill up the whole subcube prescribed by their frozen variables, like in the RSM. The same is true for k -COL (compare (E14) and (E27) in [10]) with the small difference that for every unfrozen variable only two (out of k) colors are allowed.

Note that this comparison is valid only in a finite vicinity of the satisfiability threshold, for $\gamma = \Theta_k(1)$. In particular, it does not encompass the clustering transition, which for k -SAT (resp. k -COL) occurs for constraint densities scaling as $2^k \ln 2/k$ (resp. $k \ln k/2$).

2.4 Decimation

An important contribution of statistical physics to the field of combinatorial optimization has been to exploit the information provided by message-passing algorithms to devise physics-guided decimation schemes.

Message-passing algorithms exchange information between units (variables and constraints) in order to obtain *estimates* of marginal probabilities (beliefs) or other related quantities (e.g. surveys, see below). Subsequently, this information is used to find a solution. A usual way to do this, called decimation, proceeds as follows: fix randomly the value of one⁶ variable according to its estimated belief, then re-run the message-passing algorithm on the reduced system, and loop. A trivial statement is that a perfect estimate of all marginal probabilities would always cause the decimation procedure to find a solution (if any).

In the RSM there is no underlying graph, therefore message-passing cannot be defined. However, it is possible to study decimation schemes based on exact marginal probability estimators. Although such procedures should really be viewed as thought experiments, they can be used to gain some insight on real algorithms. Here two idealized algorithms are considered:

- Belief estimator: outputs the exact marginal probabilities $\mu_i(\sigma_i) = \sum_{\sigma_{\setminus i}} \mu(\sigma)$, where $\mu(\sigma) = \mathbb{I}(\sigma \in S)/|S|$.
- Survey estimator: outputs “surveys”, i.e. marginal probabilities over the clusters: $v_i(\pi_i) = \sum_{\pi_{\setminus i}} \nu(\pi)$, where $\nu(\pi) = \sum_A \mathbb{I}(\pi = \pi^A)/[2^{N(1-\alpha)}]$.

⁵Note the difference in the logarithmic base between here and [8, 10].

⁶In practice the number of variables fixed at each step can range from one to a small fraction of the variables.

In real CSPs, belief and survey propagation arguably provide *asymptotically* accurate estimators, as long as the number of clusters dominating the measure μ (or ν for survey propagation) scales exponentially with N , and as long as the one-step replica symmetry breaking description is the correct one [10, 14]. However, belief-guided and survey-guided decimation schemes are difficult to analyze in real CSPs—see [29] for an empirical study on surveys and [30] for recent analytical study on beliefs in k -SAT.

Let us study decimation in the RSM. As long as the phase is non-condensed, the belief estimator always outputs $\mu_i(\sigma_i) \approx 1/2$ for all i in the limit $N \rightarrow \infty$. Likewise, the survey estimator will output $\nu_i(\{0\}) \approx \nu_i(\{1\}) \approx p/2$, $\nu_i(\{0, 1\}) \approx 1 - p$. In both cases, the decimation procedure is completely unbiased: it will fix a random variable i to 0 or 1, with probability 1/2. This observation remains true in the subsequent decimation steps as long as the number of clusters dominating the reduced measure μ (or ν) remains exponential. Within this assumption, after $T = tN$ ($0 \leq t \leq 1$) decimation steps, T variables will be fixed randomly and independently. We are then left with a restricted space of solutions compatible with these T fixed variables. The logarithm of the number of clusters of entropy sN is then

$$N \Sigma_t(s) = N \left[1 - \alpha + t \log_2 \left(1 - \frac{p}{2} \right) - (1 - t) D \left(\frac{s}{1 - t} \parallel 1 - p \right) \right]. \tag{21}$$

Rescaling by the number of unfixed variables $(1 - t)N$: $\Sigma_t = (1 - t)\bar{\Sigma}_t$, $s = (1 - t)\bar{s}$, we obtain

$$\bar{\Sigma}_t(\bar{s}) = 1 - \frac{\alpha - t\alpha_d}{1 - t} - D(\bar{s} \parallel 1 - p). \tag{22}$$

The parameter $\bar{\alpha}(t) := (\alpha - t\alpha_d)/(1 - t)$ now plays the same role as α in the analysis of the RSM. Consequently, the system undergoes the same condensation and unsatisfiability transitions as t is increased. Assume for example that $\alpha_d < \alpha < \alpha_c$. Fixing a fraction $t_c := (\alpha_c - \alpha)/(\alpha_c - \alpha_d)$ of the variables will cause the system to condense. The belief estimator will then be dominated by a finite number of clusters, yielding instance-dependent biases on variables. An extensive number of variables become suddenly near-frozen, i.e., $\mu_i(1) \approx 0$ or ≈ 1 , and remain so for $t > t_c$. At $t_s := (1 - \alpha)/(1 - \alpha_d)$, the total complexity goes to zero, and near-frozen variables become truly frozen, i.e. $\mu_i(1) = 0$ or 1, as all subdominant clusters disappear. By contrast, the survey estimator will output unbiased marginal probabilities as long as the total number of clusters is exponential in N , that is if $t < t_s$. At $t = t_s$, decimation concentrates on a single cluster, causing a freezing avalanche in the surveys, i.e. for each variable i , $\nu_i(\{0, 1\}) = 1$ or $\nu_i(\{0\}) = 1$ or $\nu_i(\{1\}) = 1$. For both estimators, any slight error at t_s will cause the failure of the decimation process.

Note that in real CSPs, belief propagation is not expected to be correct in the condensed phase, beyond t_c ; in fact, it will not detect condensation, nor near-frozen variables. For a recent study of the belief-propagation-guided decimation in k -SAT see [30]. Remarkably, in real CSPs, survey-guided decimation usually simplifies the problem: after a certain number of decimation steps, it outputs $\nu_i(\{0, 1\}) = 1$ for all i , and the problem can then be passed over to a simple local search algorithm [6, 29]. Although the RSM, which is intrinsically clustered, is unable to capture this property, it sheds some light on why decimation may not work in some cases.

3 Distance and Ergodicity

The geometrical organization of solutions to CSPs is thought to play an important role in setting intrinsic limits to the performance of search algorithms. In particular, the clustering phenomenon, by which the solution space is fragmented into many connected components far from each other, has been proposed by physicists as a possible explanation for the failure of most known algorithms [4–6]. Also, the role of frozen variables was recently discussed in [10]. Conversely, the success of survey propagation [6] is usually explained by the fact that it explicitly incorporates the existence of clusters. The separability of clusters has been proved in the k -SAT problem, in compliance with the predictions of statistical physics [15]. Despite this evidence, the precise relation between geometry and algorithms still lacks a rigorous base. The random-subcube model offers an excellent opportunity to study these questions in a well controlled framework.

3.1 The Dynamical Transition

Let us first argue why the clustering transition at α_d , defined by (9), actually corresponds to what is commonly referred to as the dynamical (ergodicity breaking) transition in real CSPs [8, 10, 12, 31–33]. To this end we study a uniform unbiased random walk on the space of solutions: at each step, one is allowed to move from one solution to the other by flipping only one spin. We choose this dynamics for the sake of simplicity, but most of the arguments below hold for more general dynamical rules, like for example the flipping of a sub-extensive number of spins at each step.

We have already pointed out that an arbitrary configuration belongs to $2^{(1-\alpha)N} (1 - p/2)^N$ different clusters w.h.p. if $\alpha < \alpha_d$, and to none if $\alpha > \alpha_d$. Therefore below α_d almost all configurations are solutions, and any reasonable dynamics will explore the entire phase space uniformly.

On the other hand, when $\alpha > \alpha_d$, solutions become exponentially rare. Let A be a cluster of internal entropy s . What is the probability that after t steps, a random walker ends up in another cluster B of internal entropy s' ? Let a denote the proportion of variables that are free in A and frozen in B . The probability distribution of a is given by the total number of partitions of $\{1, \dots, N\}$ into four categories: frozen in A and B , only frozen in A , only frozen in B , and frozen neither in A nor B . This probability reads:

$$q(a) = \frac{1}{\binom{N}{Ns} \binom{N}{Ns'}} \frac{N!}{(Na)! [N(s-a)]! [N(1-s'-a)]! [N(s'-s+a)]!}. \tag{23}$$

In order for the walker to reach B from A , it has to match perfectly the prescriptions (freezings) of B on these aN variables. If $t = \Theta(N^d)$, the probability of this happening is $\approx t2^{-aN}$. Additionally, variables that are frozen in both clusters must coincide, so that A and B have a non-empty intersection. For a random choice of B , this happens with probability $2^{N(s'+a-1)}$.

Consequently, the probability that the walker wanders in any other cluster of entropy s' after t steps is union-bounded by:

$$\phi(s \rightarrow s') \leq 2^{N\Sigma(s')} \sum_a q(a) t 2^{-aN} 2^{N(s'+a-1)}. \tag{24}$$

The maximum of $\frac{1}{N} \log_2 q(a)$ tends to zero when $N \rightarrow \infty$, so that:

$$\limsup_{N \rightarrow \infty} \frac{1}{N} \log_2 \phi(s \rightarrow s') \leq \Sigma(s') + s' - 1. \tag{25}$$

This quantity remains negative for all s' , as long as $\alpha > \alpha_d$. Therefore, in this regime, hopping from one cluster to the other is very unlikely in a sub-exponential number of steps, even though clusters are not disjoint. In this sense we say that ergodicity is broken. Remarkably, the space of solutions becomes non-ergodic as soon as it becomes non-trivial, at α_d .

Let us stress the importance of the entropic barriers between clusters in our analysis. In the real CSPs and optimization problems the energetic barriers are usually thought of as more important, and clusters are even sometimes described as separated (by an extensive distance). The importance of entropic barriers and the possibility of non-extensively separated clusters should, however, not be neglected in the studies of richer models.

3.2 x -Satisfiability

The notion of x -satisfiability was first introduced as a tool to study the geometrical structure of the solution space of CSPs [15]. An instance of CSP is said x -satisfiable if and only if it admits a pair of solutions separated by a Hamming distance $\sim xN$. In other words, x -satisfiability gives the distance spectrum of the solution space. This spectrum is estimated using three quantities:

- a) $d_1 = x_1(\alpha)N$: the maximum distance between two solutions inside one cluster,
- b) $d_2 = x_2(\alpha)N$: the minimum distance between two solutions from two distinct clusters,
- c) $d_3 = x_3(\alpha)N$: the maximum distance between any two solutions (presumably from two different clusters).

The first of these quantities is estimated by noting that the maximum distance between any two solutions in a given cluster, i.e. its diameter, equals its entropy s . Therefore the maximum diameter/entropy x_1 is given w.h.p. by the largest internal entropy s_M , i.e. the largest root of $\Sigma(s) = 1 - \alpha - D(s \parallel 1 - p)$.

Now take two clusters A and B at random, and consider the probability that their distance be xN . This distance is given by the number of variables which are frozen in both clusters, but in a contradictory way, such that $\pi_A(i) \neq \pi_B(i)$. This happens independently with probability $p^2/2$ for each variable, so that the number of such variables follows a binomial law of parameter $p^2/2$. Therefore, the number $\mathcal{N}(x)$ of pairs of clusters at distance xN coincides w.h.p. with its mean value:

$$\mathbb{E}[\mathcal{N}(x)] = 2^{2(1-\alpha)N} \binom{N}{Nx} \left(1 - \frac{p^2}{2}\right)^{(1-x)N} \left(\frac{p^2}{2}\right)^{xN} \asymp 2^{Ns_2(x)} \tag{26}$$

if $s_2(x) := 2(1 - \alpha) - D(x \parallel p^2/2) > 0$, and $\mathcal{N}(x) = 0$ w.h.p. if $s_2(x) < 0$. Consequently the smallest possible distance between any two clusters is given by x_2N , where $x_2(\alpha)$ is the smallest root of $s_2(x)$. A similar argument gives the largest distance between any two solutions from two different clusters: $x_3(\alpha) = 1 - x_2(\alpha)$.

To sum up, we find that a random instance is x -satisfiable w.h.p. if $\alpha < \alpha_s(x)$, and is x -unsatisfiable w.h.p. if $\alpha > \alpha_s(x)$, with:

$$\alpha_s(x) = \begin{cases} 1 & \text{if } x \in [0, 1 - p] \cup [p^2/2, 1 - p^2/2] \\ 1 - D(x \parallel 1 - p) & \text{if } x \in [1 - p, x_0] \\ 1 - \frac{1}{2}D(x \parallel p^2/2) & \text{if } x \in [x_0, p^2/2] \\ 1 - \frac{1}{2}D(1 - x \parallel p^2/2) & \text{if } x \in [1 - p^2/2, 1] \end{cases} \tag{27}$$

where x_0 is solution to $D(x \parallel p^2/2) = 2D(x \parallel 1 - p)$.




Fig. 2 (Color online) The x -satisfiability threshold is constructed from the three distances $x_1(\alpha)$, $x_2(\alpha)$ and $x_3(\alpha)$. Below a threshold α_{gap} , distance spectra fail to detect the fragmentation of the solution space (there is no more “gap” between intra and inter-cluster distances). Below another threshold α_{sep} , clusters cease to be all well separated, although ergodicity is still broken. The condensation threshold α_c is shown for information. The dynamical threshold α_d lies outside the picture, and its value is ≈ 0.070 . In this figure $p = 0.95$

Figure 2 shows how $\alpha_s(x)$ can be constructed from the three distances x_1, x_2, x_3 . We put $\alpha_{\text{sep}} := 1 + (1/2) \log_2(1 - p^2/2) > \alpha_d$, the threshold below which some pairs of clusters have a non-empty intersection. An interesting observation is that ergodicity can still be broken even below this threshold. We can also define the $\alpha_{\text{gap}} := \alpha_s(x_0) > \alpha_{\text{sep}}$, below which distances from the same cluster and distances from distinct clusters overlap. This threshold sets the limit below which the notion of x -satisfiability fails to detect clustering. The random-subcube model allows us to make a clear and intelligible distinction between the three thresholds $\alpha_d < \alpha_{\text{sep}} < \alpha_{\text{gap}}$. We expect this distinction to hold in most CSPs.⁷

4 Random Energy Landscape

4.1 Definition

The random-subcube model can be enriched by adding the notion of energy to the definition of states. The motivation for doing this is to mimic the optimization version of CSPs (where energy is defined as the number of unsatisfied constraints), but it can also be used to reproduce some properties of glassy systems.

For each energy level E_0 , we define $\mathcal{N}(E_0) = \lfloor 2^{N\Sigma(E_0/N)} \rfloor$ valleys of energy E_0 , where $\Sigma(e_0)$ is an increasing complexity function. This function can be arbitrary, but for simplicity we will restrict our examples to the form⁸ $\Sigma(e_0) = a + be_0 - ce_0 \ln(e_0)$, where $a > 0$ corresponds to a SAT phase, and $a < 0$ to an UNSAT phase.

Each valley is defined as a subcube V , chosen at random in the same way as clusters in the previous sections, cf. (1). In the following, the freezing probability p will be fixed

⁷With the notable exception of k -XORSAT, where $\alpha_d = \alpha_{\text{sep}}$. Incidentally in k -XORSAT we also have $\alpha_c = \alpha_s$.

⁸This form corresponds to the function used to fit data from the cavity method in the k -SAT problem [6].

for all energies, but one could easily generalize the model by making it energy-dependent: $p = p(e_0)$.

The number $\mathcal{N}(E_0, S_0)$ of valleys of energy $E_0 = e_0N$ and entropy $S_0 = s_0N$ is w.h.p.:

$$\begin{aligned} \mathcal{N}(E_0, S_0) &\asymp 2^{N\Sigma(e_0, s_0)} \quad \text{if } \Sigma(e_0, s_0) := \Sigma(e_0) - D(s_0 \parallel 1 - p) \geq 0 \\ &= 0, \quad \text{otherwise.} \end{aligned} \tag{28}$$

Given a configuration σ , we define its energy as a trade-off between the energy of surrounding valleys and their distance. Let us denote the energy of a valley V by $E_0(V)$. Then the energy of σ is:

$$E(\sigma) := \min_V [E_0(V) + d(\sigma, V)] \tag{29}$$

where $d(\sigma, V)$ is the distance between σ and the nearest element of V . By definition, we say that σ belongs to the basin of attraction of V if V minimizes the sum. Observe that with this definition, it may happen that some valleys are not represented at all in the energy landscape. In the following, the term ‘‘state-bottom energy’’ shall refer to the energy E_0 of the valley minimizing the sum, and the term ‘‘state-bottom entropy’’ to the entropy S_0 of that valley.

4.2 Static Description of the Energy Landscape

What is the energy of an arbitrary configuration σ ? Let us start with the typical case: for each $e_0 = E_0/N$, we compute the distance to the nearest valley with energy E_0 . Standard arguments show that the number of valleys of state-bottom energy e_0N at distance $d = \omega N$ is governed by the exponent:

$$\Sigma(e_0) - D(\omega \parallel p/2).$$

Then, the minimum distance is given, in the $N \rightarrow \infty$ limit, by $\delta[\Sigma(e_0), p/2]N$, where $\delta(x, y)$ is solution to: $x = D(\delta \parallel y)$. Then, the typical energy is obtained as the best compromise between state-bottom energy and distance:

$$e^* = E/N = \min_{e_0} \{e_0 + \delta[\Sigma(e_0), p/2]\}, \tag{30}$$

The argmin gives the typical state-bottom energy e_0^* of a random σ .

As we just saw, most configurations have roughly the same energy, and belong to valleys with the same state-bottom energy. At finite temperature however, thermodynamics will be dominated by configurations of lower energy than e^*N . We thus need to estimate the entropy function (governing the number of configurations of given energy), the Legendre transform of which shall give us the free energy. Given a valley V of energy $E_0 = e_0N$ and entropy $S_0 = s_0N$, the number of configurations of energy E belonging to this valley is:

$$\begin{aligned} \mathcal{N}_V(E) &= 2^{S_0} \binom{E - E_0}{N - S_0}, \\ &\asymp 2^{N s_V(e|e_0, s_0)}, \quad \text{with } s_V(e|e_0, s_0) := s_0 + (1 - s_0)H\left(\frac{e - e_0}{1 - s_0}\right), \end{aligned} \tag{31}$$

where $H(x) = -x \log_2 x - (1 - x) \log_2 (1 - x)$ is the entropy function. Summing up over all valleys, the total number of configurations with energy $E = eN$ is:

$$\mathcal{N}(E) = \sum_{S_0, E_0} 2^{S_0} \binom{E - E_0}{N - S_0} 2^{N\Sigma(E_0/N, S_0/N)} \asymp 2^{Ns(e)}, \tag{32}$$

$$\text{with } s(e) = \max_{\substack{e_0, s_0 \\ \Sigma(e_0, s_0) \geq 0}} [s_V(e|e_0, s_0) + \Sigma(e_0, s_0)]. \tag{33}$$

Here we have implicitly assumed that all elements in the sphere of radius $E - E_0$ and center V are in the basin of attraction of V , as long as $E < e^*N$. This is not true in general, as some configurations may in fact belong to a more favorable basin, and may thus have lower energies. However, such configurations remain exponentially rare in comparison to the total weight of the sphere. Therefore the previous estimate holds.⁹

Canceling the derivative w.r.t. s_0 in (33) yields the saddle for s_0 :

$$\tilde{s}_0 = \frac{(1 - p)(1 - e + e_0)}{1 - p/2}. \tag{34}$$

Provided that the maximum is reached in a region where $\Sigma(e_0, \tilde{s}_0) > 0$, we get:

$$s(e) = \max_{e_0} [1 - D(e - e_0 \parallel p/2) + \Sigma(e_0)] \quad \text{for } e_c < e < e^*. \tag{35}$$

This will be valid from $e = e^*$ (for which we find $s(e^*) = 1$ as expected) down to a certain condensation energy e_c . Below that energy ($e \leq e_c$), the phase is condensed: $\Sigma(e_0, \tilde{s}_0) < 0$, and the maximum in (33) is reached on the border on the definition domain, where $\Sigma(e_0, s_0) = 0$. If we denote by $s_M(e_0)$ the biggest valley of energy e_0 (i.e. the largest of root of $\Sigma(e_0, s_0) = 0$), we get:

$$s(e) = \max_{e_0} \{s_V[e|e_0, s_M(e_0)]\} \quad \text{for } e < e_c. \tag{36}$$

So far we have worked in the microcanonical ensemble, but the same arguments hold in the canonical ensemble. In particular the condensation temperature is $T_c = (\partial s / \partial e|_{e=e_c})^{-1}$. The number of dominating states in the condensed phase follows again a Poisson-Dirichlet process [25, 26] with parameter m , where m/T is the slope of the curve $\Sigma_T(f)$ at its smallest root. The function $\Sigma_T(f)$ is the canonical counterpart of $\Sigma(e_0, s_0)$, and reads:

$$\Sigma_T(f) = \max_{e_0, s_0: f_V(T|e_0, s_0)=f} \Sigma(e_0, s_0), \tag{37}$$

where $f_V(T|e_0, s_0)$ is the single-state free energy, obtained as the Legendre transform of $s_V(e|e_0, s_0)$ in (31).

In both low-temperature phases (condensed and non-condensed), equilibrium is reached for different values of (e_0, s_0) as e varies. Said differently, the states dominating the micro-canonical measure at e and at $e + \delta e$ are completely distinct. In the canonical language, we say that the system exhibits temperature chaos [34–36]: slightly changing the temperature from T to $T + \delta T$ dramatically modifies the free-energy landscape, reshuffling the ordering of states. Consequently, correlations are nonexistent between T and $T + \delta T$. This phenomenon of free-energy crossings is illustrated (in the condensed phase) by Fig. 3, where the maximum of (36) is constructed geometrically.

⁹This argument is similar to the one used in (10), where some solutions were counted several times, but with no consequence at the exponential scale.




Fig. 3 (Color online) Illustration of the temperature chaos by construction of the microcanonical entropy in the condensed phase, cf. (36), and level crossing. We have represented three entropy curves $s_V(e)$ corresponding to three extremal states V_1 , V_2 and V_3 , the envelope of all these curves is the microcanonical entropy. These states are maximally atypical; they realize a balance between low state-bottom energies and a high state-bottom entropies, which are related by $s_0 = s_M(e_0)$ (thick curve). As the temperature (or energy) is decreased, the curves $s_V(e)$ cross each other, and the system is dominated by states of lower state-bottom energies and entropies. These data were obtained for number of valleys $\Sigma(e_0) = -0.05 - 0.5e_0 \ln e_0$ and $p = 0.2$

4.3 Relation with Dynamics

We now undertake to describe the dynamical properties of this energetic landscape. To that end we shall make use of the static picture, which we know precisely from bottom-up construction.

Our reasoning proceeds in two steps. First, we study the behavior of a single spin-flip Monte-Carlo dynamics with detailed balance evolving in a *single* valley of state-bottom energy e_0^* , and state-bottom entropy s_0^* . In a second step, we argue that the same dynamics run on the full rugged energy landscape is entirely governed by this single-state-like behavior.

The thermodynamics of a single typical state is given by:

$$s(e|e_0^*, s_0^*) = s_0^* + (1 - s_0^*)H\left(\frac{e - e_0^*}{1 - s_0^*}\right) \tag{38}$$

where $s_0^* = (1 - p)(1 - e^* + e_0^*)/(1 - p/2)$ is the typical state-bottom entropy of a random configuration, computed from (34). Supposing now that the energy landscape is made of this state only, we can easily convince ourselves that the dynamics is ergodic: the energy landscape is convex, and there are neither energetic nor entropic barriers. Both the quenched and annealed¹⁰ dynamics in this simple landscape is then given by the thermodynamics.

We know from the previous paragraphs that states extend up to energy e^* , where a “crest” connects the different valleys. This crest is actually more of a plateau, as it embeds almost all configurations: $s(e^*) = 1$. Therefore, our dynamics will remain there as long as the temper-

¹⁰By quench, resp. annealing, we mean a fast, resp. slow, change of temperature from $T_1 > T_d$ to $T_2 < T_d$ and some time (a finite number of Monte Carlo sweeps) spent at T_2 afterwards.

ature does not allow configurations of lower energy. This happens at $T_d := (\partial s / \partial e|_{e=e^*})^{-1}$, where exploring valleys starts to be more favorable.

Below T_d , the system will find itself trapped in one state, since barriers between valleys are extensive. This randomly picked state has typical properties: in particular, its bottom-energy is e_0^* and its bottom-entropy is s_0^* .

We argue that below T_d everything happens as if this trapping state was put in isolation, as we have described above. The justification comes from the fact that a valley does not “see” its neighbors as long as $e < e^*$: even though many configurations of energy $e < e^*$ in the isolated state actually belong other valleys in the full landscape, their proportion is exponentially small. In other words, although there may be some directions for which the energy barrier is lower than $e^* - e$, these directions are beaten entropically.

With this reasoning, both the quenched and annealed dynamics are described by the microcanonical entropy of a single typical state: $s_{\text{dyn}}(e) = s(e|e_0^*, s_0^*)$. We note in passing:

$$T_d^{-1} = \left. \frac{\partial s}{\partial e} \right|_{e=e^*} = \left. \frac{\partial s_{\text{dyn}}}{\partial e} \right|_{e=e^*}, \tag{39}$$

as partial derivatives w.r.t. e_0 and s inside (33) cancel at the maximum. The dynamical temperature T_d is thus well defined.

Note that this analysis is exact: the fact that it purely relies on static arguments should be attributed to the simplicity of the model. In richer mean-field models this kind of arguments might also apply, but will most probably not give the full picture.

4.4 Glassy Behaviour

The analysis of the energy landscape in the RSM is summarized in the energy-temperature phase diagram of Fig. 4. As anticipated, the behaviour of the RSM resembles the one observed in glasses and spin-glasses. The two distinct glassy transitions (dynamical and condensation), as well as the phenomenon whereby the physical dynamics gets stuck in metastable states, have been described for example in the p -spin glass [28], the spherical p -spin glass [37], the Potts glass [38] and the lattice glass [39]. Several related examples of energy-temperature diagrams were derived recently in [40].

In the aforementioned mean-field models, the static behaviour is better understood than the dynamics. Static properties are usually analyzed by the replica/cavity method, with the help of Parisi’s replica symmetry breaking scheme. A satisfactory analytic treatment of the dynamics exists only for the spherical p -spin glass [37], in which all states have the same entropy. A remarkable step towards connecting the static picture and dynamical behaviour in a rather general framework was done in [31–33]. The RSM could provide a tractable playground for studying several aspects of glassy dynamics, e.g. aging and rejuvenation [41].

We would like to emphasize the freedom we can enjoy in the definition of the energy landscape. First, arbitrary numbers and sizes of valleys at each energy e_0 , $\Sigma(e_0)$ and $p(e_0)$, can be considered. Second, the definition of the configurational energy $E(\sigma)$ in (29) could be generalized to an arbitrary function of all valleys V and σ ; $E(\sigma) = \mathcal{F}(\{V\}, \sigma)$. By tuning these parameters, one could hope to reproduce the dynamics of more complex models on a quantitative level.




Fig. 4 (Color online) Energy as a function of temperature. At temperature $T > T_d$, the system is in a liquid state: the dynamics is exploring ergodically all configurations of energy e^* . Below the dynamical ergodicity is broken. The upper “dynamical” curve shows the result of a quench/annealing in temperature, whereby the systems remains stuck in a typical state of bottom energy e_0^* . For $T_c < T < T_d$, equilibrium thermodynamics is dominated by an exponential number of states (curve “glass”). Below the condensation (Kauzmann) temperature thermodynamics is dominated by a finite number of states (curve “condensed glass”). This number is given by a Poisson-Dirichlet process of parameter m (plotted in the upper part of the diagram). The dashed line shows the result of a quench/annealing starting from an equilibrium state at temperature $T_c < T < T_d$. The bottom line shows the bottom energies of the thermodynamically dominating states. These curves were obtained for $\Sigma(e_0) = -0.05 - 0.5e_0 \ln e_0$ and $p = 0.6$

5 Conclusions

The random-subcube model is a simple exactly solvable model capturing several interesting properties of random constraint satisfaction problems. Rather than an attempt to construct a new realistic model for practical instances of constraint satisfaction problems, it allows us to identify which properties of random CSPs can be reproduced by a simple probabilistic structure, and conversely, which of these properties may be intrinsically non-trivial. Examples of reproducible properties include condensation, non-monotony of the x -satisfiability threshold, temperature chaos and dynamical freezing in metastable states. From this point of view, the RSM stands just next to the random-energy model [19], the random-code model [20–22] or the random-energy random-entropy model [36].

Since the relation between the RSM and the large- k limit of random k -SAT and k -COL is based on non-rigorous results from [8, 10], it would be interesting to establish this equivalence rigorously. Further, the RSM should be helpful for understanding some properties which are too difficult to study in more realistic models, such as finite-size corrections or some aspects of glassy dynamics.

Finally, our work addresses the broad question of producing, inferring, and representing complex and rugged structures of the hypercube. Although we retained the simplest choice of subcubes for clusters, more sophisticated alternatives could be explored and used to reproduce detailed geometrical features of solution spaces in CSPs.

Acknowledgements We would like to thank Florent Krzakala and Marc Mézard for fruitful discussions about this work, and Guilhem Semerjian and Jiří Černý for critical reading of the manuscript. This work has been partially supported by EVERGROW (EU consortium FP6 IST).

References

1. Papadimitriou, C.H.: *Computational Complexity*. Addison–Wesley, Reading (1994)
2. Kirkpatrick, S., Selman, B.: Critical behavior in the satisfiability of random boolean expression. *Science* **264**, 1297–1301 (1994)
3. Monasson, R., Zecchina, R., Kirkpatrick, S., Selman, B., Troyansky, L.: Determining computational complexity from characteristic phase transitions. *Nature* **400**, 133–137 (1999)
4. Biroli, G., Monasson, R., Weigt, M.: A variational description of the ground state structure in random satisfiability problems. *Eur. Phys. J. B* **14**, 551 (2000)
5. Mézard, M., Parisi, G., Zecchina, R.: Analytic and algorithmic solution of random satisfiability problems. *Science* **297**, 812–815 (2002)
6. Mézard, M., Zecchina, R.: Random k -satisfiability problem: from an analytic solution to an efficient algorithm. *Phys. Rev. E* **66**, 056126 (2002)
7. Kschischang, F.R., Frey, B., Loeliger, H.-A.: Factor graphs and the sum-product algorithm. *IEEE Trans. Inf. Theory* **47**(2), 498–519 (2001)
8. Krzakala, F., Montanari, A., Ricci-Tersenghi, F., Semerjian, G., Zdeborová, L.: Gibbs states and the set of solutions of random constraint satisfaction problems. *Proc. Natl. Acad. Sci.* **104**, 10318 (2007)
9. Mézard, M., Palassini, M., Rivoire, O.: Landscape of solutions in constraint satisfaction problems. *Phys. Rev. Lett.* **95**, 200202 (2005)
10. Zdeborová, L., Krzakala, F.: Phase transitions in the coloring of random graphs. *Phys. Rev. E* **76**, 031131 (2007)
11. Cocco, S., Dubois, O., Mandler, J., Monasson, R.: Rigorous decimation-based construction of ground pure states for spin glass models on random lattices. *Phys. Rev. Lett.* **90**, 047205 (2003)
12. Mézard, M., Ricci-Tersenghi, F., Zecchina, R.: Alternative solutions to diluted p -spin models and XOR-SAT problems. *J. Stat. Phys.* **111**, 505 (2003)
13. Mora, T., Mézard, M.: Geometrical organization of solutions to random linear Boolean equations. *J. Stat. Mech. Theory Exp.* **10**, P10007 (2006)
14. Mézard, M., Parisi, G.: The Bethe lattice spin glass revisited. *Eur. Phys. J. B* **20**, 217 (2001)
15. Mézard, M., Mora, T., Zecchina, R.: Clustering of solutions in the random satisfiability problem. *Phys. Rev. Lett.* **94**, 197205 (2005)
16. Achlioptas, D., Ricci-Tersenghi, F.: On the solution-space geometry of random constraint satisfaction problems. In: *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pp. 130–139. ACM, New York (2006).
17. Montanari, A., Shah, D.: Counting good truth assignments of random k -sat formulae. In: *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1255–1264. ACM, New York (2007).
18. Dubois, O., Mandler, J.: The 3-xorsat threshold. In: *FOCS*, p. 769 (2002).
19. Derrida, B.: Random-energy model: Limit of a family of disordered models. *Phys. Rev. Lett.* **45**, 79–82 (1980)
20. Shannon, C.E.: A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379–423, 623–655 (1948)
21. Montanari, A.: The glassy phase of Gallager codes. *Eur. Phys. J. B* **23**, 121–136 (2001)
22. Barg, A., Forney, G.D. Jr.: Random codes: minimum distances and error exponents. *IEEE Trans. Inf. Theory* **48**, 2568–2573 (2002)
23. Semerjian, G.: On the freezing of variables in random constraint satisfaction problems. *J. Stat. Phys.* **130**, 251 (2008), [arXiv.org:0705.2147](https://arxiv.org/abs/0705.2147)
24. Mézard, M., Parisi, G., Sourlas, N., Toulouse, G., Virasoro, M.A.: Replica symmetry breaking and the nature of the spin-glass phase. *J. Phys.* **45**, 843–854 (1984)
25. Talagrand, M.: Rigorous low temperature results for the p -spin mean field spin glass model. *Probab. Theory Relat. Fields* **117**, 303–360 (2000)
26. Pitman, J., Yor, M.: The two-parameter Poisson-Dirichlet distribution derived from a stable subordinator. *Ann. Probab.* **25**, 855–900 (1997)
27. Kauzmann, W.: The nature of the glassy state and the behavior of liquids at low temperatures. *Chem. Rev.* **43**, 219 (1948)
28. Gross, D.J., Mézard, M.: The simplest spin glass. *Nucl. Phys. B* **240**, 431 (1984)

29. Parisi, G.: Some remarks on the survey decimation algorithm for k -satisfiability. [arXiv:cs/0301015](https://arxiv.org/abs/cs/0301015) (2003)
30. Montanari, A., Ricci-Tersenghi, F., Semerjian, G.: Solving constraint satisfaction problems through belief propagation-guided decimation. [arXiv:0709.1667v1](https://arxiv.org/abs/0709.1667v1) [cs.AI] (2007)
31. Montanari, A., Semerjian, G.: From large scale rearrangements to mode coupling phenomenology. *Phys. Rev. Lett.* **94**, 247201 (2005)
32. Montanari, A., Semerjian, G.: On the dynamics of the glass transition on Bethe lattices. *J. Stat. Phys.* **124**, 103–189 (2006)
33. Montanari, A., Semerjian, G.: Rigorous inequalities between length and time scales in glassy systems. *J. Stat. Phys.* **125**, 23 (2006)
34. McKay, S.R., Berker, A.N., Kirkpatrick, S.: Spin-glass behavior in frustrated Ising models with chaotic renormalization-group trajectories. *Phys. Rev. Lett.* **48**, 767–770 (1982)
35. Bray, A.J., Moore, M.A.: Chaotic nature of the spin-glass phase. *Phys. Rev. Lett.* **58**, 57–60 (1987)
36. Krzakala, F., Martin, O.C.: Chaotic temperature dependence in a model of spin glasses. *Eur. Phys. J. B* **28**, 199–208 (2002)
37. Cugliandolo, L.F., Kurchan, J.: Analytical solution of the off-equilibrium dynamics of a long-range spin glass model. *Phys. Rev. Lett.* **71**, 173 (1993)
38. Gross, D.J., Kanter, I., Sompolinsky, H.: Mean-field theory of the Potts glass. *Phys. Rev. Lett.* **55**(3), 304–307 (1985)
39. Biroli, G., Mézard, M.: Lattice glass models. *Phys. Rev. Lett.* **88**, 025501 (2002)
40. Krzakala, F., Kurchan, J.: A landscape analysis of constraint satisfaction problems. *Phys. Rev. B* **76**, 021122 (2007)
41. Barrat, J.-L., Feigelman, M.V., Kurchan, J., Dalibard, J.: In: *Slow Relaxations and Nonequilibrium Dynamics in Condensed Matter. Les Houches Session LXXVII*, 1–26 July, 2002. Springer, Berlin (2003)