A replica analysis of the minimum weight solution to the linear equations problem

1994 J. Phys. A: Math. Gen. 27 8029

(http://iopscience.iop.org/0305-4470/27/24/014)

J. Phys. A: Math. Gen. 27 (1994) 8029–8038. Printed in the UK

A replica analysis of the minimum weight solution to the linear equations problem

D M L Barbato and J F Fontanari

Instituto de Física de São Carlos, Universidade de São Paulo, Caixa Postal 369, 13560-970 São Carlos, SP, Brazil

Received 11 July 1994, in final form 30 September 1994

Abstract. We investigate, through statistical mechanics techniques, the problem of finding N-tuples $J = (J_1, J_2, \ldots, J_N)$ with a fixed fraction κ of non-vanishing entries that solve a set of $P = \alpha N$ linear equations. In particular, we determine the region in the plane (α, κ) where these solutions exist with a probability of 1 in the limit of large N. We also obtain at insight into the microscopic structure of these solutions by calculating the distribution of the probability of an arbitrary non-vanishing entry J_i. Moreover we evaluate analytically the performances of several easy-to-implement heuristic algorithms and compare them with the optimal solution, a task that is suited very well to the statistical mechanics approach.

1. Introduction

Statistical mechanics techniques developed in the study of disordered systems, namely the replica method, have been employed successfully in the analysis of combinatorial optimization problems (Mézard *et al* 1987). The main goal of these studies is to estimate the typical value of the cost of the minimal-cost (optimal) solutions. This typical value is obtained by averaging over the minimal-cost solutions of an ensemble of realizations or instances of the optimization problem. As the cost is an extensive and therefore self-averaging quantity one expects its mean value to coincide with its most probable (typical) value when the size of the system becomes very large (Binder and Young 1986). Thus, in the thermodynamic limit the cost density converges to a unique value, independently of the instance of the optimization problem. In this sense we say that the cost density of any randomly chosen instance is given by the average cost density with probability 1 in the thermodynamic limit. This is a direct consequence of the property of self-averageness of the free-energy (or cost) which, in spite of being well established numerically, has not yet been proven rigorously (Mézard *et al* 1987). It should be emphasized, however, that the procedure of averaging over instances is highly desirable when the goal is to compare the performance of heuristic algorithms in the search for optimal or near-optimal solutions of a particular combinatorial optimization problem.

In this paper we consider the problem of finding the minimum weight solution to a set of linear equations: given the set of P pairs (S^l, t^l) $l = 1, \ldots, P$ where $S^l = (S_1^l, \ldots, S_N^l)$, find the N-tuple $J = (J_1, \ldots, J_N)$ that solves the P linear equations

$$\frac{1}{\sqrt{N}} \sum_{i=1}^{N} J_i S_i^l = t^l \qquad l = 1, \ldots, P \tag{1}$$

0305-4470/94/248029+10$19.50

and possesses the minimal number of non-zero entries. This problem belongs to the *NP*-complete class which basically means that the computational cost of any known *deterministic* algorithm for finding the optimal solution grows exponentially with the number of variables N (Garey and Johnson 1979). More specifically, we assume that the entries S_i^l and t^l are Gaussian, statistically independent, random variables with a mean of zero and unit variance. The statistical independence of these random variables and the thermodynamic limit $N \to \infty$ are the main simplifying assumptions that make the problem amenable to analytical study. The practical motivation for studying these rather far-fetched randomly generated sets of linear equations is for their widespread use in the empirical evaluation of heuristic algorithms. Moreover, although real-world sets of linear equations tend to be highly structured rather than random, we think that the solution of the random problem may be useful as it can be viewed as a zeroth-order approximation around which one should be able to expand to find solutions for more realistic problems.

To tackle this problem using statistical mechanics techniques we fix the fraction of non-zero entries of J to a constant value $\kappa \in [0, 1]$ and proceed to minimize the energy or cost function

$$E_\kappa(W, c) = \frac{1}{2}\sum_{l=1}^{P}\left(\frac{1}{\sqrt{N}}\sum_{i=1}^{N} c_i W_i S_i^l - t^l\right)^2 \tag{2}$$

where $J_i = c_i W_i$ with $c_i = 0, 1$. The binary variables c_i satisfy the constraint

$$\frac{1}{N}\sum_{i}^{N} c_i = \kappa \tag{3}$$

which guarantees that the fraction of non-zero entries of J is κ. The goal is then, for fixed values of P and κ, to determine the global minima (ground states) of the cost function—equation (2). The minimum-weight solution is determined by the minimal value of κ for which there exists at least one global minimum with cost zero. In this paper we use the replica method not only to estimate the typical cost of the optimal solutions but also to obtain microscopic information about these solutions. This can be achieved by calculating the distribution of probability that a certain entry W_i of the optimal solution takes a value within the range $[W, W + \mathrm{d}W]$. We find, for instance, that there is a finite region around $W = 0$ where this probability vanishes. We illustrate the usefulness of the statistical mechanics approach as a tool for evaluating the performance of heuristic algorithms by calculating the minimal cost associated with three simple heuristic methods which attempt to minimize (2) subject to the constraint (3). The analysis presented in this paper is based on the statistical mechanics framework proposed by Gardner to study the configuration space of neural networks (Gardner 1988, Gardner and Derrida 1988).

The remainder of the paper is organized as follows. The typical cost of the optimal solutions as well as the distribution of probability of their entries are calculated analytically in section 2. The evaluation and comparison of the performances of three heuristic algorithms for finding N-tuples J with κN non-zero entries which minimize E_κ is presented in section 3. Finally, in section 4 we summarize our results and discuss their relevance to the problem of learning on a diluted or damaged neural network composed of linear neurons.

2. The optimal solution

In order to obtain a well defined thermodynamic limit, $N \to \infty$, in the following we will assume that the number of equations is extensive, i.e. $P = \alpha N$. Within the standard

canonical formulation of statistical mechanics the appropriately normalized minimal cost is given by

$$\epsilon_m = \frac{1}{\alpha} \lim_{\beta \to \infty} \frac{\partial(\beta f)}{\partial \beta} \tag{4}$$

where f is the average free-energy density

$$-\beta f = \lim_{N \to \infty} \frac{1}{N} \langle\langle \ln Z \rangle\rangle \tag{5}$$

and Z is the partition function

$$Z = \sum_{\{c\}} \delta_{Kr}\left(\sum_i^N c_i, \kappa N \right) \int_{-\infty}^{\infty} \prod_i dW_i \exp[-\beta E_\kappa(\boldsymbol{W}, \boldsymbol{c})]. \tag{6}$$

The notation $\langle\langle \cdots \rangle\rangle$ stands for the average over the Gaussian random variables S_i^l and t^l, while δ_{Kr} is the Kronecker delta. The summation over all realizations of $\boldsymbol{c}$ that satisfy constraint (3), and the integration over all values of the real entries W_i, exhaust the configuration space of possible solutions. The limit $\beta \to \infty$ then ensures that only the configurations which minimize $E_\kappa(\boldsymbol{W}, \boldsymbol{c})$ will contribute to Z. Attention must be paid, however, when carrying out the integrals over W_i so as to avoid unwanted divergences. In fact, we must impose two additional constraints,

$$Q = \frac{1}{N} \sum_i^N c_i W_i^2 \tag{7}$$

and

$$Q^0 = \frac{1}{N} \sum_i^N (1 - c_i) W_i^2 \tag{8}$$

in order to guarantee that all integrals are convergent. Clearly, our results must not be affected by the *a priori* choice of Q^0 as it is the squared norm of the subset of entries W_i that do not contribute to the cost function—equation (2). The parameter Q, however, is germane to the thermodynamical analysis of the system: it should be chosen so as to minimize ϵ_m. In the case where there are several values of Q that minimize the cost, we will choose the smallest one, which then corresponds to the solution of minimal norm, or the so-called pseudo-inverse (Kohonen 1984).

To carry out the quenched average in (5) we employ the replica method: using the identity

$$\langle\langle \ln Z \rangle\rangle = \lim_{n \to 0} \frac{1}{n} \ln \langle\langle Z^n \rangle\rangle \tag{9}$$

we first evaluate $\langle\langle Z^n \rangle\rangle$ *for integer* n *and then continue analytically to* $n = 0$. *Using standard* techniques (Gardner 1988, Gardner and Derrida 1988) we obtain, in the thermodynamic limit

$$-\beta f = \lim_{n \to 0} \operatorname{extr} \frac{1}{n} \left\{ -2 \sum_{a<b}^n q_{ab} \hat{q}_{ab} + \sum_a^n \left(\kappa \hat{c}_a + Q^0 \hat{Q}_a^0 - Q \hat{Q}_a \right) + G_0(\hat{q}_{ab}, \hat{c}_a, \hat{Q}_a, \hat{Q}_a^0) + \alpha G_I(q_{ab}) \right\} \tag{10}$$

where

$$G_0 = \ln \sum_{\{c^a = 0,1\}} \int \prod_{a=1}^n dW^a \exp\left\{ -\sum_a^n \left[\hat{c}_a c^a + \hat{Q}_a^0 (1 - c^a)(W^a)^2 - \hat{Q}_a c^a (W^a)^2 \right] + 2 \sum_{a<b}^n \hat{q}_{ab} c^a c^b W^a W^b \right\} \tag{11}$$

and

$$G_1 = \ln \int \prod_{a=1}^{n} \frac{dy_a}{\sqrt{2\pi}} \exp\left\{ -\tfrac{1}{2} \sum_{a}^{n} y_a^2 [1 + \beta(Q+1)] - \beta \sum_{a<b}^{n} y_a y_b (q_{ab} + 1) \right\}. \tag{12}$$

The extremum in (10) is taken over all saddle-point parameters $(\hat{c}_a, \hat{q}_{ab}, \hat{Q}_a, \hat{Q}_a^0, q_{ab})$. The physical order parameter

$$q_{ab} = \frac{1}{N} \sum_{i=1}^{N} c_i^a c_i^b W_i^a W_i^b \qquad a < b \tag{13}$$

measures the overlap between two different solutions J^a and J^b.

To proceed further we make the replica-symmetric ansatz, i.e. we assume that the values of the saddle-point parameters are independent of their replica indices: $q_{ab} = q$, $\hat{q}_{ab} = \hat{q}$ $\forall a < b$ and similarly for $\hat{c}_a, \hat{Q}_a$ and $\hat{Q}_a^0$. Evaluation of (11) and (12) with this ansatz is straightforward, resulting in the following expression for the replica-symmetric average free-energy density:

$$-\beta f_{rs} = q\hat{q} - Q\hat{Q} + Q^0\hat{Q}^0 - \tfrac{1}{2}(1-\kappa)\ln \hat{Q}^0 + \kappa\hat{c}' + \tfrac{1}{2}\ln\pi - \tfrac{1}{2}\alpha \ln[1+\beta(Q-q)]$$
$$- \frac{\alpha\beta}{2} \frac{q+1}{1+\beta(Q-q)} + \int Dz \ln\left\{ 1 + \frac{\exp\left[-\hat{c}' + z^2\hat{q}/2(\hat{q}-\hat{Q})\right]}{\sqrt{\hat{q}-\hat{Q}}} \right\} \tag{14}$$

where $Dz = dz/\sqrt{2\pi}\exp(-z^2/2)$ is the Gaussian measure. We have introduced the parameter $\hat{c}' = \hat{c} - \frac{1}{2}\ln\hat{Q}^0$ which allows for the total decoupling between $\hat{Q}^0$ and the remaining saddle-point parameters which are relevant to the computation of ϵ_m. The minimal cost (4) becomes

$$\epsilon_m = \lim_{\beta\to\infty} \frac{1 + Q + \beta(Q-q)^2}{2[1+\beta(Q-q)]^2}. \tag{15}$$

The replica-symmetric saddle-point parameters $(\hat{c}', \hat{q}, \hat{Q}, \hat{Q}^0, q)$ are obtained by extremizing f_{rs}, which gives rise to a set of four coupled saddle-point equations since the equation for $\hat{Q}^0$ does not involve the other parameters. In the case where $\beta(Q-q) = x$ is finite and Q is chosen so as to minimize f_{rs}, i.e. $\partial f_{rs}/\partial Q = 0$, the task of solving the saddle-point equations is extremely simplified resulting in the following expression for the squared norm of the optimal solution (7):

$$Q = x = \frac{\Lambda_\kappa}{\alpha - \Lambda_\kappa} \tag{16}$$

where

$$\Lambda_\kappa = 2\int_{\xi}^{\infty} Dz\, z^2 \tag{17}$$

and ξ is the unique solution of

$$\kappa = 2\int_{\xi}^{\infty} Dz\,. \tag{18}$$

It is clear then that a solution with finite x only exists for $\alpha > \Lambda_\kappa$. In this region the minimal cost reduces to

$$\epsilon_m = \frac{1}{2} - \frac{\Lambda_\kappa}{2\alpha} \tag{19}$$

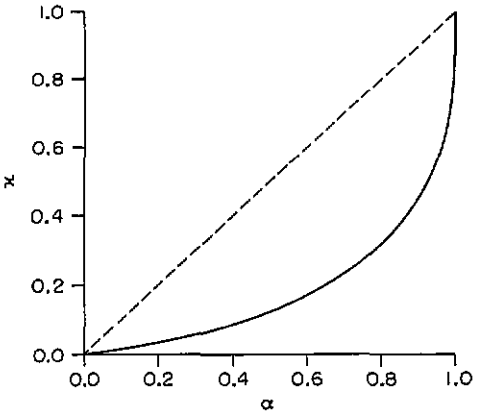

Figure 1. The full curve, $\alpha = \Lambda_\kappa$, delimits the region where a solution with cost zero exists with probability 1 (above the curve). The broken curve, $\alpha = \kappa$, shown here for comparison, is the trivial guess for this boundary.

while for $\alpha < \Lambda_\kappa$ we find $x \to \infty$, so that $\epsilon_m = 0$, in accordance with (15). In fact, the curve given by the equation $\alpha = \Lambda_\kappa$, shown in figure 1, delimits the region in the space (α, κ) where the minimal cost is non-zero. Thus, a minimum with cost zero exists with probability 1 in the thermodynamic limit provided that the condition $\alpha < \Lambda_\kappa$ is satisfied. As mentioned in the introduction, this result is a consequence of the self-averageness of the free-energy. Note that $\Lambda_\kappa = 1$ for $\kappa = 1$, as expected.

To investigate the regime $\alpha < \Lambda_\kappa$ we follow the analysis of Fontanari (1993) for the limit $\kappa = 1$ and study the solution of the saddle-point equations for Q fixed. This analysis indicates that there exists an infinity of values of Q consistent with $\epsilon_m = 0$. More specifically, any $Q > Q^P$ where

$$Q^P = \frac{\alpha}{\Lambda_\kappa - \alpha} \tag{20}$$

yields $\epsilon_m = 0$ in this regime. The choice $Q = Q^P$ corresponds to the pseudo-inverse solution (Kohonen 1984).

To obtain some insight into the microscopic structure of the optimal solution, we look at the distribution of values its entries can take on. The density of probability that a *non-vanishing* entry W_i of the optimal solution assumes a value within the range $[W, W + dW]$ is defined as (Bouten *et al* 1990)

$$\begin{aligned} P(W) &= \lim_{\beta\to\infty} \langle\!\langle \langle \delta(W_i - W)\rangle_T \rangle\!\rangle \\ &= \lim_{\beta\to\infty} \left\langle\!\!\left\langle \frac{\sum_c \int dW\, \delta(W_i - W)\, \exp[-\beta E_\kappa(W, c)]}{\sum_c \int dW\, \exp[-\beta E_\kappa(W, c)]} \right\rangle\!\!\right\rangle \end{aligned} \tag{21}$$

where $\langle\cdots\rangle_T$ stands for the thermal average. We note that $P(W)$ is normalized to κ. The averages are performed by using a standard replica trick to lift the denominator to the numerator (Bouten *et al* 1990). Assuming replica symmetry yields

$$P(W) = \lim_{\beta\to\infty} \int \frac{Dz}{\sqrt{\pi}} \frac{\exp\left[-\hat{c}' - (\hat{q} - \hat{Q})W^2 + z\sqrt{2\hat{q}}\, W\right]}{1 + \exp\left[-\hat{c}' + z^2\hat{q}/2(\hat{q} - \hat{Q})\right]/\sqrt{\hat{q} - \hat{Q}}}. \tag{22}$$

To carry out the limit $\beta \to \infty$ explicitly we must consider separately the two regimes $\alpha < \Lambda_\kappa$ where $Q = Q^P = \alpha/(\Lambda_\kappa - \alpha)$, and $\alpha > \Lambda_\kappa$ where $Q = \Lambda_\kappa/(\alpha - \Lambda_\kappa)$. In the first

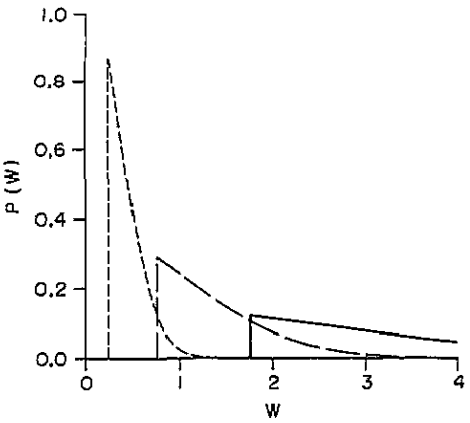

Figure 2. Distribution of probability that an arbitrary non-vanishing entry W_i takes a value within the range $[W, W + dW]$ for $\kappa = 0.5$ and $\alpha = 0.1$ (short broken curve), $\alpha = 0.5$ (long broken curve) and $\alpha = 0.8$ (full curve). Note that $\Lambda_{0.5} \approx 0.93$.

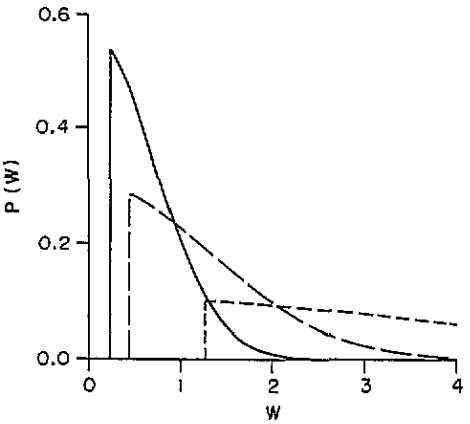

Figure 3. Same as figure 2 but for $\alpha = 1.0$ (short broken curve), $\alpha = 1.5$ (long broken curve) and $\alpha = 3.0$ (full curve).

case we find

$$P(W) = \begin{cases} 0 & \text{if} \quad |W| < \xi\sqrt{\dfrac{2\alpha}{\Lambda_\kappa(\Lambda_\kappa - \alpha)}} \\ \sqrt{\dfrac{\Lambda_\kappa(\Lambda_\kappa - \alpha)}{2\alpha\pi}} \exp\left[\dfrac{-W^2\Lambda_\kappa(\Lambda_\kappa - \alpha)}{2\alpha}\right] & \text{otherwise} \end{cases} \tag{23}$$

while in the second case we find

$$P(W) = \begin{cases} 0 & \text{if} \quad |W| < \dfrac{\xi}{\sqrt{2(\alpha - \Lambda_\kappa)}} \\ \sqrt{\dfrac{(\alpha - \Lambda_\kappa)}{2\pi}} \exp\left[\dfrac{-W^2(\alpha - \Lambda_\kappa)}{2}\right] & \text{otherwise.} \end{cases} \tag{24}$$

This distribution is depicted in figures 2 and 3 for $\kappa = 0.5$ and several values of α below and above $\Lambda_{0.5} = 0.93$. Since the distribution is symmetric we only present it for positive entries. As α approaches $\Lambda_{0.5}$ the gap increases and the non-vanishing part of the distribution becomes more uniform. Branch and bound heuristics may benefit greatly from these results as they suggest that tree searches leading to values of $|W_i|$ smaller than the intrinsic thresholds given in (23) and (24) are unlikely to be fruitful. Moreover, this result motivates the proposal of a heuristic algorithm which consists of solving the system of linear equations without imposing any constraints on the entries and then deleting the $(1-\kappa)N$ smallest entries. The performance of this algorithm will be considered in the next section.

The condition for the local stability of the replica-symmetric saddle-point (de Almeida and Thouless 1978) is given by

$$\alpha\gamma_0\gamma_1 < 1 \tag{25}$$

where γ_0 and γ_1 are the transverse eigenvalues of the matrices of second derivatives of G_0 and G_1 with respect to $\hat{q}_{ab}$ and q_{ab}, respectively. Following the analysis of Gardner and

Derrida (1988) we find that condition (25) reduces to

$$\frac{\kappa}{\alpha} < 1 \qquad \text{if} \quad \alpha > \Lambda_\kappa \tag{26}$$

and

$$\frac{\alpha\kappa}{2\Lambda_\kappa^2} < 1 \qquad \text{if} \quad \alpha < \Lambda_\kappa \tag{27}$$

which are always satisfied since $\Lambda_\kappa > \kappa$ (figure 1). We believe, then, that the results presented in this section are exact, since the replica-symmetric saddle-point is *locally* stable for all values of α and κ. The exactness of our results depends, of course, on the *global* stability of the replica-symmetric ansatz, a very difficult issue which we do not address in this paper. We mention, however, that the model we consider is a diluted version of the spherical model of a spin glass (Kosterlitz *et al* 1976), whose rigorous solution coincides with the replica-symmetric solution.

3. Heuristic algorithms

In this section we consider three easy-to-implement heuristic algorithms that generate N-tuples with κN non-vanishing entries and attempt to minimize the cost function E_κ for α and κ fixed.

In the first heuristic algorithm, termed *quenched cutting*, we set to zero $(1-\kappa)N$ arbitrary entries J_i, say $i = 1+\kappa N, \ldots, N$, and choose the remaining ones so as to minimize the cost function

$$E^q(J) = \frac{1}{2}\sum_{l=1}^{P}\left(\frac{1}{\sqrt{N}}\sum_{i=1}^{\kappa N} J_i S_i^l - t^l\right)^2. \tag{28}$$

Clearly, if $\kappa \geqslant \alpha$ this procedure is guaranteed to find a solution with cost zero. As pointed out by Bouten *et al* (1990) and Barbato and Fontanari (1993), the minimal normalized cost associated with (28) can be obtained by replacing α by α/κ in (19) evaluated at $\kappa = 1$. This yields

$$\epsilon_q = \frac{1}{2} - \frac{\kappa}{2\alpha} \tag{29}$$

for $\alpha > \kappa$ and $\epsilon_q = 0$ otherwise.

The second heuristic we consider, *selective cutting*, consists of finding the N-tuple J that minimizes the cost

$$E_{\kappa=1}(J) = \frac{1}{2}\sum_{l=1}^{P}\left(\frac{1}{\sqrt{N}}\sum_{i=1}^{N} J_i S_i^l - t^l\right)^2 \tag{30}$$

without imposing any constraint on the number of vanishing entries. Once J is known one sets to zero all entries such that $|J_i| < \lambda$ where λ is chosen so as to guarantee that the fraction of non-vanishing entries equals κ. The performance of selective cutting is measured by its normalized cost defined as

$$\epsilon_s = \frac{1}{\alpha N}\lim_{\beta\to\infty}\left\langle\!\!\left\langle\left\langle\frac{1}{2}\sum_{l}^{P}\left(\frac{1}{\sqrt{N}}\sum_{i}^{N} J_i\Theta(|J_i|-\lambda)S_i^l - t^l\right)^2\right\rangle_{\mathrm{T}}\right\rangle\!\!\right\rangle \tag{31}$$

where the thermal average $\langle\cdots\rangle_T$ is with the Boltzmann weights $\exp[-\beta E_{\kappa=1}(J)]$. Here $\Theta(x) = 1$ if $x > 0$ and 0 otherwise. The relation between λ and κ can be obtained by calculating explicitly the fraction of entries such that $|J_i| > \lambda$, i.e.

$$\kappa = \left\langle\!\!\left\langle \left\langle \frac{1}{N}\sum_i^N \Theta(|J_i| - \lambda) \right\rangle_T \right\rangle\!\!\right\rangle \tag{32}$$

which yields

$$\kappa = 2\int_{\lambda/\sqrt{Q_1}}^{\infty} Dz \tag{33}$$

where Q_1, the squared norm of the N-tuples that minimize the cost function $E_{\kappa=1}$, is obtained by setting $\kappa = 1$ in (16) and (20). Since the choice of κ determines the value of the threshold λ, the tempting strategy of equating λ to the intrinsic thresholds defined in (23) and (24), which are themselves also determined by κ, cannot be used because it gives rise to an ambiguity in the value of κ. The evaluation of the averages in (31) is rather involved but can easily be carried out following steps essentially identical to those used to calculate $P(W)$ in section 2. The final result is

$$\epsilon_s = \frac{1-\Lambda_\kappa}{2}\left(1 - \Lambda_\kappa + \frac{\alpha}{1-\alpha}\right) \qquad \alpha < 1 \tag{34}$$

and

$$\epsilon_s = \frac{\Lambda_\kappa(1-\Lambda_\kappa)y^3 + (1-\Lambda_\kappa^2)y^2 + (2-\Lambda_\kappa)y + 1}{2(1+y)^2} \qquad \alpha > 1 \tag{35}$$

where $y = 1/(\alpha - 1)$. In the limit of large α this last equation can be written as

$$\epsilon_s = \frac{1}{2} - \frac{\Lambda_\kappa}{2\alpha} + \frac{\Lambda_\kappa(1-\Lambda_\kappa)}{2\alpha^2} + O(\alpha^{-3}) \tag{36}$$

which shows that selective cutting has the same asymptotic behaviour as the optimal solution—equation(19).

Finally, we consider the third heuristic, *random cutting*, which, similarly to selective cutting, consists of modifying the N-tuple J that minimizes $E_{\kappa=1}$. Once J is known, $(1-\kappa)N$ randomly chosen entries are set to zero so as to satisfy the condition that the solution must possess κN non-vanishing entries. The performance of this heuristic algorithm is measured by the quantity

$$\epsilon_r = \frac{1}{\alpha N}\lim_{\beta\to\infty}\left\langle\!\!\left\langle\!\!\left\langle \frac{1}{2}\sum_l^P \left(\frac{1}{\sqrt{N}}\sum_i^N c_i J_i S_i^l - t^l\right)^2 \right\rangle_T \right\rangle_c \right\rangle\!\!\right\rangle \tag{37}$$

where $\langle\cdots\rangle_c$ stands for the average over the statistically independent random variables c_i distributed according to

$$p(c_i) = \kappa\delta(c_i - 1) + (1-\kappa)\delta(c_i)\,. \tag{38}$$

As $E_{\kappa=1}$ does not depend on c_i, the average over these random variables can be readily performed resulting in the following expression:

$$\epsilon_r = \frac{1}{\alpha N}\lim_{\beta\to\infty}\left\langle\!\!\left\langle\!\!\left\langle \left[\frac{1}{2}\sum_l^P \left(\frac{\kappa}{\sqrt{N}}\sum_i^N J_i S_i^l - t^l\right)^2 - \frac{\alpha\kappa(1-\kappa)}{2}\sum_i^N W_i^2 (S_i^l)^2\right]\right\rangle_T \right\rangle\!\!\right\rangle. \tag{39}$$

The evaluation of the quenched $\langle\!\langle\cdots\rangle\!\rangle$ and the thermal $\langle\cdots\rangle_T$ averages follows the same procedure mentioned before yielding

$$\epsilon_r = \frac{(1-\kappa)^2}{2} + \frac{\alpha\kappa(1-\kappa)}{2(1-\alpha)} \qquad \alpha < 1 \tag{40}$$

and

$$\epsilon_r = \frac{1}{2} + \frac{\kappa(1-\kappa)}{2(\alpha-1)} - \frac{\kappa(2-\kappa)}{2\alpha} \qquad \alpha > 1 . \tag{41}$$

The asymptotic behaviour of this cost is

$$\epsilon_r = \frac{1}{2} - \frac{\kappa}{2\alpha} + \frac{\kappa(1-\kappa)}{2\alpha^2} + \mathcal{O}(\alpha^{-3}) \tag{42}$$

which coincides with the behaviour of quenched cutting (29) for large α.

In figures 4(a) and (b) we show the normalized costs as predicted by the three heuristic algorithms presented above, together with the optimal solution, for $\kappa = 0.1$. It is remarkable that in the regime $\alpha < 1$ (figure 4(a)) there exists a range of α where selective cutting outperforms quenched cutting, though the latter should always be preferred for $\alpha < \kappa$. This region increases as κ decreases. Another rather unexpected result is the best performance of random cutting in comparison with selective cutting near $\alpha = 1$. This seems to indicate that the small entries play a very important role on the stability of a system of linear equations in the saturation regime. In fact, we have also verified that deleting the $(1-\kappa)N$ largest entries produces the lowest cost, near $\alpha = 1$, as compared with selective and random cutting. For $\alpha > 1$ (figure 4(b)) selective cutting outperforms quenched cutting for all α not too near $\alpha = 1$ and it actually approaches the optimal solution in the regime of large α.

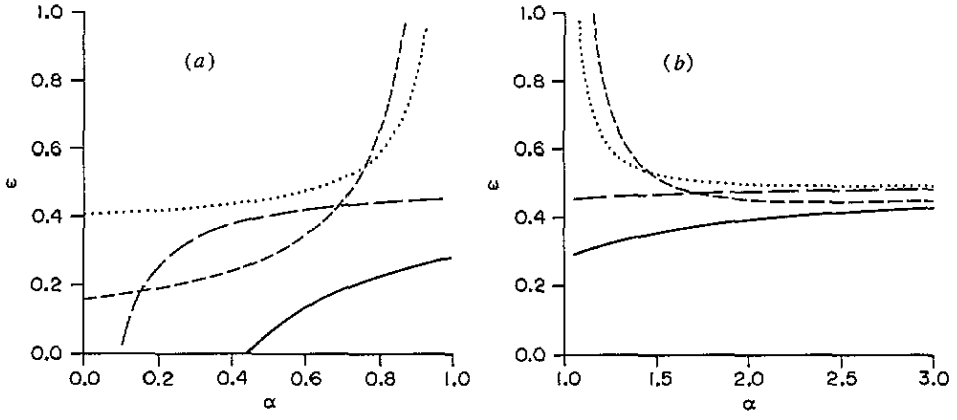

Figure 4. (a) Normalized cost for $\kappa = 0.1$ and $\alpha < 1$ as predicted by the heuristic algorithms: random cutting (dotted curve), selective cutting (short broken curve) and quenched cutting (long broken curve). The full curve is the optimal solution. (b) Same as (a) but for $\alpha > 1$.

4. Conclusion

We have shown the usefulness of the statistical mechanics approach for estimating the typical cost of the optimal solutions of a combinatorial optimization problem—the minimum-weight solution to linear equations. This approach is akin to the rather recent average-case complexity analysis (Traub and Wozniakowski 1994) which contrasts strongly with the more traditional worst-case analysis of computing theory. Actually, the usefulness of worst-case analysis for large problems is rather doubtful since the worst cases are very unlike to arise. For instance, worst-case analysis tells us that given κ and α the problem of determining whether there exists a solution to (1) is *NP*-complete (Garey and Johnson 1979). Nevertheless, the statistical mechanics approach tells us that for $\alpha < \Lambda_\kappa$ (see figure 1) a solution exists with a probability that tends to 1 in the limit of very large systems.

Moreover, the statistical mechanics approach can also provide valuable information on the microscopic structure of the optimal solutions as illustrated by figures 2 and 3 where the density of probability of a non-vanishing entry $J_i = W_i$, taking a value within the range $[W, W + dW]$, is depicted.

Any fair comparison between the performances of heuristic algorithms aimed at solving a particular optimization problem must necessarily involve averaging over a large number of instances of the problem. Thus the statistical mechanics approach is very well suited to this comparative task as shown in figures 4(*a*) and (*b*) where the minimal costs predicted by several heuristic algorithms are presented together with the optimal solution. It is interesting to note that the performance of a class of general purpose heuristic algorithms—genetic algorithms (Goldberg 1989)—has recently been analysed within the statistical mechanics framework (Prügel-Bennett and Shapiro 1994), though the optimization problem considered by those authors was much simpler than the one discussed in this paper.

A point worth emphasizing is the equivalence between the problem of finding a solution J to a set of linear equations and the problem of learning (or memorizing) a random set of input/output pairs (S^l, t^l) in a linear perceptron of synaptic weights J (Fontanari 1993). In this sense, the results presented in this paper may also be viewed as a study of the effects of different types of lesion or dilution on such a neural network. On the one hand, the optimal solution and the quenched cutting are models for lesions that take place *before* the learning process, being then analogous to the annealed and quenched dilution, respectively, discussed by Bouten *et al* (1990) and Barbato and Fontanari (1993). On the other hand, both selective and random cutting attempt to model lesions that occur *after* the learning process. In particular, the comparison between the effects of the damage caused by selective and random cutting on the performance of the network is useful to single out the role of the strength of the synaptic weights: the results shown in figures 4(*a*) and (*b*) suggest that small weights are important only when the network is near its saturation capacity ($\alpha = 1$), otherwise their deletion causes less harm than the random elimination of weights. It should also be interesting to investigate the effects of these different types of lesions on the generalization capability of neural networks.

Acknowledgment

This work was supported in part by Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq).

References

Barbato D M L and Fontanari J F 1993 *J. Phys. A: Math. Gen.* **26** 1847
Binder K and Young A P 1986 *Rev. Mod. Phys.* **58** 801
Bouten M, Engel A, Komoda A and Serneels R 1990 *J. Phys. A: Math. Gen.* **23** 4643
de Almeida J R and Thouless D J 1978 *J. Phys. A: Math. Gen.* **11** 983
Fontanari J F 1993 *J. Phys. A: Math. Gen.* **26** 6147
Gardner E 1988 *J. Phys. A: Math. Gen.* **21** 257
Gardner E and Derrida B 1988 *J. Phys. A: Math. Gen.* **21** 271
Garey M R and Johnson D S 1979 *Computers and Intractability: a Guide to the Theory of NP-Completeness* (San Francisco: Freeman)
Goldberg D E 1989 *Genetic Algorithms in Search, Optimization and Machine Learning* (Reading, MA: Addison-Wesley)
Kohonen T 1984 *Self-organization and Associative Memory* (Berlin: Springer)
Kosterlitz J M, Thouless D J and Jones R C 1976 *Phys. Rev. Lett.* **36** 1217
Mézard M, Parisi G and Virasoro M A 1987 *Spin Glass Theory and Beyond* (Singapore: World Scientific)
Prügel-Bennett A and Shapiro J L 1994 *Phys. Rev. Lett.* **72** 1305
Traub J F and Wozniakowski H 1994 *Sci. Am.* **270** 102