Pergamon

Neural Networks, Vol. 7, No. 3, pp. 477–490, 1994

0893-6080/94 $6.00 + .00

CONTRIBUTED ARTICLE

The Invisible Hand Algorithm: Solving the Assignment Problem With Statistical Physics

J. J. KOSOWSKY AND A. L. YUILLE

Harvard University

(*Received* 13 *August* 1992; *revised and accepted* 1 *October* 1993)

Abstract—*We propose a novel method for solving the assignment problem using techniques adapted from statistical physics. We derive a convex effective energy function whose unique minimum corresponds to the optimal assignment. Steepest descent results in a continuous-time dynamical system that is guaranteed to converge arbitrarily close to the optimal solution. Our algorithm has an appealing economic interpretation and has very interesting connections to the discrete auction algorithm proposed by Bertsekas. We also derive an alternative discrete algorithm for minimizing the effective energy based on a theorem by Sinkhorn.*

Keywords—Assignment problem, Statistical physics, Auction algorithm, Interior point method, Optimization, Neural networks.

1. INTRODUCTION

The assignment problem, also known as the bipartite weighted matching problem, is a classical combinatorial optimization problem. We are given N persons and N objects; each person–object pairing has an associated benefit, A_{ia}, corresponding to the benefit that the ith person associates with the ath object. The goal is to assign exactly one object to each person so that the resulting one-to-one correspondence maximizes the total benefit. Each possible assignment can be thought of as an $N \times N$ permutation matrix so that the space of all admissible assignments is isomorphic to S_N, the group of all permutations of N elements. Mathematically, we want to find the permutation $\pi \in S_N$ that *maximizes* the benefit function:

$$B_{\text{ass}}[\Pi] = \sum_{i,a=1}^{N} A_{ia}\Pi_{ia} = \sum_{a=1}^{N} A_{\pi(a)a} \qquad (1)$$

where Π indicates the permutation matrix corresponding to the permutation $\pi \in S_N$, that is, the i,jth entry of Π is $\delta_{i\pi(j)}$.[1] The best known algorithm for solving the assignment problem is the so-called *Hungarian Method* that requires $O(N^4)$ operations and was first proposed by Kuhn (1955). Over the years, many other polynomial time algorithms have been developed. In 1979, Bertsekas proposed the auction algorithm for finding the optimal assignment (Bertsekas, 1981). Both theoretical analysis and computational experience have shown the auction algorithm to be superior to all of its principal competitors, especially when the benefit matrix is sparse (Bertsekas & Castañon, 1989; Bertsekas & Eckstein, 1988; Bertsekas & Tsitsiklis, 1989).

All of these algorithms use iterative methods for finding the optimal assignment from the discrete feasible set of $N!$ possible assignments. However, we propose to use continuous-time dynamical systems, motivated by statistical physics, to solve the assignment problem.

Statistical physics techniques, and deterministic annealing in particular, have proven very useful in motivating heuristic algorithms for solving a wide range of optimization problems (Hopfield & Tank, 1985; Durbin & Willshaw, 1987; Peterson & Söderberg, 1989; Simic, 1990; Yuille, 1990). Statistical physics provides a unified framework for modeling diverse problems in vision and brain development in which many existing

Acknowledgements: We would like to acknowledge support from DARPA and the Air Force with contracts AFOSR-89-0506 and F49620-92-J-0466. We also thank the Brown, Harvard, and M.I.T. Center for Intelligent Control Systems for a United States Army Research Office Grant Number DAAL03-86-C-0171 and for providing an exceptionally stimulating working environment. We would also like to thank Roger Brockett, Leonid Faybusovich, and David Mumford for helpful conversations.

Requests for reprints should be sent to J. J. Kosowsky, Division of Applied Sciences, Harvard University, Cambridge, MA 02138.

[1] Because the group of permutation matrices is naturally isomorphic with the permutation group of N elements, we will generally not distinguish between the two.

theories appear as special cases and new interrelationships become evident (Geiger & Girosi, 1991; Yuille, 1990).

Although statistical mechanics methods have been very successful empirically, the convergence of these algorithms to the optimal solution has only been proved in simple cases. We propose a novel continuous-time deterministic algorithm that is guaranteed to converge to the optimal assignment in bounded time. Moreover, the proposed algorithm exhibits many intriguing similarities to Bertsekas' iterative auction algorithm. We also show how an iterative algorithm based on a theorem by Sinkhorn can be derived from our statistical physics framework.

In Section 2, we use statistical physics to construct an effective energy and use steepest descent to derive an algorithm for solving the assignment problem. We then prove convergence to the optimal solution and demonstrate the algorithm's practicality by performing computer simulations. In Section 3, we discuss similarities to the auction algorithm. We also suggest two alternative methods for minimizing the effective energy: temperature tracking and a discrete iterative algorithm based upon a classical theorem by Sinkhorn. Finally, in Section 4, we summarize recent results that bound how fast the algorithm will converge.

2. CONSTRUCTION AND PROOF OF THE ALGORITHM

In this section, we derive an effective energy for the assignment problem and prove that generically the energy function has a unique minimum corresponding to the optimal solution.

This effective energy, which we denote as the *P-Energy*, E_P, will be shown to have the following form:

$$E_P[P:\beta] = \frac{1}{\beta}\sum_i \log\left(\sum_a e^{\beta(A_{ia}-P_a)}\right) + \sum_a P_a. \qquad (2)$$

The P-Energy admits a very simple economic interpretation, discussed further in Section 3. The $\{A_{ia}\}$ measure the benefit that the ith person receives by owning the ath object. Each person wishes to own the object that maximizes the benefit that inevitably leads to conflicts because several people may prefer the same object. Prices are introduced for each object, so that the net benefit becomes $\{A_{ia} - P_a\}$ where P_a represents the price of the ath object. Minimizing $E_P[P:\beta]$ with respect to P corresponds to adjusting the prices, in a way reminiscent of the laws of supply and demand, so that there are no remaining conflicts. As discussed in Section 3, this minimization formulation is similar to the auction algorithm proposed by Bertsekas (1981).

In the first subsection, we use saddle point techniques to obtain the P-Energy as a function of N Lagrange multipliers $\{P_a\}$ and the temperature parameter T. We show in the second subsection that the P-Energy is weakly convex, bounded below, and hence has a unique minimum up to a constant translation invariance. In the third subsection, we derive bounds on the P-Energy and its trajectories as a function of temperature and prove that steepest descent on the P-Energy at finite nonzero temperature solves the assignment problem. Finally, we present simulation results in the final subsection.

2.1. Deriving the P-Energy

This subsection motivates our introduction of the P-Energy (2) by showing how it can be derived from the saddle point approximation to a statistical physics system. Readers willing to accept the P-Energy without this statistical physics motivation can skip to the subsequent subsections that prove that minimizing this P-Energy does indeed solve the assignment problem.

We first reformulate the assignment problem (1) by introducing an $N \times N$ decision matrix of binary variables $\{V_{ia}\}$. Let $\{A_{ia}\}$ denote the benefit matrix. We *minimize* the energy function

$$E_{ass}[V] = -B_{ass}[V] = -\sum_{i,a=1}^{N} V_{ia}A_{ia}, \qquad (3)$$

with respect to the binary decision matrix $\{V_{ia}\}$, subject to the constraint that there is a unique 1 in each row and each column. This minimization problem is clearly equivalent to our original formulation of maximizing the benefit function $\sum_{a=1}^{N} A_{\pi(a)a}$ where π is an element of S_N, the permutation group of N elements.

In statistical physics (e.g., see Parisi, 1988), one considers a system in thermal equilibrium at finite temperature. The Gibbs distribution represents the probability of the system existing in any specific state. The most probable state is the one with lowest energy; at zero temperature it occurs with probability one. Hence, minimizing the energy function is equivalent to finding the most probable state. Note that this statistical formulation is the basis for simulated annealing (Kirkpatrick, Gelatt, & Vecchi, 1983).

We define the Gibbs distribution (e.g., see Parisi, 1988),

$$P[V] = \frac{1}{Z} e^{-\beta E_{ass}[V]}, \qquad (4)$$

where $\beta = (1/T)$ is an inverse temperature. Z is a normalization constant, called the *partition function*, and is given by

$$Z = \sum_{\mathbf{V}} e^{-\beta E_{ass}[\mathbf{V}]}, \qquad (5)$$

where the sum is taken over all admissible configurations $\mathbf{V}$. For the assignment problem, admissible configurations are those that satisfy the above row and column constraints.

Let $\{\bar{V}_{ia}\}$ denote the expected value of the $\{V_{ia}\}$ with respect to the Gibbs distribution. In statistical physics, the $\{\bar{V}_{ia}\}$ are called the *mean fields.* It is straightforward to verify that

$$\bar{V}_{ia} = \frac{1}{\beta}\frac{\partial \log Z}{\partial A_{ia}}. \tag{6}$$

As $T = (1/\beta) \to 0$, the mean fields minimize the energy function and hence correspond to solutions of the assignment problem.

Unfortunately, it is usually not possible to use the above partition function to calculate the mean fields because we typically cannot compute the partition function analytically. Instead, we may derive a *saddle point approximation* to it. Recall (e.g., see Amit, 1989) that this approximation includes the dominant contribution to the partition function and is generically exact as the temperature goes to zero. For the assignment problem, this approximation is justified because we prove in Section 2.3 that the saddle point is unique and corresponds to the minimum energy state and hence the optimal assignment as the temperature goes to zero.

To obtain the saddle point approximation, we adapt a technique from Peterson and Söderberg (1989). We embed the space of binary decision matrices into $gl(N, \mathbf{R})$, the space of $N \times N$ matrices over $\mathbf{R}$, and rewrite the partition function as

$$Z = \sum_V \int [dS] e^{-\beta E_{\text{ass}}[S]} \prod_{i,a} \delta(S_{ia} - V_{ia}), \tag{7}$$

where we sum over all configurations $\{V_{ia}\}$ satisfying the following row and column constraints: (i) for all i, there exists a unique a such that $V_{ia} = 1$, (ii) for all a there exists a unique i such that $V_{ia} = 1$. These row and column conditions encapsulate the fact that the space of admissible binary decision matrices is precisely the set of permutation matrices.

Using Fourier theory, we can represent a delta function as an integral of an exponential, $\delta(x) = (1/2\pi) \int_I e^{-wx}\, dw$, where the integral is taken along the imaginary axis. We can thus replace the delta functions in eqn (7) by integrals in the Fourier variables $\{U_{ia}\}$. We will drop constant factors, such as $(1/2\pi)$ because they only scale the partition function by a constant independent of β. Imposing the column constraints by adding an additional vector of delta functions in the integral, we get that

$$Z = \sum_V \int\int_I [dS][dU] e^{-\beta E_{\text{ass}}[S]} e^{-\sum_{i,a} U_{ia}(S_{ia} - V_{ia})} \times \prod_a \delta\left(\sum_i S_{ia} - 1\right), \tag{8}$$

where we can now sum over all configurations satisfying just the row constraints. We can once again replace the delta functions by their Fourier representations, thereby introducing a new set of auxiliary variables $\{P_a\}$ and yielding

$$Z = \sum_V \int\int_I \int_I [dS][dU][dP] e^{-\beta E_{\text{ass}}[S]} e^{-\sum_{i,a} U_{ia}(S_{ia} - V_{ia})} \times e^{\sum_a P_a(\sum_i S_{ia} - 1)}. \tag{9}$$

We can now explicitly impose the row constraints by interchanging the order of summation and integration and then computing the summation over all states, V, satisfying these row constraints. The resulting summation yields

$$\begin{aligned} \sum_V e^{\sum_{i,a} U_{ia} V_{ia}} &= \sum_V \prod_i e^{\sum_a U_{ia} V_{ia}} \\ &= \prod_i \left(\sum_a e^{U_{ia}}\right) \\ &= e^{\sum_i \log(\sum_a e^{U_{ia}})}. \end{aligned} \tag{10}$$

Combining these results gives

$$Z = \int\int_I \int_I [dS][dU][dP] e^{-\beta E_{\text{eff}}[S,U,P]}, \tag{11}$$

where the effective energy E_{eff}, is then

$$E_{\text{eff}}[S, U, P : \beta] = E_{\text{ass}}[S] + \frac{1}{\beta}\sum_{i,a} S_{ia} U_{ia} - \frac{1}{\beta}\sum_i \log\left(\sum_a e^{U_{ia}}\right) + \sum_a P_a\left(\sum_i S_{ia} - 1\right). \tag{12}$$

Note that we have rescaled the variables $P_a \mapsto (1/\beta)P_a$, so that the critical values of P will scale properly with $\beta = 1/T$. In Kosowsky and Yuille (1991), we show that with this change of variables, $|P_a|$ is bounded by $\max_{i,a} A_{ia}$ as $\beta \to \infty$.

We apply saddle point approximation to the partition function (11) by determining the critical points of the effective energy.

$$0 = \frac{\partial E_{\text{eff}}}{\partial S_{ia}} = -A_{ia} + \frac{1}{\beta} U_{ia} + P_a, \tag{13a}$$

$$0 = \frac{\partial E_{\text{eff}}}{\partial U_{ia}} = \frac{1}{\beta} S_{ia} - \frac{1}{\beta}\frac{e^{U_{ia}}}{\sum_b e^{U_{ib}}}, \tag{13b}$$

$$0 = \frac{\partial E_{\text{eff}}}{\partial P_a} = \sum_i S_{ia} - 1. \tag{13c}$$

Equation (13c) suggests that we can interpret the $\{P_a\}$ as Lagrange multipliers enforcing the column constraints $\sum_i S_{ia} = 1$ for all a.

By rewriting eqn (13a) as $U_{ia} = \beta(A_{ia} - P_a)$, we can eliminate $\{U_{ia}\}$, obtaining the following simultaneous equations

$$\begin{aligned} S_{ia} &= \frac{e^{\beta(A_{ia} - P_a)}}{\sum_b e^{\beta(A_{ib} - P_b)}}, \quad \forall i, a \\ \sum_i S_{ia} &= 1, \quad \forall i. \end{aligned} \tag{14}$$

We can further eliminate the $\{S_{ia}\}$ variables, resulting in a set of N simultaneous equations for the N unknowns $\{P_a\}$

$$\sum_i \frac{e^{\beta(A_{ia}-P_a)}}{\sum_b e^{\beta(A_{ib}-P_b)}} = 1, \quad \forall a. \tag{15}$$

Given a vector $\{P_a\}$ solving the above simultaneous equations, we can use eqns (14) and (13a) to recover the corresponding values for $\{S_{ia}\}$ and $\{U_{ia}\}$. Thus, the critical points of the effective energy are completely characterized by solutions to eqn (15).

Equation (14), shows that $\{S_{ia}\}$ takes values in the range [0, 1]. We will prove in Section 2.3 that as $T = 1/\beta \to 0$, the matrix $\{S_{ia}\}$ converges to a permutation matrix representing the optimal assignment. Observe that eqns (14) and (15) are invariant under the constant translation $\{P_a\} \mapsto \{P_a + K\}$, where K is an arbitrary constant. This translational degree of freedom leaves $\{S_{ia}\}$ and the effective energy, E_{eff}, unchanged.

A general property of the saddle point approximation (e.g., see Amit, 1989) is that the mean of $\{V_{ia}\}$ is approximated by the value of $\{S_{ia}\}$ at the saddle. We will prove that in the limit $\beta \to \infty$, the matrix $\{S_{ia}\}$ converges to the solution of the assignment problem. To analyze the saddle point approximation, it is convenient to express the effective energy as a function of $\{P_a\}$ alone by using the saddle point equations (13a) and (13b) to eliminate $\{S_{ia}\}$ and $\{U_{ia}\}$. This yields

$$\begin{aligned} E_{\text{P}}[P:\beta] &\stackrel{\text{def}}{=} -E_{\text{eff}}[S(P), U(P), P:\beta] \\ &= \frac{1}{\beta}\sum_i \log\left(\sum_a e^{\beta(A_{ia}-P_a)}\right) + \sum_a P_a, \end{aligned} \tag{16}$$

which we call the P-Energy.[2] The P-Energy is thus simply the negative of the effective energy made stationary with respect to $\{S_{ia}\}$ and $\{U_{ia}\}$. Because differentiating eqn (16) with respect to $\{P_a\}$ recovers eqn (15), we conclude that the effective energy, E_{eff}, and the P-Energy, E_{P}, have equivalent critical points. Observe that the P-Energy is also invariant under the constant translation $\{P_a\} \mapsto \{P_a + K\}$, where K is any constant.

2.2. Weak Convexity of P-Energy

The following theorem shows that the P-Energy is weakly convex and bounded below; hence, it has a unique minimum that can be found by a simple steepest descent algorithm.

THEOREM 1. *The Hessian of $E_{\text{P}}[P:\beta]$ is positive semi-definite with $(N-1)$ positive eigenvalues and 1 zero eigenvalue. Furthermore, the eigenvector corresponding to the zero eigenvalue is $(1, \ldots, 1)$ and thus corresponds to the constant translation invariance of the P-Energy.*

Moreover, the P-Energy is bounded below by

$$\frac{1}{\beta} N \log N + \sum_{i,a} A_{ia} \tag{17}$$

and hence has a unique minimum up to constant translation invariance.

Proof. See Appendix 1. ■

Now, consider the steepest descent equation,

$$\begin{aligned} \frac{dP_a}{dt} &= -\frac{dE_{\text{P}}[P_a]}{dP_a} \\ &= \sum_i \frac{e^{\beta(A_{ia}-P_a)}}{\sum_c e^{\beta(A_{ic}-P_c)}} - 1 \qquad (18) \\ &= \sum_{i=1}^{N} S_{ia} - 1. \qquad (19) \end{aligned}$$

Because the P-Energy has a unique minimum up to constant translation invariance, steepest descent is guaranteed to converge to the minimum. Note that $\sum_a dP_a/dt = 0$ and hence, $\sum_a P_a$ is constant along trajectories, thereby fixing the translation invariance.

2.3. Bounds on P-Energy

We use the notation $\{\tilde{P}_a(\beta)\}$ to denote the value of $\{P_a\}$ at the P-Energy minimum at fixed inverse temperature $\beta = 1/T$. We will often abbreviate this notation as $\{\tilde{P}_a\}$ when fixed inverse temperature is implied.

Recalling that $\sum_a S_{ia} = 1$, we can rewrite the P-Energy as

$$E_{\text{P}}[P, \beta] = \frac{1}{\beta}\sum_{i,a} S_{ia}\log\left(\sum_b e^{\beta(A_{ib}-P_b)}\right) + \sum_a P_a. \tag{20}$$

Now, using eqn (14) to substitute $\sum_b e^{\beta(A_{ib}-P_b)} = e^{\beta(A_{ia}-P_a)}/S_{ia}$ and then simplifying, we get that

$$\begin{aligned} -E_{\text{P}}\left[P, T = \frac{1}{\beta}\right] &= E_{\text{ass}}[S(P)] \\ &+ T\sum_{i,a=1}^{N} S_{ia}(P)\log S_{ia}(P) + \sum_{a=1}^{N} P_a\left(\sum_{i=1}^{N} S_{ia} - 1\right). \end{aligned} \tag{21}$$

This form of the P-Energy has an intuitive physical interpretation as the negative of the free energy for the corresponding statistical physics system. The first term is the original assignment energy, the second is $-T$ times a standard entropy term, and the third imposes constraints via Lagrange multipliers.

THEOREM 2. *At a minimum of the P-Energy,*

$$0 \le E_{\text{P}}[\tilde{P}(T), T] - B_{\text{ass}}[S(\tilde{P}(T), T)] \le TN \log N. \tag{22}$$

Proof. See Appendix 2. ■

A matrix is said to be *doubly stochastic* if all of its entries are positive and if each of its rows and each of

[2] Consistent with the notation of statistical physics, we have introduced a minus sign to insure that solving the assignment problem corresponds to minimizing rather than maximizing an energy function.

its columns sum to one. A famous theorem by Birkhoff (1946), states that the space of $N \times N$ doubly stochastic matrices is equal to the *permutation polytope* formed by the convex hull of the $N!$ permutation matrices. Using eqn (14), we know that at a minimum of the P-Energy, $S_{ia}[\tilde{P}, T]$ is a doubly stochastic matrix. Because the assignment problem is a linear optimization problem and the permutation polytope is convex, we conclude that at any temperature, T,

$$B_{\text{ass}}\{S[\tilde{P}, T]\} = \sum_{i,a=1}^{N} A_{ia} S_{ia}[\tilde{P}, T] \leq B^*_{\text{ass}} \tag{23}$$

where $B^*_{\text{ass}} \overset{\text{def}}{=} \max_{\pi \in S_N} \sum_{i=1}^{N} A_{i\pi(i)}$ denotes the total benefit of the optimal assignment.

Before proceeding, we need the following technical lemma.

LEMMA 3. *For any choice of* $\{P_a\}$,

$$\lim_{T \to 0} E_{\text{P}}[P, T] = \sum_{i=1}^{N} \max_a (A_{ia} - P_a) + \sum_{a=1}^{N} P_a. \tag{24}$$

Proof. See Appendix 2. ■

We will now prove that the optimal assignment benefit, B^*_{ass} equals the zero temperature limit of the P-Energy.

THEOREM 4.

$$B^*_{\text{ass}} = E_{\text{P}}[\tilde{P}, T = 0] \tag{25}$$

$$= \min_P \left[\sum_{i=1}^{N} \max_a (A_{ia} - P_a) + \sum_{a=1}^{N} P_a \right] \tag{26}$$

where $E_{\text{P}}[P, T = 0] = \min_{P \in \mathbf{R}^N} E_{\text{P}}[P, T = 0]$.

Proof. See Appendix 2. ■

This theorem shows that in the zero temperature limit, the minimum of the P-Energy corresponds to the optimal assignment benefit. This result thus confirms the validity of our statistical physics approach and the saddle-point approximation in particular.

The following theorem puts bounds on the relationships among the minimum of the P-Energy at a fixed temperature, the corresponding assignment benefit and the optimal assignment benefit. In particular, we see how well the minimum of the P-Energy at nonzero temperature approximates the optimal assignment benefit.

THEOREM 5. *Let* $\tilde{P} = \tilde{P}(T)$ *be the P-vector minimizing P-Energy at fixed temperature, T, and let* B^*_{ass} *represent the benefit of the optimal assignment. Then,*

1. $0 \leq \dfrac{dE_{\text{P}}[\tilde{P}(T), T]}{dT} = \dfrac{1}{T}\{E_{\text{P}}[\tilde{P}(T), T] - \text{B}_{\text{ass}}[S(\tilde{P}(T), T)]\} \leq N \log N$
2. $B^*_{\text{ass}} \leq E_{\text{P}}[\tilde{P}, T] \leq B^*_{\text{ass}} + TN \log N$
3. $B^*_{\text{ass}} - TN \log N \leq B_{\text{ass}}[S(\tilde{P}, T)] \leq B^*_{\text{ass}}$.

Proof. See Appendix 2. ■

For additional inequalities bounding the P-Energy and the P-vectors, see Kosowsky and Yuille (1991) and Yuille and Kosowsky (1991).

Observe that $\{S_{ia}(\tilde{P}, T)\}$ denotes the unique doubly stochastic matrix associated with the minimum of the P-Energy at temperature T. The next theorem bounds how close this doubly stochastic matrix $\{S_{ia}(\tilde{P}, T)\}$ is to the permutation matrix, Π^*, representing the optimal solution to the assignment problem, as a function of temperature. In particular, we see that as the temperature tends to zero, $S_{ia}(\tilde{P}, T) \to \Pi^*$.

THEOREM 6. *Suppose that the assignment problem associated with the* $N \times N$ *benefit matrix* $\{A_{ia}\}$ *admits a unique optimal solution,* $\{\Pi^*_{ia}\}$. *Let* Δ *equal the difference in total benefit between the optimal solution and the second best solution. Then,*

$$\max_{i,a} |\Pi^*_{ia} - S_{ia}(\tilde{P}, T)| \leq \frac{TN \log N}{\Delta}. \tag{27}$$

So, in particular,

$$\Pi^*_{ia} = \lim_{T \to 0} S_{ia}[\tilde{P}(T), T] \overset{\text{def}}{=} S_{ia}(\tilde{P}, 0). \tag{28}$$

Proof. See Appendix 2. ■

Before proving our main result, we pause to introduce the following terminology. We will say that a square matrix is *row (column) dominant* if in every row (column), the maximum element occurs in a different column (row); furthermore, we will say that the matrix is *strictly row (column) dominant* if the maximum in each row (column) is unique. It is clear that the assignment problem is trivially solvable for a strictly row or column dominant benefit matrix. For example, in a strictly row dominant matrix, the optimal assignment corresponds to each person choosing his favorite object because there are no conflicts. Moreover, it is easily checked that the optimal assignment associated with an arbitrary benefit matrix is unchanged if we add a constant vector to each row or column of the benefit matrix. Specifically, if $\{A_{ia}\}$ is an $N \times N$ benefit matrix and $\{P_a\}$ and $\{Q_i\}$ are N-vectors, then the matrix $\{A_{ia} + Q_i - P_a\}$ yields the same optimal assignment as the original benefit matrix, although the corresponding optimal assignment benefits differ by $\sum_{j=1}^{N} (Q_j - P_j)$. Economically, P_a can be interpreted as the *price* of the ath object and Q_j can be thought of as a *wage* paid to the ith person.

Now, we can use the previous results to find bounds on the temperature required to recover the optimal assignment from the P-Energy minimum. In particular, we show that minimizing the P-Energy at sufficiently low temperature yields a vector of prices $\{P_a\}$ that converts the benefit matrix into the easily solvable row dominant form.

THEOREM 7. *Suppose that the assignment problem associated with the $N \times N$ benefit matrix $\{A_{ia}\}$ admits a unique optimal solution. Let Δ equal the difference in net benefit between the optimal solution and the second best solution. Then, rounding-off each of the entries of $S_{ia}[\tilde{P}, T]$ to the nearest integer yields the permutation matrix that solves the assignment problem whenever*

$$T < \frac{\Delta}{2N \log N}. \tag{29}$$

*In fact, the matrix $\{A_{ia} - \tilde{P}_a(T)\}$ is strictly row dominant, and hence, the i*th *person is assigned the object a that maximizes $\{A_{ia} - \tilde{P}_a(T)\}$.*

If all the entries of the benefit matrix are scaled to be integers, then $\Delta \geq 1$ and it suffices to have

$$T < \frac{1}{2N \log N}. \tag{30}$$

Proof. See Appendix 2. ■

For fixed precision matrices, the temperature required for convergence is thus $O(1/N \log N)$. Because having a unique optimum to the assignment problem is the generic case, we conclude that at sufficiently low temperature T, we can solve the generic assignment problem by performing steepest descent on the P-Energy at fixed temperature T. Simply fix an arbitrary initial condition for the $\{P_a\}$ vector and then integrate the steepest descent equation

$$\frac{dP_a}{dt} = \sum_i \frac{e^{\beta[A_{ia} - P_a(t)]}}{\sum_c e^{\beta[A_{ic} - P_c(t)]}} - 1. \tag{31}$$

We are then guaranteed to obtain the unique solution to the assignment problem by rounding off each term of the matrix $\lim_{t\to\infty}[S_{ia}(t) = e^{\beta[A_{ia} - P_a(t)]} / \sum_c e^{\beta[A_{ic} - P_c(t)]}]$ to the nearest binary value.

2.4. Simulations

In order to confirm our theoretical results, we simulated the steepest descent eqn (19) using fifth order adaptive step size Runge–Kutta (e.g., see Press, Flannery, Teukolsky, & Vetterling, 1988). Because the steepest descent equation as written above contains ratios of potentially large exponentials that will tend to introduce numerical inaccuracies, we rewrote the descent equation as follows

$$\frac{dP_a}{dt} = \sum_i \frac{1}{1 + \sum_{b \neq a} e^{\beta(A_{ib} - P_b - A_{ia} + P_a)}} - 1. \tag{32}$$

The number of simulation steps required to solve the steepest descent problem is roughly proportional to β because we know from eqn (A.1) that $\partial^2 E_P / \partial P_a \partial P_b$ scales linearly with β.

The sample benefit matrices consisted of uniformly distributed random integers between 0 and 2047. Because all the entries were integers, we set $\Delta = 1$ and used Theorem 7 to fix the temperature $T = 1/(2N \log N)$. We ran simulations on benefit matrices ranging in size between 2×2 and 150×150 with unique optimal assignments. The descent algorithm yielded the correct optimal assignment in every trial. In Figure 1a and b, we plot the evolution of the prices along the steepest descent path for a 10×10 and a 100×100 benefit matrix, respectively.

To speed up the simulations, we implemented the following *temperature scaling* descent algorithm. Start at a relatively high temperature, say $T = 1$, and begin steepest descent. Continue the descent while decreasing the temperature according to an *annealing schedule.* For example, each time the steepest descent equation begins to converge, that is, the derivative approaches zero, halve the temperature T until the critical temperature $T = \Delta/(2N \log N)$ is reached. This scaling method significantly speeds up the descent simulation because the energy landscape becomes sharper as $T \to 0$, slowing down machine simulation. Because $\tilde{P}(T)$ converges as $T \to 0$, it is beneficial to get a quick approximation of $\tilde{P}(T)$ at high temperature when the energy landscape is smoother and then refine the approximations at successively lower temperatures. The gain in performance is quite sensitive to the choice of annealing schedule. In practice, the heuristic method of halving the temperature improved performance by more than two orders of magnitude allowing us easily to solve assignment problems up to size 400×400. In Figure 2a and b, we show the evolution of the P-vector with temperature scaling as a function of Runge–Kutta steps for the same sample 10×10 and 100×100 benefit matrices as before.

There are three possible criteria that we can use to determine when to stop the steepest descent simulation and still be guaranteed that we can recover the optimal assignment. One possibility, is to use the bounds on the gradient of the P-Energy derived in Yuille and Kosowsky (1991) to determine when we are sufficiently close to the P-Energy minimum to allow rounding off to succeed. Alternatively, we could simply run the steepest descent equations for the time bound described in Section 4 and then round off to the nearest binary matrix. The final possibility is to continue the descent until the matrix $\{A_{ia} - P_a\}$ becomes strictly row dominant and then stop and solve the corresponding trivial assignment problem. This final option gives the earliest stopping time because the other criteria cannot be satisfied until $\{A_{ia} - P_a\}$ is strictly row dominant.

3. COMPARISON TO OTHER ASSIGNMENT ALGORITHMS

3.1. Bertsekas' Auction Algorithm

Our saddle point statistical physics model has many interesting parallels to the fast auction algorithm of

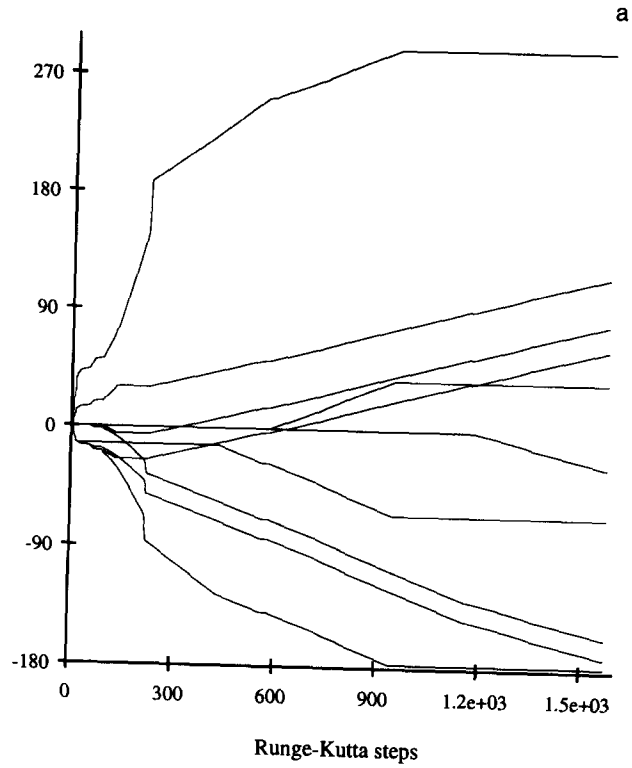

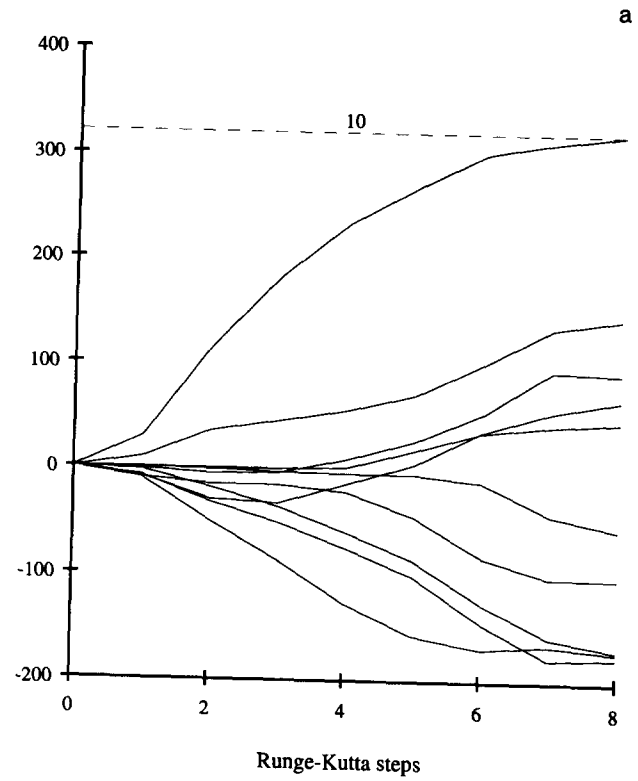

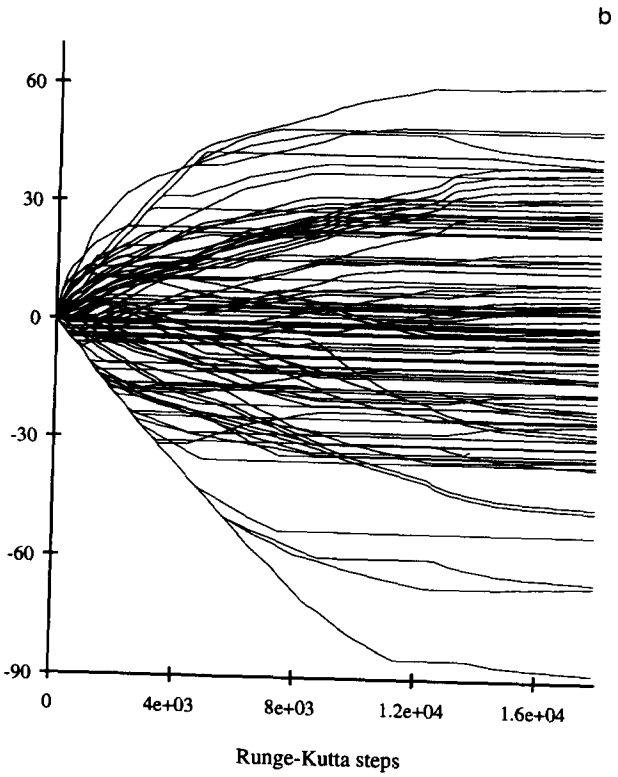

FIGURE 1. Steepest descent using Runge–Kutta *for minimizing* the P-Energy. Graph of $\{P_s\}$ versus Runge-Kutta steps at fixed low temperature bound. (a) 10 × 10 assignment ($T \approx 0.022$); (b) 100 × 100 assignment ($T \approx 0.0011$).

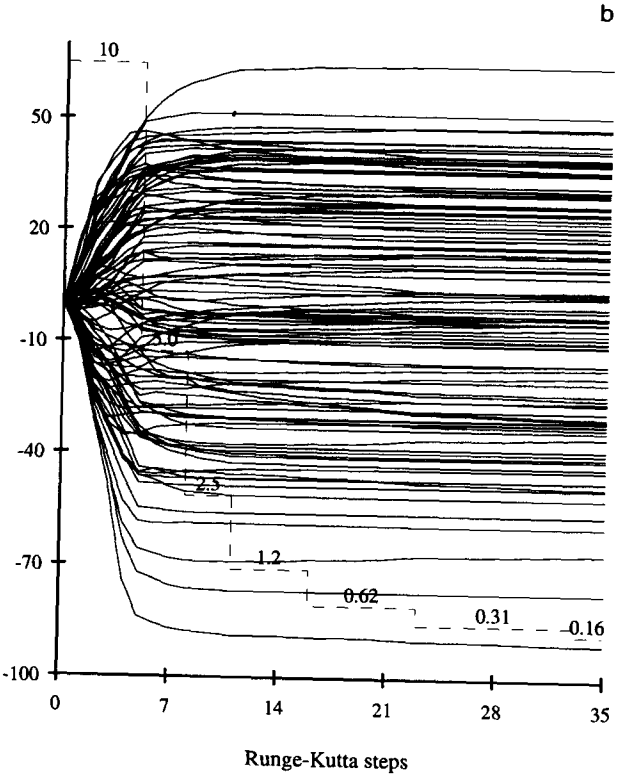

FIGURE 2. Steepest descent using Runge–Kutta with temperature scaling for minimizing the P-Energy. Graph of $\{P_s\}$ versus Runge-Kutta steps for the P-Energy steepest descent equation. The dotted line marks the temperature evolution. (a) 10 × 10 assignment; (b) 100 × 100 assignment.

Bertsekas (1981). We will first present a short summary of the auction algorithm. The auction algorithm solves the assignment problem by constructing a vector $\{P_a\}$ such that the matrix $\{A_{ia} - P_a\}$ is strictly row dominant.

According to Bertsekas' scheme, we interpret the entry A_{ia} as the *benefit* that the ith person receives by owning the ath object, and we consider P_a to be the *price* of the ath object. Person i is said to be *happy* with object b at the reigning price structure $\{P_a\}$, if

$$A_{ib} - P_b \geq \max_a (A_{ia} - P_a). \tag{33}$$

Given some fixed $\varepsilon > 0$, the person is *ε-almost happy* if

$$A_{ib} - P_b \geq \max_a (A_{ia} - P_a) - \varepsilon. \tag{34}$$

To solve the assignment problem it is clearly sufficient to find a distribution of prices such that all persons are happy. If $\varepsilon \leq \Delta/N$ whereas before Δ is the difference in total benefit between the best and second best assignments, then Bertsekas shows that it is sufficient for all persons to be ε-almost happy. This notion of *almost* optimality is called *epsilon-complementary slackness.*

The prices are adjusted by the following procedure until everyone is ε-almost happy.

1. Assign each person a unique object.
2. Fix an arbitrary price vector $\{P_a\}$.
3. Pick a person j who is not ε-almost happy. Let a_j be the object that the person values most at the current price level. Specifically,

$$a_j = \arg \max_{a \in \{1,2,\ldots,n\}} (A_{ja} - P_a). \tag{35}$$

4. Person j exchanges the object with the person currently assigned to object a_j.
5. The price P_{a_j} of object a_j is raised by the amount

$$(A_{ja_j} - P_{a_j}) - \max_{b \neq a_j} (A_{jb} - P_b) + \varepsilon. \tag{36}$$

6. Continue with step 3 until everyone is ε-almost happy.

We can view this procedure as analogous to a standard auction. Any player who can increase the benefit by more than ε by exchanging objects, switches objects and then raises the bid on the newly chosen object until the player is within ε of being indifferent between the two objects considered most valuable. If ε is greater than zero, then this procedure will always terminate. Once an object is exchanged, it will continue to be assigned to people who are ε-almost happy with the object; furthermore, because the price of an object increases by at least ε with each exchange, every object that is not ε-almost happily assigned will eventually be exchanged as the other objects become relatively more expensive.

If ε is greater than zero and less than Δ/N, then the procedure converges to an assignment that makes everyone happy and thus yields an assignment that maximizes the total benefit. To prove this assertion, Bertsekas introduces the following dual minimization problem (Bertsekas, 1990). Minimize over all $\{P_a\} \in \mathbf{R}^N$, the dual energy

$$E_{\text{dual}}[\{P_a\}] = \sum_{i=1}^{N} \max_a (A_{ia} - P_a) + \sum_{a=1}^{N} P_a \tag{37}$$

Bertsekas shows that in the limit as ε goes to zero,

$$B^*_{\text{ass}} = \min_{\{Pa\}} \left[\sum_{i=1}^{N} \max_a (A_{ia} - P_a) + \sum_{a=1}^{N} P_a \right] = \sum_{i=1}^{N} A_{i\bar{\pi}(i)} \tag{38}$$

where B^*_{ass} is the optimal assignment benefit and $\bar{\pi}$ is the permutation computed by the auction algorithm.

3.2. Comparison of Our Algorithm With the Auction Algorithm

There are many evident similarities between Bertsekas' auction algorithm and our proposed algorithm. Both methods introduce a vector $\{P_a\}$ whose underlying purpose is to transform the original benefit matrix into a row dominant matrix, thereby revealing the optimal assignment. Bertsekas motivates the introduction of the $\{P_a\}$ vector by interpreting them as prices or bids in his auction analogy. In our model, the $\{P_a\}$s were first introduced in eqn (9) and were interpreted as a vector of Lagrange multipliers enforcing the constraint that each object is assigned to exactly one person.

In the spirit of Bertsekas, if we interpret the $\{P_a\}$s in our model as *prices,* then our steepest descent algorithm can be given the following appealing economic interpretation. The entry $S_{ia} = e^{\beta(A_{ia}-P_a)}/\sum_b e^{\beta(A_{ib}-P_b)}$ can be thought of as the relative utility that person i gets by choosing object a over all the other objects. $\sum_{i=1}^{N} S_{ia}$ is then the total demand associated with object a. The total supply of each object at any price level is rigidly fixed at one. The steepest descent equation

$$\frac{dP_a}{dt} = \sum_{i=1}^{N} S_{ia} - 1 \tag{39}$$

now makes perfect sense! If total demand for object a exceeds the rigid supply, then raise its price; if supply exceeds demand, then lower the price. Equilibrium occurs when supply equals demand. Because our descent algorithm so nicely illustrates Adam Smith's (Smith, 1778) famous law of supply and demand, we will call our algorithm, the *invisible hand algorithm.*

In order to prove convergence, both algorithms convert the original problem of maximizing the assignment benefit over all $N!$ possible assignments to a dual problem of minimizing a *dual* energy function over the set of all N-vectors $\{P_a\}$. In our approach, we introduce the P-Energy and then prove that steepest descent on the P-Energy solves the assignment problem; Bertsekas uses his dual energy function to show that the auction

algorithm yields the optimal assignment. Most strikingly, referring to Lemma 3, we note that the P-Energy at zero temperature is precisely Bertsekas' dual energy function, E_{dual}! It then follows from Theorem 4 that minimizing the P-Energy at zero temperature or equivalently minimizing Bertsekas' dual energy function maximizes the assignment benefit function and solves the assignment problem. Bertsekas then uses the dual energy function and the notion of ε-complementary slackness to show that the auction algorithm gives the optimal assignment for sufficiently small positive values of ε. Similarly, in our proofs, we show that steepest descent at a sufficiently small fixed nonzero temperature suffices to solve the assignment problem. Thus, both Bertsekas' ε parameter and our temperature parameter are used to achieve an *almost optimal* solution that for sufficiently small values of the parameter corresponds to the optimal assignment.

The number of bidding rounds required for the auction algorithm to terminate is roughly inversely proportional to the value of ε because ε determines the minimum bid increment. To speed up the auction algorithm Bertsekas suggests using ε-scaling. Specifically, beginning with a large value for ε, the auction algorithm is run several times using successively smaller values of ε until $\varepsilon < \Delta/N$. The resulting prices from each cycle are used as the initial conditions for the prices in the succeeding cycle. As described in Section 2.4, a similar notion of temperature scaling applies to our invisible hand algorithm. We use the steepest descent equation to minimize the P-energy at successively lower temperatures until the temperature satisfies the bound of Theorem 7. After each descent iteration, the resulting equilibrium prices are used as the initial price distribution for the next cycle.

3.3. Tracking Down the High Temperature Solution

Instead of following the steepest descent path for the P-energy at successively lower temperatures, we could find the unique minimum of the P-Energy at $T = 1/\beta \gg 1$ and then follow the trajectory of the $\{P_a\}$s that minimize the P-Energy as the temperature decreases. Recall that at any fixed temperature, the following system of equations

$$\sum_i \frac{e^{\beta(A_{ia}-P_a)}}{\sum_b e^{\beta(A_{ib}-P_b)}} = 1, \quad \forall a, \tag{40}$$

determines the unique (up to constant translation) vector $\{P_a\}$ that minimizes the P-Energy. At high temperatures, $T = 1/\beta \gg 1$, we can solve for $\{P_a\}$ to first order in β by using a Taylor series expansion.

$$\sum_i [1 + \beta(A_{ia} - P_a)]\left[1 - \frac{\beta}{N}\sum_b (A_{ib} - P_b)\right] = N + O(\beta^2). \tag{41}$$

Hence,

$$\sum_i (A_{ia} - P_a) - \frac{1}{N}\sum_i \sum_b (A_{ib} - P_b) = 0 + O(\beta), \tag{42}$$

and

$$P_a(\beta) = \frac{1}{N}\sum_i A_{ia} - \frac{1}{N^2}\sum_{i,b} A_{ib} + O(\beta), \tag{43}$$

where we have fixed the translation invariance so that $\sum_{a=1}^{N} P_a = 0$. By fixing β to be sufficiently small, we can to any desired degree of accuracy solve explicitly for the vector $P(\beta)$ that minimizes the P-Energy, $E_{\text{P}}[P, \beta]$. Now at any fixed β, the minimum of the P-Energy is determined by solutions of

$$\frac{\partial E_{\text{P}}[P(\beta), \beta]}{\partial P_a} = 0. \tag{44}$$

We want to track the solutions as a function of β. Differentiating with respect to β gives

$$\frac{\partial^2 E_{\text{P}}}{\partial P_a \partial \beta} + \sum_b \frac{\partial^2 E_{\text{P}}}{\partial P_a \partial P_b}\frac{dP_b}{d\beta} = 0. \tag{45}$$

Because the Hessian of E_{P} is always positive semidefinite, except for the unimportant translation freedom, we can solve for $dP_b/d\beta$ once we have fixed the translation invariance by setting $\sum_a P_a$ = constant. Finally, by integrating $dP_b(\beta)/d\beta$, we can track the minimum as β increases. By Theorem 7, we know that for sufficiently large β, the matrix $\{A_{ia} - P_a(\beta)\}$ is strictly row dominant allowing the optimal assignment to be recovered easily.

This temperature tracking method is reminiscent of simulated annealing (Kirkpatrick et al., 1983) in the sense that we decrease the temperature of a statistical system to find its global minimum, but we emphasize that our method is deterministic not stochastic and that the rate of descent is not restricted by an annealing schedule. Moreover, as shown in Yuille and Kosowsky (1991), our temperature tracking approach is analogous to interior point methods of Karmakar (1990) and Faybusovich (1991). Consequently, this similarity suggests an intriguing connection between deterministic annealing and interior point methods.

In a forthcoming paper (Yuille & Kosowsky, 1994), we show that we can find $\beta_{\min} < \beta_{\max} < \infty$ such that if we begin at $\beta = \beta_{\min}$ with the fixed initial conditions $P_a(\beta_{\min}) = 0$ for all a and track the trajectory of constant P-Energy gradient, $\partial E_{\text{P}}[P_a(\beta), \beta]/\partial P_a$ = constant, as β increases to $\beta_{\max}$, then we can recover the optimal assignment by rounding off the final value of $\{S_{ia}[P(\beta), \beta]\}$ to a binary matrix.

3.4. Iterative Algorithm for Minimizing P-Energy

In both the invisible hand algorithm and in the above temperature tracking scheme, we find a solution to eqn

(15) and thus solve the assignment problem by following the trajectory of a *continuous* dynamical system. Applying a general result due to Sinkhorn (1964) on doubly stochastic matrices,[3] we can instead find an *iterative* algorithm that solves eqn (15) at any fixed temperature and thus, in particular, at sufficiently low temperature, solves the assignment problem temperature. Sinkhorn proves the following theorem:

THEOREM 8 (Sinkhorn). *Given a strictly positive $N \times N$ matrix M, there exists a unique corresponding doubly stochastic matrix $\Theta^M = EMD$, where D and E are strictly positive diagonal matrices and are themselves unique up to a multiplicative scale factor. Moreover, the iterative process of alternatively normalizing the rows and column of M to each sum to 1, converges to the corresponding doubly stochastic matrix Θ^M.*

Note that since Θ^M is doubly stochastic,

$$E_{ii} = \frac{1}{\sum_{a=1}^{N} M_{ia} D_{aa}}. \tag{46}$$

Now, because $\{e^{\beta A_{ia}}\}$ is a strictly positive matrix that we identify with $\{M_{ia}\}$, we conclude that Sinkhorn's iterative procedure yields a vector $\{P_a = -(1/\beta)\log(D_{aa})\}$, which is unique up to the constant translation invariance $\{P_a\} \rightarrow \{P_a + K\}$, such that

$$S_{ia} \overset{\text{def}}{=} \frac{e^{\beta(A_{ia} - P_a)}}{\sum_{c=1}^{N} e^{\beta(A_{ic} - P_c)}} \tag{47}$$

is doubly stochastic and hence solves eqn (15). Because this procedure works for arbitrary β, we conclude that Sinkhorn's method can be adapted to yield an iterative solution to the assignment problem. Conversely, because any positive matrix M can be expressed as $e^{\beta A_{ia}}$ where $A_{ia} = (1/\beta)\log M_{ia}$, we conclude from Theorem 1, that the steepest descent eqn (19) converges to $\{\tilde{P}_a\}$ such that $S_{ia}(\tilde{P})$ is the unique doubly stochastic matrix corresponding to M. So, $\Theta^M = EMD$, where $\{D_{aa}\} = \{e^{-\beta \tilde{P}_a}\}$ and $\{E_{ii}\} = 1/\sum_c M_{ic} e^{-\beta \tilde{P}_c}$. Hence, Theorem 1 can be interpreted as a simple new proof of Sinkhorn's theorem.

3.5. Simulations

In practice, on finite precision machines, a direct application of Sinkhorn's method leads to numerical overflow problems. The difficulty is that the exponent βA_{ia} tends to be so large that the resulting exponential overflows. Furthermore, even if we avoid overflow, we still end up alternatingly normalizing the rows and the columns by very large numbers, which is numerically unstable. In fact, using standard double precision arithmetic, this direct approach only worked when the

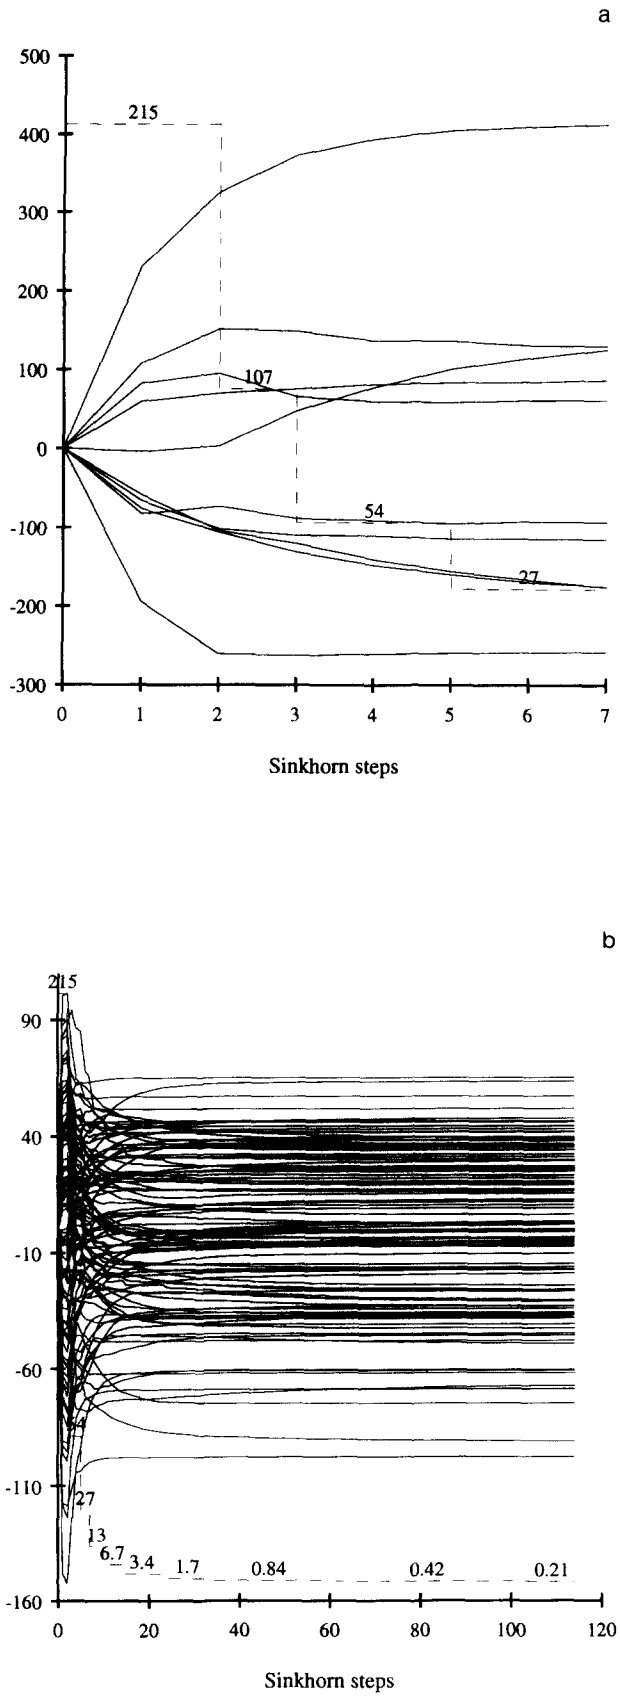

FIGURE 3. Discrete Sinkhorn method with temperature scaling. Graph of $\{P_a\}$ versus Sinkhorn iterations. The dotted line marks the temperature evolution. (a) 10 × 10 assignment; (b) 100 × 100 assignment.

[3] We would like to thank D. Mumford for bringing this result to our attention.

benefit matrix was no larger than about 7×7 and when the (integer) entries did not exceed order 10.

These numerical difficulties can be avoided by adapting the temperature scaling approach of Section 2.4 to Sinkhorn's method. To preserve numerical accuracy, at each iteration we write the normalized matrix as $e^{-\beta Q_i} e^{\beta A_{ia}} e^{-\beta P_a}$, where Q_i corresponds to normalizing the ith row and P_a corresponds to normalizing the ath column. When we normalize the rows, we additively adjust the vector $\{Q_i\}$ and when we normalize the columns, we additively adjust the vector $\{P_a\}$. The temperature scaling Sinkhorn algorithm then proceeds as follows. First choose $\beta = 1/T$ sufficiently small so that $\max_{i,a} \beta |A_{ia}|$ is of order one and initialize $Q_i = P_a = 0$. At each fixed value of β, alternatively normalize the rows and columns until the resulting matrix is almost doubly stochastic. Then, if β is less than the critical inverse temperature, $2N \log N/\Delta$, double β and proceed once again to normalize the rows and columns using the previous values for $\{Q_i\}$ and $\{P_a\}$ as the initial conditions; otherwise, we are done and we can recover the optimal assignment by rounding off the entries of the now nearly doubly stochastic matrix. Note that this method of successively halving the temperature is the exact analog of the simple annealing schedule proposed in Section 2.4. Other annealing schedules are possible, yielding somewhat different performances.

On benefit matrices consisting of uniformly distributed random integers between 0 and 2047, this temperature scaling Sinkhorn approach worked well on matrices up to about 400×400. Experimentally, this Sinkhorn-type algorithm runs more than 10 times faster than Runge–Kutta simulations of the temperature scaling steepest descent approach described in Section 2.4. In Figure 2a and b, we show the evolution of the price vector, P, as a function of Sinkhorn steps, using temperature scaling, for the same sample 10×10 and 100×100 benefit matrices as before.

4. BOUNDS ON TIME CONVERGENCE

In Section 2, we showed that for sufficiently low temperature, the optimal assignment can be recovered at a minimum of the P-Energy. We will now prove that the optimal assignment can be recovered as long as we are in a neighborhood of the P-Energy minimum or equivalently, because the P-Energy is convex, as long as the gradient of the P-Energy, $\nabla E_P = \{\partial E_P/\partial P_a\}$, is sufficiently small. In particular, this will show that the optimal assignment can be recovered by following the steepest descent trajectory for a finite predetermined amount of time. To this end, we must extend the results of the Section 2.3 to prove bounds on the P-Energy and assignment benefit function in the *neighborhood* of a minimum of the P-Energy.

THEOREM 9. *Suppose $\{A_{ia}\}$ is any $N \times N$ benefit matrix with a unique optimal assignment, Π^*, and $N \geq 4$. Let Δ equal the benefit difference between the optimal and second best assignments. If $\|\nabla E_P[P; T]\| \leq \varepsilon$, then for all i and a,*

$$|S_{ia}(P, T) - \Pi^*_{ia}| < 8 \frac{N^2 \log N}{\Delta} \left[\varepsilon \left\{ 2(\max_{i,a} A_{ia} - \min_{i,a} A_{ia}) + T \log \frac{N - 1 + \varepsilon}{1 - \varepsilon} \right\} + T \log N \right] + \varepsilon N \log(N - 1) \quad (48)$$

Proof. See Yuille and Kosowsky (1991). ■

This theorem shows that provided the temperature is sufficiently small, we do not need to get to the minimum of E_P. Instead, we can put a threshold on $\|\nabla E_P\|$ and stop the descent as soon as this threshold is reached. This will only take finite time because, as E_P is convex, we can put a lower bound on the rate of decrease of E_P. Specifically,

THEOREM 10. *Suppose $\{A_{ia}\}$ is any $N \times N$ benefit matrix with a unique optimal assignment and $N \geq 3$. Let Δ equal the benefit difference between the optimal and second best assignments. Fix temperature*

$$T < \frac{\Delta}{64N^2(\log N)^2}. \quad (49)$$

If

$$\|\nabla E_P[P; T]\| < \frac{1}{64N^2 \log N} \quad (50)$$

then, we can recover the optimal assignment by rounding off $\{S_{ia}(P, T)\}$ to the nearest binary matrix.

In particular, if we use steepest descent and start at $P = 0$, then convergence is guaranteed in time

$$t < 4096N^5(\log N)^2[T \log N + (\max_{i,a} A_{ia} - \min_{i,a} A_{ia})]. \quad (51)$$

Proof. See Appendix 3. ■

Thus, we can always find a sufficiently low temperature at which steepest descent is guaranteed to converge to the unique optimal assignment in polynomial time (in N). Note that eqn (51) refers to the time for the analog dynamical system to converge. This notion of polynomial time convergence should *not* be confused with the standard measure of complexity used in computer science or with the amount of CPU time it takes to simulate the descent equation on a digital computer.

5. CONCLUSION

In this paper, we have used statistical physics to construct a weakly convex effective energy function parameterized by temperature. We proved that for sufficiently low temperature, the unique minimum of the effective energy corresponds to an optimal solution of an associated assignment problem. The resulting continuous-time dynamical system induced by gradient descent can be used to solve this classic combinatorial

problem. Converting discrete problems into dynamical systems is desirable because this process yields inherently parallelizable algorithms suggestive of neural computation.

Moreover, this effective energy construction provides a bridge between statistical physics/neural network ideas, Sinkhorn's results about doubly stochastic matrices, and Bertsekas' discrete optimization auction algorithm. Indeed, Sinkhorn's iterative algorithm can be applied to minimize the effective energy and hence solve the assignment problem; conversely, our dynamical system can solve Sinkhorn's problem. The similarities between the effective energy construction and the auction algorithm are striking, our Lagrange multipliers correspond to Bertsekas' prices, and his dual problem is precisely our effective energy at zero temperature. Finally, steepest descent on the effective energy admits an elegant economic interpretation reminiscent of Adam Smith's invisible hand.

Although statistical physics and related neural network techniques have been successfully applied to yield good heuristic algorithms for solving combinatorial problems such as the traveling salesman problem, rigorous proofs of convergence have been lacking. In this paper, we proved convergence in the context of the assignment problem. We also presented time convergence bounds for the invisible hand algorithm that are proved in greater detail in an extension of this paper (Yuille & Kosowsky, 1991). In a forthcoming paper (Yuille & Kosowsky, 1994), we extend our results to prove convergence for a general class of optimization problems and demonstrate similarities to interior point algorithms and barrier function methods.

REFERENCES

Amit, D. J. (1989). *Modeling brain function.* Cambridge University Press, Cambridge, England.

Bertsekas, D. P. (1981). A new algorithm for the assignment problem. *Mathematical Programming,* **21,** 152–171.

Bertsekas, D. P. (1990). The auction algorithm for assignment and other network flow problems: A tutorial. *Interfaces,* **20,** 133–149.

Bertsekas, D. P., & Castañon, D. A. (1989). The auction algorithm for the minimum cost network flow problem. Laboratory for information and decision systems report LIDS-P-1925, M.I.T., Cambridge, MA.

Bertsekas, D. P., & Eckstein, J. (1988). Dual coordinate step methods for linear network flow problems. *Mathematical Programming, Series B,* **42,** 203–243.

Bertsekas, D. P., & Tsitsiklis, J. N. (1989). *Parallel and distributed computation: Numerical methods.* Englewood Cliffs, NJ: Prentice-Hall.

Birkhoff, G. (1946). Trés observaciones sobre el algebra linear. *Universidad Nacional de Tucumán Revista, Series A,* **5,** 147–151.

Durbin, R., & Willshaw, D. (1987). An analog approach to the travelling salesman problem using an elastic net method. *Nature,* **326,** 689–691.

Faybusovich, L. (1991). Interior point methods and entropy. In *Proceedings of the IEEE Conference on Decision and Control,* 2094–2095.

Geiger, D., & Girosi, F. (1991). Parallel and deterministic algorithms from MRFs: Surface reconstruction and integration. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* **13,** 401–412.

Hopfield, J. J., & Tank, D. W. (1985). Neural computation of decisions in optimization problems. *Biological Cybernetics,* **52,** 141–152.

Karmakar, N. (1990). Riemannian geometry underlying interior point methods for linear programming. Preprint.

Kirkpatrick, S., Gelatt, Jr., C., & Vecchi, M. (1983). Optimization by simulated annealing. *Science,* **220,** 671–680.

Kosowsky, J. J., & Yuille, A. L. (1991). The invisible hand algorithm: Solving the assignment problem with statistical physics. Harvard robotics laboratory technical report 91-1, Harvard University, Cambridge, MA.

Kuhn, H. W. (1955). The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly,* **2,** 83–97.

Parisi, G. (1988). *Statistical field theory.* Reading, MA: Addison-Wesley.

Peterson, C., & Söderberg, B. (1989). A new method for maping optimization problems onto neural networks. *International Journal of Neural Systems,* **1,** 3–22.

Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. (1988) *Numerical recipes in C: The art of scientific computing.* Cambridge, MA: Cambridge University Press.

Simic, P. (1990). Statistical mechanics as the underlying theory of 'elastic' and 'neural' optimization. *NETWORK: Computation in Neural Systems,* **1,** 1–15.

Sinkhorn, R. (1964). A relationship between arbitrary positive matrices and doubly stochastic matrices. *Annals of Mathematical Statistics,* **35,** 876–879.

Smith, A. (1778). *An Inquiry into the nature and causes of the wealth of nations* (2nd ed.). London: W. Strahan and T. Cadell.

Yuille, A. L. (1990). Generalized deformable models, statistical physics and matching problems. *Neural Computation,* **2,** 1–24.

Yuille, A. L., & Kosowsky, J. J. (1991). The invisible hand algorithm: Time convergence and temperature tracking. Harvard robotics laboratory technical report 91-10, Harvard University.

Yuille, A. L., & Kosowsky, J. J. (in press). Statistical physics algorithms that converge. *Neural Computation.*

APPENDIX 1

Proof of Theorem 1. We first prove that the P-Energy is weakly convex, by showing that the Hessian

$$\frac{1}{\beta}\frac{\partial^2 E_P}{\partial P_a \partial P_b} = \sum_i \frac{e^{\beta(A_{ia}-P_a)}}{\sum_c e^{\beta(A_{ic}-P_c)}}\left(\delta_{ab} - \frac{e^{\beta(A_{ib}-P_b)}}{\sum_c e^{\beta(A_{ic}-P_c)}}\right) \tag{A.1}$$

is positive semidefinite. From eqn (14), we obtain

$$\frac{1}{\beta}\frac{\partial^2 E_P}{\partial P_a \partial P_b} = \sum_i S_{ia}(\delta_{ab} - S_{ib}). \tag{A.2}$$

Let $\{x_a\}$ be an arbitrary vector, then

$$\begin{aligned}\frac{1}{\beta}\sum_{a,b}\frac{\partial^2 E_P}{\partial P_a \partial P_b}x_a x_b &= \sum_i \sum_a S_{ia}x_a^2 - \sum_i \sum_{a,b} S_{ia}S_{ib}x_a x_b \\ &= \sum_i \left[\left(\sum_a S_{ia}x_a^2\right) - \left(\sum_a S_{ia}x_a\right)^2\right].\end{aligned} \tag{A.3}$$

We will show that the summand is nonnegative for each i. Recall the Cauchy–Schwarz inequality

$$\left(\sum_\mu f_\mu g_\mu\right)^2 \le \left(\sum_\mu f_\mu^2\right)\left(\sum_\nu g_\nu^2\right), \tag{A.4}$$

where equality holds if and only if $\{f_\mu\}$ and $\{g_\nu\}$ are co-linear. We apply this for each i, setting $f_a = x_a\sqrt{(S_{ia})}$ and $g_a = \sqrt{(S_{ia})}$ and noting that the $\{S_{ia}\}$ are all positive by eqn (14). This yields

$$\left(\sum_a S_{ia}x_a\right)^2 \le \left(\sum_a S_{ia}x_a^2\right), \tag{A.5}$$

since $\sum_a S_{ia} = 1$. Hence, $(1/\beta)\sum_{a,b}(\partial^2 E_P/\partial P_a \partial P_b)x_a x_b$ is nonnegative for all vectors $\{x_a\}$, so the Hessian is indeed positive semidefinite. By Cauchy–Schwarz, equality holds only when $\{x_a\} \propto (1, \ldots, 1)$. This eigenvector corresponds to the translation invariance $\{P_a\} \mapsto \{P_a + K\}$.

We will now show that the P-Energy is bounded below. Without loss of generality, we fix the translation degree of freedom by imposing the condition $\sum_a P_a = 0$. Then E_P can be rewritten as

$$E_P = \frac{1}{\beta}\log\prod_i\left(\sum_a e^{\beta A_{ia}}e^{-\beta P_a}\right). \tag{A.6}$$

Consider the term $\sum_a e^{\beta A_{ia}}e^{-\beta P_a}$. Because $\sum_a P_a = 0$, this term is bounded below, but not above (send one P_a to minus infinity and the others to plus infinity). Using Lagrange multipliers to impose the constraint $\sum_a P_a = 0$, we calculate that the lower bound on each term is $Ne^{(\beta/N)\sum_c A_{ic}}$, which occurs when $e^{-\beta P_a} = e^{[(\beta/N)\sum_c A_{ic}]}e^{-\beta A_{ia}}$. Hence, E_P is bounded below by $(1/\beta)N\log N + \sum_{i,a} A_{ia}$.

We conclude that since E_P is weakly convex and bounded below, it has a unique minimum up to a constant translation invariance. ■

APPENDIX 2

Proof of Theorem 2. At a minimum of the P-Energy, $\sum_i S_{ia} = 1$, so

$$E_P[\tilde{P}, T] - B_{ass}[S(\tilde{P}, T)] = -T\sum_{i,a=1}^{N} S_{ia}(\tilde{P}, T)\log S_{ia}(\tilde{P}, T). \tag{A.7}$$

Using Lagrange multipliers to minimize $\sum_{i,a} S_{ia}\log S_{ia}$ subject to the constraints $\sum_a S_{ia} = 1$ for all i and $S_{ia} \ge 0$ for all i, a, we conclude that

$$0 \le -\sum_{i,a} S_{ia}\log S_{ia} \le N\log N. \quad ■ \tag{A.8}$$

Proof of Lemma 3.

$$\lim_{T\to 0} E_P[P, T] = \lim_{T\to 0}\left\{T\sum_i \log\left(\sum_a e^{(1/T)(A_{ia}-P_a)}\right) + \sum_a P_a\right\} = \sum_{i=1}^{N}\max_a (A_{ia} + P_a) + \sum_{a=1}^{N} P_a. \quad ■$$

Proof of Theorem 4. We will first prove that $B^*_{ass} \le E_P[\tilde{P}, T = 0]$. Now for any choice of $\pi \in S_N$ and for all $\{P_a\} \in \mathbf{R}^n$

$$\sum_{i=1}^{N} A_{i\pi(i)} = \sum_{i=1}^{N}(A_{i\pi(i)} - P_{\pi(i)}) + \sum_{a=1}^{N} P_a \le \left[\sum_{i=1}^{N}\max_a(A_{ia} - P_a) + \sum_{a=1}^{N} P_a\right] = E_P[P, T = 0].$$

where the final equality follows from Lemma 3. Hence,

$$B^*_{ass} \overset{\text{def}}{=} \max_{\pi\in S_N}\sum_{i=1}^{N} A_{i\pi(i)} \le \min_{\{P_a\}}\left(\sum_{i=1}^{N}\max_a(A_{ia} - P_a) + \sum_{a=1}^{N} P_a\right) = \min_P E_P[P, T = 0] = E_P[\tilde{P}, T = 0],$$

where $\tilde{P}$ denotes the vector minimizing the P-Energy at $T = 0$.

The converse inequality follows by applying eqns (A.7) and (23) at $T = 0$ to get that

$$E_P[\tilde{P}, T = 0] = \sum_{i,a} A_{ia}S_{ia}[\tilde{P}, T = 0] \le B^*_{ass}. \quad ■ \tag{A.9}$$

Proof of Theorem 5. Using the chain rule,

$$\frac{dE_P[\tilde{P}, T]}{dT} = \frac{\partial E_P}{\partial T} + \sum_a \frac{\partial E_P}{\partial P_a}\frac{dP_a}{dT} = \frac{\partial E_P}{\partial T},$$

recalling that we are at a minimum of the P-Energy. Setting $T = 1/\beta$,

$$\begin{aligned}\frac{dE_P[\tilde{P}, T]}{dT} &= -\beta^2\frac{\partial E_P[\tilde{P}, T = 1/\beta]}{\partial\beta} \\ &= \sum_i \log\left(\sum_a e^{\beta(A_{ia}-\tilde{P}_a)}\right) - \beta\sum_i \frac{\sum_a (A_{ia} - \tilde{P}_a)e^{\beta(A_{ia}-\tilde{P}_a)}}{\sum_b e^{\beta(A_{ib}-\tilde{P}_b)}} \\ &= \frac{1}{T}\left(E_P[\tilde{P}, T] - \sum_a \tilde{P}_a - \sum_{i,a}(A_{ia} - \tilde{P}_a)S_{ia}[\tilde{P}, T]\right) \\ &= \frac{1}{T}\{E_P[\tilde{P}, T] - B_{ass}[S(\tilde{P}, T)]\}\end{aligned} \tag{A.10}$$

where we have used the fact that at equilibrium, $\sum_i S_{ia} = 1$. The proof of the first inequality is then completed by referring to the bounds in Theorem 2.

The second inequality follows by integrating the first inequality, and substituting the result of Theorem 4. The third inequality follows from the second using eqns (A.7) and (A.8). ■

Proof of Theorem 6. Set $B^*_{ass} = B_{ass}[\{\Pi^*_{ia}\}]$ and let $B^{2*}_{ass} = B_{ass}[\{\Pi^{2*}_{ia}\}]$ denote the total benefit of the second best assignment, $\{\Pi^{2*}_{ia}\}$. Because we are assuming that there exists a unique optimum to the assignment problem, $\Delta = B^*_{ass} - B^{2*}_{ass}$ is strictly positive. Now $S_{ia}(\tilde{P}, T)$ is a doubly stochastic matrix and hence belongs to the permutation polytope, the convex hull of the permutation matrices. So, we can write

$$S_{ia}(\tilde{P}, T) = \sum_{\pi\in S_N} t_\pi \Pi_{ia}, \tag{A.11}$$

where, $t_\pi \ge 0$ and $\sum_{\pi\in S_N} t_\pi = 1$.

Suppose there exist $j, b \in \{1, 2, \ldots N\}$ such that $TN\log N/\Delta < |\Pi^*_{jb} - S_{jb}(\tilde{P}, T)|$. If $\Pi^*_{jb} = 1$ then

$$\frac{TN\log N}{\Delta} < 1 - \sum_{\pi\in S_N} t_\pi\Pi_{jb} = (1 - t_{\pi^*}) - \sum_{\pi\in S_N - \pi^*} t_\pi\Pi_{jb} < 1 - t_{\pi^*},$$

while if $\Pi^*_{jb} = 0$ then

$$\frac{TN\log N}{\Delta} < \sum_{\pi\in S_N} t_\pi\Pi_{jb} = \sum_{\pi\in S_N - \pi^*} t_\pi\Pi_{jb} < 1 - t_{\pi^*}.$$

Hence, in either case

$$\frac{TN\log N}{\Delta} < 1 - t_{\pi^*}. \tag{A.12}$$

Now,

$$\begin{aligned}B_{ass}[S(\tilde{P}, T)] &= \sum_{i,a} A_{ia}S_{ia}(\tilde{P}, T) \\ &= \sum_{\pi\in S_N} t_\pi\left(\sum_{i,a} A_{ia}\Pi_{ia}\right) \\ &= t_{\pi^*}B^*_{ass} + \sum_{\pi\in S_N - \{\pi^*\}} t_\pi\left(\sum_{i,a} A_{ia}\Pi_{ia}\right) \\ &\le t_{\pi^*}B^*_{ass} + (1 - t_{\pi^*})B^{2*}_{ass} = B^*_{ass} - (1 - t_{\pi^*})\Delta.\end{aligned}$$

Applying the bounds from Theorem 5, we get a contradiction

$$\frac{TN \log N}{\Delta} < 1 - t_{**} \leq \frac{TN \log N}{\Delta}. \tag{A.13}$$

Hence, we conclude that

$$\max_{i,a} |\Pi^*_{ia} - S_{ia}(\tilde{P}, T)| \leq \frac{TN \log N}{\Delta}. \tag{A.14}$$

So, the entries of $\{S_{ia}(\tilde{P}, T)\}$ converge uniformly to the optimal assignment permutation as the temperature goes to zero. ■

Proof of Theorem 7. The proof follows trivially from Theorem 6 because rounding off to the nearest integer will yield the optimal assignment whenever

$$\max_{i,a} |\Pi^*_{ia} - S_{ia}(\tilde{P}, T)| < \frac{1}{2}. \tag{A.15}$$

Because rounding off gives the optimal assignment, we know that for each row i there exists a unique a such that $S_{ia}(\tilde{P}, T) > \frac{1}{2}$, (i.e. object a is assigned to person i). By eqn (14),

$$\sum_{b \neq a} e^{\beta(A_{ib} - \tilde{P}_b - A_{ia} + \tilde{P}_a)} = \frac{1}{S_{ia}} - 1 < 1. \tag{A16}$$

Hence, for all $b \neq a$.

$$A_{ia} - \tilde{P}_a > A_{ib} - \tilde{P}_b, \tag{A.17}$$

so, $\{A_{ia} - P_a\}$ is indeed strictly row dominant. ■

APPENDIX 3

Proof of Theorem 10. As in the proof of Theorem 7, if $|S_{ia}(P, T) - \Pi^*_{ia}| < \frac{1}{2}$, then we can recover the optimal assignment by rounding off. This condition can be enforced by requiring that the four terms on the right-hand side of eqn (48) are all less $\frac{1}{8}$. For the third term we need,

$$T < \frac{\Delta}{64N^2(\log N)^2}. \tag{A.18}$$

Similarly, for the first term we require,

$$\varepsilon < \frac{\Delta}{128N^2(\log N)(\max_{i,a} A_{ia} - \min_{i,a} A_{ia})}. \tag{A.19}$$

Now, the fourth term is less than $\frac{1}{8}$ if

$$\varepsilon < \frac{1}{8N \log(N-1)}. \tag{A.20}$$

Finally, for the second term, $[8N^2(\log N)T \log(N - 1 + \varepsilon)/(1 - \varepsilon)]/\Delta < \frac{1}{8}$ will follow from eqn (A.18), provided $\varepsilon < 1/(N+1)$. Because $N \geq 3$, eqn (A.20) in turn ensures that $\varepsilon < 1/(N+1)$, so that the second relation is automatically satisfied.

Now, we can bound

$$\Delta \leq 2(\max_{i,a} A_{ia} - \min_{i,a} A_{ia}). \tag{A.21}$$

It then follows that eqns (A.19) and (A.20) are satisfied if $\varepsilon < 1/(64N^2 \log N)$. Thus, convergence is assured if,

$$T < \frac{\Delta}{64N^2(\log N)^2}$$

$$\|\nabla E_P[P; T]\| < \frac{1}{64N^2 \log N}. \tag{A.22}$$

The convergence time for this dynamical system can now be estimated. Using the chain rule, $dE_P[P(T)]/dt = \nabla E_P[P(T)] \cdot dP/dt$, and the steepest descent equation, $dP/dt = -\nabla E_P[P(T)]$, we observe that $dE_P[P(T)]/dt = -\|\nabla E_P[P(T)]\|^2$. Recalling that the P-Energy is convex, it follows that after $(E_P[P = 0, T] - E_P[\tilde{P}(T)])/\varepsilon^2$ units of time we can descend to a point where $\|\nabla E_P[P(T)]\| < \varepsilon$. Using Lagrange multipliers, it can be shown that $E_P[P = 0, T] = B_{ass}[S(0)] - T \sum_{i,a} S_{ia}(0) \log S_{ia}(0)$ is bounded above by $(TN \log N + N \max_{i,a} A_{ia})$. By Theorem 5, $E_P[\tilde{P}(T)]$ is bounded below by $N \min_{i,a} A_{ia}$. Hence,

$$E_P[P = 0, T] - E_P[\tilde{P}(T)] \leq TN \log N + N(\max_{i,a} A_{ia} - \min_{i,a} A_{ia}). \tag{A.23}$$

So, if we fix

$$T \leq \frac{\Delta}{64N^2(\log N)^2} \tag{A.24}$$

and set $\varepsilon = 1/(64N^2 \log N)$, then we have convergence in time

$$t \leq 4096N^5(\log N)^2(T \log N + (\max_{i,a} A_{ia} - \min_{i,a} A_{ia})). \blacksquare \tag{A.25}$$