

“UNLEARNING” INCREASES THE STORAGE CAPACITY OF CONTENT ADDRESSABLE MEMORIES

D. KLEINFELD AND D. B. PENDERGRAFT*

Department of Molecular Biophysics, AT&T Bell Laboratories, Murray Hill, New Jersey 07974; and

**Department of Physics, University of California at San Diego, La Jolla, California 92093*

ABSTRACT The storage and retrieval of information in networks of biological neurons can be modeled by certain types of content addressable memories (CAMs). We demonstrate numerically that the amount of information that can be stored in such CAMs is substantially increased by an unlearning algorithm. Mechanisms for the increase in capacity are identified and illustrated in terms of an energy function that describes the convergence properties of the network.

INTRODUCTION

Content addressable memories (CAMs) function by retrieving information based on partial knowledge of the contents of the desired memory, as opposed to knowledge of the location of the memory. Some of these systems exhibit many features in common with the way information is stored and retrieved in biological memories. Thus, these CAMs appear to be useful models for studying networks of biological neurons (Nakano, 1972; Cooper, 1973; Little and Shaw, 1975; Anderson et al., 1977; Hopfield, 1982, 1984; Peretto, 1984; Amit et al., 1985a) (for review, see Kohonen [1977] and Hopfield and Tank [1986]).

Crick and Mitchison (1983, 1986) proposed that the performance of biological memories may be improved by the selective “unlearning” of stored information. In support of this conjecture, Hopfield et al. (1983) used their model neural network to demonstrate that a mathematically analogous process improves the accessibility of information stored in a CAM. In this work, we examine numerically the information storage capacity of the CAM described by Hopfield (1982). We quantify, in terms of an entropic measure, the effect of unlearning on the capacity. The results are discussed in the context of an energy function (Hopfield, 1982; Cohen and Grossberg, 1983) that determines the dynamic properties of the network. Finally, we examine the effect of unlearning on the convergence properties of the network.

Description of the CAM

We consider a CAM consisting of N bits, or neurons. Each neuron is either active (+1) or quiescent (−1). A state of the network, S , is defined by a sequence of N bits that specifies the output of each neuron. For example, $S = (+1 -1 -1 \dots)$ implies that neuron 1 is active, neurons 2 and 3 are quiescent, etc. There are 2^N possible states. Memories are represented by specific states chosen from the 2^N possibilities.

Each neuron receives inputs from all other neurons via a set of pair-wise connections. The strength of each connec-

tion, T_{ij} , is determined from the stored memory states, M^r , according to the outer-product rule (Nakano, 1972; Cooper, 1973; Hopfield, 1982)

$$T_{ij} = \begin{cases} \sum_{r=1}^n M_i^r M_j^r & i \neq j \\ 0 & i = j \end{cases} \quad (1)$$

where n is the number of memories. Note that M^r and its complimentary state, $-M^r$, give identical T_{ij} 's; thus both states are stored as memories. The state of the CAM will converge from an arbitrary initial state to an unchanging state, i.e., a stable state, when the output of each neuron is updated according to the rule (e.g., Hopfield, 1982)

$$S_i \leftarrow \text{sgn} \left(\sum_{j=1}^N T_{ij} S_j \right) \quad (2)$$

where

$$\text{sgn}(z) = \begin{cases} +1 & \text{if } z > 0 \\ \text{unchanged} & \text{if } z = 0 \\ -1 & \text{if } z < 0. \end{cases} \quad (3)$$

The updating is performed asynchronously, i.e., the value of the index i is chosen at random. This procedure simulates variations in the setting times of biological neurons.

When relatively few memories are stored, i.e., $n \ll N$, the stable states of the network will correspond to the stored memory states. However, as the number of stored memories becomes large, the stable states may differ from the memory states (Hopfield, 1982). In particular, the storage algorithm (Eq. 1) produces spurious stable states that do not correspond to stored memories (Hopfield et al., 1983; Amit et al., 1985a).

Reprint requests should be addressed to D. Kleinfeld.

D. B. Pendergraft's present address is Department of Radiological Sciences, University of California at Irvine, Irvine, California 92717.

The Hamming distance between two states is defined as the number of bits that are different between the bit sequences that specify each state. This definition can be used to give a geometric interpretation to the update rule (Eq. 2). In terms of the Hamming distance, H' , between the present state of the network, S , and the ν th memory, M' , Eq. 2 becomes

$$S_i \leftarrow \text{sgn} \left[\sum_{\nu=1}^n \left(1 - \frac{H'}{N/2} \right) M'_i - \frac{n}{N} S_i \right] \quad (4)$$

When the present state is close to the ν th memory H' is close to zero. The weight for that memory is $(1 - H'[N/2]) \approx 1$. When the present state is far from the ν th memory, H' approaches the value $N/2$. The weight is now close to zero. Thus the present state should converge to the closest memory state. Note that the term $(n/N)S_i$ in Eq. 4, which results from the constraint $T_{ii} = 0$, will weaken the convergence property of the network when the number of stored memories is large, i.e., $n \approx N$.

The convergence properties of the network can be characterized in terms of an average (Hamming) radius of convergence, \bar{H}_{crit} . When the state of the CAM differs from the nearest memory by less than \bar{H}_{crit} bits, the network is most likely to converge to that memory. When the state differs from the nearest memory by a larger distance, the network preferentially converges to a different stable state. The number of states, \bar{V} , enclosed within \bar{H}_{crit} is given by

$$\bar{V} = \sum_{i=0}^{\bar{H}_{\text{crit}}} \binom{N}{i} \quad (5)$$

The value of \bar{V} cannot be larger than the average number of states available per stored state, i.e.

$$\bar{V} \leq \frac{2^N}{2 \cdot n}. \quad (6)$$

The factor of 2 in the denominator above results from the symmetry of the network. The largest possible value of \bar{H}_{crit} is found by equating Eqs. 5 and 6.

The accumulation of errors in the network, and the removal of these errors by unlearning, will be illustrated in a later section by energy diagrams. The energy of a state, S , is defined by (Hopfield, 1982).

$$E = -\frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N S_i T_{ij} S_j \quad (7)$$

In terms of the Hamming distance between S and the ν th stored memory, Eq. 7 becomes

$$E = -\frac{N^2}{2} \left[\sum_{\nu=1}^n \left(1 - \frac{H'}{N/2} \right)^2 - \frac{n}{N} \right]. \quad (8)$$

The stable states in the network, either memories or spurious states, correspond to the minima in E . When S is close to a particular memory, but far from all others, the

value of E is both large and negative. When S is far from all memories, the value of E is small.

The unlearning process is a repetitive algorithm that selectively modifies the connections in the network (Hopfield, et al., 1983; Clark et al., 1984; Clark et al., 1985). Each repeating unit in the algorithm, denoted as a trial, begins by initializing the network in a randomly chosen state. The neurons are updated until the network converges to a stable state, S' . The connection strengths are then changed to "unlearn" this state, i.e.

$$\Delta T_{ij}^{\text{unlearn}} = \begin{cases} -\epsilon S'_i S'_j & i \neq j \\ 0 & i = j \end{cases}, \quad (9)$$

where the unlearning strength, ϵ , is a small ($\epsilon \ll 1$) constant. The procedure is repeated for m trials ($m \gg 1$). The optimum choice of ϵ and m will be discussed later. Note that unlearning increases the energy of the state S' by

$$\Delta E^{\text{unlearn}} = +\epsilon \frac{N^2}{2} \left(1 - \frac{1}{N} \right). \quad (10)$$

The increase in the energy of neighboring states falls off quadratically with increasing Hamming distance from S' .

Storage Capacity

The probability of an error occurring when storing data (a spatial process) is equivalent to the probability of an error occurring when transmitting data (a temporal process). This equivalence suggests that the storage capacity of a network may be quantified in terms of the number of errors that occur when recalling memories. We define P as the (normalized) average number of bits that differ between the stored memory states and the stable states that correspond to these memories, i.e., the retrieved memories. The storage capacity, C , is defined in analogy with Shannon's (1948) entropic measure for the transmission capacity of a (binary symmetric) data channel, i.e.

$$C = nN [1 + P \log_2 P + (1 - P) \log_2 (1 - P)]. \quad (11)$$

The average number of differing bits, P , is calculated from

$$P = \frac{1}{N} \sum_{x=0}^N x p(x), \quad (12)$$

where $p(x)$ is the probability that x bits differ between the stored bit sequence and the retrieved sequence of a memory (see next section). When the retrieved memories are free of errors, $P = 0$ and the capacity scales linearly with the number of memories stored, i.e., $C = nN$. When there is no relation between the stored memories and the stable states in the network, the average number of bits in error is $P = 1/2$ and therefore $C = 0$.

METHODS

Calculations were performed primarily with networks consisting of 30 neurons. This number was large enough to serve as a useful model network but small enough to keep computation times acceptable. Unless noted otherwise, stored memories were composed of bit sequences chosen at random.

The number of bits that differed between a stored memory and the retrieved value of the memory, x , was determined by starting the network in the state corresponding to the stored memory and randomly updating the neurons (Eqs. 2 and 3). The updating was stopped when the network reached a stable state. This process was repeated for all stored memories to determine the distribution of $p(x)$'s, i.e., the probabilities for x bits differing. To decrease the statistical uncertainty in the $p(x)$'s we averaged the results from separate simulations. From the probabilities $p(x)$ and Eqs. 11 and 12 we calculated the storage capacity, C .

The probability of converging to a memory state was determined by starting the network in a state a specified Hamming distance, H , from a retrieved memory and updating the network until a stable state was reached. The probability of convergence was calculated from the number of times the network returned to the original memory, averaged over many starting states. This process was repeated for consecutively larger values of H . The radius of convergence, \bar{H}_{crit} , equaled the (interpolated) value of H for which the convergence probability was $1/2$.

Diagrams of the energy values were constructed along selected paths through the state-space of the network, where each location in state-space corresponds to one of the 2^N possible states of the network. The energy of a state was calculated from Eq. 7.

RESULTS

Storage Capacity

When four or less memories were stored in the network, the probability of a bit being in error between a stored memory and the retrieved value of the memory was, within statistical uncertainty, zero. This finding is in accord with the prediction of Amit et al. (1985b) for the behavior of asymptotically large networks. The subsequent storage of additional memories led to the onset of errors. The distribution of bit errors $p(x)$ following the storage of 8, 11, and 14 memories is shown in Fig. 1 *a*, *b*, and *c*, respectively (note the logarithmic scale). The probability of an error occurring in any bit [i.e., $1 - p(0)$] as well as the average number of bits in the error (i.e., NP), increased sharply with increasing n (compare Fig. 1, *a-c*).

The unlearning algorithm substantially reduced both the total number of bits in error and the average size of the error. This decrease is illustrated in Fig. 1, *a-c*, with $\epsilon = 1/N = 0.033$ and m as indicated. For example, with 11 memories stored (Fig. 1 *b*) only ~25% of the memories were free of incorrect bits before unlearning ($p(0) \approx 0.25$); this value increased to ~50% after 30 unlearning trials and ~95% after 120 trials.

The storage capacity was determined as a function of the number of memories stored using Eqs. 11 and 12 and calculations of $p(x)$ vs. x (e.g., Fig. 1, *a-c*); this is shown in Fig. 2. Consider first the capacity of the network before unlearning. When the number of memories was small, i.e., $n \ll N$, the probability of a bit being in error was also small and the capacity scaled linearly with the number of




FIGURE 1 The distribution of bit errors in the stable states corresponding to memories for a network of 30 neurons; note the logarithmic scale. Results are shown for (a) 8, (b) 11, and (c) 14 stored memories. Calculations were performed both before and after unlearning with $\epsilon = 0.033$ and m as indicated. Each datum represents the average of 200 simulations; the relative standard deviation of the mean for all data points was $\delta p/p \leq 0.09$.

memories stored, i.e., $C = nN$ (dashed line, Fig. 2). As the number of memories was increased toward $n \approx N$, the probability of a bit being in error increased such that C passed through a broad maximum and asymptotically reached a constant nonzero value. Note that the oscillatory behavior in C is a consequence of the small size of the network. It results from alternately summing over an odd then even number of stored states in the update rule (see Eq. 4).

The effect of unlearning on the capacity is shown in Fig. 2, with $\epsilon = 1/N = 0.033$ and m as noted. Unlearning extended the region over which C scaled linearly with n

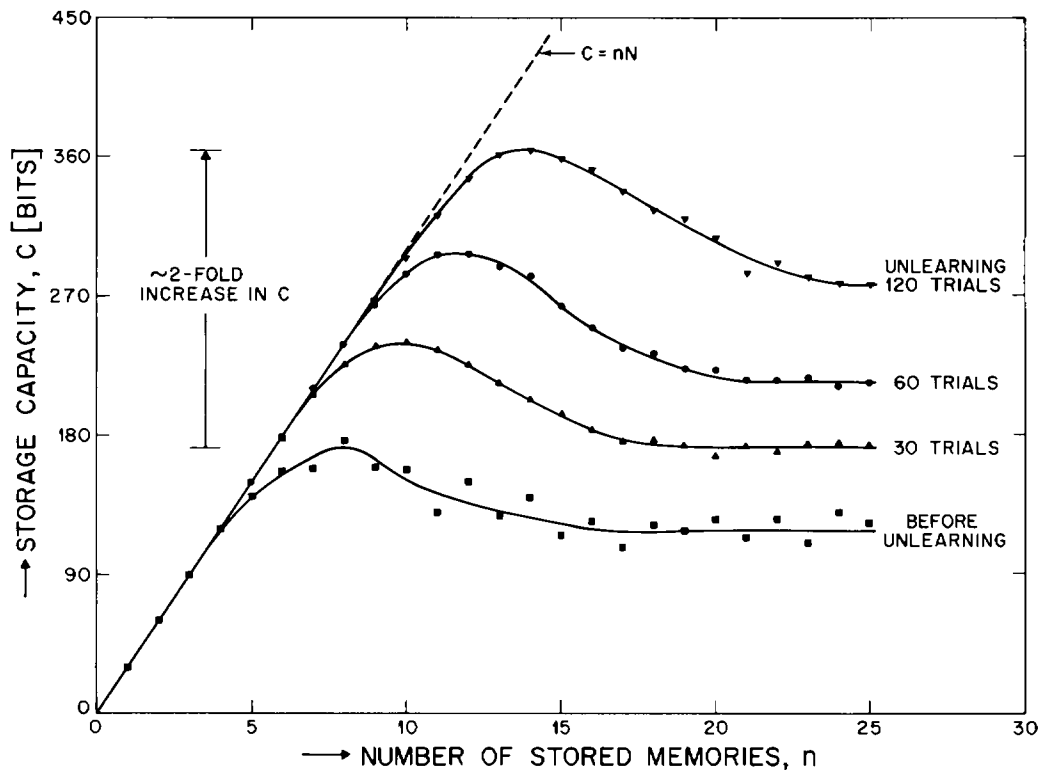


FIGURE 2 The information storage capacity, C (see Eqs. 11 and 12), as a function of the number of memories stored, n , for a network of 30 neurons. Calculations were performed before unlearning and, for $n > 4$, after unlearning with $\epsilon = 0.033$ and m as indicated. Each datum represents the average of 100 simulations; the relative standard deviation of the mean for all points was $\delta C/C \leq 0.04$. The solid lines were drawn as guides to the eye.

(dashed line, Fig. 2). This in turn increased the number of memories that could be stored subject to a maximum error rate. For example, after 120 unlearning trials C reached its maximum value after the storage of 14 memories, as opposed to 8 memories before unlearning. This improvement represents an approximately twofold increase in the capacity after unlearning, compared with the value before unlearning.

We assessed the sensitivity of the unlearning process to the size of the unlearning strength, ϵ . In these studies the product $m\epsilon$, which represents the maximum possible change in a connection strength T_{ij} , was fixed. The increase in storage capacity after unlearning was found to be independent of the size of ϵ for $\epsilon \leq 1/n$, but the extent of the increase diminished for larger values of ϵ . This is illustrated in Fig. 3 for the case $m\epsilon = 4$ and n as noted. Similar behavior was found for other values of $m\epsilon$. This result justifies the choice $\epsilon = 1/N$ used in the previous simulations (Figs. 1 and 2). It also suggests that memories must be sampled only once on average for unlearning to be effective.

The dependence of the storage capacity on the number of unlearning trials, m , was studied for fixed values of ϵ . We found that the capacity increased with successive unlearning trials, until a maximum value (for $n < N/2$) of $C \approx nN$ was reached. The maximum value of the capacity




FIGURE 3 The information storage capacity, C (see Eqs. 11 and 12), as a function of the normalized unlearning strength, $n\epsilon$, for a network of 30 neurons. For each value of n an increasingly larger value of ϵ was used in the unlearning procedure with m determined from the constraint $m\epsilon = 4$. Each datum represents the average of 225 simulations; the relative standard deviation of the mean for all points was $\delta C/C \leq 0.03$. The solid lines were drawn as guides to the eye.

remained essentially constant over the range $0.3n < m\epsilon < 0.6n$. Larger values of $m\epsilon$ caused a decrease in C concomitant with the magnitude of the T_{ij} 's approaching ϵ . This is illustrated in Fig. 4, with $\epsilon = 1/N$ and n as noted. When 14 memories were stored, unlearning increased the storage capacity by more than a factor of 3.

The previous result implies that the unlearning procedure allows nearly $n = N/2$ memories to be stored with essentially no errors in the retrieved content of those memories. To check if this result was independent of the size of the network, we extended the simulations to networks containing 100, 300, and 1,000 neurons. We found similar results for all values of N . In fact, the larger networks more closely approached the limit $C \rightarrow N^2/2$ for $n = N/2$.

Convergence Properties

We now examine the accessibility of the states stored in the CAM. The accessibility was characterized by the average radius of convergence for the retrieved memories, \bar{H}_{crit} . We found that this radius decreased linearly as the number of stored memories increased, i.e., $\bar{H}_{\text{crit}} \propto -n$. However, the value of \bar{H}_{crit} was only weakly changed by unlearning, increasing by less than one bit under optimal conditions of unlearning. This shows that the reduction of errors in the network is accompanied by only mild increases in the radius of convergence. Note, however, that a one-bit increase in \bar{H}_{crit} increases the number of states that converge to a memory by $\delta\bar{V}/\bar{V} \sim N/\bar{H}_{\text{crit}}$ (Eq. 5 with $1 \ll \bar{H}_{\text{crit}} \ll N$).

Although the radius of convergence was largely unaffected by unlearning, details of the convergence properties of the network were changed. This is illustrated for a

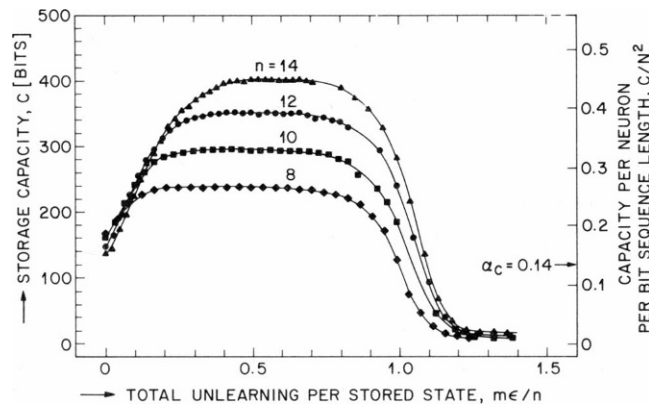


FIGURE 4 The information storage capacity, C (see Eqs. 11 and 12), as a function of the total unlearning per stored state, $m\epsilon/n$, for network of 30 neurons. For each value of n the value $\epsilon = 1/N$ and increasing values of m , as indicated, were used in the unlearning algorithm. The scale C/N^2 corresponds to the number of memories normalized to the number of neurons; the value $C/N^2 = \alpha_c = 0.14$ is the (pre-unlearning) limit of Amit et al. (1985b). Each datum represents the average of 100 simulations; the relative standard deviation of the mean for all points was $\delta C/C \leq 0.04$. The solid lines were drawn as guides to the eye.

network with nine stored memories (Fig. 5) by plotting the convergence probability as a function of the Hamming distance, H , from a memory. Unlearning skews the likelihood of convergence towards those states with $H < \bar{H}_{\text{crit}}$, in contrast to the gradual dependence of the convergence on H observed without unlearning.

The value of the \bar{H}_{crit} for all values of n , $n > 2$, was found to be substantially less than the value expected if all 2^N states in the network converged to one of the memories (Eqs. 5 and 6). For example, with nine stored memories the measured convergence radius, $\bar{H}_{\text{crit}} \sim 6\frac{1}{2}$ (Fig. 5), was substantially less than the maximum possible radius, $\bar{H}_{\text{crit}} (\text{max}) \approx 10$ (Eqs. 5 and 6 with $N = 30$ and $n = 9$). The measured radius encompasses ~ 20 times less states than maximum radius (Eq. 5). Where does the network converge to for initial states with Hamming distances $H > \bar{H}_{\text{crit}}$? We found that the network typically converged to a spurious state. Thus the state-space of the network can be qualitatively visualized as containing regions of attraction, of radius \bar{H}_{crit} , around each memory (see also McEliece et al., 1986). Outside of these regions the network is most likely to converge to a spurious state, whose number increases as 3^n (Amit et al., 1985a), as opposed to the closest memory.

Energy Diagrams

To illustrate how unlearning increases the storage capacity, we constructed a diagram of the energy values along two paths through the state-space of a network. The energies were calculated following the storage of one, four, seven, and nine memories; Fig. 6, *a-d*. The paths begin at the location of the first stored memory, M^1 , and were

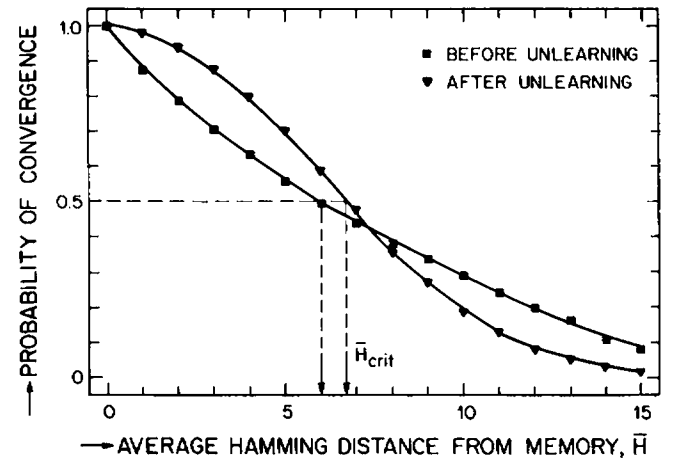


FIGURE 5 The average probability of converging to a retrieved memory state as a function of the Hamming distance, H , from that state for a network of 30 neurons with nine stored memories. Probabilities were calculated both before and after unlearning, with $\epsilon = 0.033$ and $m = 120$. Each datum represents the average of 140 simulations; the standard deviation of the mean for all points was $\delta p \leq 0.02$. The solid lines were drawn as guides to the eye.

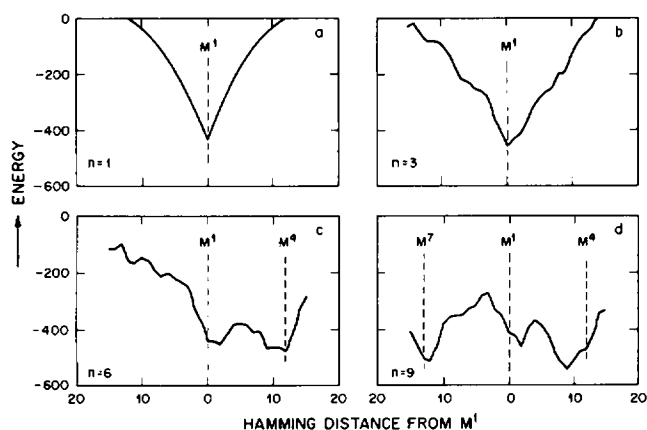


FIGURE 6 The energy values along two 16-bit paths through the state-space of a 30 neuron network. Both paths begin from the state M^1 . The energy, calculated from Eq. 7, is shown after the storage of (a) a single memory (M^1), (b) three memories, (c) six memories, and (d) nine memories. The paths, two of $[30!/(30-16)!] \approx 3 \cdot 10^{21}$ possible 16-bit paths, encompass 31 of the $2^{30} \approx 1 \cdot 10^9$ states in the network. They were constrained to pass through memory states M^1 , M^4 , and M^7 (dashed lines) and the corresponding retrieved states, i.e., minima in the energy, after the storage of nine memories.

constrained to pass through M^1 , M^4 , and M^7 and the retrieved states corresponding to these memories. With only M^1 present, the energy of states at succeeding distances away from M^1 had a quadratic dependence on the Hamming distance (Fig. 6a). This dependence follows from Eq. 7 with $\nu = 1$. As additional memories were stored, the smoothness of the energy curve was lost (cf. Fig. 6a and b) and the minima corresponding to M^1 , M^4 , and M^7 shifted from their correct position (Fig. 6d). These shifts correspond to a change in the bit sequence specifying the memory, i.e., a retrieval error.

The minima corresponding to memories are essentially restored to their correct positions after successive unlearning trials (Fig. 7). The mechanisms found to be responsible for this improvement will be illustrated by considering the effect of unlearning on two sources of retrieval errors. One source is the formation of spurious states (Hopfield et al., 1983; Amit et al., 1985a). Although the spurious states have relatively weak minima, they can shift the position of the memories. This is analogous to the apparent shift in the position of a spectroscopic absorption line caused by a second line, which is both weaker and narrower than the original line, that peaks close to the position of the original line.

The energy diagrams of Fig. 6, a-d can be used to illustrate the effect of spurious states on the position of the memories. The minimum corresponding to M^1 remained correctly positioned after three memories were stored (Fig. 6, a and b). After the sixth memory was stored a spurious state appeared near M^1 (Fig. 6c) and shifted the position of the minimum corresponding to M^1 by two bits (Fig. 6, c and d). The energy of the minimum for the spurious state alone, however, was much less than the energy of M^1 . Thus

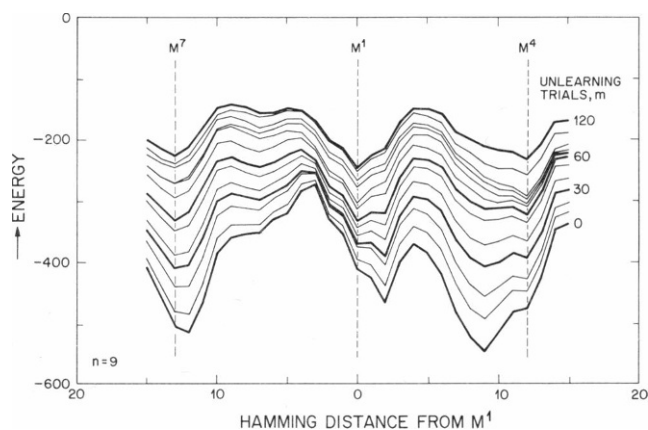


FIGURE 7 The energy values along two paths as a function of the number of unlearning trials. The paths are the same as those used in Fig. 6, a-d, with nine stored memories (Fig. 6d). The values $\epsilon = 0.033$ and m as indicated were used in the unlearning procedure. Note that the stable states corresponding to each memory (M^1 , M^4 , and M^7) gradually return to their correct positions with successive unlearning trials.

the unlearning process primarily affected, and essentially removed, the spurious state (Fig. 7). The minimum corresponding to M^1 consequently returned to its correct position.

A second source of retrieval errors was found to be the "attraction" between memories. To illustrate this point we constructed a network with four memories; a network of this size is free of errors when the memories are composed of randomly chosen bit sequences (see previous section). In this particular construction the bit sequences of the first three memories, M^1 , M^2 , and M^3 , were chosen at random while the sequence of M^4 was chosen to be a specified Hamming distance, H_{initial} , from M^3 . The distance between the stable states corresponding to M^3 and M^4 , H_{final} , was determined for each value of H_{initial} . Fig. 8 shows the shift in this distance, $H_{\text{initial}} - H_{\text{final}}$, as a function of H_{initial} . When M^3 and M^4 were close ($H_{\text{initial}} \leq 8$ bits) the minima for the two states shifted toward each other, i.e., $H_{\text{final}} < H_{\text{initial}}$. The memory with the shallower minimum was shifted the most. Unlearning equalized the energy of both memories, i.e., the depth of the two minima. This in turn reduced the shift in Hamming distance (Fig. 8). Conversely, for a fixed shift in Hamming distance (i.e., a fixed probability of bit errors) unlearning allows memories to be spaced closer together. This increases the density with which memories can be stored in the network and suggests a geometric interpretation for the increase in storage capacity that accompanies unlearning.

DISCUSSION AND CONCLUSIONS

It is interesting to consider the possible physiological significance of these results. Crick and Mitchison (1983, 1986) proposed that during REM sleep, mammals "dream in order to forget" as a means for removing spurious, undesirable memories from the neocortex. This process

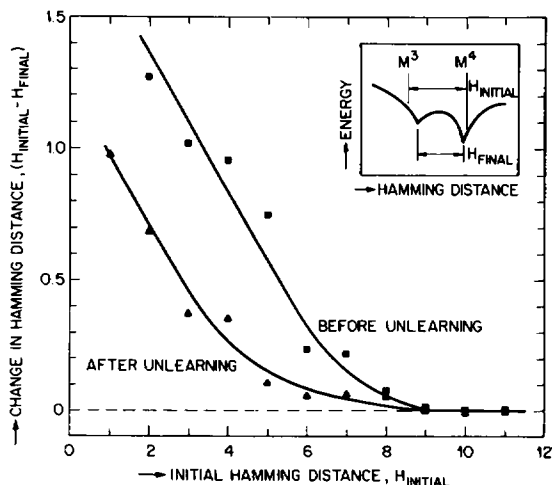


FIGURE 8 The shift in Hamming distance between the minima corresponding to adjacent memories as a function of the initial Hamming distance between these memories. The bit sequences of M^1 , M^2 , and M^3 were selected at random while the sequence of M^4 was chosen to be H_{initial} bits from M^3 . The values $\epsilon = 0.033$ and $m = 30$ were used in the unlearning procedure. Each datum represents the average of 70 simulations; the standard deviation of the mean for all points was $\delta H \leq 0.1$. The solid lines were drawn as guides to the eye. The inset illustrates the relation between H_{initial} and H_{final} on an energy diagram.

was conjectured to weaken obsessive memories and to prevent inappropriate associations. In independent work, Hopfield et al. (1983) showed that unlearning enhanced the performance of their CAM in accessing stored memories and in minimizing the presence of spurious memories. Their results appear to support Crick and Mitchison's conjecture. Based on the results found in this study, we suggest that REM sleep also increases the efficiency of cortical networks by increasing the number of memories that can be recalled correctly.

We have shown that the information storage capacity of a CAM is substantially increased by an unlearning algorithm. The mechanisms for this increase were illustrated by analyzing the energy of the states in the network. Our results imply that hardware implementations of CAMs may benefit from the application of such algorithms, particularly when a priori knowledge of the contents of the stored memories is unavailable.

We thank J. P. Allen, F. H. C. Crick, G. Feher, D. R. Fredkin, and J. J. Hopfield for useful discussions and suggestions and Mona Ludhag for technical assistance.

This work was supported in part by a grant from the National Institutes of Health (GM-13191) to G. Feher. D. B. Pendergraft was the recipient of a University of California Presidential Undergraduate Fellowship.

Received for publication 19 March 1986 and in final form 29 July 1986.

REFERENCES

- Amit, D. J., H. Gutfreund, and H. Sompolinsky. 1985a. Spin-glass models of neural networks. *Phys. Rev. A* 32:1007-1018.
- Amit, D. J., H. Gutfreund, and H. Sompolinsky. 1985b. Storing infinite number of patterns in a spin-glass model of neural networks. *Phys. Rev. Lett.* 55:1530-1533.
- Anderson, J. A., J. W. Silverstein, S. A. Ritz, and R. S. Jones. 1977. Distinctive features, categorical perception, and probability learning: some applications of a neural model. *Psychol. Rev.* 84:413-451.
- Clark, J. W., J. V. Winston, and J. Rafelski. 1984. Self-organization of neural networks. *Phys. Lett.* 102A:207-211.
- Clark, J. W., J. Rafelski, and J. V. Winston. 1985. Brain without mind: computer simulation of neural networks with modifiable neuronal interactions. *Phys. Rep.* 123:215-273.
- Cohen, M. A., and S. Grossberg. 1983. Absolute stability of global pattern formation and parallel memory storage by competitive neural networks. *IEEE Trans. Sys. Man. Cybern.* 13:815-826.
- Cooper, L. N. 1973. A possible organization of animal memory and learning. In *Proceedings of the Nobel Symposium on Collective Properties of Physical Systems*. B. Lundqvist and S. Lundqvist, editors. Academic Press, Inc. New York. 252-264.
- Crick, F., and G. Mitchison. 1983. The function of dream sleep. *Nature (Lond.)* 304:111-114.
- Crick, F., and G. Mitchison. 1986. REM sleep and neural nets. In *Cognition and Dream Research*. R. Haskell, editor. In press.
- Hopfield, J. J. 1982. Neural networks and physical systems with emergent collective properties. *Proc. Natl. Acad. Sci. USA* 79:2554-2558.
- Hopfield, J. J. 1984. Neurons with graded response have collective computational properties like those of two-state neurons. *Proc. Natl. Acad. Sci. USA* 81:3088-3092.
- Hopfield, J. J., D. I. Feinstein, and R. G. Palmer. 1983. "Unlearning" has a stabilizing effect in collective memories. *Nature (Lond.)* 304:158-159.
- Hopfield, J. J., and D. W. Tank. 1986. Computing with neural circuits: a model. *Science (Wash. DC)* 233:625-633.
- Kohonen, T. 1977. *Associative Memory: A System Theoretic Approach*. Springer-Verlag, Berlin.
- Little, W. A., and G. L. Shaw. 1975. A statistical theory of short and long term memory. *Behav. Biol.* 14:115-133.
- McEliece, R. J., E. C. Posner, E. R. Rodemich, and S. S. Venkatesh. 1986. The capacity of the Hopfield associative memory. *IEEE Trans. Info. Theory*. In press.
- Nakano, K. 1972. Associatron: a model of associative memory. *IEEE Trans. Sys. Man. Cybern.* 2:380-387.
- Peretto, P. 1984. Collective properties of neural networks: a statistical physics approach. *Biol. Cybern.* 50:51-62.
- Shannon, C. E. 1948. The mathematical theory of communication. *Bell Sys. Tech. J.* 27:379-423.