A TEMPORAL SEQUENCE LEARNING MODEL[1]

James C. Stanley
Computer and Information Science

COINS Technical Report 75C-10

University of Massachusetts
Amherst, MA 01002 U.S.A.

(December 20, 1975)

ABSTRACT

A neural model is presented relating to possible mechanisms of habituation in the dentate gyrus of the mammalian hippocampus. The model is designed such that once it has been sufficiently exposed to a sequence of regularly repeated inputs, it can regenerate that sequence upon being cued with an initial fragment of the sequence. Propagating waves of excitation generated by each input form a dynamic representation of the input, coding both the set of cells stimulated by the input and the time elapsed since those cells were excited. Results of simulations of the dynamic wave generating mechanism and the full learning model are presented.

---

[1]This paper was presented at the 1975 Winter Computer Simulation Conference, December 18-19, 1975, Sacramento, California. Research supported in part by NIH Grant No. 5R01 NS09755-04 COM awarded to M. A. Arbib, Dept. of Computer and Information Science, Univ. of Mass., Amherst, MA.

A TEMPORAL SEQUENCE LEARNING MODEL

James C. Stanley
University of Massachusetts, Amherst, Massachusetts

The ability to adapt to changes in the environment is one of the most interesting aspects of animal behavior. Neuronal plasticities of various sorts allow organisms to reshape and reorder their behavioral repertoires in order to survive in a changing world. Learning is a form of plasticity that allows an animal to appreciate new relations between its behavior and things and events in its world. The phenomenon of learning has been widely studied at both the behavioral and the neurophysiological levels. Many mathematical and computer models have been constructed in order to help understand the neurophysiological basis of learning (1, 2, 3, 4, 5). These models show how simple training rules enable circuits composed of neuron-like elements to learn and remember. While many of these learning models are based on generalized circuit anatomies, others are founded upon known structures of particular brain regions. One such model, designed by Kilmer (1), is based on the anatomy and physiology of a part of the mammalian hippocampus.

The hippocampus is a region of cerebral cortex that has been implicated in a variety of functions, many of which require that the hippocampus learn to respond in particular ways to its inputs (6, 7). The inputs to the hippocampus consist of pre-processed sensory information and motivational information, and its outputs go to motor and motivational regions of the brain. The hippocampus is made up of a number of subregions. One of the main outputs of the hippocampus stems from region CA3, which was modeled by Kilmer (1). In Kilmer's model, circuit elements representing cells that give rise to this output path are trained to respond to particular input patterns. This model shows how positive and negative feedback can be used to shape circuit response. In order to extend this work, we have considered a circuit model of another region of the hippocampal complex called the dentate gyrus. The dentate gyrus also receives sensory and motivational information, and its outputs go to the CA3 region. While the dentate gyrus may serve a variety of functions in the overall hippocampal system, we have modeled the dentate gyrus as an habituation circuit, following a suggestion of McLardy (8).

Habituation is a form of plasticity that allows an animal to ignore those things in its world that are of no immediate consequence. Habituation is defined as a reversible decrement of response to repeated stimulation, and is displayed by organisms from mollusc to man. The reversibility of habituation sets it apart from fatigue and accommodation. An habituated stimulus may once again evoke a response following presentation of a novel stimulus. This phenomenon is known as dishabituation, and rules out sensory and motor fatigue as the bases of the response decrement.

Habituation, like learning, has been studied at both the behavioral and physiological levels. Several circuit models have been designed that realize many of the detailed properties of habituation (9, 10, 11). In these models, each measurable attribute of a stimulus is coded in terms of the firing of a cell or a group of cells. Upon repeated presentation of the stimulus, the outputs of these cells are depressed by a buildup of inhibition or by a decrease in synaptic efficacy or cell sensitivity. The known anatomy and physiology of the dentate gyrus may be related to these circuit models in a straightforward way to investigate ways the dentate gyrus might display many of the properties of habituation.

One property of habituation has not been satisfactorily dealt with in terms of the detailed workings of a neuronal system, however. It is known that, following habituation to a regularly repeated stimulus, changes in stimulus duration or interstimulus interval can bring about dishabituation (12). There must be some mechanism, then, that codes the temporal qualities of the stimulus, just as the other qualities are coded. A number of temporal memory models exist that could be applied to the problem of temporal coding (9, 13, 14, 15), but none can be readily related to the structure of the dentate gyrus. Accordingly, we have designed a temporal sequence memory based on the anatomy of the dentate gyrus. The hippocampal system could use such a memory to anticipate the arrival of a regularly repeated, short duration input, as described below.

TEMPORAL MEMORY OPERATION AND THE HIPPOCAMPUS

The temporal sequence memory presented here is so designed that once it has been sufficiently exposed to a sequence of inputs seperated by particular time intervals, it can reproduce the sequence with proper timing if cued by an initial portion of the sequence. The system thus learns by rote the interstimulus intervals involved in a given repetitive sequence. Predictions generated by the memory may be used with match-mismatch circuits of the sort considered by Horn (9) as a part of a habituation system.

The outputs of the temporal match-mismatch circuits, together with the outputs of filters sensitive to other stimulus qualities, may be used as the inputs to circuits of the sort referred to above that embody the processes of habituation and dishabituation and that realize the detailed properties of habituation (see Fig. 13). If the memory is able to anticipate an input, then, the temporal qualities of that input can be habituated to, along with its other qualities. Any subsequent variation in the temporal characteristics of the habituated input will give rise to a mismatch signal, causing dishabituation of the overall system response. Inputs that do not repeat in a regular pattern, or that repeat against changing background conditions, give rise to continual mismatches, and eventually cause the habituation circuitry to come to ignore the variable properties of the stimulus in a process of generalization.

The structure of our temporal memory network reflects the structure of the dentate gyrus. We refer the reader to Kilmer (7) for a discussion of the anatomy of the hippocampus. Here we stress the fact that functionally, the hippocampus is organized into transverse slices called lamellae (16). Most fibers that enter or originate within the hippocampus remain largely within such slices, displaying little longitudinal spread. Sensory inputs excite granule and pyramidal cells, the main circuit elements of the dentate gyrus and CA3, respectively. Granule cells give rise to fibers that excite pyramidal cells in CA3. The portion of CA3 that lies closest to the dentate gyrus is called the endfolial or hilar region, and pyramidal cells there will be denoted endfolial cells. These cells give rise to fibers that spread to contact granule cells in other lamellae (17). This anatomy is represented in the structure of the model, as described below.

In simulating this system, we have expressed the operation of a region of the brain in terms of the functioning of a large network of simple elements acting in parallel. We have simulated the network itself on two levels. On the level of dynamic circuit action, each cell is represented by a set of nonlinear differential equations. Cells are coupled so as to generate waves of activity. Simulation shows how the nonlinearities may be used to tailor the waves to allow the overall system to function properly. On the second level, finite-state automata are used to represent groups of cells. These automata are coupled so as to generate waves, as before. On this level, however, the interactions between waves can be more easily studied. Simulation of the system on this level shows how such a dynamic memory can be controlled and what mechanisms are necessary for its proper operation.

We assume here, then, that a granule cell in the dentate gyrus that is excited by a system input spreads excitation

to neighbors in its lamella, causing a wave of activity to travel down the lamella. Details of the cell coupling and dynamic properties used to produce waves are discussed in section three, following presentation of the model's overall structure in section two. Results of simulation of the wave-producing dynamic system are also presented in section three. Waves generated by a given input are associated with granule cells excited by the following input to effect the memorization process, as described in section four. In section five we describe mechanisms added to the basic model to surmount problems of interference and instability. Finally, in sections six and seven we present a simulation of the full model, and discuss the results of simulation.

MODEL SYSTEM STRUCTURE

Like the hippocampus, the model is made up of a number of arrays of cells called lamellae. Each model lamella consists of four lines of neuromimes as indicated in Fig. 1. G-neuromimes receive the system inputs and their

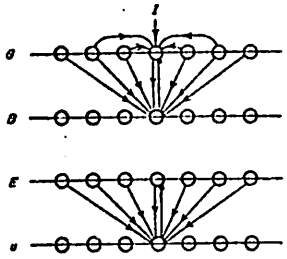

Fig. 1. Patterns of interconnections between G and B neuromimes and E and U interneuromimes within a lamella. Arrows indicate excitatory influence, dots indicate inhibitory influence.

outputs constitute the outputs of the lamellae. The G-neuromimes are coupled to one another and to B-neuromimes and E-neuromimes in their lamellae. B-neuromimes are excited by the G's and in turn inhibit them. Patterns of connection and connection strengths between G-neuromimes and B-neuromimes are fixed and identical for all G's and B's. Details of these connections are discussed in the section on wave generation. Coupling between E's and U's is similar to that between G's and B's, with the exception that E's are not coupled to one another. These connections are also fixed and identical for all E- and U-neuromimes.

Connections between G's and E's are assigned with some randomness in the model. Each E is connected to a fixed number of G's in its lamella. These neuromimes are chosen at random from a range of G's centered around the E, as indicated in Fig. 2(a). The strengths of these connections are fixed and are identical for all connections.

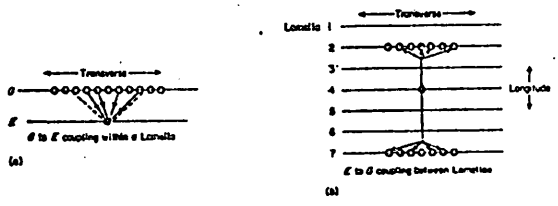

Fig. 2. Patterns of interconnection between G and E neuromimes. (a) Connections from G to E within a lamella. Dashed lines indicate range of possible contact. (b) Connections from E to G in other lamellae.

Each E in turn gives rise to two lines, called Z-lines, that run perpendicular to the lamellae, making contact with G's as indicated in Fig. 2(b). The length of each Z-line is chosen according to a specified probability distribution, F_Z. G's in each lamella the Z-line passes through are chosen for contact with fixed probability to a maximum range on either side of the Z-line. These longitudinal lines are the only connections between lamellae, and have the only variable strengths used in the model, as discussed in section four.

GENERATION AND PROPAGATION OF WAVES IN THE MODEL

The wave-supporting substrate of the model has a configuration suggested by hippocampal physiology. Pyramidal cells in CA3 and granule cells in the dentate gyrus receive excitatory influences from the hippocampal input pathways (16). These cells in turn excite a variety of interneurons whose axons play back onto other pyramidal and granule cells. In both regions, a type of interneuron called a basket cell is assumed to have a profound inhibitory effect on pyramidal and granule cells (18). Basket cells have widely branching axonal arborizations, so excitation of a basket cell by one pyramid or granule cell will depress the activity of others nearby. In CA3, it has been shown that pyramidal cells also excite neighboring pyramids (19). We assume here that granule-to-granule excitation also exists in the dentate gyrus, leading us to the circuit configuration for G's and B's shown in Fig. 1.

The generation of single impulses in nerve cells is a highly nonlinear process involving a buildup of excitatory potentials in dendritic membranes and the subsequent triggering of a propagating action potential (20). However, the rate of firing of a single cell and the average firing level of a homogeneous population of cells may reasonably be modeled in terms of simple dynamic equations (3, 21, 22). Accordingly, each of our model cells obeys a first-order differential equation, the value of which may represent the firing rate of a single neuron, or the number of active cells in a population of neurons. We refer to the model cells below in terms of output pulse rates of single cells.

We assume that the output pulse rate of a cell is given by the difference between the total excitatory influence on the cell and some function of the total inhibitory influence on the cell. An initial set of equations for the operation of the G and B cells under these assumptions is:

$$T_G \frac{dG_j}{dt} = -G_j + \sum_m \lceil G_m^0 - \theta_m \rceil^+ \cdot w_m + I_j$$

$$T_B \frac{dB_j}{dt} = -B_j + \sum_k \lceil G_k^0 - \Gamma_k \rceil^+ \cdot v_k$$

$$G_j^0 = M\lceil G_j - B_j \rfloor_m$$

where G_j represents the total excitatory influence on the G-cell,

B_j represents the B-cell output,

θ_m and Γ_k are connection thresholds,

w_m and v_k are connection weights,

I_j is the external input to the G-cell,

G_j^0 is the output pulse rate of the G-cell,

$$M\lceil x \rfloor_m = \begin{cases} M ; & x > M \\ x ; & m < x < M \\ m ; & x < m , \end{cases}$$

for $m < 0 < M$, and

$\lceil x \rceil^+$ is x if $x > 0$ and is zero otherwise.

Note that the output rate of the G-cell is constrained to lie between a maximum value greater than zero and a minimum less than zero. The maximum reflects the fact that nerve cells cannot fire more rapidly than some maximum rate. The minimum is set to a value less than zero under the assumption that the zero level in the model represents a nonzero spontaneous firing rate in the actual nerve cells. Then the minimum value in the model corresponds to a firing

rate of zero in the actual cells.

Each G and B cell is driven through weighted threshold connections by the output rates of neighboring cells. That is, neither is influenced by any neighbor to which it is connected until that neighbor begins to fire above a specified rate. The influence of the neighbor is weighted by a constant associated with the connection. Note that neighbors firing below threshold rates and at rates below the spontaneous level do not lead to depression of the firing of the cell. Hence cells that are firing below the spontaneous rate are effectively uncoupled from the system. We assume then that information processing and maintenance of a background spontaneous firing rate (the system equilibrium state) are two separable factors in our system dynamics.

A simulation of this system of equations has been carried out using an integration routine written in FORTRAN. The propagating waves illustrated in Fig. 4 were generated with the connection template shown in Fig. 1 and the associated weight and threshold profiles of Fig. 3.

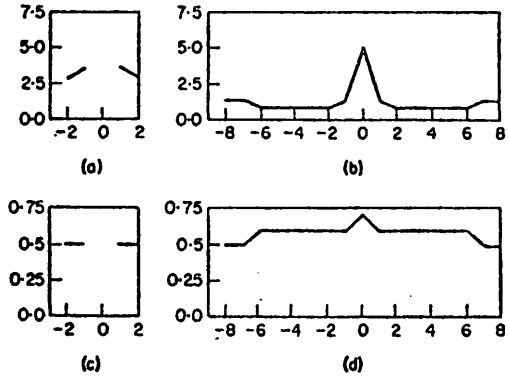

Fig. 3. Initial weight and threshold profiles. Ordinates indicate number of neighbor relative to cell. (a) G to G weights. (b) G to B weights. (c) G to G thresholds. (d) G to B thresholds.

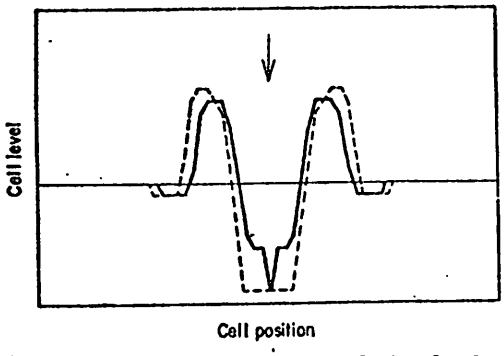

Fig. 4. Wave produced with profiles of Fig. 3. Dashed lines indicate progression of wave.

All cell values were set initially to zero, representing an undisturbed system displaying spontaneous firing. Following external stimulation at the position indicated by the arrow in Fig. 4, each sufficiently excited G-neuromime stimulated neighboring G's according to the G-G weight and threshold profiles and inhibited a wider range of G's according to the G-B profiles. The indicated weight and threshold settings allow the excitation of each G to build to a high enough level to excite further G's before being countered by a heavy and long-lasting inhibition arising from the central weight and threshold of the G-B profile. This central connection corresponds to a strong self-inhibition activated when the cell fires enough.

A wave may travel to the ends of the lamella and die there, or may die before reaching the ends. In the first case, the wave propagates with a constant shape until the end of the lamella is reached. At that point, there are no further G cells to excite and the wave dies. Because of the connection thresholds, the zero level of activity is stable, so no further activity arises in the lamella until a new input is presented. In the second case, inhibition builds that eventually stops the wave. Here, the wave's amplitude decays as it travels until further propagation is impossible. The distance traveled depends on the strength of inhibition and on the relative B time constant. The full memory model to be described below uses waves that die after going a short distance rather than those that propagate unchanged.

The strong self-inhibition employed above results in a circuit that for long periods after passage of a wave is unable to support another wave. Too long a period of depression decreases the chance that new inputs can enter the lamella, as will be described in section four, so means of shortening the depression time were sought. Two ways to shorten the depression time are to decrease the level of inhibition needed to produce waves and to decrease the inhibition time constant. In order to decrease inhibition levels, an amplitude-dependent nonlinearity is used to allow low levels of inhibition to have a greater effect on the circuit. The amplitude-dependent gain appears in the computation of the output rate as follows:

$$G_j^0 = {}^{M}\lceil G_j - (P + W \cdot F(B_j)) \cdot B_j \rfloor_m$$

where G_j^0, G_j, and B_j are as before,

P is a bias term set close to unity,

W is a weight set greater than one, and

F is the logistic function given by

$$F(x) = \frac{1}{1 + c\,e^{-d(x-x_0)}}$$

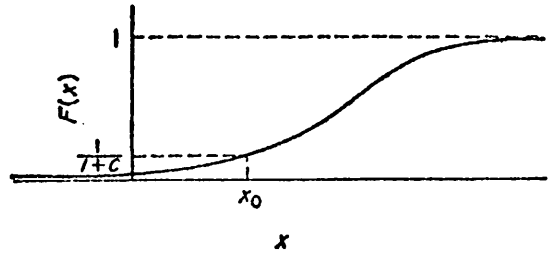

Fig. 5. Logistic function.

and illustrated in Fig. 5.

For small values of inhibition, G's output rate is approximately the difference between excitation and inhibition as before. As the level of inhibition rises, however, its weighting in the output rate calculation increases from unity to 1 + W. In this way low levels of inhibition may have a great effect on the system. The lower values of inhibition decay to zero sooner, leaving the system depressed for a shorter period of time.

A smaller inhibition time constant may be used if different weight and threshold profiles are employed, corresponding to a different way of producing waves.

With the modified weight profiles, self-inhibition begins to build when the G-cell reaches lower levels of firing, due to the smaller self-inhibitory connection threshold, but builds more slowly than before due to the decreased weight involved. In the previous scheme, waves were produced by allowing cells to build to large firing rates and then depressing them with strong inhibition. In order to produce waves, this inhibition had to last long enough for the excitatory levels to decay below the thresholds necessary for coupling to other cells. Shorter-lasting or weaker inhibition allowed the cells to remain for long periods of time at a level just above the self-inhibition threshold. In

order to produce the wave profiles of Fig. 4, in which each cell fires for a short time and then is silenced, the inhibition time constant had to be great enough to keep the cell depressed until the wave passed out of its neighborhood and its excitatory level decayed considerably. In the new scheme, however, inhibition is applied more slowly, and instead of depressing the cell immediately, pulls the output rate down gradually. Hence inhibitory levels need not remain large for too long, and the inhibition time constant may be shortened. Figure 6 shows waves generated by this

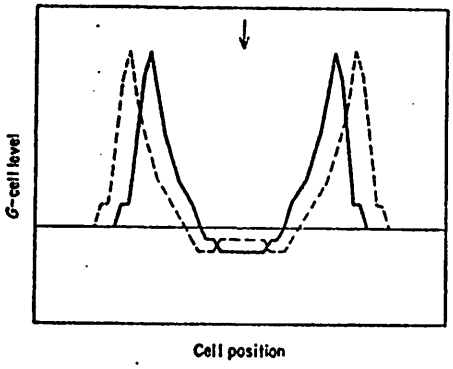

Fig. 6. Waves produced with modified profiles.

scheme. These waves leave the circuit depressed for shorter periods of time, and so are better suited for use in the overall model. Again, these waves may travel to the ends of the lamellae, or may die due to inhibitory buildup before reaching the ends.

The equations employed here are similar to those used by Ratliff (23) to investigate the dynamics of lateral inhibition, and by Grossberg (13) to explore learning and memory in a number of circuit anatomies. In neither case were these equations employed to produce traveling waves of activity. The wave generating mechanism in this one-dimensional system is similar to the ones studied by Beurle (24) and by Wilson and Cowan (22). In each case, the wave peak arises and propagates due to positive feedback between excitatory elements. In Beurle's studies this feedback is controlled by the refractory properties of the model cells, while in Wilson and Cowan's model, as in the one presented here, inhibitory buildup stops runaway positive feedback. As stated above, the activity level of a given cell in our model may represent the average firing rate of a nerve cell or the number of cells in a coupled population that are firing at a given moment. The nonlinear inhibition used here may represent a nonlinearity in the way single cells react to inhibition, or the action of a population of inhibitory cells with a distribution of thresholds.

MEMORY STORAGE AND RECALL IN THE MODEL.

The model is designed to learn and recall both the sets of cells stimulated by system inputs and the time intervals between inputs. It performs this recall by a process of association chaining, in which each input or evoked memory of an input stimulates (after the proper time interval) the set of G's normally stimulated by the next input of the sequence. This set of G's then stimulates the set associated with the next input, continuing the process of recall.

This process is indicated in Fig. 7 for the case in which the second input immediately follows the first. The model must form associations such that future presentation of the first input will cause immediate stimulation of the second input's G's. This association is handled via the longitudinal lines arising from the E's. The strength of each longitudinal connection is modeled as a first-order differential equation as

$$T_z \frac{dZ_{ij}^{km}}{dt} = -Z_{ij}^{km} + \lceil (G_{km}^0(t) - \theta_G) \rceil^+ \cdot \lceil (E_{ij}^0(t) - \theta_E) \rceil^+$$

where Z_{ij}^{km} connects the j^{th} E-cell in the i^{th} lamella

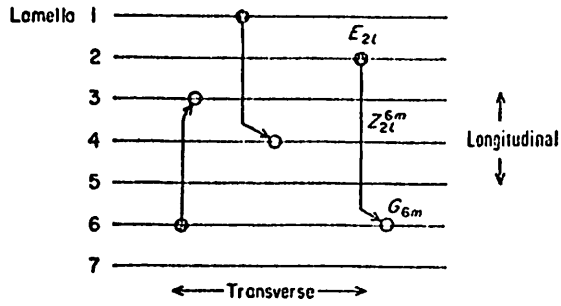

Fig. 7. Simple association of inputs. Open circles represent G neuromimes excited by present input; hatched circles represent E neuromimes excited by last input's Gs.

with the m^{th} G-cell in the k^{th} lamella,

T_z is the connection time constant,

G_{km}^0 is the firing rate of the G-cell,

E_{ij}^0 is the firing rate of the E-cell, and

$\lceil x \rceil^+$ is equal to x if $x > 0$, and is zero otherwise.

Connections are modified according to the Hebb modification rule (5). A connection is strengthened if the E and G it links are simultaneously active above the thresholds θ_G and θ_E, while an unused connection decays to zero at a rate determined by T_Z. Strengths of the longitudinal lines are initially zero, with some exceptions discussed below.

The thresholds in the modification rule allow the system to discriminate inputs that are to be learned from the waves and from the effects of noise. A system input is assumed to excite a G-cell to a firing level much greater than the maximum reached during passage of a wave. The value of θ_G is set just below this input excitation level to prevent associations from forming between G's excited by waves. Similarly, θ_E is set just below the level of excitation reached by an E-cell when a wave passes by it, so that inputs are associated only with a few E's recently excited by waves.

Thus if an E excited by one of the first input's G's gives rise to a longitudinal line that makes contact with one of the G's excited by the second input, that connection is strengthened according to the learning rule above. Later presentation of the first input will again excite that E, in turn tending to activate the second input's G via the strengthened connection. If the system is properly constructed, as discussed below, each of the second input's G's will with high probability be connected in this way to at least one of the first input's G's. Repeated presentation of the sequence will result in connections strong enough to excite the G's to the input level of excitation. Presentation of the first input will then cause immediate activation of the second input's G's, effecting recall of the sequence. This chaining process may then continue with the association between the second and third inputs, and so on.

If the next input does not occur immediately, a trace of the first input must be held in the system long enough to form associations as described above, and to code the interval between inputs. Input storage and interval coding are accomplished through the wave action of the G's. Each overall input excites a few G's throughout the system, giving rise to waves moving in the associated lamellae. These waves excite E's as they go, so that an input's G's may be associated with waves generated by the previous input, effecting both storage and interval coding.

This process is illustrated in Fig. 8. A wave generated by the first input at time T_0 in lamella 4 moves along the lamella, as indicated by the cross-hatched rectangle. Before

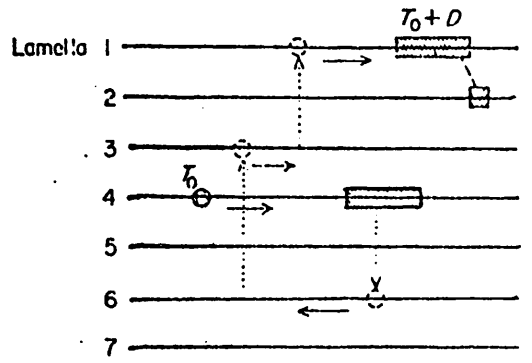

Fig. 8. Association of inputs and waves.

dieing, the wave will move a distance determined by the inhibitory parameters. In order to handle time intervals longer than those obtainable with a wave moving along a single lamella, some of the longitudinal connections are permanently strengthened initially so that a wave activating such a connection may generate a wave in another lamella. The second wave may continue after the first has died, and may generate more waves. A series of such wave generations is indicated by the dotted lines in Fig. 3 while arrows show the directions of wave motion along lamellae. At time $T_0 + D$ a wave generated in this way is at the position shown in lamella 1 when a new input excites a G-neuromime in lamella 2, as indicated by the cross-hatched square. If a longitudinal connection exists between an E excited by the wave at that time and the G excited by the new input, that connection is strengthened according to the system's learning rule. Such a connection is indicated by the dashed line in Fig. 8.

Now if at some later time the first input is again presented to the system, a wave will be generated at the same place in lamella 4. If this wave causes the same sequence of waves to be generated as before, after an interval of length D, the expected time between the two inputs, a wave will activate the learned longitudinal connection, exciting one of the second input's G's. If each of the second input's G's is associated in this way with at least one wave stemming from the first input, presentation of the first input will cause the activation of the second input's G's after about D seconds. This process may then continue, effecting recall of the sequence.

SYSTEM DESIGN CONSTRAINTS

The association system described above can fail in five main ways. The first set of problems concerns the representation of inputs using patterns of moving waves. If the waves all die, or the pattern falls into a cycle, information is lost and proper association becomes impossible. Interference between the wave pattern and the inputs to the network is the second problem. Cells that lie in the troughs of waves are inhibited, so inputs cannot excite them enough for associations to form. A third difficulty concerns the probability of association. If the system is not properly designed, the probability that a wave can become associated with a given input will be too small for reliable operation. The fourth class of problems concerns interference between wave patterns generated by successive inputs. Finally, the fifth problem arises from the effects of changes in the system's structure brought about by learning. Each of these problems is treated below. In some cases, proper setting of system parameters can alleviate the difficulty. Other problems require that further mechanisms be added to the model

Pattern Persistance and Cycling

The first problem arises because all the waves generated by a given input may die out before the next input arrives. This will happen if there are too few opportunities for waves to generate others in other lamellae. Simulations indicate that waves generally persist indefinitely in systems with more than a few prestrengthened connections from each lamella to the others. Second, a given set of waves may fall into a cyclic pattern of movement through the network, thus losing the ability to represent the exact interval between inputs. Simulation again shows that in systems with more than a few strong connections between lamellae, wave interactions are complex enough to preclude this possibility. These difficulties, then, will seldom arise in systems with enough prestrenthened connections between lamellae.

Trough Interference and Association

The second set of problems concerns the interference between a wave pattern and an input, and the probability that all of an input's G's become associated with waves. These two problems are linked through their dependence on the number of waves in the system. First, the input may stimulate a cell that lies in the trough of a wave. Such a cell is inhibited and cannot reach a firing level that allows it to become associated with other waves or to generate a wave itself. That part of the input, then, cannot enter the system and is unavailable for association with other inputs. Second, a stimulated cell may not become associated with any wave in the system. This will occur if there are too few waves or if there are too few longitudinal connections available for association. Formulas dealing with these two sources of failure are derived in the appendix. There, Table 4 shows the result of applying the formulas to a system with parameters given in Table 3, in which waves and stimulated cells are distributed at random. It may be seen that full input entry and full association in such a system are possible only when the numbers of waves and stimulated cells are kept quite low with respect to the total number of cells.

The number of waves present in the system depends on the number of cells initially stimulated and on the number of prestrengthened longitudinal connections between lamellae. With too many such connections, too many waves are spawned and input entry becomes a problem. With too few connections, however, the wave pattern may die or fall into a cycle, as discussed previously. Simulation shows that systems with enough prestrengthened connections to avoid cycling and pattern extinction generate far too many waves to ensure that none of an input's G's is inhibited. Further, the number of waves exhibits oscillations, as shown in Fig. 9 , making the system at times able to accept inputs and unable at other times. Accordingly, means were sought to

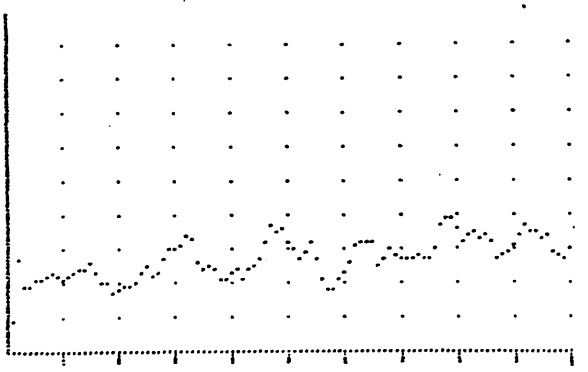

Fig. 9. Plot of number of waves produced in uncontrolled system as a function of time.

hold the number of waves at a constant low level, while still avoiding pattern extinction and cycling.

The overall number of waves in the network can be controlled through regulation of wave death. Recall that waves are assumed to travel a certain distance (called the propagation distance), then die due to inhibitory buildup. A feedback mechanism was added to control the number of waves by changing this distance. The propagation distance

may be the same at all points in the system, corresponding to an overall setting of inhibitory parameters, or may vary locally, corresponding to local parameter setting. In either case, when the number of waves in the system is small, the propagation distance is made large. This insures that waves will travel far enough to generate other waves before dieing, thus keeping the overall wave pattern from dieing. As the number of waves grows, the propagation distance is decreased to keep waves from generating too many others. In the case of global propagation distance variation, the propagation distance at each point is adjusted according to the total number of waves in the system. With local variation, the distance is set at each point according to the number of waves present in an area about that point.

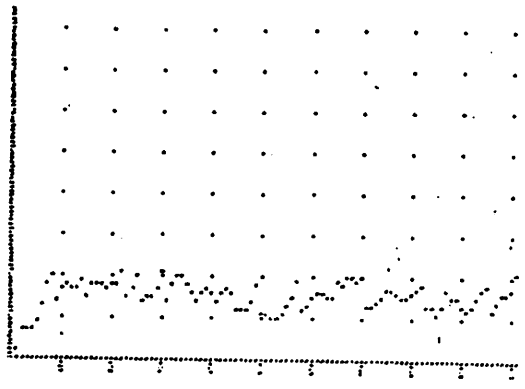

Fig. 10. Plot of number of waves produced in controlled system as a function of time.

Fig. 10 illustrates the result of applying global propagation distance variation to a system with the same parameters as those used in Fig. 9. The number of waves is considerably reduced and is more constant than before. Similar results are obtained using local variation. Local variation has the added advantage that waves are able to spread more evenly through the network, leading to increased local probabilities of association. This method of propagation distance variation thus allows the number of waves to be controlled to satisfy the input entry and association constraints and to keep the wave pattern from dieing entirely or falling into a short cycle.

Inter-sequence Interference

Another cause of difficulty is the interference between wave patterns generated by different inputs or by successive presentations of the same input. As an example, consider the sequence formed by regular repetition of a single input, and suppose that there are already waves in the system when this input is first presented. These waves will interact with those generated by the new input, so that the pattern that exists when the input is presented the second time may not be the same as the initial wave pattern. Because of this difference, the third presentation of the input is likely to become associated with still another pattern, and so on. This difficulty stems from the instability of the wave patterns, which is such that two patterns that differ only slightly give rise over time to patterns that are more and more different. Because of this instability, it is possible that a repeated input will not generate the same wave pattern twice over very many repetitions, and hence that no strong connections will form in a short time. The network would thus have stored a large number of memories of this input and would take a prohibitively long time to train. Even if a short sequence of patterns were formed, if the input was presented with a different initial pattern during recall, the learned sequence of patterns might not arise again, leading to a recall failure.

The first mechanism added to combat this interference problem is one that inhibits all waves when a peak of maximum activity is detected anywhere in the system. This corresponds to a strong blanket inhibition applied when an input is presented or when a strong association is recalled. It's assumed that the strong peaks of activity are able to survive this inhibition and to generate new waves of activity. With the wave-killing mechanism in operation, interference effects are limited to two forms. First are the effects of inhibitory troughs left behind by waves killed when the input was presented. Second is the effect of lingering longitudinal connection inhibition.

Wave troughs give rise to the input entry problem, as discussed earlier. Troughs left after waves are killed also block the propagation of waves generated by new peaks of activity. The effects of these forms of interference may be minimized by use of the proper numbers of waves. Longitudinal connection inhibition left after the wave pattern is killed can affect the development of the new pattern by blocking wave generation. As above, the instability of the wave patterns is such that further patterns generated without the waves whose generation is blocked will differ from patterns generated with those waves. Because of the differences, associations will not readily form and recall may be impossible. A way to obviate this difficulty is to assume that E's have a maximal inhibitory level, and to set all E's to this level when the wave killing mechanism is activated. This corresponds to a saturating inhibitory input applied at that time to all E's. This input effectively eradicates all traces of the E-inhibition due to the previous pattern, allowing the new pattern to develop with interference due only to the wave troughs. The two mechanisms of wave killing and E inhibition input and saturation can be used to minimize the effects of interference between wave patterns due to successive inputs.

Interference Effects of Learning

The final difficulty lies in the effects of changes in the system's structure brought about by learning. Because waves travel throughout the system between input presentations, a longitudinal line that has been strengthened through learning is likely to be activated at times when it should not be. Activation of a strong enough connection at the wrong time will cause a wave to arise in the system at a point in the development of the wave pattern where none has been before. Such extra waves will alter the development of the pattern, making proper association and recall impossible.

A way to ensure that longitudinal lines strengthened through learning are activated only at the proper times is to allow E activation only when a particular pattern of waves exists in the network. In this scheme, an E that gives rise to no strengthened connections, that is, one that has never taken part in association or that has forgotten any associations ever made, may be activated by a single wave that goes by its position, as before. When an E takes part in association, however, that E learns the pattern of waves in an area around its position in the network. Thereafter, the E can be activated only when the wave pattern in the network matches its stored pattern. If the stored pattern is composed of enough waves, it is unlikely to occur at any but the correct time.

This mechanism requires a change in the way the E's are activated. Each E must have strong connections from G's in its lamella, and weak connections from G's in other lamellae. All of these connections must be trainable, such that when a longitudinal line stemming from the E is strengthened, any connection to the E from a G active at or above the wave propagation level is also made strong. At the same time, the E's threshold of activation must increase, so that the entire set of strengthened connections must be excited to activate the E. In this way, a particular wave pattern is stored in the connections to the E from G's in its own and neighboring lamellae.

These mechanisms added to the basic model, then, minimize the effects of the problems arising from fluctuations in the number of waves in the network and from the instability of the wave patterns. The number of waves must be kept small enough to allow inputs to enter the system

freely, but large enough to allow proper association and to keep the wave pattern from dieing or cycling. The number of waves is controlled through feedback that decreases the wave propagation distance as the number of waves increases. Pattern instability gives rise to problems of interference between patterns generated by successive inputs and to disturbances due to learning itself. Interference problems are minimized using a wave inhibition mechanism and a strong longitudinal line inhibition, both of which operate when an input is presented to the network. Pattern disruption due to improper activation of longitudinal lines strengthened through learning is prevented by allowing longitudinal lines to be activated only when specific patterns of waves are present in the system. These mechanisms are related to changes in the operation of the basic model, and may be interpreted in terms of the action of particular excitatory and inhibitory mechanisms in the hippocampal system. The means of simulating this system, and results of simulation, are discussed below.

SYSTEM SIMULATION

The differential equation formulation of the model, described in section two, has the disadvantages of being difficult to control and expensive to simulate. It is difficult to know in advance what detailed characteristics the wave generating mechanism must have so that the overall system can function properly. Much work would be needed to find ways to realize a given set of desired characteristics in turns of the dynamic equations. Furthermore, a large system of such equations can be simulated only at considerable expense. In order to simplify the simulation and to make the system easier to control, simulation of the full model was carried out using cells modeled as pairs of finite-state automata, rather than as differential equations.

Each simulated cell consists of a wave-generation portion, representing the action of the G and B cells, and a longitudinal line activation portion that represents the E and U cells. The operation of each cell's wave-generation automaton reflects the generation of wave activity in the differential equation form of the model. These automata are much more easily controlled than the differential equations, however, and may easily be designed to have desired wave propagation characteristics. In this formulation, the entire wave peak is represented by a single cell in a specified state. Each cell therefore represents a number of cells in the differential equation formulation, making this simulation more economical. Similarly, all of the E's activated by a wave peak at a given position in the system are now modeled in the single longitudinal-activation automaton associated with the cell at that position. The structures of these automata are described below.

The Wave-Generating Automaton

The wave-generating automaton of each cell is shown in Fig. 11. Each state is numbered, and larger numbers represent higher levels of firing in the differential equation formulation. State Q is the resting state, in which a cell remains if undisturbed and to which it returns following perturbation. The highest state, I, represents the level of activation reached when a system input excites a quiescent cell. States W and T are the wave propagation and maximal trough states, respectively, representing the level reached as a wave peak passes the cell and the level of maximal inhibition following the peak. States between the trough state and the quiescent state are trough recovery states, through which a cell goes to quiescent following passage of a wave. The states between the quiescent state and the wave state are states of subactivation to which a cell is sent by small inputs. Finally, the states between the wave state and the state of maximal activation are superactivation states representing levels great enough to generate waves but smaller than the state to which a cell is sent by a system input.

Each cell has three sources of input. These inputs are wave-activation inputs from neighboring cells in the same lamella, system inputs, and longitudinal-activation

Fig. 11. Structure of wave-generating automaton. -+-+-+ indicates action with no input, – – – action of wave input from neighbors, ——— action of input applied to non-quiescent cell. Multiple arrows indicate that state entered depends on size of input.

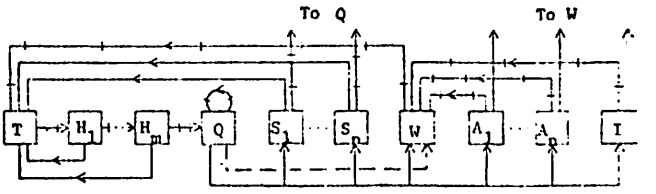

inputs from cells in other lamellae. If a cell is quiescent, that is, if the cell's wave-generating automaton is in state Q, a system input to the cell moves it to state I. From I, the cell goes in one time-step to cell W, the wave state. From state W the cell goes to the maximal trough state, T, at the next time-step, regardless of further inputs. This corresponds to a strong inhibition triggered by the cell's recent activity. If the cell receives no further inputs it moves one state at a time from T to the quiescent state. Any input to the cell while it is in one of the trough states causes it to go back to T again, representing the effect of further inhibition.

In order to propagate waves, each cell receives inputs from the two nearest neighbors in its lamella. A quiescent cell enters the wave state at a given time when one of its neighbors was in the wave state or greater at the time before. The cell next enters the lowest trough state, as above. In this way, a wave moves through a quiescent region represented by a cell in the wave state followed by cells in the trough states. The single cell in the wave state represents the entire wave peak in the differential equation form of the model. A cell in any trough state is sent to the lowest trough state at the next time-step if one of its neighbors is in the wave state or greater. Thus when two waves collide they annihilate one another, leaving behind only an inhibited region that recovers to quiescence. This action reflects the behavior of colliding waves in the dynamic formulation of the model.

The third source of inputs to a cell is activity communicated over the longitudinal lines from cells in other lamellae. The activity level of a given longitudinal line is an integer calculated in a manner described below. The activity levels of all longitudinal lines to a given cell are summed at each time-step, together with the external input to the cell, to arrive at the total extrinsic input to the cell at that time. If a cell is in any state other than the quiescent state or a preactivation state when it receives a nonzero extrinsic input, it is sent to the lowest trough state at the next time step. This corresponds to a strong inhibition activated by the cell's recent activity and reinforced by further inputs to the cell. A small input to a quiescent cell will set the cell to a subactivation state, from which it goes to quiescence if it receives no further inputs. Larger inputs set the cell to the wave state or to a superactivation state. The superactivation states decay to the wave state in one step with no further inputs. An input to a cell in a superactivated state sends it to the lowest trough state at the next time. A strong enough total input can send a quiescent cell to the maximal activation state. In this way, the activation of longitudinal lines can lead to the generation of new waves and strong peaks of activity in the system.

The Longitudinal-Activation Automaton

A longitudinal line is activated as a result of the action of both portions of the cell that gives rise to the line. The structure of the longitudinal-activation automaton , denoted LA, is shown in Fig. 12. The LA

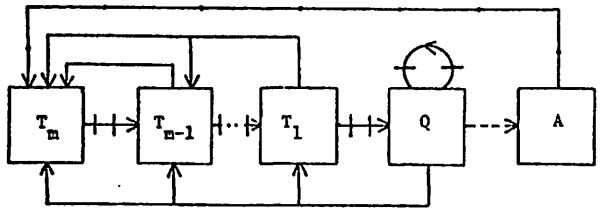

Fig. 12. Structure of the longitudinal-activation automaton. +++ indicates action with no input, - - - - indicates activating input, ——— action of inhibitory input from neighbors, ——— action of self-inhibition. Multiple arrows indicate state entered depends on size of input.

automaton has a quiescent state, Q, an activated state, A, and a set of trough states, T_1. Each LA automaton receives an activating input from the wave-generating automaton in its own cell, and inhibiting inputs from the LA automata associated with neighboring cells in its own lamella. Like the E's in the dynamic form of the model, an LA automaton with no preferred wave pattern can be activated by a wave that passes by its position in the lamella. The LA automaton is inhibited following its own firing or the firing of neighboring LA automata.

To illustrate LA activation, first consider a cell that has no preferred activation pattern, that is, one that has not recently taken part in an association. If the wave-generating portion of that cell is at or above the wave state, and the cell's LA automaton is not in an inhibited state, then the LA automaton is sent to the activated state and the longitudinal lines stemming from the cell are activated. At the same time, the LA automata of neighboring cells receive inhibitory inputs depending on their distances from the cell whose lines became active. Inhibitory inputs are additive, and the greater the total inhibitory input at a given time, the lower the inhibition state to which the LA automaton is sent at the next time. If left unperturbed, the LA automaton recovers one state at a time to quiescence. Note that an inhibited LA automaton can't activate its longitudinal lines or produce inhibition in other LA automata. In this way, a wave of activity alternately activates and inhibits LA automata as it travels along its lamella. This sequence of events is the same for cells with preferred wave patterns, except that activation can occur only if the automaton is quiescent and if the proper wave pattern exists in the network.

The level of activity transmitted from one wave-generating automaton to another through an activated longitudinal line depends on the state of the activating wave-generating automaton and on a weight value associated with the connection between the two automata. In the dynamic form of the model, these weights conform to differential equations. Here, they are modeled as difference equations with values that are increased according to the learning rule presented in section two, and that fall to zero otherwise. Note, however, that the prestrengthened connections that allow a wave in one lamella to generate waves in others have non-zero values that are fixed for the life of the network. The input to a cell in another lamella due to an activated longitudinal connection is simply the connection weight multiplied by the level above quiescence of the activating wave-generating automaton. Recall that all such inputs are summed at the receiving cell to determine its state at the next time step. So a wave in one lamella that activates a strong-enough longitudinal connection can send a cell in another lamella to its highest state, while somewhat weaker connections allow waves to generate other waves in different parts of the system.

This formulation of the model in terms of automata is an abstraction from the dynamic form that is designed to be easier to control and less costly to simulate. Enough of the operation of the differential equations has been embodied in the functioning of the automata, however, that results obtained with this form of the model can also be achieved using the dynamic system form. Results of simulating the automaton form of the full model are described below.

SIMULATION RESULTS

A simulation of the system in the automaton form described above was carried out to investigate the effects of interference and to determine the network's memory capacity. Two networks were created that differed in the details of their structure and in their wave control parameters. In order to ensure that associations were made properly and that recall errors were minimized, two further constraints were imposed on the operation of these networks. First, when an input entered the system, the network was searched to find at least two cells active at or above the wave propagation level at the time before that did not already give rise to strengthened connections. These cells were then connected to the cell excited by the input, and the connections were given non-zero values. At the same time, the wave pattern about each of the cells chosen for connection was searched until more than a fixed minimum number of waves were found. The positions of these waves were then recorded to act as the activating configurations for the cells chosen for association.

The first condition ensured that if there were any free cells activated by waves when an input entered the system, then the input would be associated with at least one wave in the network. This corresponds to a network in which each E gives rise to a large number of longitudinal lines ready for association. The second condition ensured that the activating wave configurations were large enough to be unique in the course of development of a wave pattern. The first condition may be relaxed to allow associations to form only with some probability. Likewise, the second condition may be changed by allowing a fixed-area search for waves to use in forming the activating configurations. Under the relaxed conditions, however, association failure is possible, and activating configurations too small to be unique may be formed. These sources of error were eliminated here in order to investigate memory capacity and the effects of errors due to interference.

Table 1 lists the parameters of the two simulated networks. In each case prestrengthened longitudinal connections were chosen according to an exponential probability distribution. In these simulations, the values of these and all other strengthened longitudinal connections did not decay, so that memory capacity could be evaluated. Connection strengthening parameters were chosen such that four occurrences of strengthening were required to make a connection strong enough to be able to generate a peak of maximal activity. Each system input consisted of two cells chosen from all cells of the network according to a uniform distribution. Input sequences were composed of two inputs repeated at fixed intervals chosen with uniform probability from a range of between thirty and fifty time steps.

Wave propagation distance was set according to the equation below:

$$D = \max[D_{min}, D_{max} - a_1 \cdot \lceil N/a_2 \rceil]$$

where D is wave propagation distance at a given point in the network,

Table 1. Network Parameters
P_{zw} is the probability that a longitudinal connection is selected for prestrengthening. Other parameters as named in appendix and below.

Network Number	N_L	N_C	R	L	D_{max}	D_{min}	a_1	a_2	P_z	P_{zw}
A	9	59	20	4	25	2	12	4	.4	.0045
B	9	59	30	6	25	4	6	2	.4	.0045

D_{min} and D_{max} are the minimum and maximum allowable propagation distances,

N is the number of waves in a rectangular region centered about the point in question, R cells wide and L lamellae deep, as indicated in Table 1.

a_1, a_2 are constants, and

$\lceil m \rceil$ denotes the largest integer less than or equal to m.

With this rule, the wave propagation distance decreased in steps of size a_1 for each increase in N of size a_2. In network A, wave control was established on the basis of the number of waves in a region about each cell that is smaller than the region used in network B. In B, the distance was decreased gradually as the number of waves increased, while in A the decrease was more abrupt due to the larger values of a_1 and a_2. Wave control was better in A than in B, leading to a more nearly constant number of waves in A. This difference in the wave control effects is reflected in the networks storage and recall success rates, as described below.

Both networks were trained as follows. At the start of each training session, four inputs were chosen to present to the network. Two of the inputs were to form the sequence to be learned and two were to act as "temporal context" during learning and recall. The first context input was presented to the network, and the wave pattern due to that input was allowed to develop. After a time, the first input of the sequence to be learned was presented, followed at the appropriate time by the second input, and then again by the first. This sequence was repeated several times to allow strong associations to form. To test the network's ability to recall the sequence, all waves were then erased from the system and the second context input was presented. As before, the pattern due to this input was allowed to develop for a time and then the newly-learned sequence was presented twice. If the network had learned the new sequence properly, and if interference problems did not occur, the network would continue to regenerate the sequence following cueing. After a complete training session, successful or not, the network was saved to be used in further training. In this way, the effects of storing many sequences in a single network could be evaluated.

Network training was stopped when errors occurred in three successive sequences. These errors could be failures of association or interference from previously stored sequences. After the networks were trained to capacity according to this criterion, each was tested for recall of its stored sequences in order to test for interference between stored sequences. In this testing, the context input was presented as in learning, and the sequence to be recalled was presented twice. As before, successful recall meant that the network would continue to regenerate the sequence after cueing was stopped. As in training, recall could fail due to lack of proper associations or to interference effects.

Network Number	No. Jump Conn.	No. Assoc. Conn.	No. Seq. Stored	No. Interference Failures	No. Assoc. Failures	Success Rate
A	110	262	23	5	2	70%
B	130	235	24	5	4	62%

Table 2. Results of Simulation.

Results of these simulations are shown in Table 2. Approximately the same number of sequences were stored in each of the two networks, and each formed about the same number of strong associations and wave-jump connections. In most cases, recall was established after one or two presentations of the inputs, that is, after about one presentation of the entire sequence. In each network, failures occurred due to context interference (interference from the pattern of waves generated by the previous input), stemming in each case from a context failure during learning. Similarly, association failures during training gave rise to recall failures.

Interference between stored sequences also occurred in each network. In network A, three sequences composed of similar inputs interfered with each other, leading to a recall composed of parts of all three when any was used as a cue. Two other traces in network A interfered with one another. Two pairs of sequences in network B interfered with one another. In one case, a sequence disrupted the recall of another, but was not itself disturbed. In the second case, a sequence that was not fully stored due to an association failure disrupted recall of another sequence.

In no case did the wave pattern cycle between input presentations. The wave pattern died entirely in only one case. The cause of death was blockage of the new input's waves by the troughs left behind by the previous pattern.

Counting all forms of failure, network A had a success rate of 70%, while B's success rate was 62%. Success rate is measured as the percentage, among all sequences presented to the network before the final three failures, in which there was proper storage and recall. Network A was more successful than B because wave control was better in A. Failures in B often occurred because there were too few or too many waves in the system. The control in A was such that there was more often the proper number of waves to allow association without blocking the inputs.

If failures due to interference from stored sequences are not included, the success rates are 85% and 68%, respectively. Mechanisms may be added that diminish the effects of stored-trace interference by actively erasing the older sequence's connections when interference occurs. Such mechanisms would then increase the success rate of the network, at the expense of the loss of older memories.

Association failures occur when an input enters the system at a time when all waves are at positions that already give rise to strengthened longitudinal connections. This source of error can be minimized by changes in system architecture or by allowing the network to store only up to about half its expected capacity. Passive forgetting due to decay of connection weights and spontaneous loss of stored activation patterns may be employed to erase old memories, keeping the network always at about half capacity. Again, the chance of error is minimized at the expense of the loss of old memories.

DISCUSSION

A temporal sequence memory model based on the structure of the dentate gyrus region of the mammalian hippocampus

was proposed in section two. This memory uses nonlinear propagating waves to generate a representation of its inputs to effect memorization of a sequence of inputs together with the time intervals between them. Mechanisms were added to this memory model to overcome sources of interference and instability, and the full system was simulated to test the effects of errors and to determine the memory's capacity. Although several types of errors occur in the operation of the network, success rates of 70% can be achieved using proper control of the number of waves in the network. Elimination of interference through active memory erasure yields an 85% success rate.

These error rates are based on only one network, but are supported as approximate error rate figures by similar results from the other network. Extensive simulation of the wave generating mechanism shows that most networks with a given set of parameters are similar in their wave generation properties. Since the sources of error are intimately linked with the wave pattern dynamics, most networks having proper control of the numbers of waves will show 60% to 90% success rate. Success rates can be increased through the use of mechanisms that erase old or conflicting memories to decrease the chances of failures due to lack of associations and to interference from stored sequences.

Success rates will also be changed if the forced association rule used in simulation is altered. Under this rule, two or more associations were forced to form between each input cell and the cells activated by waves at the time the input entered the system. However, each activated cell was allowed to associate with only one of the two input cells. As training continued and memory capacity was used up, the chance decreased that there were enough waves in the system to associate with the inputs according to this rule. The forced association rule could be modified to allow each activated cell to form connections to both input cells, and to make connections with some probability. This modified rule corresponds more closely to the learning rule in the original description of the model. Under the modified rule, the chance of association failure is decreased, since any activated cell that does not already have an association connection can contact either or both of the cells activated by the input. At the same time, however, failure may occur under this rule, since associations are formed with some probability. The balance between these two effects could be explored in further simulations. Note that the modified association rule also extends the memory's effective capacity, since fewer cells are likely to associate with a given input.

The memory network presented here may be used as part of an habituation system as shown in Fig. 13, following

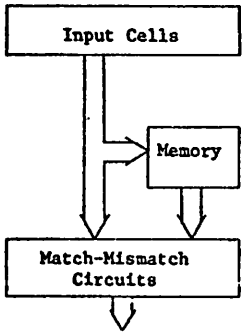

Fig. 13. Use of the model with match-mismatch circuits as part of an habituation system.

a scheme of Horn (9). The outputs of the C's form the network's overall output. The output of the memory network, together with the outputs of the cells that form its input are fed to a rank of match-mismatch circuits. These circuits may be set up initially, or may develop through a learning process, to respond when the signal from the memory cell corresponding to a particular input does not match the signal from the input cell itself. Then if the input sequence is changed after the memory has been trained, a mismatch will be signalled and dishabituation can occur as a result. The system will thus be able to dishabituate due to changes in the inputs themselves or in the times between stimuli. As stated previously, simpler habituation networks have been designed that capture many of the properties of habituation and that may be applied to the problem of habituation in the dentate gyrus or hippocampal complex. The temporal sequence memory described here allows an habituation system based on the structure and physiology of the hippocampus to duplicate known temporal properties of habituation.

This work was supported by NIH Grant No. 5RO1 NS09755-04 COM. We wish to thank Dr. Michael Arbib and Dr. William Kilmer of the University of Massachusetts and Dr. Turner McLardy for helpful discussions and critical reviews of earlier versions of this paper.

APPENDIX

Probabilities of Input Entry and Association

We calculate here the probability that an input stimulates no cell that lies in the trough of a wave, and the probability that a given cell stimulated by an input becomes associated with at least one wave in the system. The following parameters are used:

N_L = number of lamellae,

N_C = number of cells/lamella,

N_H = number of cells stimulated by an input,

N_W = number of waves in the system,

M_W = total number of cells/wave,

M_T = number of cells in a wave trough,

M_R = number of cells in another lamella that a cell in a given lamella is able to form connection to.

$F_Z(i,j)$ = probability that a longitudinal line from lamella i reaches lamella j,

P_Z = probability that a longitudinal line makes contact with a given cell (out of a total of M_R) in a lamella it reaches,

P_E = probability that a wave can activate a given longitudinal line.

To find the probability that an input will stimulate a cell in the trough of a wave, note that each wave trough uses M_T cells, so that the total number of depressed cells is $N_W \cdot M_T$. So

p(a given stimulated cell lies in a wave trough) =

$$P_{ST} = \frac{N_W \cdot M_T}{N_L N_C}$$

If waves and stimulated cells are distributed at random through the system,

p(k stimulated cells fall in wave troughs) =

$$\binom{N_H}{k} P_{ST}^{k} (1 - P_{ST})^{N_H - k}$$

So the probability that all stimulated cells fall outside wave troughs is

$$P_{FE} = (1 - P_{ST})^{N_H}$$

To find the probability that a given stimulated cell is connected with at least one wave in the system, we assume first that an input is equally likely to stimulate a cell in any of the N_L lamellae. Then

p(a stimulated cell is connected to at least one wave when there are N_W waves in the system) =

$$P_H(N_W) = \frac{1}{N_L} \sum_{i=1}^{N_L} P_C(i,N_W),$$

where

$P_C(i,N_W)$ = p(a stimulated cell in lamella i is connected to at least one wave, given N_W in the system).

The waves may be distributed through the network in many ways. Each such distribution will be called a configuration, and configurations will be denoted C_j, $j = 1,\ldots,M_C$, where M_C is the total number of configurations of N_W waves distributed in N_L lamellae. Associated with each C_j is a set of numbers $K_j = (k_{j1},\ldots,k_{jN_L})$ giving the number of waves in each lamella for that configuration.

Hence

$$P_C(i,N_W) = \sum_{j=1}^{M_C} P_C(i,N_W|C_j)P(C_j),$$

where

$P(C_j)$ = probability that configuration C_j occurs.

Now

$P_C(i,N_W|C_j)$ = p(a stimulated cell in lamella i is connected with at least one wave in the system given N_W waves in configuration C_j)
= 1 - p(a stimulated cell in lamella is connected with no wave, given N_W waves in configuration C_j)
= $1 - P_W(i,N_W|C_j)$,

where

$$P_W(i,N_W|C_j) = \prod_{m=1}^{N_L} [1 - P_A(m,i)]^{k_{jm}}$$

and

$P_A(m,i)$ = p(a given wave in lamella m makes contact with a given stimulated cell in lamella i),
= p(the wave is at a cell in lamella m that sends a z-line to the stimulated cell in lamella i) · p(the E-cell associated with the z-line can be activated),
= $\frac{M_R}{N_C} \cdot F_z(m,j) \cdot P_E$.

The first two terms in the final expression give the probability of horizontal and vertical connection from a cell in lamella m to one in lamella j, assuming that each G-cell is connected to exactly one E-cell in its lamella.

These expressions are evaluated for a system with parameters given in Table 3, with the results shown in Table 4.

$N_L = 6$
$N_C = 60$
$M_R = 30$
$M_T = 3$
$P_z = .9$
$P_E = .66$
$F_z(i,j) = 1.0;\ 0<|i-j|<3$
$0.0;$ otherwise

Table 3. Parameter values used in preparing Table 4.

N_S \ N_W	6	7	8	9	10	11	12
1	0 71 / 0 95	0 76 / 0 94	0 79 / 0 93	0 82 / 0 92	0 85 / 0 91	0 87 / 0 91	0 89 / 0 90
2	0 50 / 0 90	0 57 / 0 88	0 63 / 0 87	0 68 / 0 85	0 73 / 0 84	0 76 / 0 83	0 79 / 0 81
3	0·36 / 0 86	0 43 / 0 83	0·50 / 0 81	0 57 / 0 79	0 62 / 0 77	0 67 / 0 75	0 71 / 0 73
4	0 25 / 0 81	0·33 / 0 78	0·40 / 0 76	0·47 / 0 73	0 53 / 0 71	0 59 / 0 68	0 63 / 0 65

Table 4. Probabilities of full input entry and full association using parameter values in Table 3.

REFERENCES

1. Kilmer, W. & Olinski, M., "Model of a plausible learning scheme for CA3-hippocampus", Kybernetik, 16, 133, 1974.

2. Grossberg, S., Classical and instrumental learning by neurological networks, In Progress in Theoretical Biology, 3, Academic Press, New York, 1974, 51-141.

3. Amari, S-I, "Characteristics of random nets of analog neuron-like elements", IEEE Trans. on Syst., Man, Cyb., SMC-2, 5, 643-57, Nov., 1972.

4. Marr, D., "Simple Memory: A theory for archicortex", Proc. Roy. Soc. London, B, 262, 23-81, July, 1971.

5. Hebb, D.O., The Organizaton of Behavior, Wiley, New York, 1949.

6. Izquierdo, I., "The hippocampus and learning", Progress in Neurobiology, 5, part 1, 37-75, 1975.

7. Kilmer, W., "Biology of decisionary and learning mechanisms in mammalian CA-3-hippocampus", Int. J. Man-Machine Studies, 7, 413, 1975.

8. McLardy, T., "Habituation deficit and paucity of dentate granule-cells in some schizophrenic brains", IRCS Int. Res. Comm. System, 73-3, 16-1-1, 1973.

9. Horn, G., "Neuronal machenisms of habituation", Nature (London), 215, 708, 1967.

10. Wickelgren, B., "Habituation in spiral interneurons", J. Neurophys., 30, 1434-38, 1967.

11. Groves, P.M. & Thompson, R.F., "Habituation: a dual-process theory", Psychol. Rev., 77, 419, 1970.

12. Graham, F.K., Habituation and dishabituation of responses innervated by the autonomic nervous system, In Habituation, 1, eds. H.V.S. Peeke and M.J. Hery, Academic Press, New York, 1973, 163-218.

13. Grossberg, S., "Some networks that can learn, remember, and reproduce any number of complicated space-time patterns, I.", Journal of Math. Mech., 19, 53, 1969.

14. Spinelli, D.N., OCCAM: a computer model for a content addressable memory in the central nervous system, In Biology of Memory, eds. K.H. Pribram and E. Broadbent, Academic Press, New York, 1970, 293-306.

15. Fukushima, K., "A model of associative memory in the brain", Kybernetik, 12, 58-63, 1973.

16. Andersen, P., Bliss, T.V.P. & Skrede, K.K., "Lamellar organization of hippocampal excitatory pathways", Exp. Brain Res., 13, 222, 1971.

17. Zimmer, J., "Ipsilateral afferents to the commissural zone of the fascia dentata, demonstrated in decommissurated rats by silver impregnation", J. Comp. Neurol., 142, 393, 1973.

18. Andersen, P., Eccles, J.C. & Löyning, Y., "Pathway of postsynaptic inhibition in the hippocampus", J. Neurophysiol., 27, 608, 1964.

19. Lebovitz, R.M., Dichter, M. & Spencer, W.A., "Recurrent inhibition in the CA3 region of cat hippocampus", Intern. J. Neurosci., 2, 99, 1971.

20. Shepherd, G., The Synaptic Organizaton of the Brain, Oxford University Press, New York, 1974.

21. Stein, R.B., Leung, K.V., Mangeron, D. & Oğuztöreli, M.N., "Improved neuronal models for studying neural networks", Kybernetik, 15, 1, 1974.

22. Cowan, J.D. & Wilson, H.R., "A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue", Kybernetik, 13, 55, 1973.

23. Ratliff, F., Mach Bands, Holden-Day, San Francisco, 1965.

24. Beurle, R.L., "Properties of a mass of cells capable of regenerating impulses", Phil. Trans. Roy. Soc. Lond., B249, 55, 1956.