PERGAMON

Neural Networks xx (0000) xxx–xxx

Neural Networks

www.elsevier.com/locate/neunet

Modeling goal-directed spatial navigation in the rat based on physiological data from the hippocampal formation

Randal A. Koene[a,*], Anatoli Gorchetchnikov[b], Robert C. Cannon[c], Michael E. Hasselmo[a]

[a]*Department of Pyschology and Program in Neuroscience, Boston University, Boston, MA 02215, USA*
[b]*Department of Cognitive and Neural Systems, Boston University, Boston, MA 02215, USA*
[c]*Institute of Adaptive and Neural Computation, Division of Informatics, 5 Forrest Hill, Edinburgh EH1 2QL, UK*

Abstract

We investigated the importance of hippocampal theta oscillations and the significance of phase differences of theta modulation in the cortical regions that are involved in goal-directed spatial navigation. Our models used representations of entorhinal cortex layer III (ECIII), hippocampus and prefrontal cortex (PFC) to guide movements of a virtual rat in a virtual environment. The model encoded representations of the environment through long-term potentiation of excitatory recurrent connections between sequentially spiking place cells in ECIII and CA3. This encoding required buffering of place cell activity, which was achieved by a short-term memory (STM) in EC that was regulated by theta modulation and allowed synchronized reactivation with encoding phases in ECIII and CA3. Inhibition at a specific theta phase deactivated the oldest item in the buffer when new input was presented to a full STM buffer. A 180° phase difference separated retrieval and encoding in ECIII and CA3, which enabled us to simulate data on theta phase precession of place cells. Retrieval of known paths was elicited in ECIII by input at the retrieval phase from PFC working memory for goal location, requiring strict theta phase relationships with PFC. Known locations adjacent to the virtual rat were retrieved in CA3. Together, input from ECIII and CA3 activated predictive spiking in cells in CA1 for the next desired place on a shortest path to a goal. Consistent with data, place cell activity in CA1 and CA3 showed smaller place fields than in ECIII.

Keywords: Hippocampal theta oscillation; Spatial navigation; Functional interaction

1. Introduction

The theta rhythm consists of large amplitude 3–12 Hz oscillations recorded in the EEG of the hippocampus of different animal species during exploration and orientation to stimuli (Buzsáki, Leung, & Vanderwolf, 1983). These theta oscillations are associated with rhythmic changes in a number of physiological variables, including neuronal membrane potential (Fox, 1989), synaptic currents in different layers (Brankack, Stewart, & Fox, 1993), the magnitude of synaptic transmission (Wyble, Linster, & Hasselmo, 1997) and the spiking activity of neurons (Fox, Wolfson, & Ranck, 1986). In addition, extensive data demonstrates that long-term potentiation of synaptic strength can be induced more effectively with stimulation at the peak of theta waves, while decreases in strength or no change is observed with stimulation at the trough of the theta wave (Hasselmo, Bodelon, & Wyble, 2002; Hölscher, Anwyl, & Rowan, 1997; Huerta & Lisman, 1993; Pavlides, Greenstein, Grudman, & Winson, 1988). This supports the hypothesis that theta rhythm allows separate phases of encoding and retrieval. In rats, lesions of the fornix remove much of the amplitude of theta rhythm. This impairs the ability to learn reversal tasks, in which a previously rewarded behavior must be replaced with an opposite behavior. For example, rats have difficulty learning a right turn response after initially learning a left turn response (M'Harzi, Palacios, Monmaur, Houcine, & Delacour, 1987). Previous analytical work has shown how performance in behavioral tasks such as the reversal task could depend upon specific phase relationships between physiological variables changing during theta rhythm oscillations (Hasselmo et al., 2002). The simulations presented here further demonstrate the functional role of phase relationships in entorhinal cortex and prefrontal cortex, as well as the functional role of other physiological properties.

* Corresponding author. Tel.: +1-617-358-2769; fax: +1-617-353-1424.
E-mail addresses: randalk@bu.edu (R.A. Koene), anatoli@cns.bu.edu (A. Gorchetchnikov), robert.cannon@ed.ac.uk (R.C. Cannon), hasselmo@bu.edu (M.E. Hasselmo).

0893-6080/03/$ - see front matter © 2003 Published by Elsevier Science Ltd.
doi:10.1016/S0893-6080(03)00106-0

We have utilized spiking network models of rat entorhinal and hippocampal circuitry to analyze the functional role of theta rhythm oscillations. These models have been developed in two different software packages: CATACOMB, developed by Cannon, Koene, and Hasselmo (2002)) and KInNeSS, developed in the Hasselmo laboratory by Gorchetchnikov. The use of separate models effectively tests the robustness of theoretical hypotheses and allows analysis of individual details of function. In both of these projects, the use of spiking neurons in network simulations to guide movements of a virtual rat in a virtual environment allow the network activity to be analyzed in the same manner as spiking activity recorded from single neurons in awake behaving rats. We will first discuss research using CATACOMB.

Using CATACOMB, a spiking neuron model of rat entorhinal and hippocampal circuitry (Cannon et al., 2002) was implemented to investigate the significance of these phasic changes while performing goal-directed spatial navigation tasks in a T-maze. The roles of neuron populations (ECII, ECIII, CA3, CA1) are based on hypothesized functions of these individual regions in spatial learning and recall tasks. Instead of static data sets, environmental input to model networks changes in accordance with task-specific behaviours that arise from physical simulations driven by the network output (e.g. when exploration of the T-maze leads to the discovery of food-reward). A virtual rat is placed in a novel environment, given an opportunity to learn the environment and then presented with the task to navigate to a goal discovered in that environment.

The task was successfully learned, enabling the virtual rat to find a shortest path to food-reward. Synchronized learning in ECIII and CA3 was achieved with a STM buffer that insured ordered repetition of sequences of place cell spikes without the variability caused by differences in place field traversal times. We improved our knowledge of the connections between physiological and behavioural data through our approach: studying task-specific behavioural problems caused by the dynamic interaction of the neuronal circuitry with the experimental environment. In this manner, we demonstrated the well-known empirical phenomenon of phase precession by recording simulated hippocampal place cell firing. We also found that place fields produced in our model exhibited the typical finer grain in CA1 and less defined, coarser appearance in ECIII. Theta modulation in ECIII and CA3 provided phases of encoding and retrieval that operate in parallel without interference. Theta also regulated episodic repetition in a STM buffer. Phase differences betwen regions undergoing theta modulation ensured the correctly timed propagation of activity. Theta modulation in our model is, therefore, crucial in order to accomplish each of these three aspects of the learning and spatial navigation task.

2. Results

During initial encoding, a virtual rat follows predetermined exploratory trajectories in the environment (Fig. 1). Place cell activity is assigned with variable place field sizes (Barnes McNaghton, Mizumori, Leonard, & Lei--Huey, 1990; Eichenbaum, Dudchenko, Wood, Shapiro, & Tanila, 1999; Frank, Brown, & Wilson, 2000; McNaughton, Barnes, & O'Keefe, 1983; Muller & Kubie, 1989; Quirk, Muller, Kubie, & Ranck, 1992). Associations between sequential place cell activity along the trajectories as well as in the goal location is encoded on recurrent fibres in ECIII and CA3 by LTP. During task performance, retrieval in the spiking neuron network produces output at the CA1 population that guides the direction of movement.

During exploration, the interval between the initial arrival of place cell spikes from adjacent place fields is arbitrary. However, the window of spike intervals that elicit effective LTP is less than 40 ms (Levy & Stewart, 1983) and LTP elicited by the single initial spike input is not reliable enough to produce strong Hebbian learning. Repetition of the ordered presentation of neighbouring place cell spikes is needed within the LTP window. A STM buffer can accomplish this.

There are no initially known patterns and buffering cannot be LTP-dependent, as the establishment of LTP requires in excess of 100 ms (Bi & Poo, 1998). It is hypothesized that intrinsic mechanisms may maintain firing patterns, as initially proposed by Lisman (Lisman & Idiart, 1995). One such mechanism is ADP (calcium sensitive cation currents induced by muscarinic receptor activation (Klink & Alonso, 1997)), as shown below. Synchronized timing can be achieved by regulating repetition through theta modulation. This assures that STM reactivation coincides with an encoding phase in both ECIII and CA3 (Fig. 2). Recurrent inhibition within the buffer can separate the reactivation of sequential items to maintain order (gamma rhythm) (Jensen & Lisman, 1996; Lisman & Idiart, 1995).

2.1. Reactivation in EC at gamma intervals within a phase of theta oscillation may provide short-term memory

Spiking produced by afferent activity during the input phase of the buffer is reactivated by the ADP during subsequent repetition phases. Theta modulation of membrane potential regulates the two phases. Theta oscillations

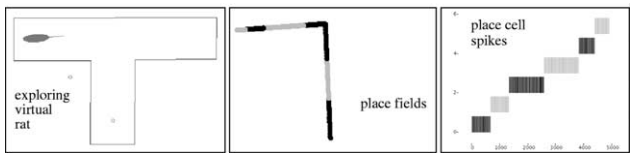

Fig. 1. Exploratory motions of a virtual rat generate place fields and corresponding spiking activity.

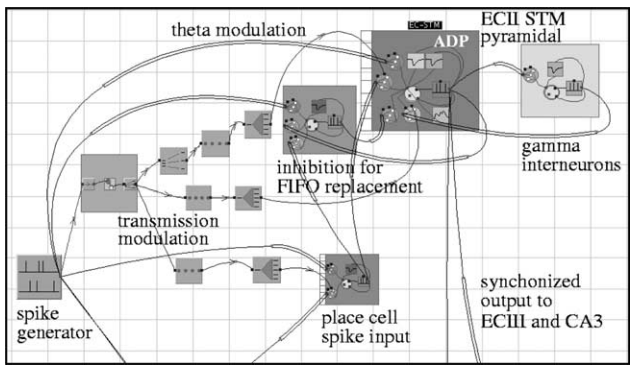

Fig. 2. The STM buffer in ECII converts arbitrary input spikes into synchronously repeated spike pairs for encoding in ECIII and CA3. *Theta modulation* of *ECII pyramidal cell* membrane potentials and phase-locked afferent and recurrent *transmission modulations* impose the repetition frequency and ensure that new input enters the buffer at the theta trough. This rhythmic modulation is regulated by a *spike generator* which simulates data on theta rhythmic firing of neurons in the septum. While the capacity is limited to two items in the current example, our STM model can accommodate up to four items that are reactivated on the depolarizing phase of each theta cycle. Items are separated by *gamma intervals* that are imposed by a population of interneurons. Another population of interneurons provides inhibition at the phase of the first item reactivation when *place cell spike input* arrives at a full STM buffer. This ensures *first-in-first-out* (*FIFO*) *replacement* of STM items.

(8 Hz) are produced by regular activity originating in septal populations (Brazhnik & Fox, 1999) believed to modulate the GABAergic inhibition of pyramidal cells via networks of interneurons (Alonso, Gaztelu, Bruno, & Garcia-Austt, 1987; Skaggs, McNaughton, Wilson, & Barnes, 1996; Stewart & Fox, 1990). Recurrent gamma inhibition produced by interneuronal networks (Bragin, Jando, Nadasdy, & Hetke, 1995) in response to each item spike imposes intervals that maintain the distinct and ordered reactivation of items (Fig. 3).

Noise and slow-AHP gradually decay STM. Prior to such decay, item replacement in the buffer is controlled by inhibition at a specific theta phase when new input is

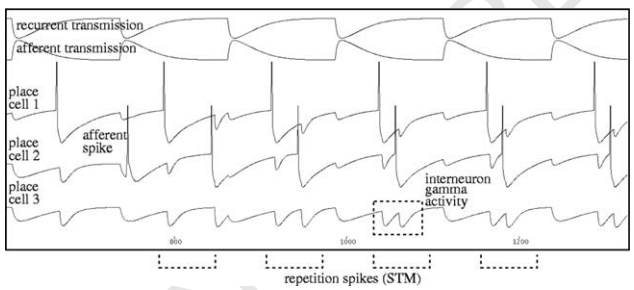

Fig. 3. New input is added to the STM buffer and the sequence is repeated. Reactivation of a buffered sequence occurs when ADP and the strongly depolarizing phase of theta modulation combine to exceed the firing threshold. Lateral inhibition maintains the separation of items in the buffer, as reactivation causes a recurrent interneuronal network to spike at gamma frequency (40 Hz). These recurrent GABAergic synaptic transmissions are enabled during the high value phase of 'recurrent transmission' modulation, while synapses receiving afferent input into the buffer are enabled during the high value phase of 'afferent transmission' modulation.

received. Replacement must retire the earliest item in the buffer (FIFO). Retirement must be rapid to avoid changes in repetition order that could become encoded in ECIII or CA3. This carefully timed inhibition triggered in response to the number of items in the buffer as well as the arrival of a new item spike (shown in Fig. 4) completes the STM model without the problematic input timing and ADP response characteristics required in the Jensen et al. model (Jensen, Idiart, & Lisman, 1996).

An interneuron population receives recurrent synaptic input from the buffer cells as well as input propagated from sources that supply afferent input to the buffer. Synaptic efficacies establish a gating function, such that N recurrent item inputs elicit subthreshold potentials in the interneuronal population, but an additional $N+1$th afferent input elicits spiking. The consequent GABAergic inhibition on the buffer is synchronized so that its hyperpolarizing effect delays reactivation of the earliest item in the buffer long enough for its ADP to subside (Fig. 4). The first item drops out of the buffer and the new item assumes the last position in the buffered sequence.

The STM buffer bridges large intervals between the appearance of successive place cell spikes and achieves extended repetition of paired spiking so that strong LTP can be established. Sequential binding is, therefore, less dependent on the speed of place field traversal and episodic learning proceeds with correct unidirectional heteroassociativity. The STM provides input that conforms to the encoding requirements of ECIII and CA3.

2.2. Encoding and retrieval can take place continuously in consecutive phases of theta

The behavioural task requires ongoing acquisition of new associations without interference from ongoing retrieval using the same synapses. We, therefore, hypothesize that learning and retrieval may take place in separate phases of theta modulation (Hasselmo et al., 2002), 180° apart to avoid interference between the two modes (Fig. 5). This hypothesis is supported by phase differences in different

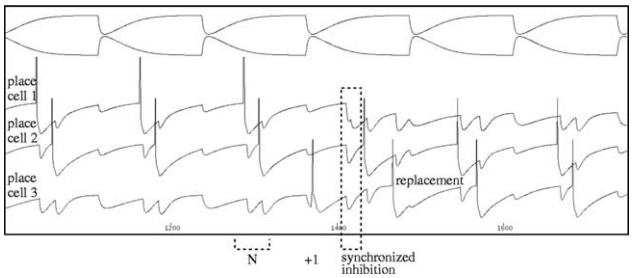

Fig. 4. The STM buffer replaces items in a first-in-first-out (FIFO) manner through targeted inhibition. When activity indicates a full buffer (N spikes), this combines with the arrival of a new item spike to trigger interneurons synchronized to spike so that the GABAergic effect on STM neurons suppresses firing of the first item until its ADP has decayed. The new item reactivates near the peak of theta modulation, thereby assuming the last position in the buffered sequence.

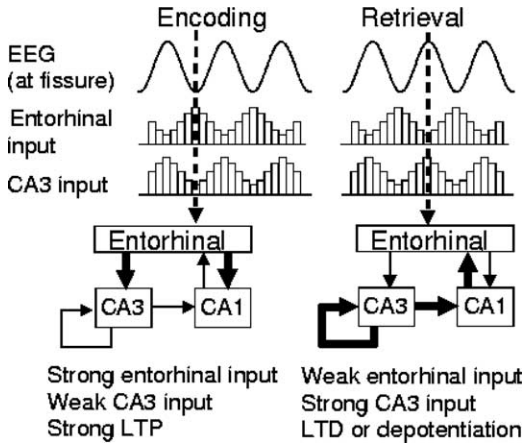

Fig. 5. When afferent synaptic input from EC to CA3 and EC to CA1 is strong encoding is favoured. When recurrent synaptic transmission in CA3 and synaptic input from CA3 to CA1 is strong retrieval is favoured.

synaptic input currents from EC and CA3 observed in current source density measures (Brankack et al., 1993; Stewart & Fox, 1990).

Phase-precession is a phenomenon recorded during in vivo experiments of spatial learning in rats (O'Keefe & Recce, 1993; Skaggs et al., 1996). Data from these experiments shows that the phase of theta at which a place cell fires gradually shifts to earlier phases as a rat traverses a known place field, as well as during learning. In the model, we measured the firing of cells relative to theta rhythm. As the virtual rat enters a place field the corresponding place cell fires at late phases of theta in CA1—predictive spiking. As the virtual rat crosses the place field, place cell firing becomes earlier in theta in CA1—encoding spiking (Fig. 6).

During encoding phases in the model we simulate suppression of recurrent fiber synaptic transmission via increased interneuron activating presynaptic $GABA_B$ channels (Hasselmo et al., 2002). This prevents interference from retrieval of existing associations (Fig. 7). Consistent with experimental data (Hölscher et al., 1997; Wyble, Hyman, Goyal, & Hasselmo, 2001), induction of LTP is

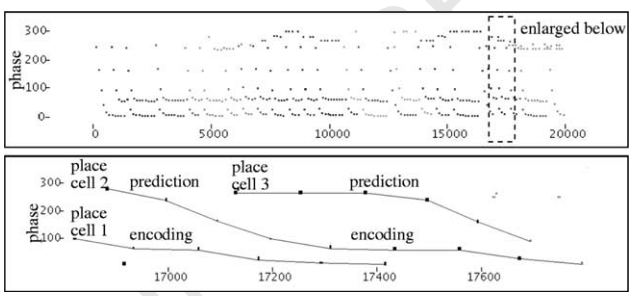

Fig. 6. Encoding and retrieval theta phases simulate theta phase precession. Model results for each place cell in CA1 show that as the corresponding place field is entered, the cell spikes in the retrieval phase of theta modulation. The retrieval allows the virtual rat to predict the next desired location on a path during navigation. As the virtual rat moves, spiking shifts to earlier phases. Spiking of the place cell during this phase allows it to participate in the episodic encoding of new paths.

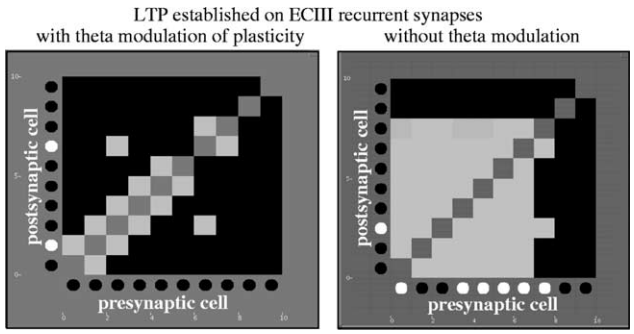

Fig. 7. The matrix of connectivity in the model demonstrates that theta rhythm is essential to the learning of correct paths as represented by ordered episodes of place cell spikes. Left: Good function with theta rhythm. The matrix on the left shows good connectivity resulting from separation of encoding and retrieval into distinct theta phases. This allows path encoding to take place without interference from retrieved spikes. The resulting patterns of synapses strengthened by LTP show pairwise associations between consecutive place cells along paths in the T-maze. Right: Poor encoding without theta. Without theta modulation the order breaks down and erroneous associations are formed between most locations.

rhythmically modulated to occur during the encoding phase, but not during the retrieval phase when recurrent synaptic transmission is strong.

Known locations adjacent to the virtual rat are retrieved in CA3 (Fig. 8), while associative retrieval in ECIII represents known paths (Fig. 9). A PFC STM buffer stores the location of goals that were discovered. PFC and hippocampus must spike synchronously.

CA1 activity (Fig. 10) depends on a combination of sub-threshold contributions from both CA3 (adjacent candidates) and ECIII (association with a path to the goal). Spiking in CA1 activates an interneuronal population that provides recurrent inhibition to CA1. A winner-take-all competition is accomplished (Fig. 11) in which one successful candidate spikes to indicate the next desired place on a shortest path to a goal.

Empirical data indicates that place fields in CA1 are significantly smaller than place fields in EC (Barnes et al., 1990; Frank et al., 2000; Quirk et al., 1992). Measurements

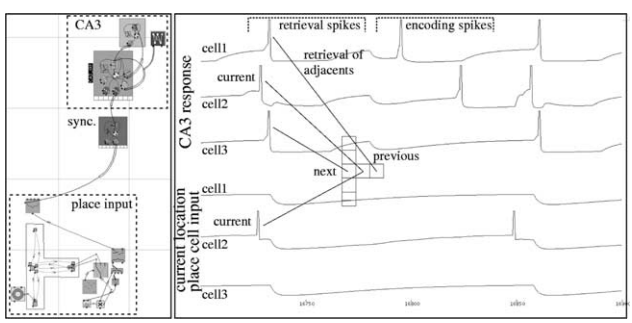

Fig. 8. Current location drives retrieval in CA3. As the virtual rat moves up the stem of the T-maze, a synchronization population ensures that spikes representing current location arrive in CA3 at its retrieval phase. Strong LTP on recurrent fibres then retrieves spikes for known adjacent locations. Lateral inhibition through interneurons in CA3 limits the spread of retrieval. Both the previous and next locations in the stem are retrieved.

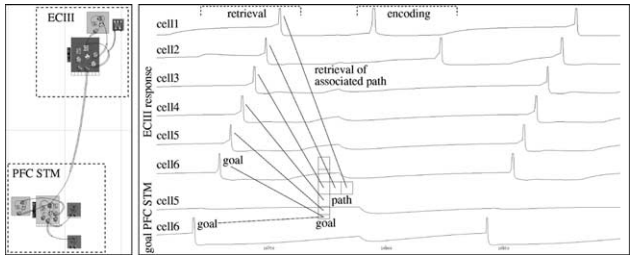

Fig. 9. Known goal locations in the T-maze, stored in PFC, cause spiking in ECIII which spreads from the goal along known paths. The Goal activity elicits spiking in ECIII at its retrieval phase in each theta cycle. Retrieval spreads away from the goal along associated known paths via recurrent fibres previously strengthened by LTP.

in the model show that activity is limited in CA3 and CA1, but broad in ECIII, consistent with that data (Fig. 12) (Cannon et al., 2002).

2.3. Phase differences in theta modulation between prefrontal and medial temporal regions enhance functional interaction

Given the need to propagate retrieval spikes from a source population (e.g. ECII STM) at the encoding phase of a target population (e.g. CA3), it is clear that specific theta phase relationships must exist between the different regions. In Fig. 13, the magnitude of modulation at theta frequency in cellular membrane potentials and synaptic transmission ratios shows the phase differences that allow PFC, ECII, ECIII, CA3 and CA1 regions to interact successfully.

2.4. An alternative model implementation provides additional support for our hypothesis of the role of theta in navigation

Our theoretical model was also evaluated in a different implementation with similar functional properties using the KIn–NeSS simulation package. We selected single-compartment integrate-and-fire neurons in the CATACOMB implementation above. In contrast, in KInNeSS each neuron consists of a dendrite compartment governed by the voltage equation and the standard approximation of the cable equation (Bower & Beeman, 1995), plus a soma compartment based on the canonical model of the type 1 neuron (Ermentrout & Kopell, 1986; Hoppensteadt & Izhikevich, 1998). The details of the cell model are provided elsewhere (Gorchetchnikov & Hasselmo, 2003), while we now emphasize the differences between the two system implementations.

Where a septal source of modulation at theta frequency is assumed in the CATACOMB simulation, septal cell activity in the KInNeSS simulation is derived from recurrent interactions with the hippocampal area via fornix connections. Similar rhythmic behaviour and synchronization at theta frequency emerges in both simulations. This synchronization affects neuronal activity within the hippocampus, as well as phase-locked spiking activity in the hippocampus and PFC. In the simulation above, ECIII learned episodes of place cell activity that represent navigable paths. In the KInNeSS implementation, this requisite knowledge for navigation is provided as sets of restrictions on certain locations. Restrictions range from absolute prohibition (an unreachable location behind a wall) to various degrees of aversion (e.g. a representation of the scent of a predator), established during the exploration of the environment. Unexplored locations remain unrestricted, so that the model can simulate path integration behaviour and attempt to take short-cuts through those locations.

The representation of space in the CATACOMB simulation is continuous in terms of the appearance of place fields during exploration and the motion directed by activity in CA1. In the KInNeSS simulation, space is mapped onto a two-dimensional grid of cells that form the centers of place fields. This approach is based on results by Trullier and Meyer (2000), who demonstrated equivalent performance with both a priori uniformly assigned place fields and place fields that were generated dynamically to

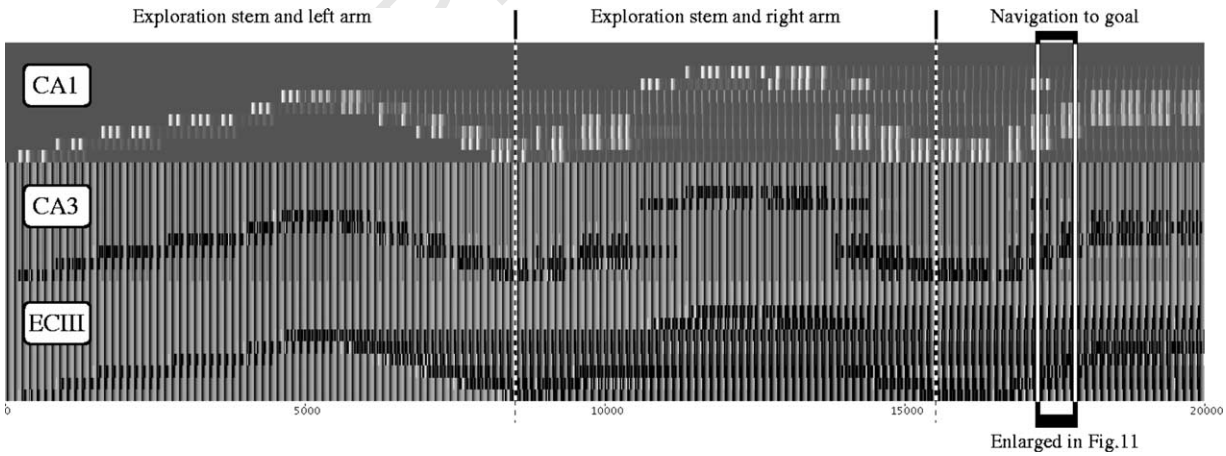

Fig. 10. Spiking in CA1 predicts desired next location on a shortest path to a goal. Retrieval in ECIII of known paths to the goal combines with retrieval in CA3 of known adjacent locations. Where the two most rapidly coincide, a spike is elicited in CA1.

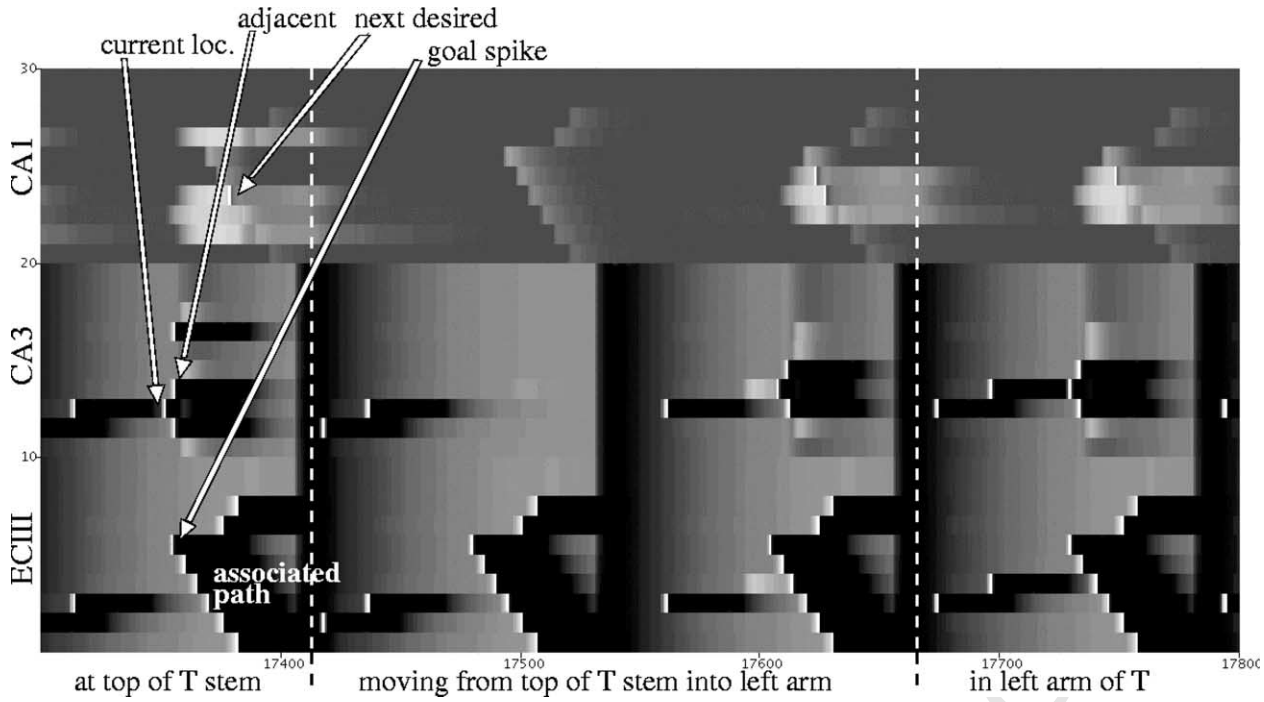

Fig. 11. Retrieved knowledge of paths in the T-maze and of adjacent locations causes predictive spiking that drives movement. The virtual rat is at the top of the T-maze stem. Place cell spiking at the current location retrieves three adjacent locations (down, left, right) Synaptic responses to these spikes in CA3 cause a slowly decaying depolarization in the corresponding cells in CA1. At the same time, a goal-spike received in ECIII elicits retrieval of known paths from the goal location at the end of the left arm of the T-maze to the bottom of the stem, as well as to the end of the right arm. This activity coincides first with the adjacency reponse in CA1 representing the adjacent location in the left arm of the T-maze. The resulting spike in CA1 predicts the next desired location and suppresses further spiking through lateral inhibition. The virtual rat consequently moves into the left arm of the T-maze. There, a new current location spike in CA3 continues the predictions that drive movement.

reflect the complexity of the environment. Similarly, the two different implementations of our theoretical model demonstrate its robust performance, regardless of the method chosen to generate place fields and their layout.

Finally, the KInNeSS simulation enables us to examine performance in the case of multiple goals in a single environment (a linear track, a maze or an open field). Recurrent inhibition in CA1 is critical to the behaviour of the model, as it underlies the competitive choice of a desired move towards the nearest goal. In a linear track or maze with narrow arms, the recurrent CA1 inhibition of the KInNeSS simulation results in optimal winner-take-all competition (Fig. 14). The situation is more complicated in the open field, as multiple CA1 cells of neighbouring locations in the general direction of the goal can spike before the onset of the competitive inhibition. This will be resolved through the inclusion of a head-direction system that takes output from CA1 and computes an average direction that can consistently drive the virtual rat's

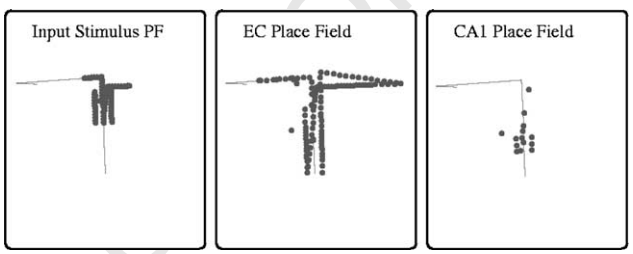

Fig. 12. The model replicates differences in place field sizes observed in experiments. These plots show the locations in the T-maze at which individual example place cells in each region were activated. The place fields generated by spiking activity in ECIII are broad and overlapping. By contrast, the place fields in CA1 are small and well defined.

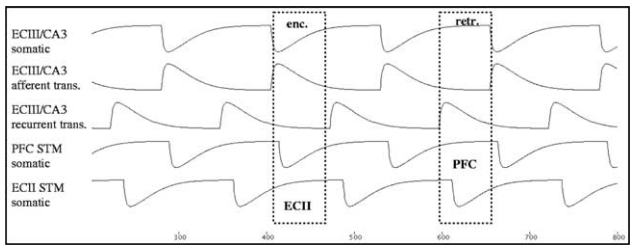

Fig. 13. The function of the model requires specific phase relationships between theta rhythmic modulation of membrane potentials and synaptic transmission in different locations. During encoding (enc.), patterns spike in the STM buffer during the depolarized membrane phase in ECII, and are transmitted to ECIII and CA3 at their encoding phases (less depolarized membrane potential) via afferent synapses in their phase of strong transmission. During retrieval (retr.), the phase of membrane depolarization in PFC activates a goal-representation that is transmitted to ECIII. During retrieval, recurrent fibre synapses are at the phase of full strength and cause spiking in ECIII and CA3 neurons, which are in their phase of somatic depolarization. This leads to the retrieval of known associations. ECIII retrieves known paths associated with the goal-location. These are combined in CA1 with associations retrieved in CA3 that indicate known place fields adjacent to current location of the virtual rat.

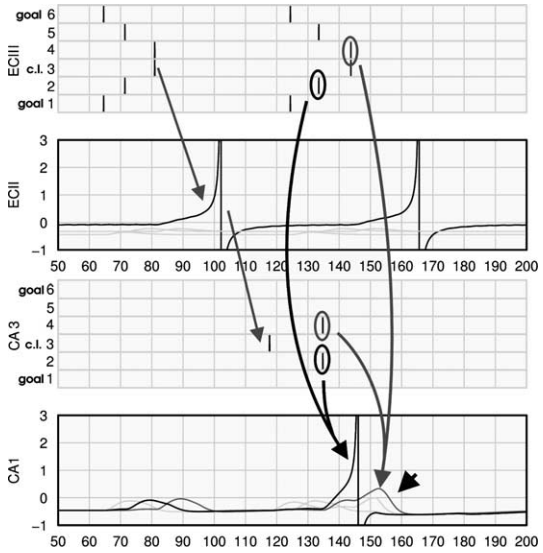

Fig. 14. Competitive decision making between two goals on a linear track with six place fields. ECIII and CA3 panels show raster plots of cells according to their locations on the track and their activity over time. The two goals and current location (c.l.) are indicated. ECII and CA1 panels show the recorded cell activity in respective areas. The spikes in ECIII and CA3 marked with black circles almost coincide in time and produce a rapid response in CA1—a correct choice. The spikes in the gray circles would produce an incorrect choice. They are further separated in time, causing a delayed response in CA1 that is suppressed by recurrent inhibition (black arrowhead).

movements. Empirical measurements support the existence of such head-direction cells in regions receiving input from CA1, such as the parasubiculum (Taube, 1995).

3. Conclusion

The model successfully enabled a virtual rat to learn its environment and to navigate on a shortest path to a food-reward goal encountered during exploration. By modeling with a focus on the behavioural task and interactions with the environment rather than individual regional functions, new insight was gained into significant requirements imposed by the dynamics involved. The manifestation of problems specific to the behaviour allowed us to link behavioural data with physiological data.

These insights support the hypothesis that theta modulation is a functional necessity in an integrative model of a realistic spatial navigation task: (1) It can impose distinct phases of ongoing encoding and retrieval. (2) It can maintain a STM buffer crucial to episodic learning. (3) It synchronizes spiking so that specific phase relationships can be achieved for the propagation of activity between cortical regions. The inspection of simulated spiking neurons resembles recording of place cells from hippocampus in awake behaving rats, allowing phenomena such as phase precession to be investigated and evaluated in terms of existing data. The implementation of a STM buffer made learning less dependent on the speed of motion and the timing of place cell spiking, while it maintained the order in which place cell spikes appeared for sequence learning. Finally, place fields in model CA1 were found to be smaller and more defined than those in ECIII.

In our upcoming work, we are adding a system of neocortical 'minicolumns' that will enable the learning of conditional relationships and task rules. This system will interact with the hippocampus for the retrieval of episodic memory items. A model that extends beyond the hippocampus will allow explicit separation of the components of behavior requiring the hippocampus from those components which are spared after hippocampal lesions.

Acknowledgements

The CATACOMB simulations described here and information about CATACOMB are available on our Computational Neurophysiology website at http://askja.bu.edu.

Supported by NIH R01 grants DA16454, MH60013 and MH61492 to Hasselmo and by Conte Center Grant MH60450.

References

Alonso, A., Gaztelu, J., Bruno, W., Jr, & Garcia-Austt, E. (1987). Crosscorrelation analysis of septohippocampal neurons during theta-rhythm. *Brain Research*, *413*, 135–146.

Barnes, C., McNaughton, B., Mizumori, S., Leonard, B., & Lei-Huey, L. (1990). Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. *Progress in Brain Research*, *83*, 287–299.

Bi, G., & Poo, M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience*, *18*(24), 10464–10472.

Bower, J., & Beeman, D. (1995). *The Book of GENESIS. Exploring realistic neural models with the GEneral Neural SImulation System*. New York: Springer.

Bragin, A., Jando, G., Nadasdy, Z., & Hetke, J. (1995). Gamma (40–100 Hz) oscillation in the hippocampus of the behaving rat. *Journal of Neuroscience*, *15*, 47–60.

Brankack, J., Stewart, M., & Fox, S. (1993). Current source density analysis of the hippocampal theta rhythm: associated sustained potentials and candidate synaptic generators. *Brain Research*, *615*(2), 310–327.

Brazhnik, E., & Fox, S. (1999). Action potentials and relations to the theta rhythm of septohippocampal neurons in vivo. *Experimental Brain Research*, *127*, 244–258.

Buzsáki, G., Leung, L., & Vanderwolf, C. (1983). Cellular bases of hippocampal EEG in the behaving rat. *Brain Research*, *287*, 139–171.

Cannon, R., Hasselmo, M., & Koene, R. (2002). From biophysics to behaviour: Catacomb2 and the design of biologically plausable models for spatial navigation. *Neuroinformatics*, *1*(1), 3–42.

Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M., & Tanila, H. (1999). The hippocampus, memory, and place cells: Is it spatial memory or a memory space? *Neuron*, *23*, 209–226.

Ermentrout, G., & Kopell, N. (1986). Parabolic bursting in an excitable system coupled with slow oscillation. *SIAM Journal of Applied Mathematics*, *46*, 233–252.

Fox, S. (1989). Membrane potential and impedance changes in hippocampal pyramidal cells during theta rhythm. *Experimental Brain Research*, *77*, 283–294.

Fox, S., Wolfson, S., & Ranck, J. (1986). Hippocampal theta rhythm and the firing of neurons in walking and urethane anesthetized rats. *Experimental Brain Research*, *62*, 495–508.

Frank, L., Brown, E., & Wilson, M. (2000). Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron*, *27*(1), 169–178.

Gorchetchnikov, A., & Hasselmo, M. (2003). Timing of consecutive traveling pulses in a model of entorhinal cortex. *Proceedings of the International Joint Conference on Neural Networks*, in press.

Hasselmo, M., Bodelon, C., & Wyble, B. (2002). A proposed function for hippocampal theta rhythm: Separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Computation*, *14*(4), 793–817.

Hölscher, C., Anwyl, R., & Rowan, M. (1997). Stimulation on the positive phase of hippocampal theta rhythm induces long-term potentiation that can be depotentiated by stimulation on the negative phase in area CA1 in vivo. *Journal of Neuroscience*, *17*(16), 6470–6477.

Hoppensteadt, F., & Izhikevich, E. (1998). *Weakly connected neural networks*. New York: Springer.

Huerta, P., & Lisman, J. (1993). Heightened synaptic plasticity of hippocampal ca1 neurons during a cholinergically induced rhythmic state. *Nature*, *364*, 723–725.

Jensen, O., Idiart, M., & Lisman, J. (1996). Physiologically realistic formation of autoassociative memory in networks with theta/gamma oscillations: Role of fast NMDA channels. *Learning & Memory*, *3*, 243–256.

Jensen, O., & Lisman, J. (1996). Novel lists of 7 ± 2 known items can be reliably stored in an oscillatory short-term memory network: Interaction with long-term memory. *Learning & Memory*, *3*, 257–263.

Klink, R., & Alonso, A. (1997). Morphological characteristics of layer ii projection neurons in the rat medial entorhinal cortex. *Hippocampus*, *7*, 571–583.

Levy, W., & Stewart, D. (1983). Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience*, *8*(4), 791–797.

Lisman, J., & Idiart, M. (1995). Storage of 7 ± 2 short-term memories in oscillatory subcylces. *Science*, *267*, 1512–1515.

McNaughton, B., Barnes, C., & O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research*, *52*(1), 41–49.

M'Harzi, M., Palacios, A., Monmaur, P., Houcine, O., & Delacour, J. (1987). Effects of selective lesions of hippocampal connections on learning set in the rat. *Physiology and Behavior*, *40*, 181–188.

Muller, R. U., & Kubie, J. (1989). The firing of hippocampal place cells predicts the future position of freely moving rats. *Journal of Neuroscience*, *9*, 4101–4110.

O'Keefe, J., & Recce, M. (1993). Phase relationship between hippocampal place units and the hippocampal theta rhythm. *Hippocampus*, *3*, 317–330.

Pavlides, C., Greenstein, Y., Grudman, M., & Winson, J. (1988). Long-term potentiation in the dentate gyrus is induced preferentially on the positive phase of theta-rhythm. *Brain Research*, *439*(1/2), 383–387.

Quirk, G., Muller, R., Kubie, J., & Ranck, J., Jr (1992). The positional firing properties of medial entorhinal neurons: Description and comparison with hippocampal place cells. *The Journal of Neuroscience*, *12*, 1945–1963.

Skaggs, W., McNaughton, B., Wilson, M., & Barnes, C. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, *6*, 149–172.

Stewart, M., & Fox, S. (1990). Do septal neurons pace the hippocampal theta rhythm? *Neuron*, *13*, 163–168.

Taube, J. (1995). Place cells recorded in the parasubiculum of freely moving rat. *Hippocampus*, *5*, 569–583.

Trullier, O., & Meyer, J.-A. (2000). Animat navigation using a cognitive graph. *Biological Cybernetics*, *83*, 271–285.

Wyble, B., Hyman, J., Goyal, V., & Hasselmo, M. (2001). Phase relationship of ltp induction and behavior to theta rhythm in the rat hippocampus. *Society for Neuroscience Abstracts*, *27*, 53719.

Wyble, B., Linster, C., & Hasselmo, M. (1997). Evoked synaptic potential size depends on the phase of theta rhythm in rat hippocampus. *Society of Neuroscience Abstracts*, *23*, 508.