# THE BOLTZMANN PERCEPTRON NETWORK: A MULTI-LAYERED FEED-FORWARD NETWORK EQUIVALENT TO THE BOLTZMANN MACHINE 

Eyal Yair and Allen Gersho<br>Center for Information Processing Research Department of Electrical \& Computer Engineering University of California, Santa Barbara, CA 93106


#### Abstract

The concept of the stochastic Boltzmann machine (BM) is attractive for decision making and pattern classification purposes since the probability of attaining the network states is a function of the network energy. Hence, the probability of attaining particular energy minima may be associated with the probabilities of making certain decisions (or classifications). However, because of its stochastic nature, the complexity of the BM is fairly high and therefore such networks are not very likely to be used in practice. In this paper we suggest a way to alleviate this drawback by converting the stochastic BM into a deterministic network which we call the Boltzmann Perceptron Network (BPN). The BPN is functionally equivalent to the BM but has a feed-forward structure and low complexity. No annealing is required. The conditions under which such a conversion is feasible are given. A learning algorithm for the BPN based on the conjugate gradient method is also provided which is somewhat akin to the backpropagation algorithm.


## INTRODUCTION

In decision-making applications, it is desirable to have a network which computes the probabilities of deciding upon each of $M$ possible propositions for any given input pattern. In principle, the Boltzmann machine (BM) (Hinton, Sejnowski and Ackley, 1984) can provide such a capability. The network is composed of a set of binary units connected through symmetric connection links. The units are randomly and asynchronously changing their values in $\{0,1\}$ according to a stochastic transition rule. The transition rule used by Hinton et. al. defines the probability of a unit to be in the 'on' state as the logistic function of the energy change resulting by changing the value of that unit. The BM can be described by an ergodic Markov chain in which the thermal equilibrium probability of attaining each state obeys the Boltzmann distribution which is a function of only the energy. By associating the set of possible propositions with subsets of network states, the probability of deciding upon each of these propositions can be measured by the probability of attaining the corresponding set of states. This probability is also affected by the temperature. As the temperature

[^0]increases, the Boltzmann probability distribution become more uniform and thus the decision made is 'vague'. The lower the temperature the greater is the probability of attaining states with lower energy thereby leading to more 'distinctive' decisions.
This approach, while very attractive in principle, has two major drawbacks which make the complexity of the computations become non-feasible for nontrivial problems. The first is the need for thermal equilibrium in order to obtain the Boltzmann distribution. To make distinctive decisions a low temperature is required. This implies slower convergence towards thermal equilibrium. Generally, the method used to reach thermal equilibrium is simulated annealing (SA) (Kirkpatrick et. al., 1983) in which the temperature starts from a high value and is gradually reduced to the desired final value. In order to avoid 'freezing' of the network, the cooling schedule should be fairly slow. SA is thus time consuming and computationally expensive. The second drawback is due to the stochastic nature of the computation. Since the network state is a random vector, the desired probabilities have to be estimated by accumulating statistics of the network behavior for only a finite period of time. Hence, a trade-off between speed and accuracy is unavoidable.
In this paper, we propose a mechanism to alleviate the above computational drawbacks by converting the stochastic BM into a functionally equivalent deterministic network, which we call the Boltzmann Perceptron Network (BPN). The BPN circumvents the need for a Monte Carlo type of computation and instead evaluates the desired probabilities using a multilayer perceptron-like network. The very time consuming learning process for the BM is similarly replaced by a deterministic learning scheme, somewhat akin to the backpropagation algorithm, which is computationally affordable. The similarity between the learning algorithm of a BM having a layered structure and that of a two-layer perceptron has been recently pointed out by Hopfield (1987). In this paper we further elaborate on such an equivalence between the BM and the new perceptron-like network, and give the conditions under which the conversion of the stochastic BM into the deterministic BPN is possible. Unlike the original BM, the BPN is virtually always in thermal equilibrium and thus SA is no longer required. Nevertheless, the temperature still plays the same role and thus varying it may be beneficial to control the 'softness' of the decisions made by the BPN. Using the BPN as a soft classifier is described in details in (Yair and Gersho, 1989).

## THE BOLTZMANN PERCEPTRON NETWORK

Suppose we have a network of $K$ units connected through symmetric connection links with no self-feedback, so that the connection matrix $\Gamma$ is symmetric and zero-diagonal. Let us categorize the units into three different types: input, output and hidden units. The input pattern will be supplied to the network by clamping the input units, denoted by $\underline{\mathbf{x}}=\left(x_{1}, \ldots, x_{i}, . ., x_{l}\right)^{T}$, with this pattern. $\underline{\mathrm{x}}$ is a real-valued vector in $R^{I}$. The output of the network will be observed on the output units $y=\left(y_{1}, \ldots, y_{m}, \ldots, y_{M}\right)^{T}$, which is a binary vector. The remaining units, denoted $\mathbf{v}=\left(v_{1}, \ldots, v_{j}, \ldots, v_{J}\right)^{T}$, are the hidden units, which are also binary-valued. The hidden and output units are asynchronously and randomly changing their binary values in $\{0,1\}$ according to inputs they receive from other units.
The state of the network will be denoted by the vector $\underline{\underline{u}}$ which is partitioned as follows: $\underline{\underline{u}}^{T}=\left(\underline{\mathbf{x}}^{T}, \underline{\underline{v}}^{T}, \mathbf{\underline { y }}^{T}\right)$. The energy associated with state $\underline{\underline{u}}$ is denoted by $E_{\boldsymbol{u}}$ and is given by:

$$
\begin{equation*}
-E_{u}=1 / 2 \underline{\mathbf{u}}^{T} \Gamma \underline{\mathbf{u}}+\underline{\mathbf{u}}^{T} \underline{\delta} \tag{1}
\end{equation*}
$$

where $\underline{\delta}$ is a vector of bias values, partitioned to comply with the partition of $\underline{\underline{u}}$ as follows: $\underline{\boldsymbol{\delta}}^{\boldsymbol{T}}=\left(\underline{\mathbf{f}}^{\boldsymbol{T}}, \underline{\underline{c}}^{\boldsymbol{T}}, \underline{\mathbf{S}}^{\boldsymbol{T}}\right)$.

The transition from one state to another is performed by selecting each time one unit, say unit $k$, at random and determine its output value according to the following stochastic rule: set the output of the unit to 1 with probability $p_{k}$, and to 0 with a probability $1-p_{k}$. The parameter $p_{k}$ is determined locally by the $k$-th unit as a function of the energy change $\Delta E_{k}$ in the following fashion:

$$
\begin{equation*}
p_{k}=g\left(\Delta E_{k}\right) ; \quad g(x) \triangleq \frac{1}{1+e^{-\beta x}} \tag{2}
\end{equation*}
$$

$\Delta E_{k}=\left\{E_{u}\right.$ (unit $k$ is off) $-E_{u}$ (unit $k$ is on) $\}$, and $\beta=1 / T$ is a control parameter. $T$ is called the temperature and $g(\cdot)$ is the logistic function. With this transition rule the thermal equilibrium probability $P_{u}$ of attaining a state $\underline{\underline{u}}$ obeys the Boltzmann distribution:

$$
\begin{equation*}
P_{k}=\frac{1}{Z_{x}} e^{-\beta E_{u}} \tag{3}
\end{equation*}
$$

where $Z_{x}$, called the partition function, is a normalization factor (independent of $\mathbf{v}$ and $\mathbf{y}$ ) such that the sum of $P_{u}$ over all the $2^{J+M}$ possible states will sum to unity.
In order to use the network in a deterministic fashion rather than accumulate statistics while observing its random behavior, we should be able to explicitly compute the probability of attaining a certain vector $\mathbf{y}$ on the output units while $\mathbf{x}$ is clamped on the input units. This probability, denoted by $P_{y \mid x}$, can be expressed as:

$$
\begin{equation*}
P_{y \mid x}=\sum_{\underline{v} \in B_{J}} P_{v, y \mid x}=\frac{1}{Z_{x}} \sum_{\underline{v} \in B_{J}} e^{-\beta E_{v y \mid x}} \tag{4}
\end{equation*}
$$

where $B_{J}$ is the set of all binary vectors of length $J$, and $v, y \mid x$ denotes a state $\underline{\underline{u}}$ in which a specific input vector $\mathbf{x}$ is clamped. The explicit evaluation of the desired probabilities therefore involves the computation of the partition function for which the number of operations grows exponentially with the number of units. That is, the complexity is $O\left(2^{J+M}\right)$. Obviously, this is computationally unacceptable. Nevertheless, we shall see that under a certain restriction on the connection matrix $\Gamma$ the explicit computation of the desired probabilities becomes possible with a complexity of $O(J M)$ which is computationally feasible.
Let us assume that for each input pattern we have $M$ possible propositions which are associated with the $M$ output vectors: $\boldsymbol{I}_{M}=\left\{\underline{\mathbf{I}}_{1}, \ldots \mathbf{I}_{m}, \ldots \mathbf{I}_{M}\right\}$, where $\underline{\mathbf{I}}_{m}$ is the $\boldsymbol{m}$-th column of the $M \times M$ identity matrix. Any state of the network having output vector $\mathbf{y}=\mathbf{I}_{m}$ (for any $m$ ) will be denoted by $v, m \mid x$ and will be called a feasible state. All other state vectors $v, y \mid x$ for $\mathbf{Y} \neq \mathbf{I}_{m}$ will be considered as intermediate steps between two feasible states. This redefinition of the network states is equivalent to redefining the states of the underlying Markov model of the network and thus conserves the equilibrium Boltzmann distribution. The probability of proposition $m$ for a given input $\underline{x}$, denoted by $P_{m \mid x}$, will be taken as the probability of obtaining output vector $\mathbf{y}=\mathbf{I}_{\boldsymbol{m}}$ given that the output is one of the feasible values. That is,

$$
\begin{equation*}
P_{m \mid x}=\operatorname{Pr}\left\{\mathbf{y}=\mathbf{I}_{m} \mid \underline{\mathbf{x}}, \mathbf{y} \in \boldsymbol{I}_{M}\right\} \tag{5}
\end{equation*}
$$

which can be computed from (4) by restricting the state space to the $2^{J} M$ feasible state vectors and by setting $\mathbf{y}=\mathbf{I}_{m}$. The partition function, conditioned on restricting $\mathbf{y}$ to lie in the set of feasible outputs, $I_{M}$, is denoted by $\bar{Z}_{x}$ and is given by:

$$
\begin{equation*}
\bar{Z}_{x}=\sum_{y \in I_{M}} \sum_{v \in B_{J}} e^{-\beta E_{v y} \mid x} \tag{6}
\end{equation*}
$$

Let us now partition the connection matrix $\Gamma$ to comply with the partition of the state vector and rewrite the energy for the feasible state $v, m \mid x$ as:

$$
\begin{equation*}
-E_{v, m \mid x}=\underline{\mathbf{y}}^{T}\left(\mathbf{R} \underline{\mathbf{x}}+\mathbf{Q}_{m}+1 / 2 \mathbf{D}_{2} \mathbf{\underline { \mathbf { v } }}+\boldsymbol{\varepsilon}\right)+\mathbf{I}_{m}^{T}\left(\mathbf{W}_{\underline{\mathbf{x}}}+1 / 2 \mathbf{D}_{3} \mathbf{I}_{m}+\underline{\mathbf{s}}\right)+\underline{\mathbf{x}}^{T}\left(1 / 2 \mathbf{D}_{1} \underline{\mathbf{x}}+\mathbf{f}\right) . \tag{7}
\end{equation*}
$$

Since $\underline{x}$ is clamped on the input units, the last term in the energy expression serves only as a bias term for the energy which is independent of the binary units $\mathbf{y}$ and $\mathbf{y}$. Therefore, without any loss of generality it may be assumed that $\mathbf{D}_{1}=0$ and $\boldsymbol{\mathcal { I }}=\underline{\mathbf{Q}}$. The second term, denoted by $T_{m \mid x}$, can be simplified since $\mathrm{D}_{3}$ has a zero diagonal. Hence,

$$
\begin{equation*}
T_{m \mid x}=\sum_{i=1}^{I} w_{m i} x_{i}+s_{m} \tag{8}
\end{equation*}
$$

The absence of the matrix $\mathbf{D}_{3}$ in the energy expression means that interconnections between output units have no effect on the probability of attaining output vectors $\mathbf{y} \in \boldsymbol{I}_{\boldsymbol{M}}$, and may be assumed absent without any loss of generality.
Defining $L_{m}(\underline{x})$ to be:
in which $\mathbf{q}_{m}$ is the $m$-th column of $\mathbf{Q}$, the desired probabilities, $P_{m \mid x}$, for $m=1, \ldots, M$ are obtained using (4) and (7) as a function of these quantities as follows:

$$
\begin{equation*}
P_{m \mid x}=\frac{1}{\bar{Z}_{x}} e^{L_{m}(X)} \quad \text { with: } \quad \bar{Z}_{x}=\sum_{m=1}^{M} e^{L_{m}(X)} . \tag{10}
\end{equation*}
$$

The complexity of evaluating the desired probabilities $P_{m \mid x}$ is still exponential with the number of hidden units $J$ due to the sum in (9). However, if we impose the restriction that $\mathbf{D}_{\mathbf{2}}=\mathbf{0}$, namely, the hidden units are not directly interconnected, then this sum becomes separable in the components $v_{j}$ and thus can be decomposed into the product of only $J$ terms. This restricted connectivity of course imposes some restrictions on the capability of the network compared to that of a fully connected network. On the other hand, it allows the computation of the desired probabilities in a deterministic way with the atractive complexity of only $O(J M)$ operations. The tedious estimation of conditional expectations commonly required by the learning algorithm for a BM and the annealing procedure are avoided, and an accurate and computationally affordable scheme becomes possible. We thus suggest a trade-off in which the operation and learning of the BM are tremendously simplified and the exact decision probabilities are computed (rather than their statistical estimates) at the expense of a restricted connectivity, namely, no interconnections are allowed between the hidden units. Hence, in our scheme, the connection matrix, $\Gamma$, becomes zero block-diagonal, meaning that the network has connections only between units of different categories. This structure is shown schematically in Figure 1.


Figure 1. Schematic architecture of the stochastic BM.


Figure 2. Block diagram of the corresponding deterministic BPN.

By applying the property $\mathrm{D}_{2}=0$ to (9), the sum over the space of hidden units, which can be explicitly written as the sum over all the $J$ components of $\mathbf{v}$, can be decomposed using the separability of the different $\boldsymbol{v}_{\boldsymbol{j}}$ components into a sum of only $J$ terms as follows:

$$
\begin{equation*}
L_{m}(\underline{x})=\beta T_{m \mid x}+\sum_{j=1}^{J} f\left(V_{j}^{m \mid x}\right) \tag{11a}
\end{equation*}
$$

where:

$$
\begin{equation*}
f(x) \triangleq \ln \left(1+e^{\beta x}\right) \quad \text { and } \quad V_{j}^{m \mid x} \triangleq \sum_{i=1}^{I} r_{j i} x_{i}+c_{j}+q_{j m} \tag{11b}
\end{equation*}
$$

$f(\cdot)$ is called the activation function. Note that as $\beta$ is increased $f(\cdot)$ approaches the linear threshold function in which a zero response is obtained for a negative input and a linear one (with slope $\beta$ ) for a positive input.
Finally, the desired probabilities $P_{m \mid x}$ can be expressed as a function of the $L_{m}(\underline{x})$ ( $m=1, \ldots, M$ ) in an expression which can be regarded as the generalization of the logistic function to $M$ inputs:

$$
\begin{equation*}
P_{m \mid x}=\left[1+\sum_{\substack{n=1 \\ n \neq m}}^{M} e^{-L_{m, n}(\underline{x})}\right]^{-1} \quad ; \quad \text { where: } \quad L_{n, m}(\underline{x})=L_{n}(\underline{x})-L_{m}(\underline{x}) \tag{12}
\end{equation*}
$$

Eqs. (8) and (11) describe a two-layer feed-forward perceptron-like subnetwork which uses the nonlinearity $f(\cdot)$. It evaluates the quantity $L_{m}(\underline{x})$ which we call the score of the $m$-th proposition. Eq. (12) describes a competition between the scores $L_{m}(\underline{x})$ generated by the $M$ subnetworks ( $m=1, . ., M$ ) which we call a soft competition with lateral inhibition. That is, If several scores receive relatively high values compared to the others, they will share, according to their relative strengths, the total amount (unity) of probability, while inhibiting the remaining probabilities to approach zero. For example, if one of the scores, say $L_{k}(\underline{x})$, is large compared to all the other scores, then the exponentiation of the pairwise score differences will result in $P_{k \mid x} \approx 1$ while the remaining probabilities will approach zero. Specifically, for any $n \neq k, P_{n \mid x} \approx \exp \left(-L_{k, n}(x)\right)$, which is essentially zero if $L_{k}(x)$ is sufficiently high. In other words, by being large compared to the others, $L_{k}(x)$ won the competition so that the corresponding probability $P_{k \mid x}$ approaches unity, while all the remaining probabilities have been attenuated by the high value of $L_{k}(\underline{x})$ to approach zero.
Let us examine the effect of the gain $\beta$ on this competition. When $\beta$ is increased, the slope of the activation function $f(\cdot)$ is increased thereby accentuating the differences between the $M$ contenders. In the limit when $\beta \rightarrow \infty$, one of the $L_{m}(\underline{x})$ will always be sufficiently large compared to the others, and thus only one proposition will win. The competition then becomes a winner-take-all competition. In this case, the network becomes a maximum a posteriori (MAP) decision scheme in which the $L_{m}(\mathbf{x})$ play the role of nonlinear discriminant functions and the most probable proposition for the given input pattern is chosen:

$$
\begin{equation*}
P_{k \mid x}=1 \quad \text { for } \quad k=\underset{m}{\operatorname{argmax}}\left\{L_{m}(\underline{x})\right\} \quad \text { and } \quad P_{n \mid x}=0 \quad \text { for } n \neq k \tag{13}
\end{equation*}
$$

This results coincides with our earlier observation that the temperature controls the 'softness' of the decision. The lower the temperature, the 'harder' is the competition and the more distinctive are the decisions. However, in contrast to the stochastic network, there is no need to gradually 'cool' the network to achieve a desired (low) temperature. Any desired value of $\beta$ is directly applicable in the BPN scheme. The above notion of soft competition has its merits in a wider scope of applications apart from its role in the BPN classifier. In many competitive schemes a soft competition between a set of contenders has
a substantial benefit in comparison to the winner-take-all paradigm. The above competition scheme which can be implemented by a two-layer feed-forward network thus offers a valuable scheme for such purposes.
The block diagram of the BPN is depicted in Figure 2. The BPN is thus a four-layer feedforward deterministic network. It is comprised of a two-layer perceptron-like network followed by a two-layer competition network. The competition can be 'hard' (winner-take-all) or 'soft' (graded decision) and is governed by a single gain parameter $\beta$.

## THE LEARNING ALGORITHM

Let us denote the BPN output by the $M$-dimensional probability vector $\underline{\mathbf{P}}_{x}$, where: $\underline{\mathbf{P}}_{x}=\left(P_{1 \mid x}, \ldots, P_{m \mid x}, \ldots, P_{M \mid x}\right)^{T}$. For any given set of weights $\underline{\theta}$, the BPN realizes some deterministic mapping $\Psi: \mathrm{R}^{I} \rightarrow[0,1]^{M}$ so that $\underline{\mathbf{P}}_{x}=\Psi(\underline{x})$. The objective of learning is to determine the mapping $\Psi$ (by estimating the set of parameters $\underline{\theta}$ ) which 'best' explains a finite set of examples given from the desired mapping and are called the training set. The training set is specified by a set of $N$ patterns $\left\{\underline{\mathbf{x}}_{1}, \ldots, \mathbf{x}_{n}, \ldots, \mathbf{x}_{N}\right\}$ (denoted for simplicity by ( $\underline{\mathbf{x}}\}$ ), the a priori probability for each training pattern $\underline{\mathbf{x}}: \underline{Q}(\mathbf{x})$, and the desired mapping for each pattern $\underline{x}: \underline{Q}_{x}=\left(Q_{1 \mid x}, \ldots, Q_{m \mid x}, \ldots, Q_{m \mid x}\right)^{T}$, where $Q_{m \mid x}=\operatorname{Pr}$ (proposition $m \mid \underline{x}$ ) is the desired probability. For each input pattern presented to the BPN, the actual output probability vector $\mathbf{P}_{x}$ is, in general, different from the desired one, $\mathbf{Q}_{\boldsymbol{x}}$. We denote by $G_{\boldsymbol{x}}$ the distortion between the actual $\underline{\mathbf{P}}_{\boldsymbol{x}}$ and the desired $\underline{\mathbf{Q}}_{x}$ for an input pattern $\underline{\mathbf{x}}$. Thus, our task is to determine the network weights (and the bias values) $\underline{\theta}$ so that, on the average, the distortion over the whole set of training patterns will be minimized. Adopting the original distortion measure suggested for Boltzmann machines, the average distortion, $\boldsymbol{G}(\theta)$, is given by:

$$
\begin{equation*}
G(\theta)=\sum_{\mathbf{x}} Q(\underline{x}) G_{x}(\theta) \quad ; \quad G_{x}(\theta)=\sum_{m=1}^{M} Q_{m \mid x} \ln \left[Q_{m \mid x} / P_{m \mid x}(\theta)\right] \tag{14}
\end{equation*}
$$

which is always non-negative since $\underline{P}_{x}$ and $\underline{Q}_{x}$ are probability vectors. To minimize $\boldsymbol{G}(\boldsymbol{\theta})$ a gradient based minimization search is used. Specifically, a Partial Conjugate Gradient (PCG) search (Fletcher and Reeves, 1964; Luenberger, 1984) was found to be significantly more efficient than the ordinary steepest descent approach which is so widely used in multilayer perceptrons. A further discussion supporting this finding is given in (Yair and Gersho, 1989). For each set of weights we thus have to be able to compute the gradient $\mathbf{g}=\nabla_{\theta} G$ of the cost function $G(\boldsymbol{\theta})$. Let us denote the components of the 'instantaneous' gradient by $G_{m}^{s \mid x}=\partial G_{x} / \partial s_{m}, \quad G_{j m}^{q \mid x}=\partial G_{x} / \partial q_{j m}, \quad G_{j}^{c \mid x}=\partial G_{x} / \partial c_{j}, \quad G_{m i}^{w \mid x}=\partial G_{x} / \partial w_{m i}$, $G_{j i}^{r \mid x}=\partial G_{x} / \partial r_{j i}$. To get the full gradient, the instantaneous components should be accumulated while the input patterns are presented (one at a time) to the network, until one full cycle through the whole training set is completed.
It is straightforward to show that the gradient may be evaluated in a recursive manner, in a fashion somewhat similar to the evaluation of the gradient by the backpropagation algorithm used for feed-forward networks (Rumelhart et. al., 1986). The evaluation of the gradient is accomplished by propagating the errors $e_{m \mid x}=Q_{m \mid x}-P_{m \mid x}$ through a linear network, termed the Error Propagation Network (EPN), as follows:

$$
\begin{array}{ll}
G_{m}^{s \mid x}=-\beta e_{m \mid x} & ; \quad G_{m i}^{w \mid x}=x_{i} G_{m}^{s \mid x} \\
G_{j m}^{q \mid x}=b_{j}^{m \mid x} G_{m}^{s \mid x} \quad ; \quad G_{j}^{c \mid x}=\sum_{m=1}^{M} G_{j m}^{q \mid x} \quad ; \quad G_{j i}^{r \mid x}=x_{i} G_{j}^{c \mid x} \tag{15}
\end{array}
$$

The only new variables required are the $b_{j}^{m \mid x}$, given by $b_{j}^{m \mid x}=g\left(V_{j}^{m \mid x}\right)$, which can be
easily obtained by applying the logistic nonlinearity to $V_{j}^{m \mid x}$. The above error propagation scheme can also be written in a matrix form. Let us define the following notation: $\mathbf{e}_{x}=\left(e_{1 \mid x}, \ldots, e_{m \mid x} \ldots, e_{M \mid x}\right)^{T} . \mathbf{E}_{x}$ will be a diagonal $M \times M$ matrix whose diagonal is $\mathbf{e}_{x}$. $\mathbf{B}_{x}=\left[b_{j}^{m \mid x}\right]$, a $J \times M$ matrix. Let $1 v$ denote a column vector of length $M$ whose components are all 1 's. Similarly we will define the vectors: $G^{\lambda \mid x}=\left(\ldots, G_{\varepsilon}^{\lambda \mid x} \ldots . .\right)^{T}$, and the matrices: $\mathbf{G}^{\lambda \mid x}=\left[G_{\varepsilon \eta}^{\lambda \mid x}\right]$ with the appropriate dimensions (for any $\lambda, \varepsilon$ and $\eta$ ). Hence, the error propagation can be written as:

$$
\begin{align*}
& \underline{G}^{2 \mid x}=-\beta \underline{\mathrm{e}}_{\boldsymbol{x}} \quad ; \quad \mathbf{G}^{w \mid x}=\underline{G}^{2 \mid x} \underline{\mathbf{x}}^{T} \tag{16}
\end{align*}
$$

The EPN is depicted in Figure 3. This is a linear system composed of inner and outer products between matrices, which can be efficiently implemented using a neural network. The gradient $\mathbf{g}$ is used in the PCG update formula in which a new set of weights is created and is used for the next update iteration. The leaming scheme of the BPN is given in Figure 4.


Figure 3. The Error Propagation Network (EPN). Diag' is a diagonalization operator. 'right' and '7eft' are right and left multipliers, respectively.


Figure 4. The learning scheme. The BPN outputs, $\underline{P}_{x}$, are compared with the desired probabilities, $\underline{Q}_{\boldsymbol{x}}$. The resulted errors, $\underline{e}_{x}$, propagate through the EPN to form the gradient g which is used in the PCG alg. to create the new weights.

## SIMULATION RESULTS

We now present several simulation results of two-class classification problems with Gaussian sources. That is, we have two propositions represented by class 0 or class 1 . Suppose there are $L$ random sources ( $i=1, ., L$ ) over the input space, some of them are atributed to class 0 , and the others to class 1 . Each time, a source is chosen according to an a priori probability $P(i)$. When chosen, the $i$-th source then emits a pattern $\underline{x}$ according to a probability density $Q_{x \mid i}(\underline{\mathbf{x}})$. Measuring a pattern $\underline{\underline{x}}$ it is desired to decide upon the most probable origin class - in the binary decision problem (MAP classifier), or obtain some estimate to $Q_{11 x}$, the probability that class 1 emitted this pattern - in the soft classification problem. In the learning phase, a training set of size $N$ was presented to the network, and the weights were iteratively modified by the learning scheme (Figure 4) until convergence. The final weights were used to construct a BPN classifier, which was then tested for new input patterns. The output classification probability of the BPN, $P_{1 \mid x}(x)$, was compared with the true (Bayesian) conditional probability, $Q_{11 x}(\underline{x})$ which was computed analytically. Results are shown in Figures 5-7. In Figure 5, two symmetric equi-probable Gaussian sources with substantial overlap were used, one for each class. The network was trained on $N=8$ patterns with gain $\beta=1$. Figure 5 b shows how the BPN performs for the problem given in Figure 5 a . For $\beta=1$, i.e., when the same gain is used in both the training and classification phases, there is an almost perfect match between the BPN output, $P_{11 x}(x)$, denoted in the figures by ' $\beta=1$ ', and the true curve, $Q_{11 x}(x)$. For $\beta=10$, the high gain winner-take-all
competition is taking place and the classifier becomes, practically, a binary (yes/no) decision network. In Fig. 6 disconnected decision regions were formed by using four sources, two of which were attributed to each class. Again, a nearly perfect match can be seen between actual ( $\beta=1$ ) and desired ( $Q_{11 x}$ ) outputs. Also, the simplicity of making 'hard' classification decisions by increasing the gain is again illustrated ( $\beta=10$ ). In Fig. 7 the classifier was required to find the boundary (expressed by $Q_{11 x}=0.5$ ) between two 2D classes.


Figure 5: Classification for Gaussian sources. (5.a) The two sources. (5.b) 'Soft' ( $\beta=1$ ) and 'hard' $(\beta=10)$ classifications versus $Q_{11 x} . J$ indicates the number of hidden units used.
Figure 6: Classification for disconnected decision regions. (6.a) The sources used: dashed lines indicate class 0 and solid lines - class 1. (6.b) Soft $(\beta=1)$ and hard $(\beta=10)$ classifications versus $Q_{11 x}$.
Figure 7: Classification in a 2D space. (7.a) The two classes and the true boundary indicated by $Q_{11 x}=0.5$. (7.b) The boundary found by the BPN, marked by $P_{11 x}=0.5$, versus the true one.

## References

Fletcher, R., Reeves, C.M. (1964). Function minimization by conjugate gradients. Computer J., 7, 149-154.
Hinton, G.E., Sejnowski T.R., \& Ackley D.H. (1984). Boltzmann machines: constraint satisfaction networks that leam. Carnegie-Melon Technical Report, CMU-CS-84-119.

Hopfield, J.J. (1987). Leaming algorithms and probability distributions in feed-forward and feed-back networks. Proc. Natl. Acad. Sci. USA, 84, 8429-8433.

Kirkpatrick, S., Gelatt, C.D., \& Vecchi M.P. (1983). Optimization by simulated amnealing. Science, 220, 671-680.
Luenberger, D.G. (1984). Linear and nonlinear programming, Addison-Wesley, Reading, Mass.
Rumelhart, D.E., Hinton, G.E., \& Williams RJ. (1986). Learning intemal representations by error propagation. In D.E. Rumelhart \& J.L. MoClelland (Eds.), Parallel Distributed Processing., MIT Press/Bradford Books.

Yair, E., \& Gersho, A. (1989). The Boltzmann perecptron network: a soft classifier. Submitted to the Jowrnal of Neural Networks, December, 1988.


[^0]:    This work was supported by the Weizmann Foundation for scientific research, by the University of Califomia MICRO program, and by Bell Communications Research, Inc.

