clinical development include novel opiates, adenosine kinase inhibitors, nicotinic agonists and various high-tech approaches, such as spinal implantation of analgesic-producing chromaffin cells. Since most conceptual advances have occurred in the last decade there is every reason to be optimistic that, after many years of limited progress, analgesic drug development programmes will lead to truly novel, and better, chronic pain therapies.

Selected references

1 **Bennett, G.J. and Xie, Y.K.** (1988) *Pain* 33, 87–107
2 **Kim, S.H. and Chung, J.M.** (1992) *Pain* 50, 355–363
3 **Novakovic, S.D. *et al.*** (1998) *J. Neurosci.* 18, 2174–2187
4 **Gold, M.S. *et al.*** (1996) *Proc. Natl. Acad. Sci. U. S. A.* 93, 108–1112
5 **Black, J.A. *et al.*** (1997) *NeuroReport* 8, 2331–2335
6 **Field, M.J. *et al.*** (1997) *Br. J. Pharmacol.* 121, 1513–1522
7 **Caterina, M.J. *et al.*** (1997) *Nature* 389, 816–824
8 **Szallasi, A. and Blumberg, M.** (1996) *Pain* 68, 195–208
9 **Waldmann, R. *et al.*** (1997) *J. Biol. Chem.* 272, 20975–20978
10 **Mantyh, P.W. *et al.*** (1995) *Science* 268, 1629–1632
11 **Mantyh, P.W. *et al.*** (1997) *Science* 278, 275–279
12 **Cao, Y.Q. *et al.*** (1998) *Nature* 392, 390–394
13 **De Felipe C. *et al.*** (1998) *Nature* 392, 394–397
14 **Rupniak, N.M. *et al.*** (1996) *Pain* 67, 189–195
15 **Walpole, C.S.J. *et al.*** (1998) *Br. J. Pharmacol.* 124, 183–192
16 **Zhang, X. *et al.*** (1998) *Neuroscience* 82, 1225–1242
17 **Su, X., Sengupta, J.N. and Gebhart, G.F.** (1997) *J. Neurophysiol.* 78, 1003–1012
18 **Sorkin, L.S. *et al.*** (1997) *Neuroscience* 81, 255–262
19 **Davidson, E.M., Coggeshall, R.E. and Carlton, S.M.** (1997) *NeuroReport* 8, 941–946

VIEWPOINT

The hippocampus as an associator of discontiguous events

Gene V. Wallenstein, Howard Eichenbaum and Michael E. Hasselmo

The hippocampus has long been thought to be an important cortical region for associative learning and memory. After several decades of experimental and theoretical studies, a picture is emerging slowly of the generic types of learning tasks that this neural structure might be essential for solving. Recently, there have been attempts to unify electrophysiological and behavioral observations from rodents performing spatial learning tasks with data from primates performing various tests of conditional and discrimination learning. Most of these theoretical frameworks have rested primarily on behavioral observations. Complementing these perspectives, we ask the question: given certain physiological constraints at the neuronal and cortical level, what class of learning problems is the hippocampus, in particular, most suited to solve? From a computational point of view, we argue that this structure is involved most critically in learning and memory tasks in which discontiguous items must be associated, in terms of their temporal or spatial positioning, or both.

Trends Neurosci. (1998) 21, 317–323

IT HAS BEEN almost 40 years since the first report of a severe memory impairment in patient HM, following a bilateral temporal lobectomy for treatment of a seizure disorder[1]. This initial observation has been followed by numerous lesion studies in rodents and nonhuman primates, in attempts to replicate the result and to delineate further the anatomical substrates for learning and memory[2–4]. From these data, it has become clear that several medial temporal lobe areas, particularly the hippocampus (Ammon's horn and dentate gyrus) and adjacent parahippocampal regions (for example, entorhinal, perirhinal and parahippocampal cortices), play an important role in certain forms of memory. However, a greater challenge has been to distinguish between the specific types of memory representations that are hippocampal dependent and those that are not and, in addition, to determine the neural mechanisms that support such phenomena.

Interestingly, both amnesia patients and animals with hippocampal system damage exhibit 'time-dependent impairments' in behavioral tasks generally described as associative or relational in nature[5,6]. HM and patients with medial temporal lobe damage show intact short-term working memory, although they exhibit severe deficits in the recall of events that were encountered only a brief period (~1 min or longer) before testing[7] (see Ref. 8 for a review of retrograde impairments). This deficit in intermediate-term memory is also seen in hippocampal-damaged animals performing tests of associative learning, such as time interval discrimination. In this paradigm, subjects must respond to indicate the end of an estimated time interval in the presence of a stimulus cue[9]. Rats with fimbria–fornix lesions (which disconnect the hippocampus from subcortical innervation) and normal rats show similar response rates and accuracy using this standard protocol. However, when the experiment is

Gene V. Wallenstein and Michael E. Hasselmo are at the Dept of Psychology, Program in Neuroscience, Harvard University, Cambridge, MA 02138, USA, and Howard Eichenbaum is at the Laboratory of Cognitive Neurobiology, Dept of Psychology, Boston University, Boston, MA 02215, USA.

 0166 - 2236/98/$19.00 PII: S0166-2236(97)01220-4

constructed such that the time period to be estimated is interrupted briefly (for example, if the stimulus cue is removed), normal rats estimate the total interval by adding together the (estimated) time periods before and after the interruption[9]. In contrast, lesioned animals ignore the estimated time period before the interruption and initiate timing after the interruption is removed[9].

Rawlins[5] has suggested that many experimental tasks that are sensitive to hippocampal dysfunction have an inherent temporal 'discontiguity' in their design. By discontiguity, we mean that the events that must be associated together do not overlap. Another example of this can be found in classic conditioning of the nictitating-membrane response in rabbits. During training, a conditioned stimulus (CS; tone) is presented before an unconditioned stimulus (UCS; air-puff to eye). When the two stimuli overlap in time, animals with hippocampal lesions and normal animals both learn the response equally well[10]. However, when there is a brief delay period (0.5 s) between the two stimuli, hippocampal-lesioned animals show marked impairment in learning compared with normal animals[10]. Similar observations using a classical conditioning procedure with rats have been reported[11]. Taken together, these data suggest that animals with hippocampal impairment have difficulty in associating temporally discontiguous events.

Spatial discontiguity

There is another class of tasks, however, that hippocampal-damaged animals are particularly poor at performing relative to controls, in which a discontiguity in the 'spatial positioning' of related items (rather than being temporal in nature) must be bridged. Previous research has shown that rats can use extramaze cues to guide performance during spatial learning tasks[12,13]. Using a radial arm maze, Suzuki and colleagues found that rotating a set of extramaze cues or the maze itself failed to disrupt task performance[13]. However, transposing the cues, which alters the spatial relationship among stimuli, impaired performance significantly. These results are consistent with the cognitive mapping theory of hippocampal function, which postulates that navigational behavior is guided by learning the spatial relationships among a constellation of stimuli in an environment rather than by any one specific cue[2]. Indeed, previous research has shown that the spatial learning performance of hippocampal-damaged rats depends critically on the placement of extramaze cues[14,15]. For example, when visual stimuli used to guide performance are situated in close proximity to one another, forming a compound cue, both hippocampal-damaged rats and controls learn maze tasks equally well. However, when the same cues are distributed around the maze, requiring the capacity for forming associations among spatially disparate items to guide performance, hippocampal-damaged rats typically show significant impairment relative to controls[14]. Thus, the idea that the hippocampus creates a cognitive map by storing relationships among spatially distributed stimuli during navigation[2] is predicated on the assumption that this neural substrate can support associations among discontiguous cues.

When one considers the spatial and temporal data together, these behavioral observations suggest that one general contribution of the hippocampal system is that it provides an associative link between discontiguous events. As shown above, the nature of the discontiguity can occur in different domains. What then is special in the neural structure of the hippocampal system that enables this particular contribution to learning?

Context and the hippocampus

Part of the answer to this question might be found in the way contextual information is represented neurally in the hippocampus. Within our present framework, 'context' can be thought of as a stimulus environment that is changing much more slowly than the specific stimuli being learned. To illustrate how cellular activity related to context might be important for learning discontiguous information, we return to the example of associating temporally disparate events. At the cellular level, it has been shown that concurrent activation of a presynaptic and postsynaptic cell provides the physiological basis for the induction of a long-term potentiation (LTP) of the postsynaptic response[16,17]. Thus, the synaptic efficacy or association between two cells can increase with their co-activation. Detailed investigations of the receptor kinetics underlying this phenomenon have indicated that a time window exists, whereby the postsynaptic cell must be sufficiently depolarized within ~100 ms of presynaptic excitation for LTP induction to occur[18,19]. If one imagines a group of cells that encodes one feature in a learning environment while another group of cells encodes a different feature, then these features might be 'bound' together neurally by the co-activation of their cellular representations using an LTP-like process. A problem arises, however, if the features being bound together are perceived at different points in time, and if the inter-event time exceeds the time window normally associated with LTP induction (~100 ms). Starting with the known biophysical mechanisms of LTP induction; how are temporally discontiguous events (having an inter-event time greater than ~100 ms) associated?

Earlier theoretical and computational models of associative memory showed that static patterns of information can be stored in simplified networks of cells using an LTP-like learning rule[20–23]. These models are capable of learning input patterns and recalling them during completion tasks in which a degraded version of the full pattern is presented[24]. However, until recently, no theoretical model has been able to learn and recall time-dependent patterns of events (sequence information) without relying on mechanisms unsupported by existing physiological observations (for example, Refs 25,26). Primarily, this is because experimental observations of hippocampal cellular behavior have not pointed to a clear mechanism that could support such a process. In the last year, however, several physiologically realistic models have emerged that learn and recall sequence information and, thus, are capable of solving the time-spanning problem alluded to above[27–30]. All of these models are based on hippocampal region CA3, as its pronounced anatomical pattern of recurrent excitatory connections among pyramidal cells results in an extensive network of potential associations[31].

During the course of learning a sequence of items separated in time, these models develop what can be

called 'context fields' as illustrated in Fig. 1 (Refs 27,29,30,32). In simulations of sequence learning, a single item is most often represented by a unique group of cells firing at the same time (five cells in this figure). With respect to hippocampal physiology, this pattern of input can be considered analogous to the afferent activation of a subset of CA3 pyramidal cells via the perforant path, owing to a unique pattern of sensory events. When a cell that encodes a sequence item fires closely in time to a synaptically connected cell (that might be firing at that particular moment as a result of recurrent excitation from another cell), the association between the two cells is increased in proportion to their temporal overlap. That is, of course, assuming that they fire within ~100 ms or less of each other, given the temporal constraints on LTP induction[18,19]. This simple Hebbian-type learning rule used to approximate LTP in several of the models results in a gradient of different synaptic strengths between cells that encode the items in a sequence and cells that are firing at random in the background, owing to recurrent activation from other pyramidal cells. With repeated exposure to the same sequence of items, this gradient of synaptic potentiation from cells encoding the sequence causes background cells to fire in a sustained manner for a period of time that is proportional to the potentiation strength. This marked change in background firing can be seen most readily by comparing Fig. 1A, which shows a sample of simulated pyramidal cell activity in our own model[29,30,33] during the first learning trial, with Fig. 1B, which shows the fourth learning trial. In this figure, each item (13 in total) in the sequence is represented by the co-activation of a unique set of five simulated pyramidal cells shown at the top of both graphs. As can be seen, the initially random background activity (Fig. 1A) becomes organized with learning (Fig. 1B) to the point where certain cells begin to respond to contiguous segments of the longer sequence. These simulated context-sensitive cells have similar characteristics and development to hippocampal place cells that fire selectively when an organism is in a specific location of a test environment (place field) during a spatial learning task[27,29,32]. In this case, one could think of a navigational path being learned as a series of specific locations in a particular order. Thus, the full path can be represented by a serially ordered sequence of items.

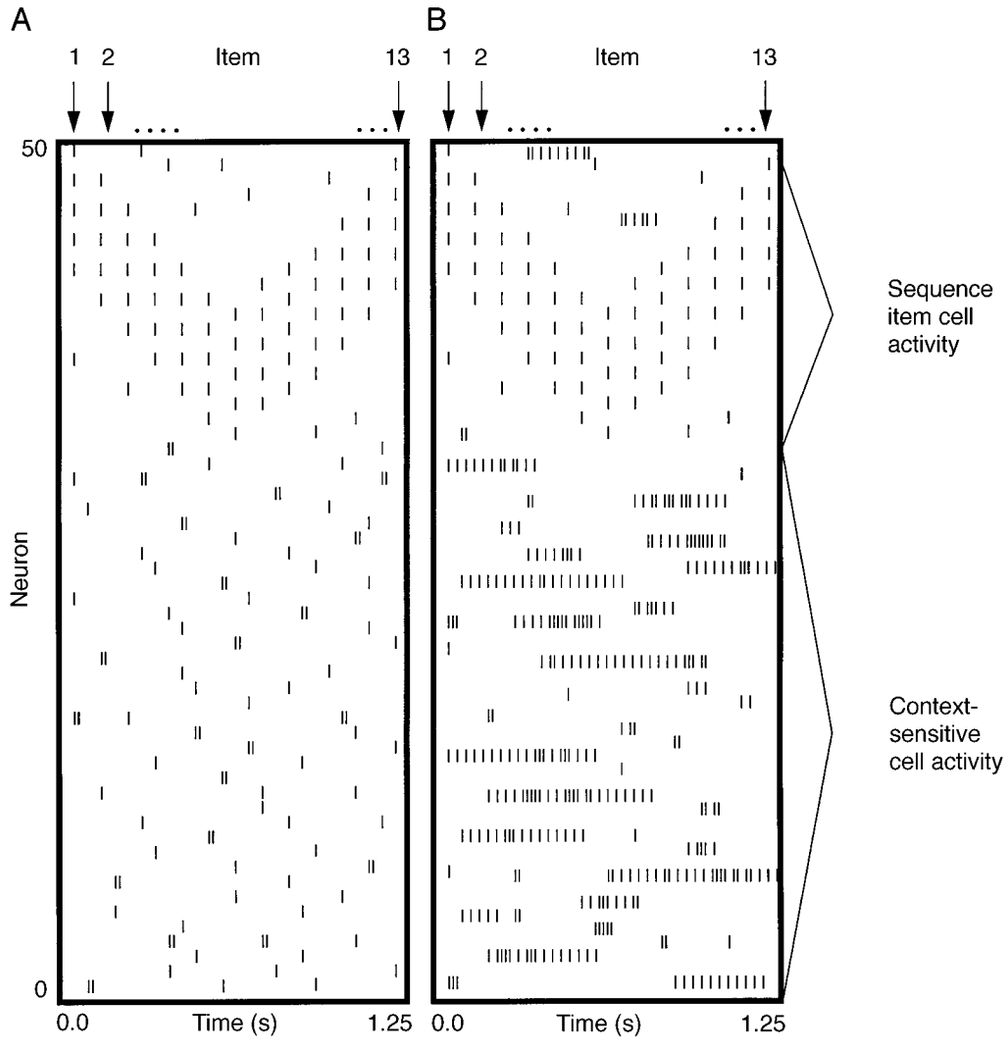

Fig. 1. ***Context field development****. Each rectangle shows a subset of 50 simulated pyramidal cells firing across time. For clarity, each action potential is represented as a vertical line. A unique group of five cells firing at the same time encodes a single item in a 13-item sequence. These items are encoded by cells firing in the top portion of each rectangle. This pattern of input can be thought of as representing the afferent activation of a subset of hippocampal CA3 pyramidal cells via the perforant path owing to a specific pattern of sensory events.* **(A)** *Notice the presence of unorganized background firing in pyramidal cells during the first learning trial that do not encode sequence items directly. This firing stems from activity at recurrent excitatory synapses. Repeated exposure to the same sequence can lead to enhanced synaptic potentiation between cells firing in the background and cells that encode sequence items if they fire closely to each other in time. This is due to a simple Hebbian-type learning rule used to approximate conditions that result in the long-term potentiation (LTP) of a synapse (see text for details).* **(B)** *After the fourth learning trial, this repeated potentiation leads to a condition where background cells begin to respond to the appearance of contiguous segments of the entire sequence. The portion of the full sequence to which the cell responds is called the 'context field' of the cell. This is somewhat analogous to the formation of place fields, where a cell fires selectively when an organism is in a specific location comprising a longer navigational path. Because the context fields overlap, the entire sequence can be reconstructed by interdigitating them in the proper order. Context fields develop in a variety of neural models that use different single-cell approximations and different learning rules as long as certain critical factors regarding network topology and activity levels are satisfied (see text).*

In these computational models, context-sensitive cell firing serves to 'glue' together subsequent items in a sequence, even when the time period between them is an order of magnitude longer than the time frame for LTP induction (~100 ms). This is because the i^{th} and i^{th} + 1 items in a sequence are associated with each other indirectly through their mutual association with one or more context-sensitive cells that fire at that particular portion of the full sequence. Indeed, because many of the context fields overlap, items with a very large temporal disparity (that is, several orders of magnitude greater than the time frame for LTP induction) can also be associated using this general

mechanism, where one context-sensitive cell potentiates the firing of a second one[29]. Thus, context fields can be thought of as a biophysical realization within CA3 of earlier behavioral theories positing a role for the hippocampus in associating temporally discontinuous stimuli[5,34,35].

A physiological recipe for context field development

It is important to note that while these models are based on hippocampal physiology, the essential components for context field development might also exist in other brain areas, particularly in the prefrontal cortex[36] and perhaps in the parahippocampal region[37,38]. An interesting feature of context fields is that they develop in a diversity of different neural models, ranging from those that use simple integrate-and-fire representations[27,32] to those that incorporate multi-compartmental approximations of the different cell types known to exist in region CA3 (Refs 29,30,33). Thus, their development is robust with regard to specific parameters at the single-cell level and, as will be noted below, with respect to the particular learning rule implemented in the model. The critical factors that determine whether or not context fields develop reside at the network level and have all been seen experimentally in the hippocampus. These factors include the following:

(1) Asymmetric connections among excitatory synapses allow items occurring at different times to be associated together in an order-preserving manner. If connections are symmetric in a model (that is, cell A contacts cell B and cell B in turn contacts cell A), the i^{th} item encoded by a group of cells will potentiate cells encoding the i^{th} + 1 and i^{th} – 1 items equally. This symmetry condition must be broken to maintain the correct order of items. In the hippocampus, it has been estimated from *in vitro* recordings that a single CA3 pyramidal cell contacts ~1–5% of its neighboring pyramidal cells[39,40]. This sparse connectivity among pyramidal cells leads to a highly asymmetric network, assuming a random distribution of connections.

(2) A low level of diffuse background pyramidal cell firing during learning provides a source of cellular activity for forming potentiated associations with cells encoding the items to be remembered (Fig. 1A). It is the sustained firing by these cells that results in context field development. Sources of background activity include firing owing to recurrent excitation within region CA3 as well as diffuse innervation from cholinergic cells in the basal forebrain[41]. Background cells firing randomly are associated, through an LTP-like process, with cells encoding an input item by varying degrees, depending on how closely in time they are activated relative to one another. This gradient of different synaptic association strengths becomes more robust with repeated presentation of the same sequence of items, and can lead progressively to sustained firing of background cells with the appearance of specific input items (formation of context fields). The size of the context field (measured as the number of contiguous items in the sequence to which the cell responds with active firing) varies proportionately to the average level of synaptic potentiation in each cell. Hippocampal recordings from rodents suggest that only a small fraction of pyramidal cells (10–20%) are active at any given point in time *in vitro*[42] and *in vivo*[43–45]. This range of activity levels is, in fact, consistent with model parameters found to produce the most stable context fields during learning[27,29,30].

(3) An LTP-like (Hebbian) learning mechanism at excitatory synapses of pyramidal cells provides a basis for making potentiated associations between different cells. It appears that context field formation is fairly robust with respect to this factor, occurring in a variety of models using very different learning rules[27,29]. Repeated experimentation has shown such a physiological mechanism to exist in the hippocampus[16,17].

(4) Periodic switching between afferent and intrinsic sources of synaptic activity driving context-sensitive cell firing enables the system to update current sensory information regularly from regions extrinsic to CA3, while allowing intrinsic associations (at recurrent collaterals between different context-sensitive cells) to be potentiated further by this information[29]. This means that a context field arises through a process where afferent synaptic transmission stemming from sensory events (the items in a sequence) is dominant for a brief period of time over synaptic transmission from other context-sensitive cells, allowing potentiated associations to form between the context-sensitive cell and cells encoding the sequence items. This is followed by an equally brief period of time in which synaptic transmission between different context-sensitive cells (intrinsic fibers) within region CA3 dominates over synaptic transmission from afferent sources, in order to potentiate associations between these cells without interference from competing sensory information.

This fourth mechanism has a basis in recent *in vitro* observations that the γ-aminobutyric acid-B ($GABA_B$) receptor agonist baclofen selectively decreases the amplitude of excitatory postsynaptic potentials (EPSPs) in hippocampal CA1 pyramidal cells induced by Schaffer collateral stimulation (stratum radiatum), while showing no significant effect on perforant path transmission[46,47]. Thus, communication between pyramidal cells within the hippocampus (intrinsic and association fibers) is suppressed, in part, by $GABA_B$ receptor activation, while sensory signals arising from afferent fibers outside this cortical area are relatively unaffected. Consequently, the relative contribution of information from inside and outside the hippocampus is shaped by this modulation. Recent biophysical modeling has shown that such GABAergic modulation might occur in synchrony with the endogenous 4–10 Hz theta rhythm[29,30], a prominent field oscillation present throughout the entorhinal–hippocampal system during exploratory behavior and learning[2,48]. Previous reports have suggested that the frequency of the theta oscillation might be optimal for inducing LTP using patterned stimuli[49]. Additional investigation has shown that LTP induction at theta frequency stimulation is dependent on $GABA_B$ receptor activation *in vitro*[50]. This suggests that $GABA_B$ receptor-activated modulation of synaptic transmission might play a role in certain forms of hippocampal plasticity during behaviors that foster theta activity.

Computational modeling in our laboratory and others suggests that either an increase in context field size or an increase in the degree of overlap between fields is required to associate temporally discontiguous items with very large inter-event times (that is, several orders of magnitude larger than the time frame for

LTP induction). Context fields overlap with associations made between different context-sensitive cells only when synaptic transmission and potentiation between them does not have to compete continually with afferent activity from cells encoding the input pattern items[29,30]. Thus, some modulation must occur that enables a period of time to occur within a theta cycle in which increased potentiation between context-sensitive cells is encouraged, without interference from cells encoding the afferent input pattern. In our particular implementation of the critical factors listed in this section, rhythmic variation in network $GABA_B$-receptor-mediated conductance levels provides a source of switching between periods (within a theta cycle) where potentiation between afferent cells and context-sensitive cells dominates (afferent dynamics) to periods where potentiation between different context-sensitive cells dominates (intrinsic dynamics)[29,30]. Numerous simulations have shown that context fields fail to develop enough overlap between them to learn and recall accurately input items separated by as little as 200 ms without such switching between afferent and intrinsic dynamics during theta oscillations[29,30].

Task classification via neural mechanisms

Levy[27] has remarked that models capable of sequence learning demonstrate a capacity for solving a wide variety of tasks thought to be hippocampal dependent. Thus, in addition to being able to solve time delay[29,30] and sequence completion tasks[27,29,30], models with the general architecture described above that develop context fields can also solve other classes of tasks that involve discontiguities in the serial positioning of related items. For example, it has been shown that rats with neurotoxic lesions of the hippocampus or fimbria–fornix perform comparably to control rats in odor-guided versions of paired-associate learning. However, when the same items are used later in tests of transitive inference (that is, satisfying the relationship: if item A is associated with item B and item B is associated with item C, then item A is inferred to be associated with item C), lesioned rats score at chance levels, while control rats display a significant capacity for solving the task[51,52]. Models that form context fields can solve tasks that involve transitive inference learning. Figure 2A shows the experimental design used by Bunsey and Eichenbaum in their study of transitive inference[51]. We have adopted this protocol in a computer simulation of the task using a detailed, biophysical model based on region CA3 (Refs 29,30,33). Each of the four factors (outlined above) found to be critical for context field development has been included in the model. We have also included realistic approximations of the intrinsic ionic conductances known to exist in real hippocampal pyramidal cells and interneurons, the synaptic kinetics governing their interactions, conduction velocities and connectivity parameters (Fig. 2A)[29,30,33]. First, we trained the model to learn simple paired-associates (such as, A–B, X–Y, B–C and Y–Z) and then tested for transitivity by priming it with either item A or item X and determining the completion accuracy. Two versions of the same model were tested; one in which normal context field development occurred during learning (normal); and another in which context fields failed to develop (damaged). As mentioned earlier, context field development depends critically on the periodic switching between afferent and intrinsic sources of synaptic activity dominating the population dynamics. This is mediated in our model by rhythmic (entrained to the endogenous population oscillation – theta rhythm) presynaptic GABAergic suppression of recurrent EPSPs (Refs 29,30). Context field development is curtailed substantially when this

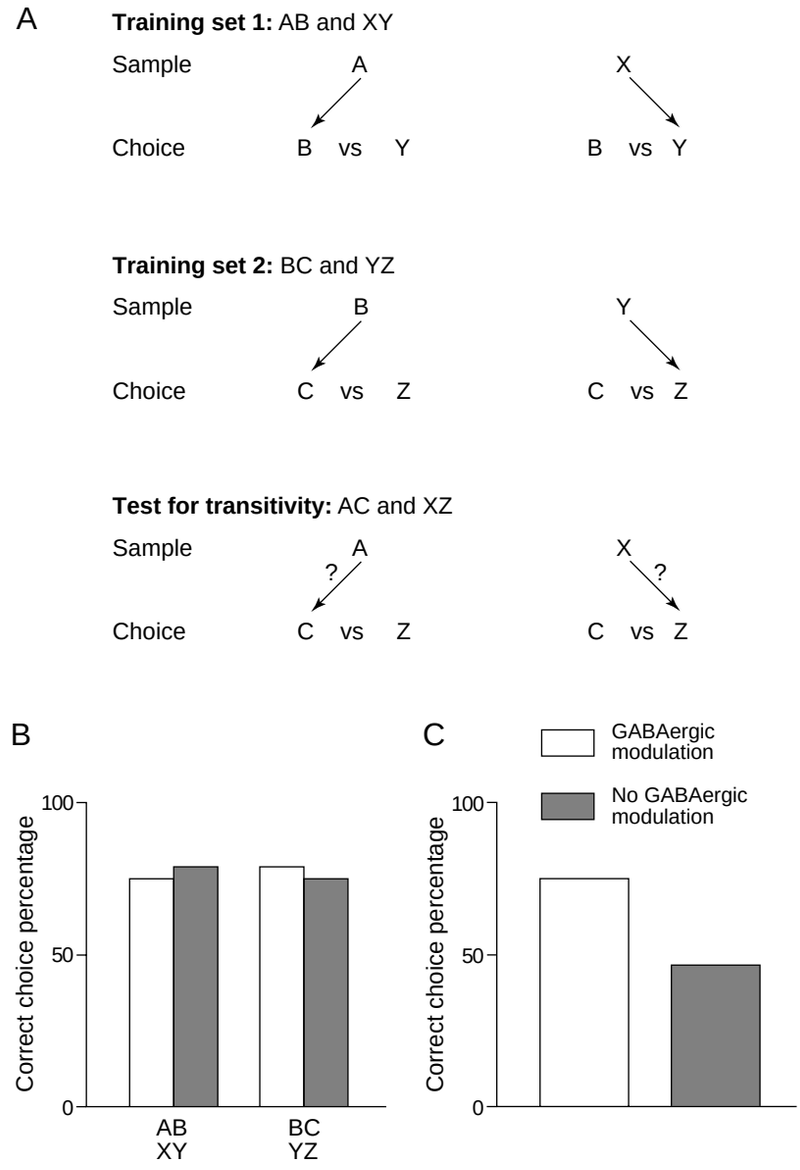

Fig. 2. ***Transitive inference.*** *Models that form context fields can also perform learning tasks that involve discontiguities in the serial positioning of items.* **(A)** *Using the Bunsey and Eichenbaum[51] design, we simulated this task using a biophysically detailed model of hippocampal region CA3 (see Refs 29,30,33 for details on model structure). We tested two versions of the same model in a simulation of transitive inference learning. One version of the model (normal) showed typical context field development similar to that illustrated in Fig. 1 and another version (damaged) was altered such that normal field development did not occur (see text for details).* **(B)** *We first tested the ability of each version of the model to learn simple paired-associates. During a five-trial learning session, we presented different pairs of items in sequence (for example, A–B, X–Y, B–C and Y–Z) together. Each item within a pair was presented 250 ms apart and 3 s passed between the presentation of each distinct pair. We next tested for accurate completion by presenting the model with the first item in each pair and determining which paired item was recalled. Both the normal and damaged versions of the model demonstrated accurate completion of these pairs at approximately the same percentage values.* **(C)** *We next tested for transitivity by presenting either item A or item X to both versions of the model and determined which item (C or Z) was recalled. In this case, only the normal version of the model showed a capacity for transitive inference learning. The damaged version of the model performed at near chance levels.*

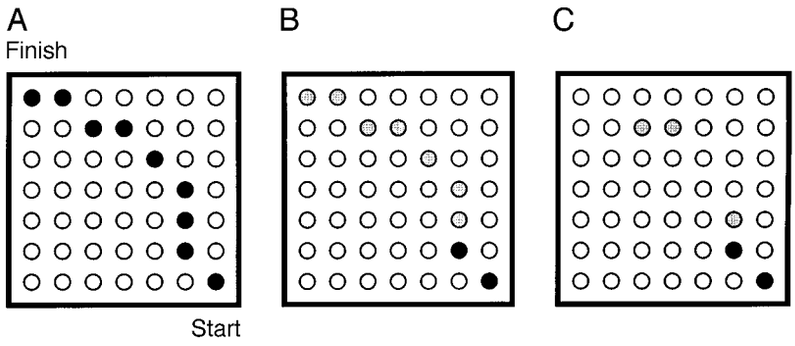

Fig. 3. ***Models that develop context fields can also perform spatial learning tasks.*** *Using the same two versions of the model described previously (see text and Refs 29,30,33 for details), we simulated simple path learning in a two-dimensional environment.* **(A)** *Each circle represents a specific 'location' in a simulated two-dimensional space. Each sequence of connected locations represents a spatial path. This simulation consisted of two learning trials in which a path was 'traversed' from a start position in the lower right-hand corner of the environment to the goal location in the top-left hand corner. A different location was reached every 100 ms. We tested recall of the path by priming each version of the model with the first two locations (black circles) in the path to give it a general vicinity and direction.* **(B)** *The full path was recalled accurately (grey circles) in the correct order by the normal version of the model.* **(C)** *However, the correct sequence order was not preserved in the damaged version of the model in which context fields failed to develop. In this example, the third location was recalled immediately, followed by locations six and seven. This suggests that the 'predictive' value of context fields might be an important component of path learning during navigational behavior.*

modulation is removed from our model[29,30]. As can be seen in Fig. 2B, both versions of the model displayed a capacity for learning simple paired-associates. However, only the model that showed context field development (normal) was able to learn the transitivity condition (Fig. 2C), as the damaged version of the model performed at near chance levels. This suggests that the neural mechanisms that support the formation of context fields might be critical elements in allowing the hippocampus to bridge discontiguous events together across time or serial positioning, or both.

There are other classes of tasks that might, in part, also be reduced to the problem of bridging together distinct events across temporal or spatial discontiguities, or some combination of the two. Undoubtedly, spatial tasks such as maze learning are dependent on numerous brain regions for successful completion. However, the hippocampal contribution to such phenomena can be framed as a problem of linking together different spatial locations into a specific serial order. Thus, hippocampal models that learn and recall sequence information using context fields should, in principle, have the capacity for solving certain spatial tasks. Figure 3 shows an example of this, where the normal model was able to recall a previously learned spatial path (in the correct sequence order) when information about the first two locations only was provided (Fig. 3B). Without context field formation, the damaged version of this model failed to recall sequence locations in the correct order (Fig. 3C). Thus, in both transitive inference and spatial learning, accurate performance depended on context field development to preserve sequence order relationships. This suggests that tasks which appear quite different with regard to overt behavior might, in fact, have common neural mechanisms supporting components of learning. These results also suggest that context fields should not be thought of as only being place fields. As has been shown, context fields can be used in computational models to solve a variety of tasks involving time delays, pattern completion and disambiguation, transitive inverse and transverse patterning, as well as spatial learning[27–30,32]. Furthermore, several studies have now emerged showing that hippocampal place cells might encode much more than just spatial location. For instance, hippocampal cells in rodents have been shown to encode temporal[53–55] and positional[55,56] relationships among non-spatial cues in different non-spatial memory tasks. Pyramidal cells in region CA1 have also been shown to fire selectively when a rodent is in close proximity to specific goals or landmarks in a learning environment, regardless of their actual spatial location[57]. Thus, it seems that robust, sustained firing of hippocampal cells might encode a diversity of features in any given task. A still unresolved set of questions concerns how such fields arise in different learning situations and if there are any generic organizing principles controlling their development.

An additional point of note is that some forms of associations between discontiguous events might depend on specific mechanisms other than those outlined above[58–60]. For instance, conditioned taste aversion (CTA) involves associating an illness with a particular food eaten some time in the past. Clearly, this association involves a long temporal discontiguity, yet CTA persists even with hippocampal lesions[59,60]. Thus, this form of association is not accounted for readily within the present framework. It remains to be determined experimentally whether the components advanced in the present model actually exist in brain regions thought to be important for CTA, such as the medial thalamus and the lateral nuclear group of the amygdala[60].

What then is the best strategy for understanding the role of the hippocampus in learning and memory? Several attempts have been made to determine a general class of learning tasks to which this brain region might contribute. From a biophysical perspective, it is still not clear what features are encoded in the hippocampus. However, by incorporating physiologically realistic approximations and behaviorally relevant tasks into neural models, one can generate a better understanding of how the biophysical machinery of the hippocampus constrains the types of information processing in which it participates. Based on a confluence of experimental observations and computational modeling, our position is that one important contribution of the hippocampus to learning and memory is the association of discontiguous events. Behavioral experiments that vary the degree to which discontiguities appear between items might be used to assess this dependency. Furthermore, specific physiological experiments can be performed to better determine how context fields arise in different task environments and their pharmacological signatures. For example, this work predicts that the local application of $GABA_B$ receptor antagonists in region CA3 will disrupt context field formation. It is known that temporary blockade of GABAergic and cholinergic innervation from the medial septum does impair place field activity in CA3 (Ref. 61). However, additional pharmacological manipulations must be performed to determine which class of receptors is involved most directly in this effect, where they reside and how they modulate relational memory *in vivo*.

Selected references
1 **Scoville, W.B. and Milner, B.** (1957) *J. Neurol. Neurosurg. Psychiatry* 20, 11–21
2 **O'Keefe, J. and Nadel, L.** (1978) *The Hippocampus as a Cognitive Map*, Oxford University Press
3 **Squire, L.R.** (1992) *Psychol. Rev.* 99, 195–231
4 **Eichenbaum, H. *et al.*** (1994) *Behav. Brain Sci.* 17, 449–518
5 **Rawlins, J.N.P.** (1985) *Behav. Brain Sci.* 8, 479–496
6 **Cohen. N.J. and Eichenbaum, H.** (1993) *Memory, Amnesia, and the Hippocampal System*, MIT Press
7 **Squire, L.R.** (1987) *Memory and Brain*, Oxford University Press
8 **Nadel, L. and Moscovitch, M.** (1997) *Curr. Opin. Neurobiol.* 7, 217–227
9 **Meck, W.H. *et al.*** (1984) *Behav. Neurosci.* 98, 3–22
10 **Thompson, R.F. *et al.*** (1982) in *Conditioning: Representation of Involved Neural Functions* (Woody, C.D., ed.), pp. 115–129, Plenum Press
11 **Ross, R.T. *et al.*** (1984) *Behav. Neurosci.* 98, 211–225
12 **Zoladek, L. and Roberts, W.A.** (1978) *Learn. Behav.* 6, 77–81
13 **Suzuki, S. *et al.*** (1980) *Learn. Motiv.* 11, 1–18
14 **O'Keefe, J. and Conway, D.H.** (1980) *Physiol. Psychol.* 2, 229–238
15 **Eichenbaum, H. *et al.*** (1990) *J. Neurosci.* 10, 331–339
16 **Bliss, T.V.P. and Lomo, T.** (1973) *J. Physiol.* 232, 331–356
17 **Bliss, T.V.P. and Collingridge, G.L.** (1993) *Nature* 361, 31–39
18 **Levy, W.B. and Steward, O.** (1983) *Neuroscience* 8, 791–797
19 **Larson, J. and Lynch, G.** (1989) *Brain Res.* 489, 49–58
20 **Marr, D.** (1971) *Philos. Trans. R. Soc. London Ser. B* 262, 23–81
21 **McNaughton, B.L. and Morris, R.G.M.** (1987) *Trends Neurosci.* 10, 408–415
22 **Hasselmo, M.E. and Bower, J.M.** (1993) *Trends Neurosci.* 16, 218–222
23 **Hasselmo, M.E.** (1993) *Neural Comput.* 5, 32–44
24 **Hasselmo, M.E.** (1995) *Behav. Brain Res.* 67, 1–27
25 **Kleinfeld, D.** (1986) *Proc. Natl. Acad. Sci. U. S. A.* 83, 9469–9473
26 **Kleinfeld, D. and Sompolinsky, H.** (1988) *Biophys. J.* 54, 1039–1051
27 **Levy, W.B.** (1996) *Hippocampus* 6, 579–590
28 **Jensen, O. and Lisman, J.E.** (1996) *Learn. Mem.* 3, 279–287
29 **Wallenstein, G.V. and Hasselmo, M.E.** (1997) *J. Neurophysiol.* 78, 393–408
30 **Wallenstein, G.V. and Hasselmo, M.E.** (1997) in *Advances in Computational Neuroscience* (Bower, J.M., ed.), pp. 547–552, Plenum Press
31 **Amaral, D.G. and Witter, M.P.** (1989) *Neuroscience* 31, 571–591
32 **Levy, W.B.** (1989) in *Computational Models of Learning in Simple Neural Systems* (Hawkins, R.D. and Bower, G.H., eds), pp. 243–305, Academic Press
33 **Wallenstein, G.V. and Hasselmo, M.E.** (1997) *Brain Res. Bull.* 43, 485–493
34 **Nadel, L. and Willner, J.** (1980) *Physiol. Psychol.* 8, 218–228
35 **Nadel, L. *et al.*** (1985) in *Context and Learning* (Balsam, P., Tomie, A. and Hillsdale, N.J., eds), pp. 385–406, Erlbaum
36 **Fuster, J.M.** (1989) *The Prefrontal Cortex: Anatomy, Physiology, and Neuropsychology of the Frontal Lobe*, Raven Press
37 **Kohler, C.** (1989) *J. Comp. Neurol.* 246, 149–169
38 **Burwell, R.D. *et al.*** (1995) *Hippocampus* 5, 390–408
39 **MacVicar, B.A. and Dudek, F.E.** (1980) *Brain Res.* 184, 220–223
40 **Traub, R.D. and Miles, R.** (1991) *Neuronal Networks of the Hippocampus*, Cambridge University Press
41 **Frotsher, M. and Leranth, C.** (1985) *J. Comp. Neurol.* 239, 237–246
42 **Traub, R.D. *et al.*** (1989) *Science* 243, 1319–1325
43 **Buzsáki, G. *et al.*** (1983) *Brain Res. Rev.* 6, 139–171
44 **Thompson, L.T. and Best, P.J.** (1989) *J. Neurosci.* 9, 2382–2390
45 **Ylinen, A. *et al.*** (1995) *Hippocampus* 5, 78–90
46 **Ault, B. and Nadler, J.V.** (1982) *J. Pharmacol. Exp. Ther.* 223, 291–297
47 **Colbert, C.M. and Levy, W.B.** (1992) *J. Neurophysiol.* 63, 1–8
48 **Stewart, M. and Fox, S.E.** (1990) *Trends Neurosci.* 13, 163–168
49 **Larson, D. *et al.*** (1986) *Brain Res.* 368, 347–350
50 **Mott, D.D. and Lewis, D.V.** (1991) *Science* 252, 1718–1720
51 **Bunsey, M. and Eichenbaum, H.** (1996) *Nature* 379, 255–257
52 **Dusek, J.A. and Eichenbaum, H.** (1997) *Proc. Natl. Acad. Sci. U. S. A.* 94, 7109–7114
53 **Otto, T. and Eichenbaum, H.** (1992) *Hippocampus* 2, 323–334
54 **Zhu, X.O. *et al.*** (1995) *Eur. J. Neurosci.* 7, 753–765
55 **Hampson, R.E. *et al.*** (1993) *Behav. Neurosci.* 107, 715–739
56 **Deadwyler, S.A. *et al.*** (1996) *J. Neurosci.* 16, 354–372
57 **Gothard, K.M. *et al.*** (1996) *J. Neurosci.* 16, 823–835
58 **Stevens, R. and Cowey, A.** (1972) *Brain Res.* 46, 1–22
59 **Reilly, S. *et al.*** (1993) *Behav. Neurosci.* 107, 996–1004
60 **Yamamoto, T. *et al.*** (1995) *Neurosci. Res.* 22, 31–49
61 **Mizumori, S.J.Y. *et al.*** (1989) *J. Neurosci.* 9, 3915–3928

Acknowledgements
This work was supported in part by the Human Frontier Science Program and the National Institutes of Mental Health (MH-52732-04).

What is the amygdala?

Larry W. Swanson and Gorica D. Petrovich

'Amygdala' and 'amygdalar complex' are terms that now refer to a highly differentiated region near the temporal pole of the mammalian cerebral hemisphere. Cell groups within it appear to be differentiated parts of the traditional cortex, the claustrum, or the striatum, and these parts belong to four obvious functional systems – accessory olfactory, main olfactory, autonomic and frontotemporal cortical. In rats, the central nucleus is a specialized autonomic-projecting motor region of the striatum, whereas the lateral and anterior basolateral nuclei together are a ventromedial extension of the claustrum for major regions of the temporal and frontal lobes. The rest of the amygdala forms association parts of the olfactory system (accessory and main), with cortical, claustral and striatal parts. Terms such as 'amygdala' and 'lenticular nucleus' combine cell groups arbitrarily rather than according to the structural and functional units to which they now seem to belong. The amygdala is neither a structural nor a functional unit.

Trends Neurosci. (1998) 21, 323–331

Larry W. Swanson and Gorica D. Petrovich are part of the Neuroscience Program at the University of Southern California, Los Angeles, CA 90089-2520, USA.

SLICES THROUGH THE TEMPORAL POLE of the human cerebral hemispheres reveal an almond-shaped mass of gray matter that Burdach[1] discovered and called the amygdalar nucleus in the early 19th century (Fig. 1). Starting about 50 years later, the microscopic examination of histological tissue sections began to reveal more and more structural differentiation in the amygdala; and the extent of its outer border, and number and classification of its subdivisions, remain controversial today. In this article, we present a model of amygdalar architecture based on recent embryological, neurotransmitter, connectional and functional data. When placed in the context of cerebral hemisphere architecture as a whole,

 0166 - 2236/98/$19.00 PII: S0166-2236(98)01265-X