The Recall of Events through the Learning of Associations between their Parts

A. R. Gardner-Medwin

Proceedings of the Royal Society of London. Series B, Biological Sciences, Volume 194, Issue 1116 (Oct. 29, 1976), 375-402.

Proc. R. Soc. Lond. B. **194**, 375–402 (1976)
Printed in Great Britain

The recall of events through the learning of associations between their parts

BY A. R. GARDNER-MEDWIN
Department of Physiology, University College London, London WC1E 6BT

(*Communicated by D. R. Wilkie, F.R.S. – Received* 10 *February* 1976)

A form of recall is discussed in which some elements of an event serve to initiate recall of the whole set of elements which comprised an event. The conditions which enable learned associations between pairs of elements to be used for such recall are quantitatively assessed for recall procedures which use only direct associations from the initial elements ('simple recall') and which use both direct and indirect associations ('progressive recall'). The theory is equivalent to the analysis of a neural model in which randomly arranged excitatory connections between a set of cells can be strengthened by simultaneous activity in their presynaptic and postsynaptic elements. The problems involved in a direct neural implementation of the model are discussed, and it is seen that some of the mechanisms resemble known physiological mechanisms in the cerebral cortex. The relationship between recall and recognition learning is discussed and the evidence relating these functions to the hippocampal formation is assessed.

1. INTRODUCTION

The ability to recall events in some detail is an obvious capability of the human brain. Recall can be triggered off either by the naming of an event or by the subsequent encounter of some of the elements of experience which made up the event. This paper is concerned with the situation in which a part of an event triggers off recall. This is a kind of recall for which the human brain seems particularly adept. For illustration, one can consider what would be involved in designing a computerized reference system which could, as quickly as the reader, furnish the details of a well known story which concerns a girl, a bear, and a bowl of porridge. The juxtaposition of a small number of details may be enough to identify one event unambiguously from a very large number of events in the individual's experience; and it can rapidly lead to the recall of many more of the details which were associated with the event. Not all the details may be recalled. Some details may be missed even though they have in fact been learned and could be recalled by some different procedure for initiating recall; and some details may be recalled which are spurious and were never part of the event. Though performance is imperfect, this is a useful and probably a major way in which we use our long term memory.

In this paper an 'event' will be taken to consist of a particular combination of what are called 'elements of experience'. Each such element may occur in a large number of events. The argument of the paper is that recall can occur through the development in the brain of associations between pairs of elements which occur together. Successive elements are recalled through their associations from some of the elements initially presented, or from elements already recalled. Recall in which only the direct associations from the initial elements are used is referred to as 'simple recall', while a procedure in which the associations from newly recalled elements become available for inducing further recall is described as 'progressive recall'. It is an important qualitative feature of such recall that the system does not have to store information about all the pairs of elements which have occurred together; information about a small proportion of them can, within certain constraints, permit nearly perfect recall.

Other mechanisms can exist for recalling events. For example, a commonly used card index system can perform the same task in a different way. In such a system holes are punched on a card to correspond to different circumstances under which access to the card may be required. The punched pattern can be regarded as an event consisting of a number of elements, stored on a particular card. This event can subsequently be recalled (separately from other events stored in a stack of cards) by the use of a probe at a few sites corresponding to elements of the required event. This system does not operate through the development of associations between pairs of elements, but it works perfectly satisfactorily within certain quantitative constraints which are different from the constraints for the system described in this paper. The recall produced with index cards has however some qualitative features which differ from human recall. Examples are the all-or-none emergence of the required elements with the requisite card, and a complete failure of recall when a single wrong element is used for initiation. It is also not easy to see how the logical operations involved in such recall could be performed with simple configurations of neurons.

Recall which takes place through the development of associations between pairs of elements turns out to have strong resemblances to human recall. Furthermore it can be implemented in a direct way by using elements similar to neurons with reasonably plausible properties. The analysis in the paper deals in fact with such an implementation. It describes a neural network with principal cells corresponding to elements of experience, some of which are active during an event. The building up of associations between elements corresponds to the strengthening of excitatory synapses between pairs of cells whenever they are both active together. The analysis of the performance of the network can be interpreted on two levels. On one level the neurones are simply convenient symbols with which to examine how recall could occur through the development of associations between pairs of elements. Each element might in this interpretation correspond to some complex pattern of activity within the brain. On the second level it can be suggested that the activation of a single cell in the

model is similar to whatever occurs in the brain as an element of experience, and that the existence of synapses in the brain with the properties described does in fact account for the learning of associations and the recall of events. On this level the theory is related to an effect in Marr's (1971) theory of archicortex.

If the implications of the analysis were to be defended rigorously it would have to be on the basis of the first interpretation described above, which could perhaps yield predictions of relevance to the psychologist, but little of relevance to the physiologist or anatomist. On the other hand the second interpretation is so specific as to be extremely vulnerable. Immediate objections come from the oversimplified assumptions about neurone behaviour, from the fact that extensive self-re-exciting networks of the sort employed are not for certain known to exist in the nervous system and from the fact that the modifiable synapses are of a type which has not yet been demonstrated, though it was first proposed by Hebb in 1949. None of these objections has a firm base however, so perhaps the theory may prove stimulating.

Summaries of some of the more mathematical sections of the paper will be found at the ends of the relevant sections.

2. The network

The network consists of a set of N principal cells connected to each other through modifiable synapses. There are also some fixed extrinsic connections (figure 1*a*).

Each of the N principal cells makes one anatomically defined synaptic connection onto each of R of the other N-1 cells. Which of the cells receive synapses is determined at random and independently for each cell. There are thus NR synapses arising from intrinsic connections within the network. All the intrinsic synapses are physiologically ineffective before the experience begins; but each synapse becomes effective as an excitatory synapse of fixed strength if and when during the experience of the network both its presynaptic and its postsynaptic cells are active together. Once rendered effective, a synapse remains effective for all time. The intrinsic excitatory synapses are thus extreme versions of the type of modifiable synapse first postulated by Hebb (1949).

One or more inhibitory cells exist, with connections onto each of the N principal cells. These cells can set the threshold of each of the principal cells to the same value T. A principal cell with a threshold T will become active if it is receiving the excitatory influence of T or more than T effective active synapses from other principal cells.

The principal cells of the network can also be made to fire by means of excitatory projections from outside the network. These extrinsic projections can over-ride any inhibition experienced by the cells and are capable of activating any subset of the principal cells. Activation by the extrinsic projections of a set of W of the N principal cells will be described as an 'event' of size W (figure 1*b*). Activation

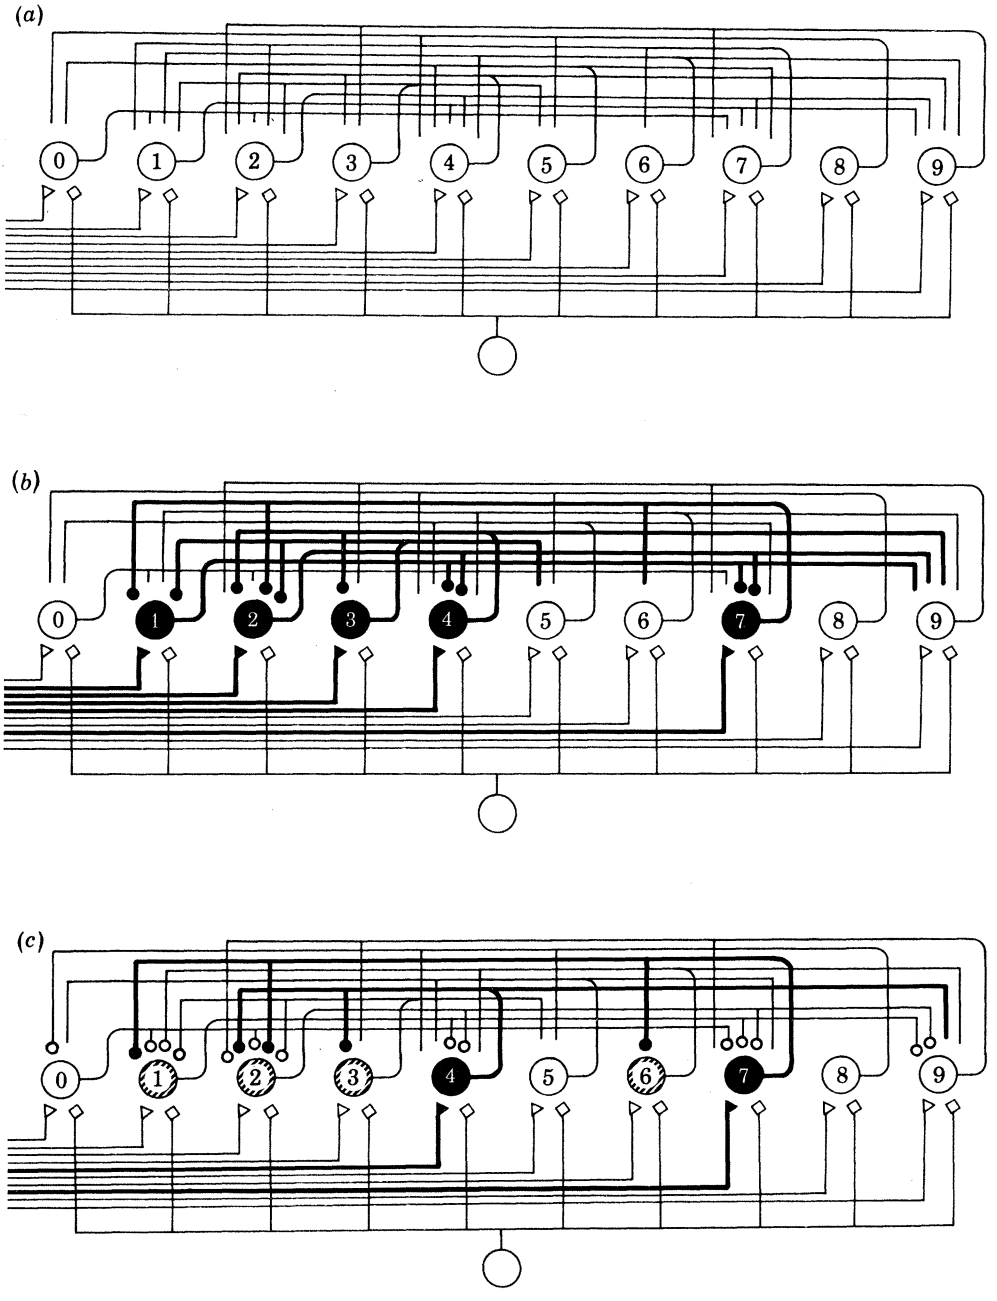

FIGURE 1. For description see opposite.

by the same means of some proportion of the cells in an event is used subsequently for the initiation of recall (figure 1*c* and §5). In figure 1 the fibres of the extrinsic projection are for simplicity drawn as having a one to one relationship with the principal cells. This is one way in which the required patterns of activity could be set up from outside the network; but it is not the only way.

Some new modifiable synapses will generally become effective for the first time during each newly experienced event. The overall proportion of the modifiable synapses which are effective at any one time as a result of the experience of the network is denoted by ρ.

Summary. A population of cells experiences 'events' consisting of activity in specific combinations of the cells. Potentially effective synapses exist from each cell onto some but not all of the others. These synapses become strengthened permanently (as excitatory synapses) whenever their cells of origin and of destination are active together in an event. Systems giving uniform inhibition of the cells and allowing extrinsic activation are described and the quantitative parameters of the system are defined.

3. Notation

N	total number of principal cells
R	number of synaptic connections made by each cell
ρ	proportion of the modifiable synapses which have been rendered effective
W	number of cells active during an event
w	number of cells constituting part of an event which are active during recall
w_0	number of cells made active in order to initiate recall of an event
M	total number of events experienced by the network
T	threshold of the principal cells at any one time

Figure 1. A simple illustrative network. (*a*) *The network before experience.* Ten numbered principal cells ($N=10$) receive excitatory synapses (—◁) and inhibitory synapses (—◇) from other cells which are under extrinsic control. Axon collaterals from the principal cells terminate on other principal cells in Hebb synapses which are potentially excitatory synapses though they are initially ineffective (terminated lines). Each axon gives rise to three synapses distributed at random ($R=3$). Note that different principal cells receive varying numbers of synapses. (*b*) *Learning.* The active cells and axons during an event (12347; $W=5$) are drawn in heavy black. Active terminals making contact with active cells become permanently strengthened (—●, or —○ when inactive). (*c*) *Simple recall.* The network is shown after synapses have been strengthened by three events (12347, 01289, 01367; $M=3$). Recall of the first event is elicited by activation of two of its cells (47; $w_0=2$). When the threshold $T = 1$ all the correct cells are recruited (hatched) together with one spurious cell (no. 6). With $T = 2$, one correct cell would be recruited with no spurious cells.

A $= WR/N$; a measure of the interaction between cells during an event of size W

a $= wR/N$

a_0 $= w_0 R/N$

P_{spur} the probability that any individual incorrect (spurious) cell will be activated during recall

$P(X)$ probability of occurrence of X

$\mathscr{E}(x)$ expectation of a variable x; this is the average of all possible values of x, weighted with their individual probabilities; it is thus the average value which a variable would attain in a large number of independent determinations

$\pi(\lambda,\theta)$ $= P(x \geqslant \theta)$, given that x is distributed according to a Poisson distribution with $\mathscr{E}(x) = \lambda$

$C(x,y)$ $= y!/(x!(y-x)!)$: the number of ways of selecting x items out of a total of y

4. The learning process

During the presentation of a new event some synapses in the network will generally become effective (figure 1*b*). If the event is identical to a previously experienced event or substantially similar, then either no synapses or a small number of synapses may become modified. In order to simplify the treatment it will be assumed that all the events experienced by the network are statistically independent of each other and consist of active cells distributed at random throughout the population. Furthermore, all events will be taken to be of the same size (W). These assumptions may be unrealistic in relation to real learning systems; but they serve to simplify the mathematics without altering the fundamental nature of the recall problem.

Consider the number of synapses which satisfy the modification conditions during a single event:

Number of presynaptic terminals active during an event $= WR$;

expected proportion of these which synapse upon active cells $= W/N$.

Therefore the expectation value of the total number of synapses satisfying the modifying conditions $= W^2R/N$,

and the probability that any one synapse in the network satisfies the modification conditions $= W^2/N^2$.

When the network has experienced a total of M independent and equally sized events we have:

Probability that any one synapse has ever satisfied the modification conditions $= 1-(1-W^2/N^2)^M$.

This expression gives us the expectation value for the proportion of synapses which have become modified:

$$\left.\begin{aligned} \mathscr{E}(\rho) = 1-(1-W^2/N^2)^M &\simeq 1-\exp(-MW^2/N^2) \quad \text{if} \quad W/N \ll 1, \\ &\simeq MW^2/N^2 \qquad\qquad\quad \text{if} \quad MW^2/N^2 \ll 1. \end{aligned}\right\} \tag{1}$$

In order to reach any given value of ρ, the number of events required (M) is critically dependent on the proportion of the cells which are active during an event (W/N). Thus to modify half of the synapses in the network ($\rho = 0.5$) requires 70 events if $W/N = 10^{-1}$ and 7×10^5 events if $W/N = 10^{-3}$. If these numbers of events are exceeded the efficiency of the network as a learning device is likely to deteriorate, since further events decrease rather than increase the specificity of the pattern of modified synapses. In the extreme case in which ρ approaches 1, all the information stored in the network about the individual events is lost, since the same end result would have been achieved after any sufficiently large set of events.

Summary. The proportion of synapses in the network which will be modified after experience of a given number of events is calculated. Different events are assumed to involve the same number of active cells and to be statistically independent. Networks which experience too many events retain little information about individual events; but the acceptable number rises sharply if the proportion of cells active during any one event is made smaller.

5. Simple recall

In a recall trial the network is required to activate as many as possible of the cells which were active in a particular event, without at the same time activating other (spurious) cells. The event to be recalled is identified by activating within the network a subset (size w_0) of the cells which constitute the whole event.

In the simplest procedure for obtaining recall the threshold (T) of the cells is set to a level at which some of the required cells are activated through the projections from the w_0 cells initiating recall (figure 1*c*). This is called 'simple recall' to distinguish it from 'progressive' or 'multi-stage' recall discussed in the next section.

The total number of active synaptic terminals within the network is $w_0 R$. The average number received by individual cells is denoted by a_0 and is equal to ($w_0 R/N$). All the active terminals onto cells of the required event will have been rendered effective, while an expected proportion (ρ) of the terminals onto spurious cells will be effective. The numbers of active effective synapses onto particular cells will be subject to a binomial distribution in which the probability of a cell receiving (r) synapses is:

$$\left.\begin{aligned} &C(r,w_0)\left(\frac{R}{N}\right)^r\left(1-\frac{R}{N}\right)^{w_0-r} \quad \text{for correct cells,} \\ \text{and} \quad &C(r,w_0)\left(\rho\frac{R}{N}\right)^r\left(1-\rho\frac{R}{N}\right)^{w_0-r} \quad \text{for spurious cells.} \end{aligned}\right\} \tag{2}$$

If $r \ll w_0$, and in nearly all cases if $R/N \ll 1$, the distributions are close to Poisson distributions in which the mean numbers of active effective synapses are a_0 for correct cells and ρa_0 for spurious cells. Poisson distributions will be assumed in the following analysis, and the small corrections arising from this assumption will be discussed later in §7.

The threshold, T, for simple recall must be set high enough so that there is a sufficiently low probability of spurious cells being activated. In general this means that only a proportion and not all of the required cells will be activated.

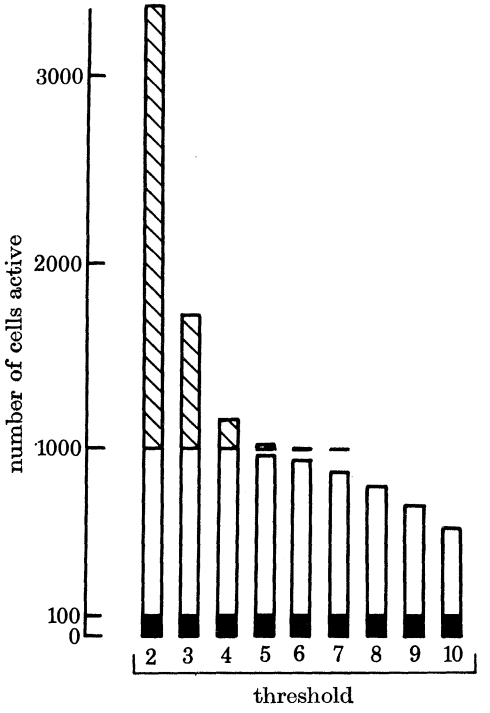

FIGURE 2. Simple recall. The expectations of the numbers of correct cells and spurious cells which will be activated with various choices of threshold are shown for a network of 10^4 cells which has learned 10 events of size 10^3. 100 cells are assumed active to initiate recall. $N = 10^4$, $W = 10^3$, $w_0 = 10^2$, $M = 10$, $\rho = 0.1$, $R = 10^3$. ■, Cells initiating recall; □, correct cells; ▨, spurious cells.

The expected numbers of correct and spurious cells which would be activated with different choices of threshold can be calculated for any particular network situation. In figure 2 these numbers are given for a set of arbitrary parameters set out in the figure caption. In this case a threshold $T = 7$ might be judged to give satisfactory recall, in which a final expected 883 of the 1000 cells in the event are activated while on average 0.7 spurious cells may be expected. A larger proportion of the correct cells may be activated if a few more spurious cells can be tolerated: for example 940 correct cells with an expected 5.3 spurious cells ($T = 6$), or 974 correct cells with 33 spurious cells ($T = 5$).

Recall in the example described above starts with $w_0 = 100$. If this figure were reduced or if the proportion of modified synapses ($\rho = 0.1$) were increased, then the quality of the final recall would deteriorate. This will be seen when we examine the quality of simple recall performance in the general conditions, determined by the following equations:

$$\left.\begin{aligned} &\text{expected value of the number of spurious cells} &&= (N-W)\pi(\rho a_0, T),\\ &\text{expected value of the final number of correct cells} &&= w_0+(W-w_0)\pi(a_0, T).\end{aligned}\right\} \quad (3)$$

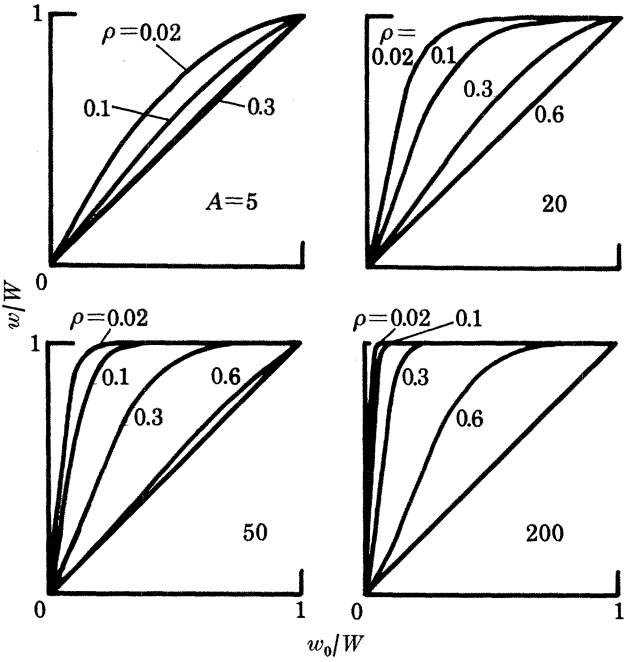

FIGURE 3. Simple recall. The performance of networks with various values of $A\,(=WR/N)$. The proportion of the required cells which are active after recall is plotted against the proportion made active to initiate recall. The probability that any one spurious cell will be activated is taken as $P_{\text{spur}} = 10^{-4}$ throughout. Each set of curves shows the performance for various values of ρ (the proportion of the synapses which are modified), which is closely proportional to the number of events experienced.

The acceptable number of spurious cells is to some extent arbitrary: it depends on how the recall output is to be used. This can be easily seen in the case of everyday human recall, where it is generally true that spurious details are more acceptable in reminiscences of a holiday than in evidence from a witness box. A limit will be set (P_{spur}) on the probability that any one of the $(N-W)$ possible spurious cells will be activated. The expected number of spurious cells ($=(N-W)P_{\text{spur}}$) depends largely on N for any given value of P_{spur}, since we have already seen (§4) that $W/N \ll 1$ if the network is to cope with a large number (M) of events.

The quality of recall performance is most easily understood if we determine

how the final activated proportion of the cells in an event (w/W) depends on the initial activated proportion (w_0/W). When expressed in these terms, the equations become:

$$\mathscr{E}\left(\frac{w_{\text{final}}}{W}\right) = \frac{w_0}{W} + \left(1 - \frac{w_0}{W}\right)\pi\left(\frac{w_0}{W}A, T\right), \tag{4}$$

where T is chosen such that:

$$\pi\left(\frac{w_0}{W}\rho A, T\right) = P_{\text{spur}}.$$

The solutions for these equations depend on three parameters: P_{spur}, A, ρ. In figure 3 a single fixed value of P_{spur} is chosen ($=10^{-4}$). A set of curves is plotted

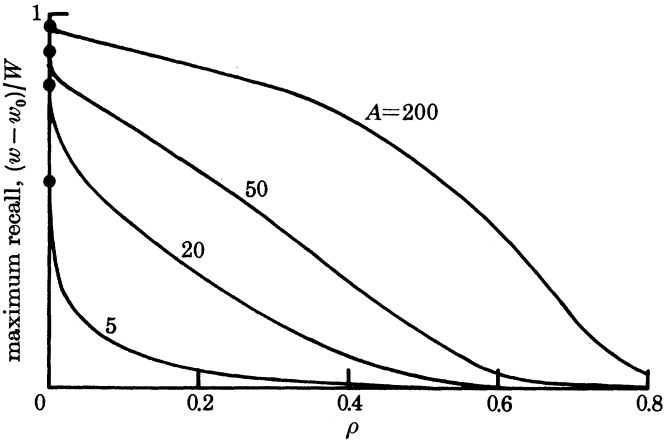

FIGURE 4. Simple recall. The maximum proportion of the required cells of an event which can be recalled (excluding those initially active) with an optimal choice of (w_0), plotted as a function of the proportion of the synapses which have been modified. Various values of $A\,(=WR/N)$ are assumed, with $P_{\text{spur}} = 10^{-4}$ throughout. The graphs show how the performance on individual tasks deteriorates as the network's total experience increases.

for each of several values of A, and within each set an individual curve is calculated for a given value of ρ. Smooth curves have been drawn through the points calculated for integer values of T.

When the vertical heights of the curves in figure 3 reach 1.0, then virtually all the cells of an event will be activated. It can be seen that in most cases this requires that a substantial fraction of the cells must be used to initiate recall. When these initially active cells are subtracted, only the proportion of the cells represented by the height of the curves above the diagonal lines ($w/W = w_0/W$) can be said to be genuinely recalled. This proportion has a maximum value for any one curve which in general occurs for low values of w_0/W when ρ is small and for $w_0/W \simeq 0.5$ when ρ becomes larger. The maximum fraction of cells which can be recalled is plotted in figure 4 as a function of ρ, for various values of A.

Figures 3 and 4 show how the recall performance deteriorates as the proportion of synapses in the network which are modified (ρ) increases. The value of ρ increases steadily with the number of events (M) experienced by the network (equation (1)). Thus the performance of the network on individual recall tasks deteriorates as its experience increases.

Summary. In the recall situation a few of the cells of an event are initially activated from outside the network. A suitable level of threshold (set via the inhibitory connections) then permits some or all of the remaining cells of the event to be activated through their strengthened excitatory connections from the initially active cells. The threshold must not be set too low or else spurious cells will be activated which were not part of the required event. The extent to which recall can take place is derived for various combinations of parameters.

6. Progressive recall

In the mechanism for simple recall described in the last section, the cells which are identified are those which receive the greatest direct activation from the set of cells used to initiate recall. We have seen how in some cases it is only possible to activate by this means a small proportion of the required cells. In the mechanism for progressive recall this small number of extra cells is added to the initially active set to form a larger group which is itself able to activate a still larger group of correct cells on repetition of the procedure for simple recall. Further repetition takes place until no additional correct cells are recruited. Progressive recall thus makes use of the indirect projections of the initially active cells as well as their direct projections. It involves a multi-stage repetition of the procedure for simple recall. The physiological means by which this could occur will be discussed in §11.

The progressive recall procedure will be illustrated firstly for a network which has learned only one event. Since no possibility exists of activating spurious cells during recall of this event, the threshold, T, may be set to 1. Suppose that the number of cells active after r stages of recall is w_r. Then the expectation value for the number of those $W-w_0$ cells initially required, which are by now receiving excitation through effective synapses is

$$(W-w_0)\pi(a_r, 1) \quad \text{where} \quad a_r = w_r R/N.$$

This is the expectation value for the number of these cells which will be active after the next stage of recall. Thus

$$\mathscr{E}(w_{r+1}) = w_0+(W-w_0)\pi(a_r, 1),$$

and

$$\mathscr{E}(a_{r+1}) = a_0+(A_0-a_0)\pi(a_r, 1).$$

If we suppose that at every stage the number of cells recruited is equal to its expectation value, then we can trace the successive increments of a on a graph

in which $\pi(a,1)$ is plotted as a function of a (figure 5). The final stable value of $a(=a_\infty)$ is the value at which the curve $y=\pi(a,1)$ intersects the straight line $y=(a-a_0)/(A-a_0)$. This final value is substantially higher than the value obtained with simple recall $(=a_1)$, and the number of cells activated (w_∞) is correspondingly a larger proportion of the total number of cells in the event.

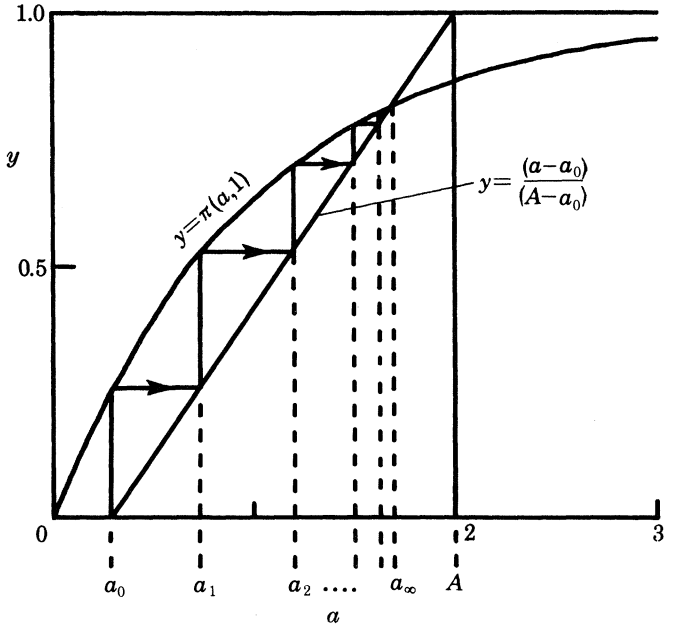

FIGURE 5. Progressive recall illustrated in a network which has only learned one event. $A = 2$; $a_0 = 0.3$. Values of a are derived for each stage of recall (a_1, a_2, etc.) from the graphs $y = \pi(a,1)$ and $y = (a-a_0)/(A-a_0)$.

When a network has experienced more than one event it will generally be necessary to use a higher threshold in order to prevent the activation of too many spurious cells. The same sort of considerations apply as for simple recall, except that the conditions must now be applied throughout all the stages of recall. If the threshold after r stages of recall is T_r, then we have

$$\pi(\rho a_r, T_r) < P_{\text{spur}}.$$

If we keep the threshold constant for all the stages of progressive recall then the analysis of performance is entirely analogous to the simple case when only one event has been experienced, with the substitution of $\pi(a_r, T)$ for $\pi(a_r, 1)$. The equation governing the steps of growth is:

$$\mathscr{E}(a_{r+1}) = a_0 + (A-a_0)\pi(a_r, T).$$

The final value of a $(=a_\infty)$ is the value at which the curve $y=\pi(a,T)$ intersects the line $y=(a-a_0)/(A-a_0)$, and the whole recall procedure will result in an

acceptably small number of spurious cells provided that:

$$\pi(\rho a_\infty, T) < P_{\text{spur}}.$$

This condition sets a higher threshold than is necessary for simple recall, since the threshold must be adequate to prevent the activation of spurious cells in the

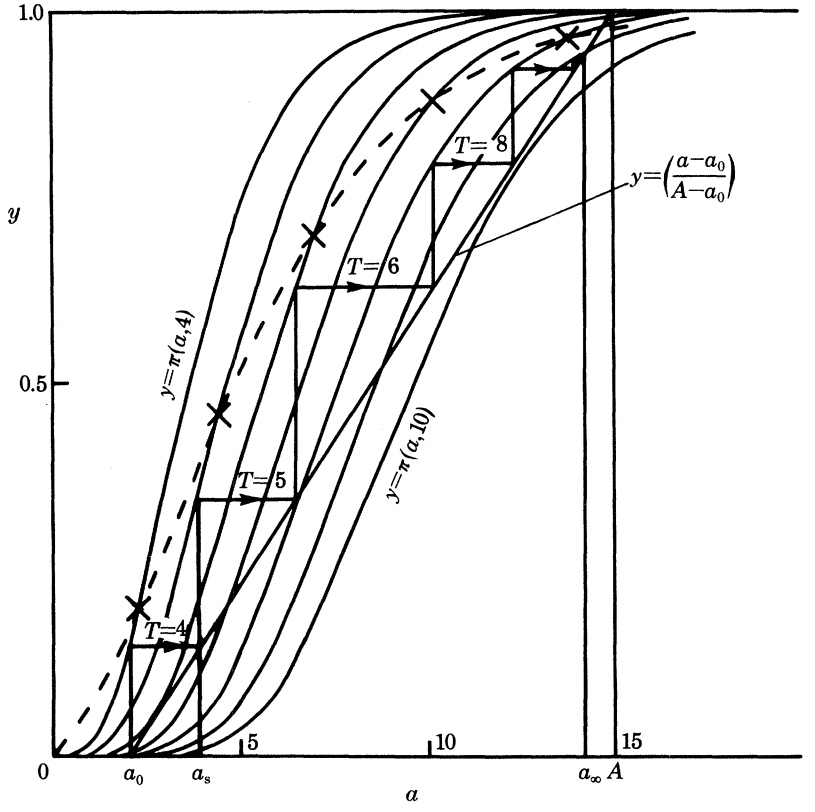

FIGURE 6. Progressive recall in a network which has learned more than one event ($\rho = 0.1$). Crosses marked on the curves $y = \pi(a,T)$ indicate the maximum values of a for which each threshold, T, may be permitted if the proportion of activated spurious cells is to be kept low ($P_{\text{spur}} \leqslant 10^{-4}$). The thresholds used for successive steps of progressive recall are shown, together with the final value of $a\,(=a_\infty)$ and the value of a obtained with simple recall $(=a_s)$. $A = 15$, $a_0 = 2$.

late stages of recall, as well as in the first stage. This allows satisfactory recall in some situations in which the proportion of modified synapses in the network (ρ) is very small. But in situations of greater interest it is often not possible to choose any single value of the threshold which gives both satisfactory recall and an acceptably small number of spurious cells.

This problem is avoided if we are able to start progressive recall with a low threshold and raise it during successive stages of recall. This is illustrated for

a particular case in figure 6. The curves of $\pi(a,T)$ are drawn for the relevant values of T, and on these curves are marked with crosses the maximum permissible values of a, determined by the condition

$$\pi(\rho a,T) < P_{\text{spur}}.$$

In this case ρ is taken as 0.1 and P_{spur} as 10^{-4}. The successive values of a_r are determined graphically in the same way as for figure 5, according to the equation

$$\mathscr{E}(a_{r+1}) = a_0+(A-a_0)\pi(a_r,T_r).$$

It is assumed once again that at each stage the value of a_{r+1} is equal to its expectation. The values of T_r are chosen to be the smallest values for which the

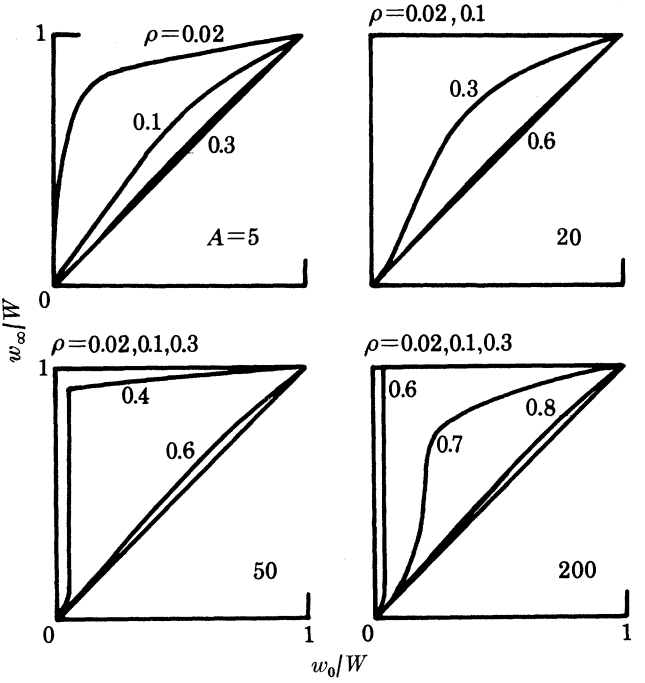

FIGURE 7. Progressive recall. The performance of networks with various values of A, plotted in the same way as figure 3. $P_{\text{spur}} = 10^{-4}$.

maximum acceptable value of a, marked by the cross, is higher than the current value. The initial threshold is 4 and the threshold for the final stages of recall is 9. It should be noted that if the initial threshold had been 9, hardly any recall would have occurred at all, because the first intersection of the curve $y=\pi(a,9)$ with the straight line $y=(a-a_0)/(A-a_0)$ occurs practically at the point $(a=a_0, y=0)$.

An approximation to the value of a_∞ attainable with progressive recall is the value of a for which the dashed curve joining the crosses in figure 6 intersects the straight line $y=(a-a_0)/(A-a_0)$. If we could allow T to take non-integral

values and interpolate between the curves for integral values, then this value of a_∞ would always be attainable, subject to the assumption that at all stages $a_r = \mathscr{E}(a_r)$. Since T must always be an integer in the strictly defined model, the performance will in some cases be slightly less good than that calculated.

We can now plot in exactly the same way as for simple recall (figures 3 and 4) the performance of progressive recall for various values of A and of ρ, assuming

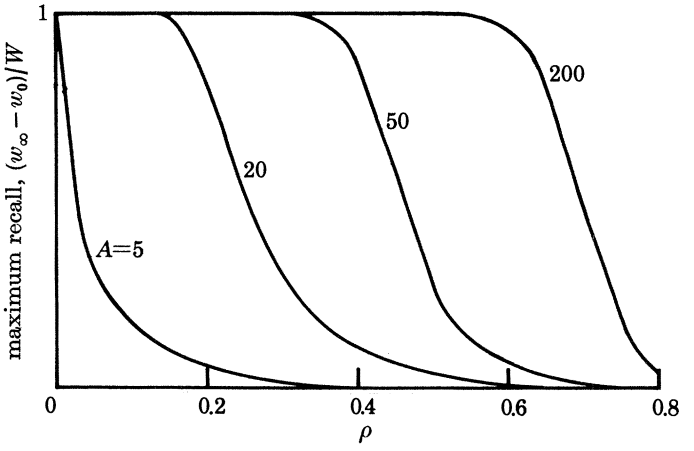

FIGURE 8. Progressive recall. The maximum proportion of cells which can be recalled, plotted as a function of ρ, in the same way as figure 4. $P_{\text{spur}} = 10^{-4}$.

a fixed value (10^{-4}) for P_{spur}. The results are given in figures 7 and 8. Comparison with figures 3 and 4 shows that a substantial improvement on simple recall is achieved. For values of A less than about 10 the improvement is not dramatic; but for values $A \geqslant 20$ nearly 100% progressive recall is possible over a range of values of ρ for which the performance of simple recall deteriorates steadily. Reference to figure 7 shows that efficient progressive recall is possible in many cases with a small proportion of initially active cells which, with simple recall, would produce only a fraction of the required cells.

Summary. Progressive recall is a procedure which achieves better performance than simple recall. It involves reiteration of the procedure for simple recall, adding the newly activated cells at each stage to the set of cells which initiate the next stage. A progressively increasing threshold is required to maintain a low probability of activation of spurious cells.

7. THE NUMBER OF ACTIVE CELLS REQUIRED TO INITIATE RECALL

The analysis of progressive recall has shown that in some cases nearly perfect recall performance may occur with a very small proportion (w_0/W) of the cells initially active. It is obvious however that recall could not take place with $w_0 < 1$, since w_0 can only take integer values. In some situations we can even say that

higher values of w_0 will not be adequate, because they are insufficient to identify unambiguously the event to be recalled out of all the events in the experience of the network. We consider first this general limitation on w_0, which is independent of any assumptions about the mechanism for recall.

The set of cells initiating recall must be sufficiently large that it is unlikely to be a subset of any event in the experience of the network other than the required event. The probability that any arbitrary event of size W contains a particular set of w_0 cells is

$$\frac{\text{no. of subsets of size } w_0 \text{ in } W \text{ cells}}{\text{no. of subsets of size } w_0 \text{ in } N \text{ cells}} = \frac{W!(N-w_0)!}{N!(W-w_0)!}$$

The expectation of the number of incorrect events in the experience of the network which are compatible with the set of w_0 cells is therefore

$$(M-1)\frac{W!(N-w_0)!}{N!(W-w_0)!} \simeq M\left(\frac{W}{N}\right)^{w_0}, \tag{5}$$

making approximations valid for $(M, W, N, W/w_0 \gg 1)$. This quantity must be very much less than one for unambiguous identification. Using equation (1) and substituting the condition that $\rho < 0.5$, we have $MW^2/N^2 < \ln(0.5) \simeq 0.7$, whence, from equation (5):

$$\text{expectation value of the number of experienced incorrect events compatible with the initial set of size } w_0 < 0.7\left(\frac{W}{N}\right)^{w_0-2}$$

Since $W/N \ll 1$ for the network to cope with a large number of events, it follows that $w_0 \geqslant 3$ will normally be enough to ensure that the initial set of cells is capable of identifying unambiguously the event required for recall.

In order to determine accurately the minimum number of cells required to initiate progressive recall, two modifications sometimes have to be made to the kind of analysis given in the last section. The first is a requirement that the initial value of a_0 must be large enough that the expected number of cells recruited in the first stage of recall must be at least one. If the expected number is much less than one then there will be a substantial probability that no cells will be recruited, and that consequently no recall will occur. Normally the expected numbers of cells recruited in subsequent stages will be larger until near the end of recall, and no problems will arise if cells are recruited in the first stage. This limitation corresponds to the first vertical step in the sequence of steps in figure 6 being larger than $1/W$. In practice this makes very little difference to the calculated performance unless W is small.

The second modification is the replacement of the Poisson curves used for the calculation with appropriate binomial distribution curves (e.g. as in figure 6). The appropriate binomial curves are different for different values of R/N, but

are hardly significantly different from the Poisson curves for $R/N \lesssim 0.2$. Where R/N is larger (i.e. where each cell projects to more than a fifth of the rest of the population), the binomial curves differ in having a more pronounced sigmoid shape near the origin. The dashed curve in figure 6 is consequently also more sigmoid, and larger minimum acceptable values of w_0 are predicted.

Summary. Some special factors are considered which sometimes affect the minimum number of cells required to initiate recall.

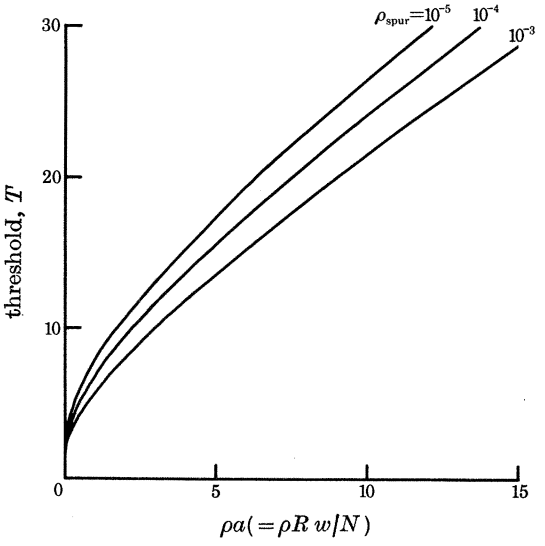

FIGURE 9. Minimum permissible thresholds. The lowest values of threshold compatible with three different limits on the probability of activation of spurious cells ($P_{spur} = 10^{-5} - 10^{-3}$), plotted as a function of ρa. The abscissa is proportional to the fraction of the principal cells in the network which are active at any one time (w/N).

8. THE VARIATION OF THRESHOLD DURING RECALL

For simple recall only a single threshold need be set, while for progressive recall a series of increasing values of threshold are needed. In both cases the relation between the required threshold and the number of cells active in the network (w) is the same, governed by the condition

$$\pi\left(\rho\frac{R}{N}w, T\right) \leqslant P_{spur}.$$

T is plotted as a function of $\rho Rw/N$ for various values of P_{spur} in figure 9. The required threshold rises rapidly at first, and then approximately linearly with the activity in the network.

The appropriate threshold for any given value of w is not fixed for all time, since it depends on ρ, which itself increases with the number of events experienced by the network. It should be noted that the number of spurious cells produced

is very sensitive to the choice of threshold. Under conditions when a threshold $T=25$ is appropriate, a 10% error in threshold will lead to a tenfold change in the expected number of spurious cells. In a physiological model this puts rather stringent requirements on a mechanism for recurrent inhibition, which is in principle capable of controlling threshold (§11).

Summary. The level of inhibition in the network must vary in a rather precise relation to the proportion of cells in the network which are active.

9. The efficiency of the network as a memory device

The number of modifiable synapses in the network is NR. Since each synapse is supposed to have just two states (effective and ineffective), the maximum amount of information which the network is in principle capable of storing is

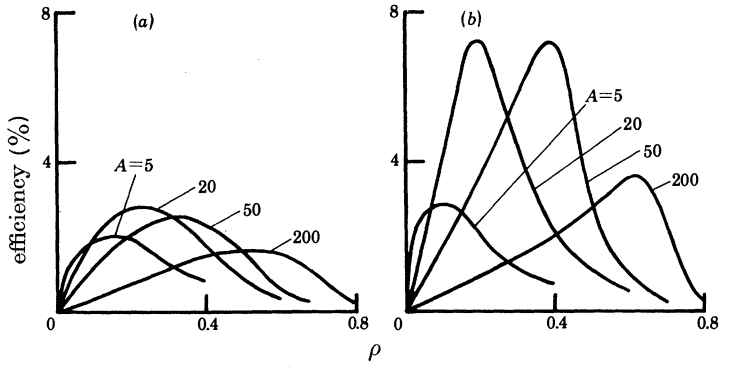

Figure 10. Efficiency (η) obtained by using (a) simple and (b) progressive recall. Values of the overall efficiency (defined in the text) are plotted as functions of ρ, which is approximately proportional to the number of experienced events. Figures are calculated from the data of figures 4 and 8, according to the formula $\eta = \frac{1}{A}\left(\frac{\Delta w}{W}\right)_{\max} \log_2\left(\frac{eN}{W}\right) \ln\left(\frac{1}{1-\rho}\right)$. N/W is taken as 10^2; $P_{\text{spur}} = 10^{-4}$.

NR bits. The network efficiency is defined as the total amount of information which can be recalled about events in the experience of the network, expressed as a fraction of this theoretical maximum.

If an event be fully recalled with the initial activation of only a small proportion of its cells, then the information about the event which is gained from the network is $\log_2 C(W,N) \simeq W \log_2(eN/W)$ bits. If only a fraction of the required cells is recalled, then the information gained is (for small W/N) closely proportional to the number of cells recalled. Thus the total recoverable information is the product of $MW\log_2(eN/W)$ with the fraction of the cells which can be recalled, already plotted in figure 4 for simple recall and figure 8 for progressive recall. The resulting values of efficiency for simple and progressive recall are plotted

in figure 10 as functions of ρ, for $W/N = 10^{-2}$. The greatest efficiency (approximately 7 %) is achieved with progressive recall, for situations in which $A \simeq 20$–50.

These calculations of efficiency yield underestimates in two ways. Firstly, more information about an event may be stored in the network than can be recalled in a single, even optimally chosen, recall trial. Secondly, the acceptable probability of activation of spurious cells (P_{spur}) has been set to 10^{-4}. If $W/N = 10^{-2}$, as in figure 10, this means that the expected number of activated spurious cells is 0.01 W, or 1 % of the number of cells in the event. It is probable that the highest figures for the retrievable information would be achieved for recall in which a larger proportion of spurious cells was permitted. Recall procedures which sacrificed the quality of recall in this way for higher storage efficiency might not be acceptable, however. Even though the results of figure 10 are based on assumptions which are to some extent arbitrary, it can be concluded that storage efficiencies of the order of 5–10 % are compatible with fairly good recall quality.

Summary. The efficiency of the network is derived by calculating the amount of information which can be recalled about the whole set of experienced events. This is expressed as a percentage of the theoretical maximum amount of information which could be stored in any device by using the same number of two-state modifiable elements.

10. General discussion

If an implementation of the recall model with real physiological neurons is to be considered, further consideration of the mechanism is necessary, as dealt with in the next section. But since the strengthening of a synapse in the network may be taken simply as a symbolic analogue of the development of an association between two elements of experience, some general discussion is appropriate at this stage. The analysis of performance has already revealed the conditions under which satisfactory recall of large combinations of elements can result from the development of associations between pairs of elements.

The suggestion that an association can be built up between two elements of experience implies that experience itself can be regarded as a combination of discrete elements. In some ways this is obviously false, since we can experience a whole continuum of different colours, sounds or emotions. But any experience which we can describe in words can be expressed in terms of discrete elements. This is true even if we deprive language of its logical relationships or syntax, since we can describe many events in a form of cryptic poetry in which a group of nouns are arranged in random order (e.g. son uncertainty station homecoming husband war waiting smoke telegrams cold). Introspection suggests that it is much in this form that recall of events takes place. Once recall has occurred we can and normally do organize the elements into logical relationships before expressing them in language. The analysis of this paper is not concerned with

the nature of these logical processes. But it is clear that deductive processes may assist the associative processes in recall, either by generating correct details (it was cold – it must have been winter – we must have been hungry) or by eliminating spurious details (it was the war – my husband died before the war – my recollection of his presence must be spurious).

It is interesting to note that dreams have many of the characteristics of recall in which the constraints of logic have been removed. Perhaps in this way the associations between elements of experience may be maintained strong through constant re-use, in a way which would not occur if recall were always to be censored by reasoning processes and related to the stringent demands of waking life.

The main quantitative conclusions which can be drawn about a recall system working in the way described concern the parameters W/N and R. If information is to be retained about a large number (M) of events, then the proportion of the total population of possible elements of experience which are involved in any one event must be small ($W/N < ca.\ 0.7\ M^{-\frac{1}{2}}$). In order for recall of a substantial proportion of the elements of experience comprising an event to take place, the number of potential associations (R) which can be made with any one element must be such that the parameter A ($= WR/N$) is greater than about 10 and, in order to avoid unnecessary redundancy, not much greater than 100. Combining these two conclusions we arrive at a value of R which is optimally of the order of $10M^{\frac{1}{2}}$–$100M^{\frac{1}{2}}$. Note that this value is independent of the actual size of the recall system, provided that the total number of elements (N) is substantially larger than the required value of R. Thus in a very large system each element needs to be able to make associations with the same number of other elements, and consequently with a smaller proportion of the total.

In order to illustrate these conclusions we can consider somewhat arbitrarily a system able to learn and recall a number of events equal to the number of days in a man's life ($M = 2.5 \times 10^4$). This requires $W/N \lesssim 0.005$ and R in the range 2×10^3–2×10^4. If individual elements of experience can be at all directly related to the activity of neurons we might relate these figures to the fraction of cells which are active in the brain at any one time and to the connectivity of cells in the brain. It is not very easy however to get figures for the proportion of cells which exceed any particular firing frequency at any one time; and indeed the appropriate firing frequency which should constitute a criterion for an 'active' cell must depend on the unknown temporal characteristics of the kinds of modifiable mechanisms responsible for learning. But we can note that the required active proportion of cells (0.005) is low. In a sensory system this active proportion may be quite high at the receptor level, approaching 1.0; but we know that mechanisms exist for reducing it at successive stages in sensory processing (e.g. lateral inhibition, trigger feature detection), and such mechanisms have been proposed elsewhere in the brain (e.g. the cerebellar granule cells: Marr 1969). The average anatomical connectivity of cells in the cerebral cortex

has been studied. From the work of Cragg (1967) the number of synaptic terminals per cell exceeds 10^4 in many regions of the cortex, though the nature of many of these connections is not known.

It is hard to take such quantitative considerations any farther. The associations between elements in human experience are by no means random and equiprobable in the way assumed for the analysis in this paper. Similar events may occur repeatedly; and different events may arise through a process of gradual transformation instead of discrete steps. But the mechanism described is probably capable of dealing with tasks of an order of difficulty which is comparable to human recall, without making any grossly implausible physiological assumptions.

11. The physiological mechanisms for recall

In considering an implementation of the recall model with real physiological neurons we have to deal with two problems. The major problem is the mechanism for controlling threshold, while related to this are the time courses of the firing patterns, the excitation and the inhibition.

In both kinds of recall there is a small number of cells active initially and a larger number later on. The threshold must be set initially to a value low enough to initiate recall and it must be raised subsequently. For progressive recall the threshold needs to be raised in an accurately controlled way at each stage of recall, according to the relation summarized in figure 9. For simple recall the threshold needs to be raised after the single stage of recall sufficiently that the newly recruited cells will not result in the activation of spurious cells. In practice the number of cells recruited during recall (of either sort) will not be easily predictable, since it will depend on chance factors and on whether or not the initial set of cells belongs to an event which has been learned. Thus if the threshold is to be related to the number of recruited cells, as in progressive recall, a feedback mechanism will be necessary in which the inhibition is itself controlled by the number of active cells. In simple recall feedback is unnecessary, since any sufficiently large increase of threshold will be adequate to prevent the activation of spurious cells.

Simple recall can occur if the threshold is brought to a low level and then rapidly raised in a predetermined way. The firing of the recruited cells will only occur in a brief burst, however, since it will be terminated by the rise in threshold. If such firing is to be useful it is likely to have to be repeated again and again. The variations in threshold would thus need to be oscillatory, with the troughs of low threshold corresponding to the values required for recall. A satisfactory procedure for simple recall would be to lower gradually the mean level of the oscillating threshold until the required number of recruited cells were firing in repetitive bursts. In order that the recall obtained should be simple recall, and to avoid the necessity of raising the threshold of successive troughs, it is necessary that the frequency of the threshold modulations should be low enough that

the e.p.s.ps produced by recruited cells have subsided between cycles. Otherwise a build-up of excitation will occur, amounting to progressive recall.

The cells initiating recall need to be activated from outside the network. It was assumed in the analysis that extrinsic excitation would overcome any inhibition in the cells. This could be the case if the extrinsic synapses are either powerful or close to the zone of the cell where action potentials are initiated. Some degree of modulation of the firing of the extrinsically activated cells is consistent with recall, provided that activation occurs sufficiently early before the troughs of threshold that the induced excitation in the cells to be recruited will coincide with the troughs.

We have already seen that raising the frequency at which the threshold is modulated provides one way of implementing progressive recall, through the build-up of excitation from recruited cells. If this is done, then connections providing recurrent inhibition are needed to raise the mean level of the oscillating threshold as cells are recruited, thus preventing the activation of spurious cells. For progressive recall it may not be necessary, however, to have an oscillating threshold. If the threshold is slowly lowered until recall is initiated, and if recurrent inhibition then serves to raise the threshold the correct amount as cells are recruited, then satisfactory progressive recall would occur. In either of these mechanisms for progressive recall, too large a rise of threshold could lead to inhibition of the cells already recruited, and the recall could temporarily subside to levels which were less good even than simple recall. We have already seen (figure 9) that the increase of spurious cells which results from too small a rise in threshold may be very severe. Thus in practice the control of threshold may be very critical. A physiological implementation of the recall mechanisms may face the alternatives of a relatively unstable but potentially efficient progressive recall mechanism and a stable but less efficient simple recall mechanism.

12. Discussion in relation to structural and functional features of the cerebral cortex

Many of the mechanisms which would be required for a direct physiological implementation of recall are known to exist in the mammalian brain. In discussing these features an attempt will be made to draw attention to the information which is lacking as much as to that which is present.

The principal anatomical features of the recall network amount to pathways for recurrent excitation and inhibition. Of these two mechanisms, recurrent inhibition is much better understood within the nervous system. Within structures of the cerebral cortex recurrent inhibition has been demonstrated in many sites including onto the dentate granule cells and pyramidal cells of the hippocampal formation (Andersen, Eccles & Løyning 1964), onto superficial and deep pyramidal cells of the prepyriform cortex (Biedenbach & Stevens 1969; Haberly 1973) and onto Betz cells in the motor cortex (Phillips 1959). In at least the first two

examples there is evidence that the inhibition may act through the relatively small number of basket or stellate cells in these regions, with widely distributed axon terminals ending on the cell bodies or the proximal dendrites of the principal cells. Such pathways could produce the widespread uniform inhibition assumed in the model. It is largely unclear however what kind of extrinsic control may be exerted on the inhibitory systems, and there is no quantitative analysis which could serve for comparison with the graph in figure 9, giving the relation of the required strength of inhibition to the activity in the network.

Recurrent excitation in which axon collaterals terminate directly on other cells of their own population is not so well established. For a discussion of much of the evidence of this and recurrent inhibition see Shepherd (1975). Fairly clear evidence for a direct recurrent excitatory pathway exists for the region of CA3 pyramidal cells in the hippocampus. These cells project to each other via the longitudinal association path of Lorente de Nó (1934). This pathway runs in both directions perpendicular to the lamellar organization of many other pathways of the hippocampus and terminates in different dendritic regions of the pyramidal cells in the two directions (Hjorth-Simonsen 1973). It has not been well studied physiologically, but it is likely to be excitatory since the Schäffer projections from the CA3 to CA1 pyramidal cells are known to be excitatory (Andersen, Blackstad & Lømo 1966). There is evidence for recurrent excitation also in the prepyriform cortex (Haberly & Shepherd 1973) and in the motor cortex (Phillips 1959; Takahashi, Kubota & Uno 1967). Pyramidal cells in the neocortex have sometimes very widespread fields of termination of their axon collaterals (Scheibel & Scheibel 1970), though the cells receiving these terminals and the extents to which they account for recurrent excitation and inhibition are not clear. The Scheibels (1970) state that the rôle of these collaterals is largely inhibitory. It is perhaps worth noting that if the pyramidal cells of the neocortex were connected in the manner required of the principal cells in this paper, then this is the result which would be expected despite the widespread excitatory connections. For the effect of electrical stimulation (e.g. by using antidromic activation) would presumably be to activate a set of cells which was randomly related to the combinations of cells which had been active in the experience of the network. Thus all the remaining cells would have the status of potentially 'spurious' cells and would be subject to inhibition strong enough (as determined in figure 9) that they would have a very small likelihood of being discharged. Obvious ways to approach this problem would be to settle anatomically whether the pyramidal cell collaterals terminate upon other pyramidal cells, and to attempt physiologically to incorporate the pattern of cells activated by an electrical stimulus into the experience of the network.

In this paper the intrinsic excitatory synapses within the network are modifiable under the conditions postulated by Hebb (1949). Many types of hypothetical synapses could provide the basis for learning in neural networks (Brindley 1967); and with suitable assumptions about the temporal patterns of firing and modifiability

all types of pre- or post-synaptic modifications could be adequate (Gardner-Medwin 1969). Nevertheless the Hebb modification conditions (strengthening when there is a combination of pre- and post-synaptic activation or depolarization) permit economical and simple solutions for many learning problems where other synapses would require extra circuitry or many more cells (Gardner-Medwin 1969; Brindley 1969). In the present model a Hebb synapse does in a direct physiological sense what corresponds to the development of an association between two elements of an event which occur together. It thus provides a particularly direct physiological implementation of the model. Hebb synapses have not been shown to exist anywhere. A number of the relatively long-term modifications of the cerebral cortex which have been produced by physiological stimulation could have Hebb synapses as the underlying mechanism (see, for example, Bindman, Lippold & Redfearn 1964; Bliss, Burns & Uttley 1968; Bliss & Lømo 1973; Bliss & Gardner-Medwin 1973; Douglas & Goddard 1975; Schwartzkroin & Webster 1975; Horne & Potter 1971); but in none of them has the mechanism yet been clearly established.

In the procedure for controlling threshold during recall it was noted that oscillatory variations of threshold provided a useful trick for controllably activating some of the cells required for recall. Though this may not be a necessary procedure it arises naturally out of the theory. It is therefore interesting to note that mechanisms exist in the thalamus for providing oscillatory variations of threshold (8–12 Hz) in the neocortex during spindle activity and probably alpha rhythm (for review see Andersen & Andersson 1968) and in the septum and hippocampus for providing somewhat slower oscillatory variations (3–8 Hz) in the limbic system during theta rhythm (see, for example, Petsch, Gogolak & van Zwieten 1965).

13. Discussion in relation to behavioural learning and amnesia

The recall performance discussed in this paper is of a type which has not been the subject of much direct behavioural study. During recall performance a large amount of information is produced as output in response to a relatively small amount fed in as input. We know from introspection and from the use of language that such recall occurs in humans. But in animal studies an experimenter has only a very limited information channel for studying output from the brain, provided by a sequence of overt actions. It is as hard to decide whether an animal can recall events as it is to decide whether a man can recall events if he refuses to tell us about them. Nevertheless, even the most taciturn individual uses recall of events in planning his actions and solving problems, and so may animals.

Animal behavioural studies thus usually involve situations with a larger information input to the animal than output from it. This is commonly but not

always true of human experiments too. Among studies of this type, recognition performance may have most in common with the recall of events. Both types of performance require the storage of a large amount of information about the particular combination of details which constituted an event or stimulus pattern. In a recognition test the question is posed: 'Is this particular combination of details familiar?', while in the situation for recall of an event the question is: 'What were the other details which went with these few?'. These tasks are distinct from reinforcement learning situations in which an animal has to classify a stimulus combination, for example as being associated with food or not; such situations require an animal to have available an appropriate classificatory procedure and then to store one bit of information about the stimulus (e.g. that a coloured corridor is good or bad for finding food). Gaffan (1972, 1974, 1975) has shown that animals are capable of recognition performance independently of their capability for reward association. He has also shown that animals with fornix lesions (transecting some of the inputs and outputs of the hippocampus) are deficient in recognition tasks though not in reward association. He is led to the conclusion that both in animals with fornix lesions and in human amnesics there is primarily a deficit in the discrimination of the familiarity of stimulus combinations, and that this discrimination is the function of the hippocampus.

This paper does not describe recognition performance. But it is easy to see how the network might be used for discriminating the familiarity of an event. The activation of a set of cells constituting a familiar event would be self perpetuating through the strengthened synapses between active cells in a way in which an unfamiliar event would not be. While testing the familiarity of an event the synapses within the network must not be permitted to strengthen: a requirement which could be met either by turning off the modifiability in some neural or chemical way, or by ensuring that the testing took only such a short time that minimal strengthening occurred. Thus recognition performance would require some extra tricks in handling the network; but it is clear that the correct sort of information is stored to make it possible. Since the same sort of information is required for recognition and for recall it seems unlikely that the brain would duplicate the storage for the two purposes.

A natural corollary of Gaffan's conclusion that recognition performance is lost with lesions of the hippocampal formation is that the recall of events may be lost too. In animals this loss of recall performance would be hard to assess because of the lack of a means of assessing recall performance in normal animals. In human amnesics with damage to the temporal lobes or hippocampal formation the evidence suggests that both recognition and recall of events are impaired. Gaffan (1975) argues differently, suggesting that amnesics may not be impaired on recall tasks unless the task is such that familiarity discrimination can aid performance. This is often the case if the required recall is for a recent event, when familiarity discrimination can provide a means of rejecting the recall of earlier wrong events. He cites the experiments of Warrington & Weiskrantz

(1970) which show that amnesics can perform as well as normal subjects if they are given partial information during testing, in the form of fragmented images or the initial letters of words. Such partial information would, in Gaffan's view, replace the familiarity discrimination as a means of eliminating spurious recall. A more straightforward interpretation of these results could be simply that in amnesics the first attempt at recall is more likely to be spurious, either because amnesics are poor at forgetting or suppressing old memories as Warrington & Weiskrantz themselves suggest, or because a recall mechanism such as is described in the present paper is functioning poorly and is not excluding spurious details. Thus it seems reasonable to maintain that amnesics really are poor at recall as well as at recognition. This conclusion is certainly supported by a commonsense interpretation of case histories in which amnesics may be straightforwardly unable to recount details of events without a tendency to produce confabulated or confused accounts (see, for example, Milner 1966).

We thus arrive at a fairly simple conclusion that the recognition of familiar events and the recall of events may use the same kind of stored information. The deficiencies of human amnesics and the effects of lesions of the fornix in experimental animals can be explained by damage to this information store or to the mechanisms for handling it. The arguments of Gaffan (1972, 1975) are important here, since he shows that a deficit in recognition performance can account for at least some of the earlier known characteristics of animals with hippocampal or fornix lesions, such as perseveration and extinction deficits. These conclusions thus support the suggestion made by Marr (1971) that the hippocampal formation may store information about the combinations of details in recent events and stimulus patterns. Marr's detailed theory proposes how such storage may take place in the hippocampal formation as well as how a suitable representation of sensory information may be encoded for optimal storage conditions and how stored information may ultimately be transferred to the neocortex (Marr 1970, 1971). The analysis of the present paper when applied directly to the hippocampal pyramidal cells is a treatment of what he calls the 'collateral effect' (Marr 1971).

The importance of the hippocampus in recall and recognition is strongly suggested by both the human and animal evidence. But human amnesics with severe damage to this system are capable of recalling events experienced before the time of the injury (Milner 1966). The hippocampus is therefore not the only structure involved in recall. The analysis in this paper provides a model of the logical processes which may be involved quite generally in the brain's recall performance. The suggested direct neuronal implementations may exist in the brain. But if not, the same logical principles could nevertheless apply to the systems which do exist.

References

Andersen, P. & Andersson, S. A. 1968 *Physiological basis of the alpha rhythm.* New York: Appleton-Century-Crofts.

Andersen, P., Blackstad, T. W. & Lømo, T. 1966 Location and identification of excitatory synapses on hippocampal pyramidal cells. *Expl Brain Res.* **1**, 236–248.

Andersen, P., Eccles, J. C. & Løyning, Y. 1964 Location of postsynaptic inhibitory synapses on hippocampal pyramidal cells. *J. Neurophysiol.* **27**, 592–607.

Biedenbach, M. A. & Stevens, C. F. 1969 Synaptic organisation of cat olfactory cortex as revealed by intracellular recording. *J. Neurophysiol.* **32**, 204–214.

Bindman, L. J., Lippold, O. C. J. & Redfearn, J. W. T. 1964 The action of brief polarizing currents on the cerebral cortex of the rat (1) during current flow and (2) in the production of long-lasting after-affects *J. Physiol., Lond.* **172**, 369–382.

Bliss, T. V. P., Burns, B. D. & Uttley, A. M. 1968 Factors affecting the conductivity of pathways in the cerebral cortex. *J. Physiol., Lond.* **195**, 339–367.

Bliss, T. V. P. & Gardner-Medwin, A. R. 1973 Long lasting potentiation of synaptic transmission in the dentate area of the unanaesthetised rabbit following stimulation of the perforant path. *J. Physiol., Lond.* **232**, 357–374.

Bliss, T. V. P. & Lømo, T. 1973 Long lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J. Physiol., Lond.* **232**, 331–356.

Brindley, G. S. 1967 The classification of modifiable synapses and their use in models for conditioning. *Proc. R. Soc. Lond.* B **168**, 361–376.

Brindley, G. S. 1969 Nerve net models of plausible size that perform many simple learning tasks. *Proc. R. Soc. Lond.* B **174**, 173–191.

Cragg, B. G. 1967 The density of synapses and neurones in the motor and visual area of the cerebral cortex. *J. Anat.* **101**, 639–654.

Douglas, R. M. & Goddard, G. V. 1975 Long-term potentiation of the perforant path granule cell synapse in the rat hippocampus. *Brain Res.* **86**, 205–215.

Gaffan, D. 1972 Loss of recognition memory in rats with lesions of the fornix. *Neuropsychologia* **10**, 327–341.

Gaffan, D. 1974 Recognition impaired and association intact in the memory of monkeys after transection of the fornix. *J. comp. physiol. Psychol.* **86**, 1100–1109.

Gaffan, D. 1975 Recognition memory in animals. In *Recognition and recall* (ed. Brown, J.), London: Wiley.

Gardner-Medwin, A. R. 1969 Modifiable synapses necessary for learning. *Nature, Lond.* **223**, 916–919.

Haberly, L. B. 1973 Unitary analysis of opossum prepyriform cortex. *J. Neurophysiol.* **36**, 762–787.

Haberly, L. B. & Shepherd, G. M. 1973 Current density analysis of summed potentials in opossum prepyriform cortex. *J. Neurophysiol.* **36**, 789–802.

Hebb, D. O. 1949 *The organization of behaviour.* New York: Wiley.

Hjorth-Simonsen, A. 1973 Some intrinsic connections of the hippocampus in the rat: an experimental analysis. *J. comp. Neurol.* **147**, 145–162.

Horne, M. & Porter, R. 1971 The effects of prolonged cortical stimulation on the natural activity of neurones in the precentral gyrus of conscious monkeys. *Brain Res.* **34**, 115–126.

Lorente de Nó, R. 1934 Studies on the structure of the cerebral cortex. II. Continuation of the study of the Ammonic system. *J. Psychol. Neurol.* **46**, 113–177.

Marr, D. 1969 A theory of cerebellar cortex. *J. Physiol., Lond.* **202**, 437–470.

Marr, D. 1970 A theory for cerebral neocortex. *Proc. R. Soc. Lond.* B **176**, 161–234.

Marr, D. 1971 Simple memory: a theory for archicortex. *Phil. Trans. R. Soc. Lond.* B **262**, 24–81.

Milner, B. 1966 Amnesia following operation on the temporal lobes. In *Amnesia* (ed. C. W. M. Whitty & O. Zangwill). London: Butterworths.

Petsche, H., Gogolak, G. & Van Zwieten, P. A. 1965 Rhythmicity of septal cell discharges at various levels of reticular excitation. *Electroenceph. Clin. Neurophysiol.* **19**, 25–33.

Phillips, C. G. 1959 Actions of antidromic pyramidal volleys on single Betz cells in the cat. *Qu. J. Exptl Physiol.* **44**, 1–25.

Scheibel, M. E. & Scheibel, A. B. 1970 Elementary processes in selected thalamic and cortical subsystems – the structural substrates. In *The neurosciences. Second Study Program* (ed. F. O. Schmitt). New York: Rockefeller.

Schwartzkroin, P. A. & Webster, K. 1975 Long-lasting facilitation of a synaptic potential following tetanization of the in vitro hippocampal slice. *Brain Res.* **89**, 107–119.

Shepherd, G. M. 1975 *The synaptic organization of the brain: an introduction.* New York: Oxford University Press.

Takahashi, K., Kubota, K. & Uno, M. 1967 Recurrent facilitation in cat pyramidal tract cells. *J. Neurophysiol.* **30**, 22–34.

Warrington, E. K. & Weiskrantz, L. 1970 Amnesic syndrome: consolidation or retrieval? *Nature, Lond.* **228**, 628–630.