LETTER Communicated by William Skaggs

Changes in $GABA_B$ Modulation During a Theta Cycle May Be Analogous to the Fall of Temperature During Annealing

Vikaas S. Sohal
Michael E. Hasselmo
Department of Psychology and Program in Neuroscience, Harvard University, Cambridge, MA 02138, U.S.A.

Changes in $GABA_B$ modulation may underlie experimentally observed changes in the strength of synaptic transmission at different phases of the theta rhythm (Wyble, Linster, & Hasselmo, 1997). Analysis demonstrates that these changes improve sequence disambiguation by a neural network model of CA3. We show that in the framework of Hopfield and Tank (1985), changes in $GABA_B$ suppression correspond to changes in the effective temperature and the relative energy of data terms and constraints of an analog network. These results suggest that phasic changes in the activity of inhibitory interneurons during a theta cycle may produce dynamics that resemble annealing. These dynamics may underlie a role for the theta cycle in improving sequence retrieval for spatial navigation.

1 Introduction

Manipulations that abolish the hippocampal theta rhythm impair spatial navigation (Winson, 1978), but the computational function of theta is not known. Previous experiments have found that the amplitude of evoked field potentials due to neuronal spiking activity changes during different phases of the theta rhythm (Rudell, Fox, & Rank, 1980; Buszaki, Grastyan, Czopf, Kellenyi, & Prohaska, 1981; Rudell & Fox, 1984). This could result from changes in presynaptic and/or postsynaptic inhibition during a theta cycle. More recent evidence suggests that the strength of synaptic transmission also depends on the phase of theta rhythm (Wyble et al., 1997), another possible indication of phasic changes in presynaptic inhibition.

Phasic changes in presynaptic inhibition could result from the rhythmic activation of presynaptic $GABA_B$ receptors, which selectively suppress recurrent synaptic transmission in piriform cortex (Tang & Hasselmo, 1994) and hippocampus (Ault & Nadler, 1982; Colbert & Levy, 1992). Stewart and Fox (1990) have hypothesized, and in vivo recordings have shown (Buszaki & Eidelberg, 1983; Fox et al., 1986; Skaggs, McNaughton, Wilson, & Barnes, 1996) that some interneurons tend to fire near a preferred phase of theta. Phasic firing by even a subset of interneurons may rhythmically activate presynaptic $GABA_B$ receptors, producing the changes in synaptic strength

observed in vivo (Wyble et al., 1997). Paired pulse depression, thought to result from activation of presynaptic $GABA_B$ receptors on interneurons, has a rise time and fast component of decay (Otis, Dekoninck, & Mody, 1993) that are compatible with theta frequency oscillations in $GABA_B$-mediated synaptic suppression.

These rhythmic changes could be important for retrieval of sequences of activity in the hippocampus. Evidence suggests that the hippocampus may represent a series of adjacent locations by a sequence of successively firing place cells (Tsodyks, Skaggs, Sejnowski, & McNaughton, 1996; Jensen & Lisman, 1996; Wallenstein & Hasselmo, 1997). This evidence includes the phenomenon of theta phase precession: during a theta cycle, place cells fire in the order in which their place fields are encountered along a path (O'Keefe & Recce, 1993; Skaggs et al., 1996).

In a neural network model of hippocampal region CA3, we find that changes in $GABA_B$ suppression during a theta cycle improve sequence disambiguation: the retrieval of sequences with identical starting points. We show that these changes in the level of $GABA_B$ modulation (within a theta cycle) correspond to changes in (1) the effective temperature and (2) the relative energy of data terms and constraints of an analog network (using the terminology of Hopfield & Tank, 1985). Data terms and constraints correspond to afferent input and recurrent connections, respectively, in our model.

2 Analysis

2.1 The Problem of Sequence Disambiguation. When two sequences with the same starting components, $\{\mathbf{X_1}, \ldots, \mathbf{X_m}\}$ and $\{\mathbf{X_1}, \ldots, \mathbf{X_i}, \mathbf{Y_{i+1}}, \ldots, \mathbf{Y_m}\}$, are stored in a network, it is ambiguous whether $\mathbf{X_i}$ should be followed by $\mathbf{X_{i+1}}$ or $\mathbf{Y_{i+1}}$. The only information that can resolve this ambiguity is knowledge of the desired end point, $\mathbf{X_m}$ or $\mathbf{Y_m}$. If the patterns of activity are neural representations of locations, then disambiguating forked sequences using knowledge of the desired end point corresponds to the everyday problem of deciding which way to turn at an intersection based on one's destination.

2.2 A Model of Sequence Disambiguation. The simplest neural realization of this problem, shown in Figure 1, contains two sequences: $\{a_1, a_2\}$ and $\{a_1, a_3\}$. When a_1 becomes active, the pattern of activity that should follow a_1 is ambiguous. We assume that during disambiguation of the sequences $\{a_1, a_2\}$ and $\{a_1, a_3\}$, a_2 receives a small amount of afferent input. This input could arrive in CA3 from the dentate gyrus via the mossy fibers or from the entorhinal cortex via the perforant path. It represents the knowledge that a_2 is the desired goal and should bias the network so that it completes the sequence $\{a_1, a_2\}$ rather than $\{a_1, a_3\}$. This assumption predicts that place cells representing the location of a goal should receive some biasing input

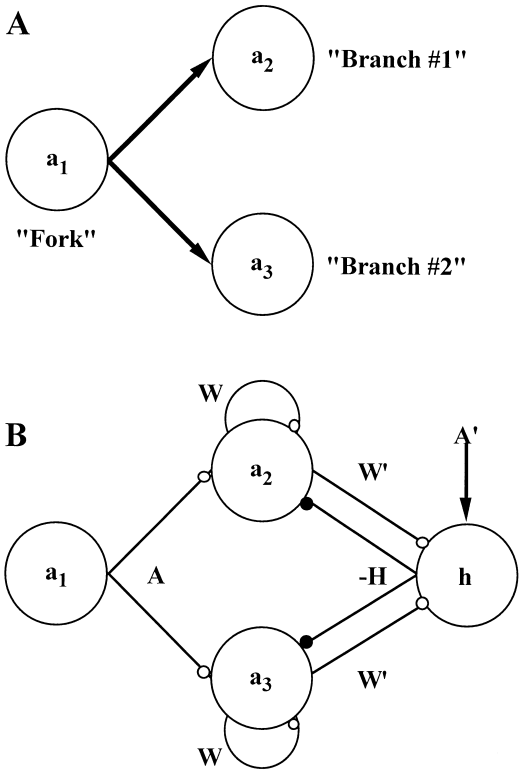

Figure 1: (A) The simplest neural representation of a forked sequence, composed of the sequences $\{a_1, a_2\}$ and $\{a_2, a_3\}$. a_1 represents the point of the "fork," and a_2 and a_3 represent the "branches" of the forked sequence. (B) The connectivity in a simple network that has stored this forked sequence.

while a rat searches for that goal. With strong enough biasing input to a_2, the network can easily disambiguate $\{a_1, a_2\}$ from $\{a_1, a_3\}$ when a_2 is given a biasing input. Here we focus on the optimal characteristics of a network that can perform this disambiguation with weak bias.

We study such a simplified network because the decision selectively to activate a_2 but not a_3 corresponds to the decision to activate neurons representing X_{i+1} but not those representing Y_{i+1} during the previously described sequence disambiguation problem. Elsewhere (Sohal & Hasselmo, 1997) we show that the conditions that optimize the decision to activate a_2 selectively in this simplified network can both be produced by $GABA_B$ receptors with

biologically plausible dynamics and significantly improve disambiguation of multiple sequences, composed of many items each, in more complex networks.

The network parameters, shown in Figure 1B, are: A, the strength of excitation from a_1 to a_2 and a_3; A_{bias}, afferent input to a_2; A', afferent input to the inhibitory interneuron; W, the strength of recurrent excitatory connections from a_2 to a_2 and a_3 to a_3; W', the strength of excitatory connections from a_2 and a_3 to the interneuron; $-H$, the strength of the inhibitory connections from this interneuron to a_2 and a_3; h, activation of the model interneuron; η, the rate of passive decay for the activities, a_2, a_3, and h; and θ, the threshold of the neurons represented by a_2, a_3, and h.

The parameters $k(t)$ and $k'(t)$ represent the $GABA_B$ suppression of recurrent synaptic transmission. All recurrent excitatory connections in the network (A, W, and W') are multiplied by $k(t)$. So $[1 - k(t)]$ represents the suppression of recurrent excitation at time t. Similarly, k' multiplies recurrent inhibition in the network (H), so $[1 - k'(t)]$ represents the suppression of recurrent inhibition.

The activation of each model neuron is a continuous variable representing the mean membrane potential of a pool of real neurons. The equations for the evolution of a_2, a_3, and h have been derived by averaging over short timescales and many neurons, and are similar to those used elsewhere (Wilson & Cowan, 1972):

$$\begin{aligned}
\dot{a}_2 &= -\eta a_2 + kA + A_{bias} + kW(a_2 - \theta)_+ - k'H(h - \theta)_+ \\
\dot{a}_3 &= -\eta a_3 + kA + kW(a_3 - \theta)_+ - k'H(h - \theta)_+ \\
\dot{h} &= -\eta h + A' + kW'(a_2 - \theta)_+ + kW'(a_3 - \theta)_+,
\end{aligned} \tag{2.1}$$

where $(x)_+ = x$ if $x > 0$ and 0 otherwise. As long as $A_{bias} > 0$, a_2 will cross threshold before a_3, giving $a_2 > \theta$ and $a_3 < \theta$. Assume the interneuron will be activated before $a_3 > \theta$ (we have found that this is the case for networks that perform sequence disambiguation). Then the activities are in the regime $a_2 > \theta$, $a_3 < \theta$, and $h > \theta$, in which network dynamics reduce to the linear system:

$$\begin{aligned}
\dot{a}_2 &= (kW - \eta)a_2 - k'Hh + kA + A_{bias} + k'H\theta \\
\dot{a}_3 &= -\eta a_3 - k'Hh + kA + k'H\theta \\
\dot{h} &= -\eta h + kW'a_2 + A' - kW'\theta.
\end{aligned}$$

We measured the performance of sequence disambiguation by this system as $a_2(t_{final}) - a_3(t_{final})$, where t_{final} is the time at which sequence disambiguation ends. Thus, $a_2(t_{final}) - a_3(t_{final})$ measures the differential completion of the two sequences, $\{a_1, a_2\}$ and $\{a_2, a_3\}$, at the end of sequence disambiguation.

2.3 Optimal Network Parameters During Sequence Disambiguation. We found the $k(t)$ and $k'(t)$ that maximized the quantity $a_2(t_{final}) - a_3(t_{final})$ using the maximum principle (Pontryagin, Boltyanskii, Gamkrelidze, & Mishchenko, 1962). The maximum principle determines necessary conditions on the functions $k(t)$ and $k'(t)$ to maximize the difference $a_2(t_{final}) - a_3(t_{final})$. These conditions uniquely specify $k(t)$ and $k'(t)$:

$$k = \begin{Bmatrix} k_{\min} \text{ if } t < t_{final} - t_1, \\ k_{\max} \text{ if } t > t_{final} - t_1 \end{Bmatrix}$$

$$\text{where } t_1 = \frac{\pi}{\sqrt{4HW'k'_{\max}k_{\max} - W^2k^2_{\max}}}$$

$$k' = \begin{Bmatrix} k'_{\min} \text{ if } t < t_{final} - t_1 - t_2, \\ k'_{\max} \text{ if } t > t_{final} - t_1 - t_2 \end{Bmatrix}$$

$$\text{where } t_2 = \frac{\pi}{\sqrt{4HW'k'_{\max}k_{\min} - W^2k^2_{\min}}}.$$

This result is true as long as:

$$t_{final} > 0$$

$$0 \le k_{\min} \le k(t) \le k_{\max} \le 1$$

$$0 \le k'_{\min} \le k'(t) \le k'_{\max} \le 1.$$

(The quantities inside the square roots are negative if and only if all eigenvalues of the linearized system are real. We assume that damped oscillations, produced by complex eigenvalues, represent the typical behavior of hippocampal neurons and therefore confine our analysis to this regime.) Thus, optimal sequence disambiguation always occurs when the suppression of recurrent excitation and inhibition are step functions that decrease from their maximum to their minimum values during the sequence disambiguation task. The maximum principle also shows that A_{bias} should be fixed at its maximum value for optimal sequence disambiguation. Note that this optimality does not depend on the specific values of $k_{\max}$, $k_{\min}$, or t_{final}.

2.4 Linearly Decreasing $GABA_B$ Suppression Significantly Improves Sequence Disambiguation. Because step function decreases in the suppression of synaptic transmission seem biologically unrealistic, we studied whether more plausible decreases in the suppression of synaptic transmission might also significantly improve sequence disambiguation in our simplified network. Any sufficiently smooth $k(t)$, $k'(t)$, such as the changes in synaptic strength observed in vivo (Wyble et al., 1997), may be locally approximated by linear functions. Therefore, we solved for the final state

of the system described above and studied how that final state changed if the suppression of recurrent synaptic transmission decreased linearly while that system disambiguated two sequences.

$k(t)$ and $k'(t)$ took the form:

$$k(t) = k'(t) = \alpha t + \beta$$

where $\alpha(100 \text{ msec}) + \beta = 1$ (100 msec represents about one-half of a theta cycle, and α measures the amount by which the suppression of recurrent synaptic transmission decreases during the sequence disambiguation task).

If a_1 is active and a_2 is receiving biasing input, then the system has only two possible final states: for low A_{bias}, both a_2 and a_3 are active, whereas for A_{bias} sufficiently large, only a_2 is active in the final state, so sequence disambiguation is successful. We found the minimum A_{bias} necessary for successful sequence disambiguation as a function of α. The results, shown in Figure 2, demonstrate that increasing α, corresponding to a greater change in the suppression of recurrent synaptic transmission, allows successful sequence disambiguation using weaker biasing inputs, that is, using less information about the desired goal. This confirms that biologically plausible changes, such as linear decreases, in the suppression of synaptic transmission can also significantly improve sequence disambiguation.

In simulations (Sohal & Hasselmo, 1997), theta frequency septal input to interneurons entrains oscillations in $GABA_B$ receptor activation such that $GABA_B$ suppression falls during the phase of pyramidal cell activity. Consistent with the preceding analysis, this decrease improves sequence disambiguation. Furthermore, the tendency toward preferred phases of interneuron firing in these simulations is similar to those observed in vivo (Skaggs et al., 1996). Thus, oscillations in $GABA_B$ suppression, which may underlie experimentally observed changes in the strength of synaptic transmission (Wyble et al., 1997), may also optimize sequence disambiguation. Alternatively, synaptic depression due to previous transmission (Markram & Tsodyks, 1996) may underlie the observed changes in synaptic strength, although this form of depression has been found only in neocortex, and hippocampal pyramidal cells are more likely to show facilitation than depression.

3 Relation to Annealing

In our model, $GABA_B$ suppression is initially high during sequence disambiguation. As a result, recurrent excitation and inhibition are weak, so neurons representing multiple sequences are active. However, as $GABA_B$ suppression falls, recurrent excitation and inhibition become stronger, so that fewer neurons, representing only one sequence, are active while the remaining neurons are inhibited. Thus, as the level of $GABA_B$ modulation falls, the network shifts from sampling many possible states to selecting the

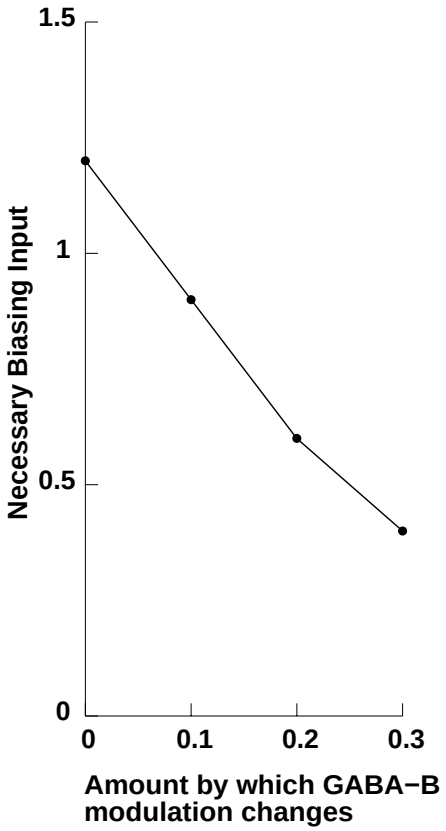

Figure 2: The minimum biasing input, A_{bias}, required for successful sequence disambiguation as a function of α. In every case, the level of GABAergic modulation reaches the same value at the end of the sequence disambiguation task, but during the task, the level of GABAergic modulation decreases with slope α. Thus, α measures the rate at which the level of GABAergic modulation, and hence the suppression of recurrent synaptic transmission, declines during the sequence disambiguation task. This figure shows that greater rates of decrease in GABAergic modulation permit successful sequence disambiguation with weaker biasing inputs.

optimal state. A similar shift also occurs during annealing as the temperature of a system decreases. The system samples states of varying energy at high temperatures, but selects the minimum energy state at low temperatures. This suggests that a falling level of $GABA_B$ modulation may improve

sequence disambiguation in the same way that a decreasing temperature drives annealing.

Hopfield and Tank (1985) showed that the equilibrium solution in an analog network equals the effective field solution for a Boltzmann machine (Hinton & Sejnowski, 1986) at a temperature determined by the gain width. Similarly, the equilibrium solution to our network is the same as the effective field solution in a Boltzmann machine, and the fall in $GABA_B$ modulation that occurs in our network is equivalent to decreasing both the effective temperature and the energy of data terms relative to constraints. Consider the energy function:

$$E = \rho_{\text{exc}}\left[-\frac{W}{2}((a_2-\theta)_+^2+(a_3-\theta)_+^2)-A(a_2+a_3)\right] + \rho_{inh}W'H(a_2a_3)-\rho_{aff}A_{bias}a_2, \tag{3.1}$$

where ρ_{exc}, ρ_{inh}, and ρ_{aff} measure the relative contributions of recurrent excitation, recurrent inhibition, and afferent input, respectively, to the energy. Following Hopfield and Tank (1985), we find the effective field solution for a Boltzmann machine in which the a_i are Ising spins restricted to the values $a_i = a^{\max}$ or $a^{\min}$, the temperature is τ, and the energy, E, is given by equation 3.1. Rescale variables so that $a^{\min} = 0 < \theta < a^{\max} = 1$. Then the expected value for a_i is:

$$\langle a_i\rangle = \frac{e^{H_i/\tau}}{1+e^{H_i/\tau}}, i = 2 \text{ or } 3, \tag{3.2}$$

where:

$$H_2 = \rho_{exc}\lfloor -A - W(a_2-\theta)_+\rfloor + \rho_{inh}HW'a_3 - \rho_{aff}A_{bias}$$

$$H_3 = \rho_{exc}[-A - W(a_3-\theta)_+] + \rho_{inh}HW'a_2.$$

Equation 3.2 reduces to:

$$\langle a_i\rangle \cong \frac{H_i}{\tau}, \tag{3.3}$$

if

$$H_i \ll \tau. \tag{3.4}$$

The condition in equation 3.4 is true when the $a^{\max}$ state is far from the saturating regime of the equation, that is, when the mean activity $\langle a_i\rangle$ still increases linearly with synaptic input. This is true (Barkai & Hasselmo, 1994) over the range of firing frequencies observed in vivo (O'Keefe & Recce, 1993; Skaggs et al., 1996).

Now consider a simplification of the neural network shown in Figure 1 in which the interneuron is replaced by linear inhibition:

$$\begin{aligned} \dot{a}_2 &= -\eta a_2 + kA + A_{bias} + kW(a_2 - \theta)_+ - kk'HW'a_3 \\ \dot{a}_3 &= -\eta a_3 + kA + kW(a_3 - \theta)_+ - kk'HW'a_2. \end{aligned} \tag{3.5}$$

The steady-state solution to equation 3.5 is:

$$\begin{aligned} a_2 &= \frac{k}{\eta}\left[\frac{1}{k}A_{bias} + A + W(a_2 - \theta)_+ - k'HW'a_3\right] \\ a_3 &= \frac{k}{\eta}[A + W(a_3 - \theta)_+ - k'HW'a_2]. \end{aligned} \tag{3.6}$$

The steady-state solution for this network (see equation 3.6) equals the effective field solution (see equation 3.3) for which the temperature, τ_{eff}, equals η/k and the relative energies of afferent input, recurrent excitation, and recurrent inhibition are given by:

$$\begin{aligned} \rho_{aff} &= \frac{1}{k} = \eta\tau_{eff} \\ \rho_{exc} &= 1 \\ \rho_{inh} &= k' = \frac{1}{\eta\tau_{eff}} \text{ (assuming } k' \approx k). \end{aligned}$$

Falling $GABA_B$ modulation produces a rise in k and is therefore equivalent to annealing, in which there is a simultaneous fall in both the temperature and the energy of afferent input relative to recurrent excitation and inhibition.

3.1 Why Do the Relative Energies of Afferent and Recurrent Inputs Change? Because activation of $GABA_B$ receptors selectively suppresses recurrent but not afferent connections (Ault & Nadler, 1982; Colbert & Levy, 1992), the energy of afferent input relative to recurrent excitation and inhibition decreases as the level of $GABA_B$ modulation falls. Two analyses showed that the selectivity of $GABA_B$ suppression optimizes network performance. First, by applying the maximum principle, we found that the strength of recurrent synapses should increase during the sequence disambiguation task, but afferent input should remain maximal throughout. Second, suppose that the relative energies of afferent input and recurrent excitation and inhibition were fixed. To model this, we multiplied afferent input by the same factor, $\alpha t + \beta$, as recurrent connections. Then the necessary biasing input, shown in Figure 2, decreases less with increasing α; that is, a fall in the level of $GABA_B$ modulation produces less of an improvement in sequence disambiguation. (However, note that sequence disambiguation still improves when the strengths of both afferent input and recurrent connections change. Thus, phasic changes in postsynaptic excitability could

improve sequence disambiguation, but this improvement would be smaller than the improvement caused by changes in $GABA_B$ suppression.)

The decrease in the energy of afferent input relative to recurrent inputs can be understood intuitively. Afferent input contributes what Hopfield and Tank (1985) call "data terms" to the energy function. These represent the locations of the starting point and goal. The energy contributions of recurrent excitation and inhibition represent the constraints that one continuous sequence should be active and only one sequence should be active, respectively. Data terms determine the general location of the global minimum of the energy function, whereas constraints refine the global minimum but also introduce local minima (Lin & Lee, 1995). The level of $GABA_B$ modulation is initially high, so that afferent input and the corresponding data terms dominate the energy function, causing the network to converge to the neighborhood of the global energy minimum. As the level of $GABA_B$ modulation falls, recurrent connections become stronger, enforcing constraints, so that the network zeros in on the global energy minimum in which exactly one continuous sequence is active. Because constraints are enforced only after the network has converged to the neighborhood of the global energy minimum, the network is less prone to become trapped in local minima.

Thus, the correspondence between annealing and a fall in $GABA_B$ suppression elucidates specific mechanisms by which $GABA_B$ modulation improves sequence dismabiguation: falling $GABA_B$ suppression both shifts the network from sampling multiple states to selecting the best one, and increases the strength of constraints relative to data terms.

Figure 3 shows how these mechanisms improve the performance of the network described by equation 3.5. When the network includes linearly decreasing $GABA_B$ suppression, the initially strong $GABA_B$ suppression slows the subthreshold rise in activity. As a result, a_2 rises well above threshold, but recurrent inhibition prevents a_3 from crossing threshold before the end of the task ($t_{final} = 50$ msec). In contrast, in the absence of $GABA_B$ suppression, both a_2 and a_3 rise above threshold. When $GABA_B$ suppression is present but does not decrease, a_3 does remain subthreshold, but a_2 barely crosses threshold.

Figure 3 does not indicate whether there exists a constant, intermediate level of $GABA_B$ suppression that produces the same final state, $(a_2(t), a_3(t))$, as the linearly decreasing $GABA_B$ suppression, shown in Figure 3A. In fact, one such constant level of $GABA_B$ suppression does exist. However, because this constant, intermediate level of $GABA_B$ suppression does not produce the slow initial activity rise that results from falling $GABA_B$ suppression, it leads to successful sequence disambiguation over a smaller range of network parameters. For example, after strengthening recurrent excitation, W, while holding all other parameters fixed, only a_2 crosses threshold in the network with the linearly decreasing $GABA_B$ suppression, whereas both a_2 and a_3 cross threshold in the network with the constant, intermediate level of $GABA_B$ suppression. Thus, consistent with the results of optimization and

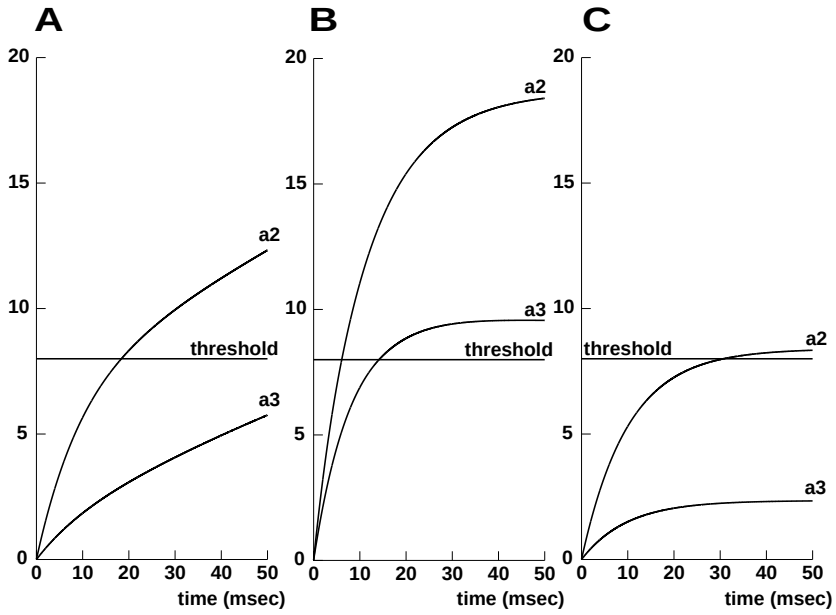

Figure 3: Activity in the network described by equation 3.5 when a_2 is receiving biasing input. (A) The network includes linearly decreasing $GABA_B$ suppression. Under these conditions, a_2 rises well above threshold. However, the initially strong $GABA_B$ suppression slows the subthreshold rise in activity, enabling recurrent inhibition to prevent a_3 from crossing the threshold before the end of the task ($t_{final} = 50$ msec). (B) $GABA_B$ suppression is absent, so syna= pses remain at their maximum strengths. As a result, both a_2 and a_3 rise above threshold. (C) $GABA_B$ suppression is present but does not decrease, so that synapses remain at their minimum strengths. In this case, although a_3 remains below threshold, synapses are so weak that a_2 barely rises above threshold.

the analogy with annealing, falling $GABA_B$ suppression leads to successful performance over a broader range of network parameters than do constant levels of $GABA_B$ suppression. We have found analogous results in simulations of more complex networks (Sohal & Hasselmo, 1997). While $GABA_B$ suppression falls, neurons that belong to a sequence whose end point receives biasing input (and which therefore correspond to a_2) become active, while neurons belonging to competing sequences (and corresponding to a_3) do not. However, when $GABA_B$ suppression is fixed at an intermediate value, both sets of neurons remain active.

4 Relation to Behavior

Spatial navigation may utilize hippocampal representations of location (Touretzky & Redish, 1996; Tsodyks et al., 1996). We have shown that the dynamics of sequence disambiguation may resemble those of annealing, and other studies have used simulated annealing by a Boltzmann machine for robot navigation (Lin & Lee, 1995). Thus, cycles of GABAergic modulation may contribute to making choices between multiple competing paths through the environment. Retrieval of specific temporal sequences among competing alternatives in the hippocampus may also be important for nonspatial tasks requiring relational processing. For example, Bunsey and Eichenbaum (1996) found that both control rats and rats with hippocampal lesions learn sample-associate pairs equally well, but that only control rats display transitivity. A network that represents sample-associate pairs as sequences can produce transitivity. In such a network, changes in $GABA_B$ would have a similar role for enhancing retrieval of the optimal relational representation from multiple possible representations.

These results suggest that one function of the theta rhythm may be to cause rhythmic changes in $GABA_B$ suppression of synaptic transmission. Our analysis demonstrates that such changes in $GABA_B$ suppression may produce dynamics that resemble annealing and improve sequence disambiguation in the hippocampus.

References

Ault, B., & Nadler, J. V. (1982). Baclofen selectively inhibits transmission at synapses made by axons of CA3 pyramidal cells in the hippocampal slice. *J Pharmacol Exp Ther, 223*, 291–297.

Barkai, E., & Hasselmo, M. E. (1994). Modulation of the input/output function of rat piriform cortex pyramidal cells. *J Neurophys, 72*, 644–658.

Bunsey, M., & Eichenbaum, H. (1996). Conservation of hippocampal memory function in rats and humans. *Nature, 379*, 255–257.

Buszaki, G., & Eidelberg, E. (1983). Phase relations of hippocampal projection cells and interneurons to theta activity in the anesthetized rat. *Brain Res, 266*, 334–339.

Buszaki, G., Grastyan, E., Czopf, J., Kellenyi, L., & Prohaska, O. (1981). Changes in neuronal transmission in the rat hippocampus during behavior. *Brain Res, 225*, 235–247.

Colbert, C. M., & Levy, W. B. (1992). Electrophysiological and pharmacological characterization of perforant path synapses on CA1: Mediation by glutamate receptors. *J Neurophys, 68*, 1–8.

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

Hinton, G. E., & Sejnowski, T. J. (1986). Learning in Boltzmann machines. In

D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 282–317). Cambridge, MA: MIT Press.

Hopfield, J. J., & Tank, D. W. (1985). Neural computation of decisions in optimization problems. *Biol Cybern, 52*, 141–152.

Jensen, O., & Lisman, J. E. (1996). Theta/gamma networks with slow NMDA channels learn sequences and encode episodic memory: Role of NMDA channels in recall. *Learning and Memory, 3*, 264–278.

Lin, C. T., & Lee, C. G. S. (1995). A multi-valued Boltzmann machine. *IEEE Trans Syst Man Cyber, 25*, 660–669.

Markram, H., & Tsodyks, M. (1996). Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature, 382*, 807–810.

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

Otis, T. S., Dekoninck, Y., & Mody, I. (1993). Characterization of synaptically elicited $GABA_B$ responses using patch-clamp recordings in rat hippocampal slices. *J Physiol, 463*, 391–407.

Pontryagin, L. S., Boltyanskii, V. G., Gamkrelidze, R. V., & Mishchenko, E. F. (1962). *The mathematical theory of optimal processes*. New York: Interscience Publishers.

Rudell, A. P., & Fox, S. E. (1984). Hippocampal excitability related to the phase of the theta rhythm in urethanized rats. *Brain Res, 294*, 350–353.

Rudell, A. P., Fox, S. E., & Rank, J. B. (1980). Hippocampal excitability phase-locked to the theta rhythm in walking rats. *Exp Neurology, 68*, 87–96.

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

Sohal, V. S., & Hasselmo, M. E. (1997). $GABA_B$ modulation improves sequence disambiguation in computational models of hippocampal region CA3. *Hippocampus*, in press.

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

Tang, A. C., & Hasselmo, M. E. (1994). Selective suppression of intrinsic but not afferent fiber synaptic transmission by baclofen in the piriform (olfactory) cortex. *Brain Res, 659*, 75–81.

Touretzky, D. S., & Redish, A. D. (1996). Theory of rodent navigation based on interacting representations of space. *Hippocampus, 6*, 247–270.

Tsodyks, M. V., Skaggs, W. E., Sejnowski, T. J., & McNaughton, B. L. (1996). Population dynamics and theta rhythm phase precession of hippocampal place cell firing: A spiking neuron model. *Hippocampus, 6*, 271–280.

Wallenstein, G. W., & Hasselmo, M. E. (1997). GABAergic modulation of hippocampal population activity: Sequence learning, place field development, and the phase precession effect. *J Neurophysiol, 78*, 393–408.

Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys J, 12*, 1–24.

Winson, J. (1978). Loss of hippocampal theta rhythm results in spatial memory deficit in the rat. *Science, 201*, 160–163.

Wyble, B. P., Linster, C., & Hasselmo, M. E. (1997). Evoked synaptic potential size depends on the phase of theta rhythm in rat hippocampus. *Soc Neurosci Abstr, 23*, 508.

Received June 2, 1997; accepted October 2, 1997.