



Temporal Sequence Compression by an Integrate-and-Fire Model of Hippocampal Area CA3

D.A. AUGUST

University of Virginia Health Sciences Center, Charlottesville, VA 22908

august@virginia.edu

W.B. LEVY

*Departments of Neurosurgery and Psychology, University of Virginia Health Sciences Center,
Charlottesville, VA 22908*

wbl@virginia.edu

Received April 28, 1997; Revised April 16, 1998; Accepted May 19, 1998

Action Editor:

Abstract. Cells in the rat hippocampus fire as a function of the animal's location in space. Thus, a rat moving through the world produces a statistically reproducible sequence of "place cell" firings. With this perspective, spatial navigation can be viewed as a sequence learning problem for the hippocampus. That is, learning entails associating the relationships among a sequence of places that are represented by a sequence of place cell firing. Recent experiments by McNaughton and colleagues suggest the hippocampus can recall a sequence of place cell firings at a faster rate than it was experienced. This speedup, which occurs during slow-wave sleep, is called *temporal compression*. Here, we show that a simplified model of hippocampal area CA3, based on integrate-and-fire cells and unsupervised Hebbian learning, reproduces this temporal compression. The amount of compression is proportional to the activity level during recall and to the relative timespan of associativity during learning. Compression seems to arise from an alteration of network dynamics between learning and recall. During learning, the dynamics are paced by external input and slowed by a low overall level of activity. During recall, however, external input is absent, and the dynamics are controlled by intrinsic network properties. Raising the activity level by lowering inhibition increases the rate at which the network can transition between previously learned states and thereby produces temporal compression. The tendency for speeding up future activations, however, is limited by the temporal range of associations that were present during learning.

Keywords: hippocampus, sequence learning, compression, chunking

1. Introduction

This paper concerns the learning and recall of sequence information by a computational model of hippocampal area CA3. We have chosen to discuss sequence learning from the perspective of spatial navigation. Experimental evidence shows that when a rat is in a given

location in space, certain "place cells" in the hippocampus fire in a specific and statistically reproducible manner (Ranck, 1975; O'Keefe and Conway, 1978; Muller and Kubie, 1989; Quirk et al., 1992; Thompson and Best, 1989). Thus, as the rat navigates through the environment, it encounters a sequence of places, which is represented by a sequence of place cell firing. From

this perspective, spatial navigation can be thought of as a sequence-learning task.

Many computational models of spatial navigation involve sequence learning (Blum and Abbott, 1995, 1996; Levy, 1994; Tsodyks et al., 1996). Perhaps the most recent example is a sequence-learning network investigated by Wallenstein and Hasselmo (1997). This CA3-like network is much more biophysically detailed than the model presented here, with multicompartmental cells, an extrinsic theta rhythm, and separate inhibitory GABA-A and GABA-B conductances. However, the simulations are still quite similar to the present study, with analogs of place fields developing from an initially random network configuration without prior assumptions about their existence or form. Unlike our model, however, the biophysical network depends on the interplay between an extrinsic theta-like input and a long, GABA-B-like timecourse of inhibition to obtain sequence learning. While these features are present in biology, here it is suggested that a simpler model is still sufficient for sequence learning.

There are two caveats to this approach of studying sequence learning from the perspective of spatial navigation. First, spatial navigation is more complex than just sequence learning. For example, spatial learning includes the phenomenon of path integration (dead reckoning), which involves traveling between two points but then returning by a different, more direct, route even in the relative absence of sensory stimuli (Redish and Touretzky, 1997; McNaughton et al., 1996; Whishaw and Jarrard, 1996). Second, in the hippocampus, sequence learning can be used to describe more than just spatial navigation. Because the firing of place cells marks a conjunction of a particular set of environmental features, these firings may not be correlated with a "place" per se (Bunsey and Eichenbaum, 1995; Wiener et al., 1989) but may be thought of in more general terms. For example, the classical conditioning paradigm of stimulus-response-reinforcement can be thought of as a very simple (three-member) sequence of cell firings. With this perspective, such hippocampally-mediated tasks as transverse patterning, transitive inference, and trace conditioning of the nictitating membrane response can be modeled as sequence learning tasks as well (Levy et al., 1995, 1996; Levy and Sederberg, 1997; Wu and Levy, 1996). Despite these two caveats, we choose sequence learning as a particularly simple spatial paradigm that might be equivalent, to the experimental situation (described below) of a rat running in one direction around a track.

Within the realm of spatial sequence learning, this article is concerned with a phenomenon called *temporal compression*, defined as the ability to recall a sequence of neural events faster than it would take to experience the sequence once during learning. Temporal compression may be relevant to a variety of different phenomena, from prediction and planning (e.g., Levy, 1989) to memory consolidation (e.g., Buzsaki, 1996). A variety of compression-like effects have been shown in several modeling studies (Jensen and Lisman, 1996; Lisman and Idiart, 1995; Prepscius and Levy, 1994; Tsodyks et al., 1996; Blum and Abbott, 1996; Wallenstein and Hasselmo, 1997) and will be described more fully in Section 4.

Recent work by McNaughton and colleagues has begun to uncover the neurophysiology of temporal compression (Wilson and McNaughton, 1994; Skaggs and McNaughton, 1996). In these experiments, rats ran around a track looking for food while multiunit recordings were made from hippocampal area CA1. By comparing this data to recordings made during subsequent episodes of slow-wave sleep (SWS), the authors found that cells that were correlated during the spatial behavior (e.g., cells with neighboring place fields) were correlated during sleep. Additionally, the temporal structure of firing was preserved during sleep. The width of the correlation peak, however, was significantly narrower during sleep than during the behavior. Taken together, these results suggest that the hippocampus recalls a temporally compressed version of recently acquired spatial sequence information during slow-wave sleep.

The firing of hippocampal cells during SWS is significantly different than the firing during waking behavior. This phase of sleep is identified on the hippocampal electroencephalogram by the appearance of numerous sharp wave (SPW) events (Ranck, 1973; Suzuki and Smith, 1987; Buzsaki, 1986). The most obvious feature of the SPW is a generalized increase in neural firing that begins in hippocampal area CA3 and is transmitted to CA1, the subiculum, and finally into the deep layers of the entorhinal cortex (Chrobak and Buzsaki, 1994; Ylinen et al., 1995). It is those CA3 cells that have recently experienced stimulation that are most likely to initiate SPWs (Buzsaki, 1984b, 1989). SPWs also include a high-frequency ripple of the field potential (≈ 100 Hz in area CA3) caused by firing of inhibitory interneurons (Buzsaki et al., 1992; Buzsaki and Chrobak, 1995). Sharp waves typically last for ≈ 100 ms and occur unpredictably at anywhere between

0.1 and 5 Hz (Buzsaki, 1986; Suzuki and Smith, 1987).

What is the mechanism of temporal compression? The simplest explanation is that the network is capable of operating at two different characteristic timescales. During learning, the timescale is driven by the external inputs. The time-constant of synaptic modification limits the temporal range over which events can be associated. As will be shown in Section 3.1, these associations are encoded by the network as *local context units*, which are hypothesized as analogous to place cells. Local context units span different portions of the sequence, thereby determining the *potential* amount of compression. During recall, there is no external input, and the timescale is set by the internal dynamics of the network. The speed at which the network can make transitions between sets of local context units encoding previously driven states determines the *actual* amount of potential compression that is expressed. Preliminary results have already appeared showing activity-dependent temporal compression (August and Levy, 1996; Levy et al., 1997).

In this article, we investigate various aspects of the network and the environment that control temporal compression. Our main findings are that the compression increases with higher levels of activity during recall and also when the duration of patterns in the input sequence is appropriately matched to the timespan of synaptic modification.

2. Methods

This section describes the general simulation strategy, the CA3-like network model, the input to this model, the data analysis procedure, and some of the computational issues involved in the simulations.

2.1. Simulation Protocol

Each simulation has two phases, *training* (learning) and *testing* (recall). To make an analogy to experiment, training represents an animal engaged in spatial behavior, and testing represents an animal during sleep or awake immobility. During training, a sequence of patterns is presented to the network repeatedly, and synaptic weights are modified according to an unsupervised Hebbian learning rule. After the synaptic weights have stabilized, associative synaptic modification is switched off and testing begins. While it may seem unrealistic to disable synaptic modification during

recall, there are data showing that synaptic plasticity indeed diminishes during slow-wave sleep (Leonard et al., 1987).

Two different testing paradigms are used. In *prompted sequence completion*, the network is given the initial part of the sequence as a test probe (or “cue”) and then is allowed to run on its own. In *spontaneous sequence completion*, the network is continually stimulated with random input activity. While spontaneous sequence completion is more like the situation during slow-wave sleep, prompted sequence completion is much easier to control and therefore to study.

2.2. Network Architecture

This model uses a two-layer network architecture to model hippocampal area CA3 and its inputs (see Fig. 1). The first layer consists of excitatory, nonmodifiable synapses projecting onto the main layer. Although CA3 actually has many sources of external input, including entorhinal cortex (EC), dentate gyrus (DG), contralateral CA3, and subcortical afferents (Amaral, 1987; Steward and Scoville, 1976; Witter, 1993), here we lump the inputs into a single EC/DG layer, noting that this procedure may ignore important computational differences among inputs (Levy et al., 1995; Treves and Rolls, 1992). The EC/DG layer projects to the main, CA3-like layer, which is characterized by sparse recurrent (feedback) connectivity. During training, these recurrent connections are modified by an unsupervised Hebbian rule governing associative synaptic modification. Activity is controlled by a feedback and a feedforward inhibitory interneuron.

2.3. EC/DG Layer Input

Unlike information arriving at a primary sensory cortex, information reaching the hippocampus is multimodal and highly processed (Amaral, 1987; Deacon et al., 1983). The information that triggers the firing of a place cell depends in a complex way on many environmental and internal cues. Therefore, rather than attempting to mimic a particular sensory feature, the model uses arbitrary patterns of EC/DG activity as the input. The input sequence consists of 100 patterns, each of which can be characterized generically by its duration, the number of cells it activates, the shift between successive patterns, and the EC/DG cell firing rate. In order to model the behavioral task of a rat running in




Figure 1. Network Schematic. Each CA3 pyramidal cell (black) receives one excitatory connection from the EC/DG layer (upper horizontal lines) and sparse recurrent excitation from within CA3 (lower horizontal lines). In addition, CA3 cells receive inhibitory projections from both a feedforward interneuron (upper gray cell) and a feedback interneuron (lower gray cell). The feedforward interneuron gets input from all the cells in the EC/DG layer, while the feedback cell receives input from the entire CA3 layer. Sparse recurrent connectivity is illustrated here by the few connections that are actually made compared to the number of possible connections. Note that the extended dendrites in this figure are only included for clarity. Each cell in the model is actually a “morpho-less” soma.

one direction around a racetrack, a circular sequence is used. That is, the (nominal) first and last patterns overlap, so that each pattern has both a predecessor and a successor.

The CA3 layer has 1,000 cells, but for any given sequence, not all of the cells are activated directly by input from the EC/DG layer. In order to make the externally driven cells of the CA3 layer easy to visualize, the sequence of patterns is formed by successively shifting each group of 10 neighboring EC/DG units by one cell. For example, the first pattern drives cells 1–10 of the CA3 layer, the second drives cells 2–11, and so forth. The sequence used in these simulations has 100 patterns, and thus only the first 100 CA3 cells are driven by the EC/DG input directly. The duration of each pattern is 20 ms, but because of the extensive overlap between patterns, each CA3 cell actually is driven by EC/DG input for $T = 200$ ms.

2.4. CA3 Layer Model

The CA3 layer is made of 1,000 integrate-and-fire elements. The somatic voltage, $V_j(t)$, of the j th cell is updated as follows:

$$\tau_m \frac{dV_j(t)}{dt} = I_j(t) - V_j(t), \quad (1)$$

where $I_j(t)$ is the synaptic current and $\tau_m = 20$ ms is the membrane time constant. When the voltage exceeds a threshold of 0.0033, the cell fires, after which the threshold is subtracted, and a 2 ms deadtime is imposed.

The synaptic excitation arriving at the j th cell is

$$I_j^{ex}(t) = K_1 x_j(t) + K_2 \sum_i c_{ij} w_{ij}(t) z_i(t - \delta_{ij}), \quad (2)$$

where $x_j(t)$ is the binary EC/DG firing state, $z_i(t)$ is the binary firing state of the i th CA3 unit, c_{ij} is a 0/1

variable representing the absence/presence of a connection from cell i to cell j , w_{ij} is the synaptic weight, and δ_{ij} is the corresponding axonal delay, randomly set between 1 and 2 ms. The scaling factors, $K_1 = K_2 = 4$, control the relative magnitudes of external and feedback excitation.

In the hippocampus, connectivity is higher locally and decays over distance (Bernard and Wheal, 1994), so that the overall average connectivity is on the order of a few percent (Amaral et al., 1990; Ishizuka et al., 1990). However, because our 1,000-cell model is much smaller than the 300,000-cell CA3, it only represents a thin slice of tissue, within which the effects of distance on connectivity can be ignored. Thus, we set the connectivity pattern, c_{ij} , randomly and uniformly, so that each CA3 cell receives connections from 10% of the other cells.

Activity is controlled by a feedforward and a feedback inhibitory interneuron. Feedforward inhibition is the running average of the EC/DG layer activity, $\bar{s}(t)$, scaled by factor K_{FFI} . Feedback inhibition is the running average of CA3 activity one millisecond in the past, $\bar{m}(t)$, scaled by K_{FBI} . The time-constant of both of these inhibitory running averages is 2 ms. Thus, the inhibitory current is

$$I_{\text{in}}(t) = K_0 + K_{\text{FFI}}\bar{s}(t) + K_{\text{FBI}}\bar{m}(t), \quad (3)$$

where $K_0 = 1$ is a constant inhibitory term.

Inhibition exerts a shunting (divisive) effect on CA3 as follows:

$$\frac{dI_j(t)}{dt} = \frac{I_j^{\text{ex}}(t)}{I_j^{\text{ex}}(t) + I_{\text{in}}(t - \delta)} - \frac{I_j(t)}{\tau_s}, \quad (4)$$

where $\tau_s = 2$ ms is the synaptic time-constant, and $\delta = 1$ ms is the axonal delay. Shunting makes the synaptic current into a *ratio* of excitation and inhibition rather than a difference.

Synaptic weights change during the course of learning according to a Hebbian rule:

$$\Delta w_{ij}(t) = 0.1 \cdot z_j(t)(\bar{z}_i(t) - w_{ij}(t)), \quad (5)$$

where $\bar{z}_i(t)$ is a running average of presynaptic activity. Each time the presynaptic cell i fires, its effect on the running average rises to a peak at 8 ms (Levy and Steward, 1983) and decays with a time-constant of

$\tau_A = 150$ ms. That is,

$$\bar{z}_i(t) = \sum_k e^{-(t-t_k^i)/\tau_A} - e^{-(t-t_k^i)/\tau_R}, \quad (6)$$

where t_k^i represents the k th firing of cell i , and $\tau_R = 0.001785$ is chosen to make the function reach a maximum at 8 ms.

The value of τ_A is meant to reflect the deactivation kinetics of the N-methyl-D-aspartate (NMDA) receptor (Holmes and Levy, 1990), which is known to participate in synaptic plasticity in the hippocampus. However, τ_A is not meant to be taken as a biophysical model of the NMDA receptor but rather a simple representation of the functional consequences of the NMDA receptor's ability to compute a relatively long-term average of presynaptic activity. Thus, as a value of τ_A , the functional duration of 150 ms is used, as suggested by experimental measurements of the timecourse of decay of associative long-term potentiation (LTP) in the hippocampus (Gustafsson and Wigstrom, 1986; Levy and Steward, 1983). In a more realistic model, τ_A could be chosen according to observed biophysical properties of the NMDA receptor, rather than its second-order effect on LTP. In this case, the value of τ_A is debated in the literature. Experimental observations show a double exponential decay of deactivation, with a value for the first time-constant ranging from lower values like 25 ms (Staley and Mody, 1992) and 70 ms (McBain and Dingledine, 1992) to higher values like 200 ms (Spruston et al., 1995), consistent with τ_A used here.

Notice that postsynaptic activity ($z_j(t) = 1$) is required for synaptic modification, while the strength and sign of modification are controlled by the average presynaptic activity ($\bar{z}_i(t)$) in relation to the current synaptic weight. This is consistent with experimental observations in the hippocampal slice preparation (Gustafsson and Wigstrom, 1986; Levy and Steward, 1983), as well as with biophysical simulations (Holmes and Levy, 1994).

With a learning rate of 0.1, only about five presentations of the sequence are required for learning. In these simulations, however, we use 10 learning trials to move the average synaptic weight closer to steady-state.

2.5. Controlling for Activity Levels

Previous experience with this and similar network models indicates that the average activity level is an important global variable, influencing network behavior

Table 1. Changing inhibition to match activity levels.

Parameter	Value	K_{FBI}	Activity (Hz)
τ_A	200 ms	420	6.8 ± 0.12
	150 ms	375	6.6 ± 0.02
τ_A	100 ms	340	6.1 ± 0.00
	150 ms	410	6.0 ± 0.09
T	300 ms	600	5.4 ± 0.00
	200 ms	440	5.6 ± 0.14
T	150 ms	340	8.4 ± 0.18
	200 ms	440	8.3 ± 0.11

In order to control for unwanted changes in activity when any model parameter is varied, inhibition must be adjusted in a compensatory manner. In this table, each row represents the parameter settings for one simulation, and rows are ordered into pairs. The first row of each pair gives the “experimental” parameter settings, in which the parameter of interest has been changed and inhibition adjusted to whatever level is required to obtain learning. The second row gives the “control” settings, in which inhibition has been adjusted to match the activity level of the previous row. For example, rows 1 and 2 compare the effect of raising τ_A from 150 to 200 ms, while rows 3 and 4 compare the effect of lowering τ_A from 150 to 100 ms. For each row, the first column lists the parameter under study, the second gives the actual parameter values, the third shows the level of feedback inhibition, and the fourth gives the average activity level \pm S.E. for the network during the final (tenth) trial of learning. Symbols used: K_{FBI} , feedback inhibition; τ_A , time-constant of synaptic modification rule; T , pattern duration.

in a variety of ways (August and Levy, 1996; Minai and Levy, 1993a, 1993c; Smith and Levy, 1998; Wu and Levy, 1995). Thus, activity is a potentially confounding variable when studying the effect of network parameters on compression. If changing a parameter does affect temporal compression, is this effect due to the parameter itself, or is it secondary to the effect of this parameter on activity levels? In order to avoid this uncertainty, we compensate for changes in activity by varying the inhibition in the opposite way, so as to match activity levels to a comparison case. For example, if we wish to study the effect on compression of increasing τ_A , but discover that increasing τ_A reduces activity, then we must compare to a case of similarly low activity obtained by increasing the inhibition. Table 1 shows each of the parameters that have been changed, along with the required compensatory changes in inhibition.

2.6. Data Analysis

The external input sequence only drives the first 100 CA3 cells directly. The remaining 900 cells develop

a characteristic firing pattern, each firing over a brief period of time. These cells, which we call *local context units*, are described more fully in Fig. 2. The collection of local context units are the CA3 network’s own internally generated code for the input sequence. More specifically, the code for pattern i of the sequence consists of all the local context units that are coactive with pattern i during learning.

In order to quantify network behavior during recall, we must not only examine those CA3 cells that were driven directly by the EC/DG layer during learning but also the pattern of local context unit firing exhibited by the remaining CA3 cells. Because the fidelity of the firing pattern in the 900 cells may not be obvious by inspection, some type of decoder is necessary. Arguably the best decoder would be a model of hippocampal area CA1, the primary target of CA3 output in the brain. However, because our main interest is CA3 itself, we instead choose a simple decoder, rather than trying to improve the performance of CA3 with an additional network layer. In particular, our decoder is a mathematically convenient function (described below) that translates the network state at any given time into an integer between 1 and the number of patterns in the sequence.

2.6.1. Decoding the CA3 State. Decoding is done in two steps. First, we calculate a *similarity matrix*, \mathbf{S} , in which each element is a number between 0 and 1 representing the similarity of the network during recall to one of the codewords developed during learning. For example, S_{ij} is the similarity of the network state at timestep j of recall to the codeword representing pattern i . Mathematically, the similarity matrix is a normalized product of two other matrices. The first matrix, \mathbf{A} , is a binary (rastergram) matrix representing the activity of the network during the final trial of learning, when still driven by EC/DG input. This matrix has 1,000 rows and 100 columns, reflecting the number of cells in the network and the number of patterns in the input sequence, respectively. The k th row of \mathbf{A} corresponds to the activity of cell k averaged into successive bins of duration T , the length of time each individual EC/DG cell is activated in the input sequence. The second matrix, \mathbf{B} , represents activity during testing as a binary matrix with 1,000 rows and a number of columns equal to the number of 1 ms bins during testing. The similarity matrix is then

$$\mathbf{S} = \frac{\mathbf{A}'\mathbf{B}}{|\mathbf{A}||\mathbf{B}|}. \quad (7)$$

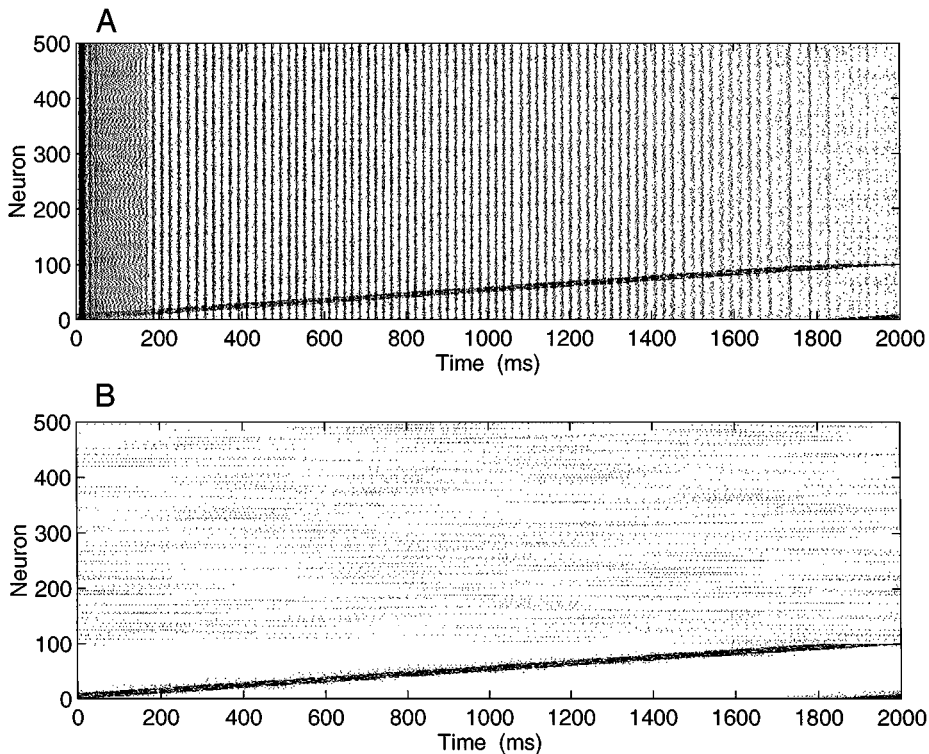


Figure 2. The development of context units. A: Network state during the first trial of training. Neural activity is presented as a rastergram, in which time proceeds along the abscissa, neurons are along the ordinate, and each dot stands for a spike. Each row represents the firing of one cell over time. For clarity, only the first 500 cells of the 1,000-cell network are shown here. The input sequence from the EC/DG layer directly excites only the first 100 CA3-layer cells, as shown by the dark diagonal band of firing. The remaining 900 cells (of which only 400 are shown) are excited by feedback (recurrent) connections from within CA3 only. At the beginning of the simulation, the network oscillates with a frequency that decreases over time and eventually dies out. B: Network state during the final trial of training. By the tenth trial of learning, the average synaptic weight has stabilized. Now, however, due to the effects of synaptic modification, the recurrently driven cells no longer fire in a simple oscillatory pattern. Instead, these cells fire over brief periods of time spanning several patterns of the sequence. We call these cells *local context units* and propose them as a rough analogy to hippocampal place cells. The set of local context units that are activated during each pattern constitute the network's own internally generated codeword for that pattern.

The similarity matrix can be conveniently viewed on a grayscale from white (least similar) to black (most similar).

The second step of decoding is to calculate the *winner*, $W(j)$, at each timestep j of testing. The winner is the largest value of each column of the similarity matrix—the pattern to which the network state during recall is most similar. The sequence of winners is the final decoding of the CA3 network state over time.

2.6.2. Measuring Temporal Compression. Temporal compression is measured as the *compression ratio*, CR , where $CR \geq 1$. The calculation of CR is based on the autocorrelation of spike train data. If we represent the spike times for neuron n by the sequence $\{t_i^n\}$, $i = 1, 2, 3, \dots$, then we may define a counting function, $X_n(\tau)$, which is incremented by one whenever $t_i^n - t_j^n = \tau$ for some i and j . The autocorrelations,

$X(\tau)$, are the average of the counting functions for 100 cells. In these simulations, the autocorrelation is measured with a 1 ms time bin, using spike train data from the first 100 local context units (actually cells 101 to 200 of the CA3 layer). If the sequence is repeatedly recalled, the autocorrelation will show a series of peaks of diminishing height. The time at which the first nonzero peak occurs, τ_1 , corresponds to the duration of the compressed sequence. Thus, the compression ratio is

$$CR = \frac{\text{sequence duration}}{\tau_1}. \quad (8)$$

One advantage of measuring compression by the autocorrelation is that, in principle, all of the necessary information can be obtained experimentally. The autocorrelation itself requires the spike trains, which are obtained using multielectrode recording. The

sequence duration can be estimated from the average time it takes a rat to complete one loop—for example, by averaging the intervals between crossings of an electric eye-beam placed at one point along the track.

2.7. Computational Considerations

These simulations use a 1,000-cell network, which is nearly 300-fold smaller than hippocampal area CA3 in the rat (Amaral et al., 1990). However, even this small network still requires a considerable computational effort to simulate. Because each synapse must be modified at every timestep, the computational burden of synaptic modification is proportional to the square of the network size. To make these simulations run efficiently, we use the Cray C-90 computing platform, combined with custom written code employing vectorized techniques. A forward Euler method is used to solve the differential equations governing synaptic modification and cellular voltage update. Using a timestep of $\Delta t = 0.00025$ s, a typical run requires about 1 sec of CPU time per simulated second.

3. Results

3.1. Development of Local Context Units

Figure 2A shows the network state during the first trial of learning. In this rastergram, each row depicts the firing of one cell over time. The dark diagonal band of firing represents the first 100 cells, which are driven directly EC/DG input. The remaining 900 cells (of which only 400 are shown) are driven by recurrent connections only. The firing pattern depends on how the initial synaptic weight distribution is chosen at the beginning of the simulation. Larger average weights result in oscillatory behavior, whereas smaller average weights result in more random, noisy firing. In this particular illustration, the initial synaptic weights are chosen from an exponential distribution with mean of 0.05. Figure 2B shows the same network during the tenth training trial, by which time the average synaptic weight has stopped changing significantly. Note the appearance of local context units—recurrently driven cells with discrete firing episodes spanning a few patterns. Learning transforms the initially random firing pattern of recurrently driven cells into organized local context unit activity. The collection of local context units active at any time constitutes the *codeword* generated by the network for that portion of the sequence.

3.2. Spontaneous Recall

Recall is first examined in the spontaneous sequence completion task, which is meant to resemble slow-wave sleep. After training, inhibition is lowered, and the network is stimulated with low levels (on average, 1 Hz across all neurons) of random EC/DG firing. Figure 3A shows the network during this random input. Three episodes of spontaneous recall are evident, starting at roughly 250, 500, and 1100 ms. No explicit triggering events are required to obtain these episodes of recall. The network exhibits a significant amount of temporal compression during these recall events, requiring only about 100 ms to replay the original 2,000 ms sequence. Note that recall happens during SPW-like periods of high-frequency firing, as shown by the timecourse of average network activity in Fig. 3B.

The similarity matrix in Fig. 3C illustrates the network state during periods of successful recall. The values from 0 to 100 on the ordinate do not correspond to the first 100 cells of the network but rather to the similarity of the entire network state at any instant of time during recall to each of the 100 codewords representing the 100 patterns in the input sequence. The dark diagonal bands demonstrate that it is not just the first 100 cells shown above (e.g., those cells driven by EC/DG activity during training) that reproduce their firing pattern during testing but the remaining 900 cells as well. We emphasize that in the simple sequence learning task shown here, looking at the similarity matrix does not add much compared to looking at the rastergram in Fig. 3A. However, in a more realistic situation in which multiple input sequences from EC/DG were scattered throughout CA3 in a nonorderly fashion—rather than the artificial situation here in which the sequence progresses in order through the first 100 cells—the similarity matrix would add a vital understanding of the fidelity of recall. The sequence of winners, which are the patterns with the highest similarity score, confirm the impression of successful recall shown in the similarity matrix (see Fig. 3D).

3.3. Compression and Activity Levels

It is difficult to study temporal compression during episodes of spontaneous recall because of their unpredictable nature. Therefore, we instead study recall during the prompted sequence completion paradigm, in which the network is given the first pattern for 50 ms as a test probe and then is allowed to relax on its own. Conceptually, this is somewhat like studying

compression during a very long sharp wave. An example of prompted sequence completion is shown in Fig. 4A, during which the network recalls the 2,000 ms sequence in ≈ 120 ms, so that $CR = 16.53$. The dark diagonal bands of high similarity, shown in Fig. 4B, demonstrate successful recall. That is, it is not just the first 100 cells that were driven by EC/DG activity during training that reproduce their firing pattern during testing but the remaining 900 cells as well.

To determine the relationship between compression and the average activity level of the network during

testing, the network is tested at two different levels of feedback inhibition. As K_{FBI} is lowered from 44 to 18, the average activity during testing increases from 42.6 to 145.3 Hz, and at the same time, CR increases from 16.53 to 25.97 (see Figs. 5A and B). To quantify more accurately the relationship between CR and activity, we test the network at 20 different values of inhibition and repeat this entire process with three random starting conditions. When the resulting compression ratios are plotted against the activity during recall, a monotonically increasing relationship is apparent (see Fig. 5C).




Figure 3. Spontaneous sequence completion. A: Network state during testing. During recall, associative synaptic modification is disabled, and EC/DG units are activated randomly at an average rate of 1 Hz. Notice the emergence of three distinct episodes of recall, two of which contain multiple repetitions of the sequence. The recalled sequence is temporally compressed by approximately 16-fold, replaying the original 2,000 ms sequence in only about 125 ms. B: Average network activity. Although the network usually remains at low levels of activity, spontaneous recall occurs during brief periods of high-frequency firing, reminiscent of sharp waves. The ordinate shows summed cell firing over the network in Hz. C: Similarity matrix. The similarity matrix quantifies how close the network state is during recall to the codeword generated for each pattern during learning. Similarity is useful because, although the fidelity of recall may seem obvious by inspection of the first 100 (externally driven) cells, the firing pattern of the remaining 900 cells is not as clear by inspection. The similarity matrix allows visualization of the entire network state at any instant of time. Specifically, the $(i, j)^{\text{th}}$ element of this matrix is the similarity between the network state at timestep j during testing and the codeword for pattern i during learning. Similarity, as defined in Section 2.6, is plotted on a scale from white (least similar) to black (most similar). Note, however, that the maximum similarity shown in this figure is only ≈ 0.3 , indicating that recall is still imperfect. D: Winners. The winner at each timestep during testing is the pattern with the highest similarity score. The sequence of winners, which is the final decoding of the network state, resembles a compressed version of the original sequence.

(Continued on next page.)




Figure 3. (Continued).

3.4. Compression and the Time-Constant of Synaptic Modification

Temporal compression should be related to the timespan over which synapses can make associations. For example, if synaptic modification during pattern 50 is only influenced as far back as pattern 49, there will be much less potential for compression than if the state is influenced as far back as pattern 1. In our model, the timespan of associativity is controlled by τ_A , the time-constant of the running averager of presynaptic activity in the synaptic modification rule. As shown in Fig. 6, increasing τ_A from 100 to 150 to 200 ms results in a dramatic increase in CR .

3.5. Compression and Pattern Duration

The compression ratio, CR , is a relative measure. Therefore, we hypothesize that decreasing the pattern duration for a fixed τ_A has a similar effect as increasing τ_A for a fixed pattern duration. Either change allows the

network to associate patterns that occur farther apart in the sequence. To test this intuition, we use pattern durations of 15, 20, and 30 ms. Because each pattern consists of 10 EC/DG cells and overlaps its neighboring patterns by nine cells, these pattern durations correspond to $T = 150, 200,$ and 300 ms as defined in Section 2.3. For the $T = 150$ ms case, the network is tested with K_{FBI} values from 9 to 20. For $T = 200$, K_{FBI} is increased from 8 to 34 during testing, and for $T = 300$, we increase K_{FBI} from 7 to 18 during testing. As shown in Fig. 7, the shorter the pattern, the more the temporal compression.

3.6. The Relative Timespan of Associativity

As just observed, there is a similarity between increasing τ_A for fixed T and decreasing T for fixed τ_A . Here, we examine directly compression as a function of the ratio τ_A/T , which we call the *relative timespan of associativity*. We expect that as τ_A/T increases, synaptic modification will be able to span more of the sequence, resulting in more compression. To examine this

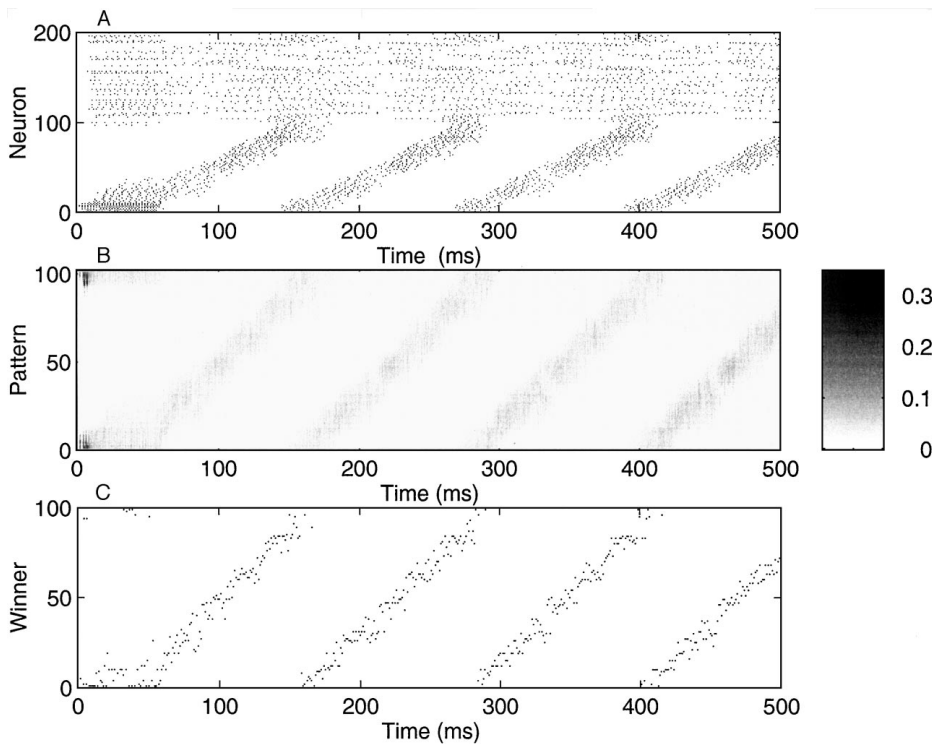


Figure 4. Prompted sequence completion. A: Network state during testing. During testing, synaptic modification is disabled and the network is allowed to freely recall the sequence after experiencing the first pattern for 50 ms as a test probe (lower lefthand corner of upper panel). After the test probe ends, the network recalls an approximate version of the original (circular) sequence several times, with a significant amount of temporal compression. That is, while the original sequence lasts for 2000 ms (see Fig. 2), each repetition of the recalled sequence lasts for only ≈ 150 ms, resulting in a compression ratio of $CR \approx 13.3$. B: Similarity matrix. The similarity matrix, described in Section 2.6, shows bands of gray corresponding to recall of the sequence. As in the previous figure, recall is not perfect, with a maximum similarity still < 0.3 . C: Winners. The sequence of patterns with the maximum similarity values over time mimics the original sequence.

relationship, data are combined from previous studies with $\tau_A = 100, 150,$ and 200 ms and $T = 150, 200,$ and 300 ms. Because these data all have somewhat different ranges of CR -values, only the average compression for activity levels between 90 and 110 Hz are compared. Figure 8 illustrates a strong positive correlation between τ_A/T and compression ($r = 0.98$). Thus, from a conceptual standpoint, τ_A/T seems to be an appropriate variable for studying compression.

The amount by which the parameter τ_A/T can be varied is limited by the fact that the network does not learn sequences that change too slowly or too rapidly compared to the timespan of the associative modification rule. For example, as shown in Fig. 9, when $\tau_A/T \geq 1.5$, the associative timespan is too small to span multiple patterns, and the network gets stuck at attractors within the sequence during recall. Conversely, when $\tau_A/T \leq 0.375$, the sequence changes too much

within a given timespan of associative modification, and there is not enough pattern-to-pattern strengthening to allow definitive sustained sequence recall.

3.7. Compression and Context Lifetime

Besides investigating how temporal compression is influenced by individual model parameters, we are also interested in the influence of global values that can be affected by many different parameters. One such global property is the average duration of local context unit firing (or *average local context lifetime*), $E[L]$ (Levy and Wu, 1996). Basically, $E[L]$ measures the average length of time each context unit (cells 101 to 1,000 in this network) fires in a contiguous manner, defined as more than three spikes within a short period of time. Because the network uses local context units to span the

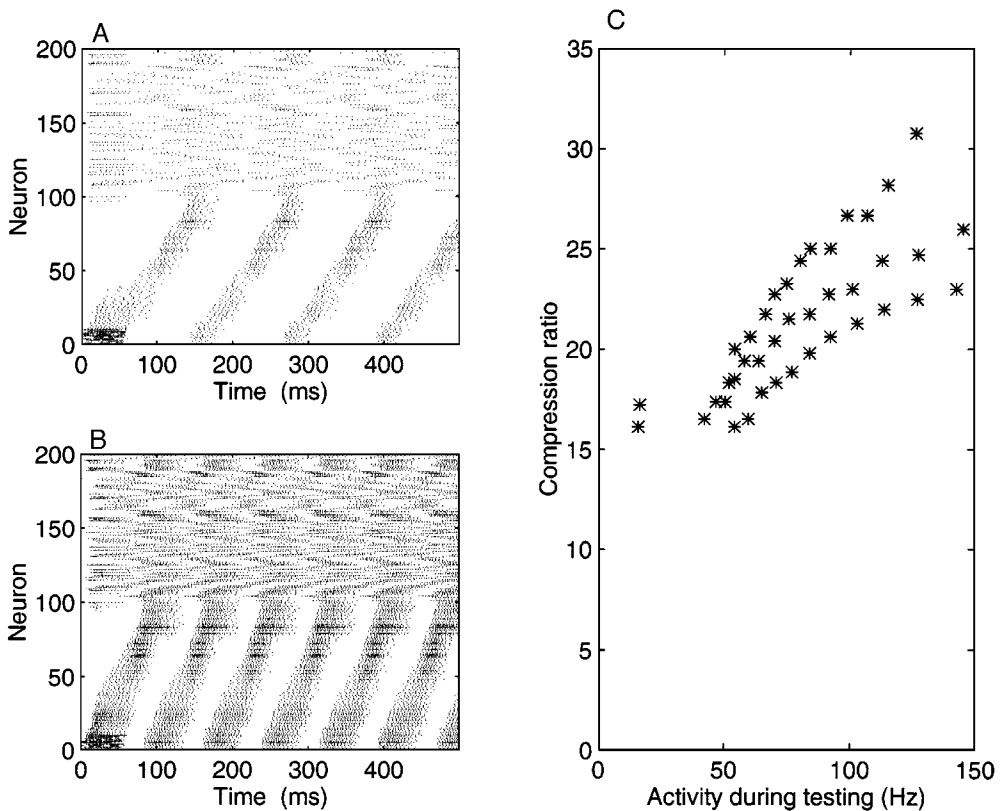


Figure 5. Compression increases with activity during testing. A: Testing at low levels of activity. In this case, the network is tested with the feedback inhibition set to $K_{\text{FBI}} = 44$, which produces about four repetitions of the sequence in 500 ms, resulting in a compression ratio of 16.5. B: Testing at high levels of activity. Here, K_{FBI} is reduced to 18, thereby raising the activity level from about 40 Hz to about 150 Hz. In this case, the network recalls the sequence about six times within the same 500 ms period, raising the CR to 26. C: Compression as a function of average activity. The network is tested over a wide range of activity levels, produced by varying K_{FBI} over 20 different values. The abscissa represents the average firing rate per unit during the 500 ms testing period. The combined data from three different simulations shows that as the average activity level during testing increases, so does the amount of temporal compression. These three simulations are produced by varying the initial random seed, which in turn controls the connectivity matrix and initial weights.

sequence, we expect that longer local context lifetimes should lead to more compression. However, there is only a weak positive correlation ($r = 0.58$) between CR and the normalized context lifetime, $E[L]/T$.

One reason for this weak correlation is that the EC/DG input is strong enough to “force” most of the local context units to have duration approximately T , rather than spreading out over many different timescales (see below). Thus, we also examine the correlation between compression and the context lifetime coefficient of variance, $CV[L]$. More specifically, data are combined from simulations using input patterns with $T = 150, 200,$ and 300 ms and simulations with associative time-constants of $\tau_A = 100, 150,$ and 200 ms, using several different random initial seeds for

each parameter value. Because the CR values for this data covers a wide range of activity levels, only those values of CR measured between 90 and 100 Hz are included in the comparison. In this case, a highly significant correlation of $r = 0.97$ emerges, indicating that the variation in the length of local context units, more than their mean length, affects the degree of temporal compression.

More insight into this relationship between compression and context lifetime can be obtained by examining the distribution of context lifetimes. Figure 10 compares the histogram of local context lifetimes for simulations with several values of τ_A , the time-constant of synaptic modification. The abscissa shows L/T , the local context lifetime normalized to the pattern




Figure 6. Compression increases with the timespan of synaptic modification. This figure shows the effect of changing τ_A , the time-constant of the running averager of presynaptic activity in the synaptic modification rule, for networks tested over a range of different activity levels. Shown here are networks with $\tau_A = 100$ (\times), 150 (\circ), and 200 ms ($*$). As τ_A increases, there is a dramatic increase in compression, as well as an enhancement of the activity dependence. To obtain this range of activity levels for the case of $\tau_A = 100$, K_{FBI} was varied from 4 to 9. For $\tau_A = 150$, K_{FBI} was varied from 11 to 34, and for $\tau_A = 200$, K_{FBI} was varied from 22 to 50. The abscissa represents the average firing rate per unit during the 500 ms testing period.

duration. Most of the context lifetimes do indeed cluster near the pattern duration, so that each histogram has a peak at $L/T = 1$. However, as τ_A increases, the number of longer duration context lifetimes increases and so does the amount of temporal compression. Thus, our original intuition that compression depends on local context lifetime is essentially correct but must be limited to local context units with $L > T$.

4. Discussion

The CA3 network model described here performs temporal compression—recalling a sequence in less time than it would take to present the sequence during a single learning trial. Temporal compression is proportional to the average network activity during recall, which we control by varying the level of feedback inhibition. Temporal compression is also influenced by

the relationship between the time-constant governing synaptic modification and the pattern duration. These latter effects are related to the development of longer local context units during training.

4.1. Spontaneous Sequence Recall

When inhibition is lowered and the network is stimulated by random input activity during testing, we observe spontaneous bursts of previously experienced sequences. These bursts occur as brief, intermittent episodes of high network activity, reminiscent of the sharp waves observed in the hippocampal electroencephalogram during slow-wave sleep and awake inactivity (Ranck, 1973). Thus, the simulation result is broadly consistent with the experimental observation that cell-specific correlations learned during awake maze-running behavior are reproduced during




Figure 7. Compression decreases with increasing pattern duration. This figure shows the amount of compression that can be performed by a network with an associative time-constant of $\tau_A = 150$ ms when stimulated by input sequences with different pattern durations. Here we compare the case of $T = 150$ (*), $T = 200$ (o), and $T = 300$ ms (x). As the pattern duration decreases, the amount of compression increases. The abscissa represents the average firing rate per unit during the 500 ms testing period.

slow-wave sleep (Wilson and McNaughton, 1994). In our model, each SPW-like event is likely terminated by background noise in the EC/DG layer. In the brain however, SPWs may be terminated by the action of inhibitory interneurons (see, for example, Buhl et al., 1994; Buzsaki, 1984a; Gulyas et al., 1993; Miles, 1990) or calcium-gated potassium channels (Benington and Heller, 1995).

The present observation of spontaneous sequence recall during episodes of high-activity extends a previous observation based on a McCulloch-Pitts network (Minai and Levy, 1993b). In those simulations, however, the completeness of the recalled sequence was not evident, nor was it possible to match the time constants to physiological values. The present results are also consistent with recent simulations demonstrating that a hippocampal network model switches between periods of relative quiescence and periods of high activity (Shen and McNaughton, 1996). However, the Shen and McNaughton study shows that the appropriate cells

(i.e., those with strong synaptic interconnections) fire during SPW-like episodes, whereas our study reflects this but also demonstrates that the temporal ordering of these firings is approximately preserved.

4.2. What Causes Temporal Compression?

From the present results, the primary mechanism of temporal compression seems to be the ability to switch from externally driven dynamics during learning to internally driven dynamics during recall. During learning, the time-constant of synaptic modification, τ_A , in relation to the pattern duration, T , controls how far apart in time events can be associated. The network encodes these associations in the form of local context units. As τ_A/T increases, the activity of individual neurons, \bar{z}_i , can span more of the sequence because individual synaptic associations span greater fractions of the whole sequence. Longer context lifetimes result,




Figure 8. Compression and the relative timespan of associativity. The individual effects of the time-constant of associative modification, τ_A , and of the pattern duration, T , are illustrated in Figs. 6 and 7, respectively. Here, we examine the relationship between compression and the ratio of τ_A/T , also known as the relative timespan of associativity. Because this figure combines data from experiments with many different CR values, compression is expressed as CR_{100} , the average CR for activity levels between 90 and 110 Hz only. There is a significant correlation ($r = 0.98$) between τ_A/T and compression. Note the the network fails to learn sequences with $\tau_A/T = 1.5$ and 0.375 .

with a greater potential for compression. During recall, the amount of compression produced depends on how fast the network can transition between these previously driven states encoding different patterns. By lowering inhibition and thereby raising activity levels, the amount of compression expressed by the network increases.

At present, we do not have an analytical theory that can quantitatively predict the compression ratio as a function of the average network activity level. However, analytical solutions are available for the average activity level in a simpler network without capacitance (Minai and Levy, 1993a, 1993c; Smith and Levy, 1998). Qualitatively, however, the speedup at higher activity levels seems straightforward. Let us denote by $S(n)$ the set of local context units active during pattern n . By virtue of the temporally asymmetric rule governing synaptic modification, the neurons in $S(n)$ will excite those neurons in the set $S(n+1)$ by the greatest amount, with less excitation for cells in $S(n)$,

and even less for neurons in $S(n-1)$. Now let us consider increasing the average network activity level as a lowering of cell firing threshold. Lowering the threshold by a small amount increases the number of cells in both $S(n+1)$ and $S(n-1)$. However, because of temporal asymmetry in the synaptic modification rule, relatively more active neurons are added from $S(n+1)$ than from $S(n-1)$. If the threshold is lowered further, the same thing happens again, but now the sets $S(n+2)$ and $S(n-2)$ are also included. Again, relatively speaking, more of the excitation from $S(n-1)$ goes to the neurons $S(n)$ than to $S(n-2)$, and more excitation goes from $S(n+1)$ to $S(n+2)$ than to $S(n)$. This process of adding relatively more cells associated with patterns in the future continues as threshold is lowered. As activity is increased in this manner, the network finds (i.e., represents) a multiplicity of states simultaneously. Because relatively more future states are represented than present or past states, there is a stronger pull into the more distant future.

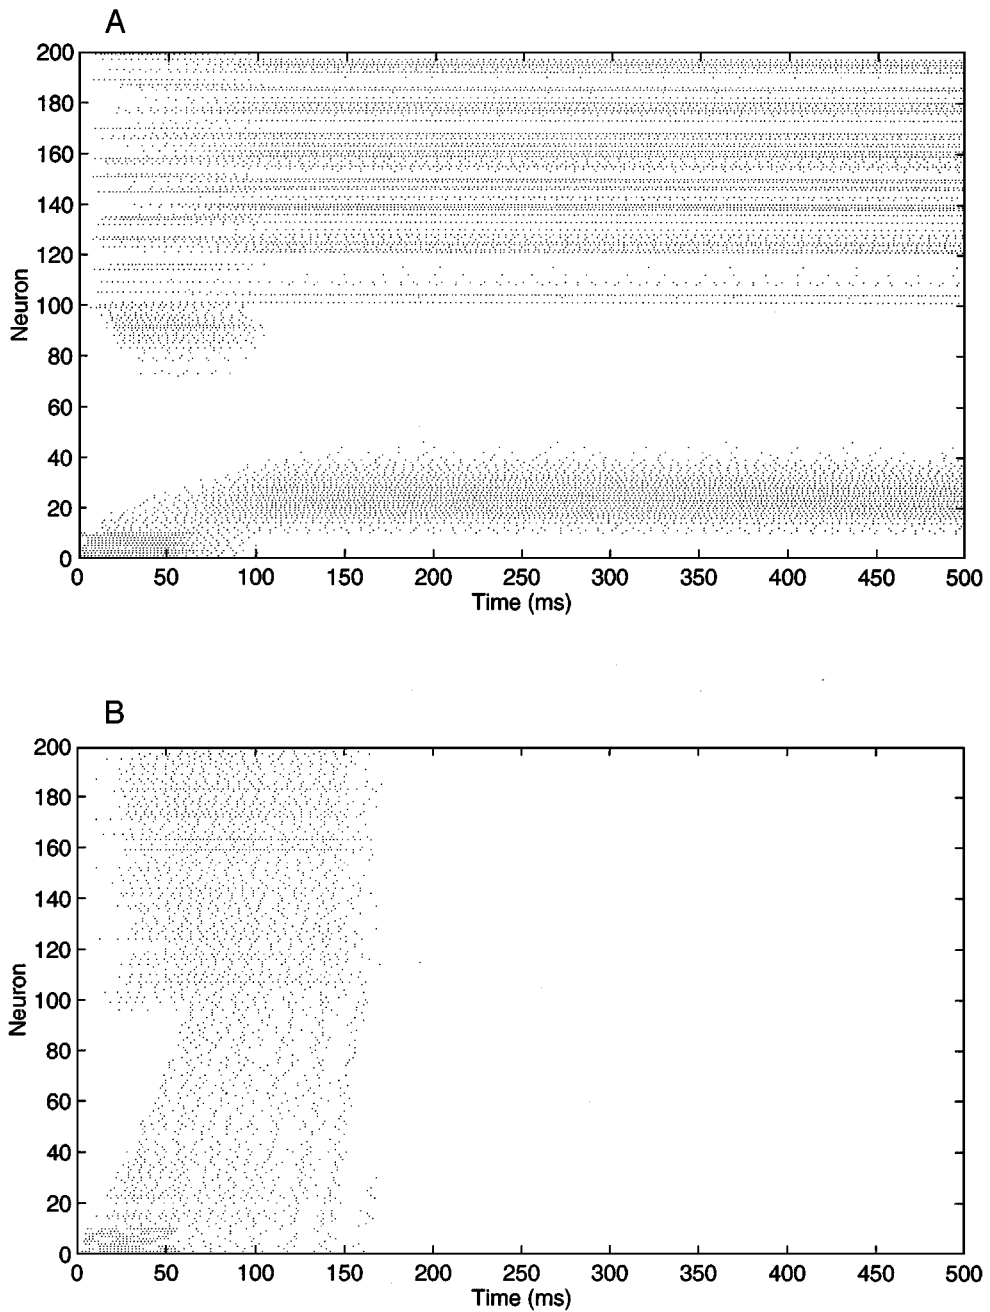


Figure 9. Failure to recall sequences with very short or long patterns. A: Recall of a sequence with prolonged patterns. The network is trained with a sequence in which each EC/DG unit is active for $T = 400$ ms instead of 200 ms. During testing, the network falls into an attractor state near pattern 30 and can not break free to recall the rest of the sequence. Attractor formation occurred in 10 out of 10 cases with different random initial seeds. B: Recall of a sequence with shortened patterns. In this case, each EC/DG unit is only active for 100 ms. During testing, network activity dies out. This effect is consistent in five out of five trials with different random seeds. In both of these simulations, the time-constant of synaptic modification is $\tau_A = 150$ ms.




Figure 10. Distribution of local context unit lifetimes. Shown here is the relative frequency of local context unit lifetimes L , normalized to the pattern duration T , for simulations with a 100 (solid) and a 200 (dash) ms time-constant of synaptic modification. Although the peak of these two distributions is similar ($L/T \approx 1$), there are more long-duration context lifetimes and a higher average CR when the associative time-constant, τ_A , is 200 ms.

4.3. Compression by Other Models

Several modeling studies have demonstrated faster than real-time sequence prediction associated with network oscillations. For example, it has been suggested (Skaggs and McNaughton, 1995, 1996; Skaggs et al., 1996) that compression arises from the so-called O'Keefe-Reece effect—the ability of place cells to shift their firing according to the phase of theta (O'Keefe and Reece, 1993). Indeed, various computational models of theta-phase shifting have recently been proposed (Shen and McNaughton, 1996; Wallenstein and Hasselmo, 1997). Another type of temporal compression arises from the interplay between higher-frequency gamma oscillation and lower-frequency theta oscillation, which allows approximately seven different events can be compressed into a single theta cycle (Jensen and Lisman, 1996; Lisman and Idiart, 1995).

In the present results, however, temporal compression is obtained without any theta-like process or phase

shifting. The temporal compression demonstrated here requires only a time-spanning learning rule, not a particular type of network oscillation. However, the O'Keefe-Reece theta-phase shifting effect nevertheless works in concert with the temporal compression described here to facilitate sequence learning. Specifically, theta-phase shifting brings temporally distant events close enough together to be associated by a learning rule with a biologically relevant timespan. Phase shifting allows place cell firings from two nearby places to be brought within a single theta-cycle of ≈ 120 ms and to remain in the correct temporal order. With these two tasks accomplished, the learning mechanism described here would associate the two events across a timespan of $\tau_A \approx 150$ ms, and the network would subsequently be capable of temporal compression during recall. The magnitude of this temporal compression would multiply the relatively smaller compressive effect of theta-phase shifting, significantly extending CR .

Besides the model presented here, another sequence learning model without oscillations is that of Blum and Abbot (1996). These investigators show that a network with a time-spanning synaptic modification rule similar to the rule used here is capable of short-term sequence prediction. However, these results are more related to continuous sequence prediction during exploration, rather than the high-speed recall of an entire learned sequence in the absence of external input demonstrated here.

A similar CA3 model based on McCulloch-Pitts neurons (Prepscius and Levy, 1994) also shows that temporal compression can be controlled by the level of inhibition during recall. This finding, with a much simpler network, emphasizes the generality of the present finding that CR depends on the activity level during testing (Fig. 5). However, the McCulloch-Pitts model shows jump-ahead compression, in which patterns in the middle of the sequence are skipped during recall, whereas the integrate-and-fire network exhibits speedup compression, in which each pattern is recalled for a shorter period of time than it lasts in the driven sequence. One reason for this qualitative difference in compression is that each pattern of the McCulloch-Pitts network lasts for only a single timestep. Thus, simple speedup is not possible because recalled patterns can not last for a fraction of a timestep. Another reason is that the study of the McCulloch-Pitts network used noncircular sequences, which form a weak attractor at the sequence end. The presence of this attractor may have increased the likelihood of jump-ahead behavior during recall.

Our laboratory has also recently observed speedup-type temporal compression in a different CA3-like model (Levy et al., 1997). This hybrid network combines the simplicity of the McCulloch-Pitts model with some (but not all) of the complexity of the full integrate-and-fire model. More specifically, the hybrid model lacks a membrane capacitance but still can associate patterns across multiple timesteps. Thus, the hybrid model represents time better than a network of McCulloch-Pitts cells but not as well as the present model. In agreement with our present results, CR is also proportional to τ_A in the hybrid model and, to a lesser extent, to activity levels.

4.4. *Functional Significance of Temporal Compression*

It is tempting to speculate on some possible functions of temporal compression with regard to memory

consolidation. Consolidation is the process by which certain memories are transferred from temporary storage in the hippocampus to their longer-term storage sites, presumably in the neocortex (McGaugh and Petrinovich, 1959; Squire and Zola-Morgan, 1991). During consolidation, the hippocampus may “teach” the cortex by repeatedly recalling its own memory traces while hippocampal-cortical synapses undergo associative modification (Alvarez and Squire, 1994; Buzsaki, 1996). However, the cortex may be a slower learner than the hippocampus (Marr, 1971; McClelland et al., 1994), requiring many more exposures to a sequence for synapses to modify. Temporal compression may facilitate this “hippocampal-neocortical dialogue” by allowing more teaching repetitions within a given period of time (Buzsaki, 1996).

Temporal compression may also facilitate chunking, the process by which several different events are combined into a single pattern of neural activity. For example, suppose that a network with an associative time constant of 160 ms can perform 20-fold compression of a sequence in which each pattern also lasts 160 ms. Thus, the network recalls this sequence as though each pattern only lasted for 8 ms. If the subsequent network layer also has an associative timespan of 160 ms, then it would be capable of associating events that were originally 3.2 seconds apart in time. Longer associative time-constants in the cortex will increase this effect proportionally, as will feeding compressed chunks rapidly back into the hippocampus.

4.5. *Experimental Predictions*

Based on our results with the τ_A/T parameter, several experimental predictions follow. First, as patterns become shorter, there will be more compression during subsequent slow-wave sleep. However, if patterns change too quickly or too slowly, learning will not occur. In theory, pattern duration can be changed by moving the rat around the track at different speeds, for example on a motorized cart. It has already been shown that rats exhibit robust place cell firing in this type of situation, so such experiments may be possible. Finally, GABA-ergic agents, by controlling activity, would be expected to change the observed amount of temporal compression. For example, adding GABA antagonists during slow wave sleep should increase temporal compression, while GABA agonists would decrease compression.

Acknowledgments

This work was supported in part by NIH MH10702 to DAA, by NIH MH48161, MH00622, EPRI RP8030-08, and Pittsburgh Supercomputing Center Grant BNS950001 to WBL, and by the Department of Neurosurgery, John A. Jane, chairman.

References

- Alvarez P, Squire LR (1994) Memory consolidation and the medial temporal lobe: A simple network model. *Proc. Natl. Acad. Sci. USA* 91:7041–7045.
- Amaral DG (1987) Memory: Anatomical organization of candidate brain regions. In: F Plum, ed. *Handbook of Physiology, Section I: The Nervous System V*. Oxford University Press, New York. pp. 211–293.
- Amaral DG, Ishizuka N, Claiborne B (1990) Neurons, numbers, and the hippocampal network. In: J Storm-Mathisen, J Zimmer, OP Ottersen, eds. *Understanding the Brain Through the Hippocampus: The Hippocampal Region as a Model for Studying Structure and Function*. Elsevier, pp. 1–11.
- August DA, Levy WB (1996) Temporal sequence compression by a hippocampal network model. In: INNS World Congress on Neural Networks. Erlbaum, Mahwah, NJ. pp. 1299–1303.
- Benington JH, Heller HC (1995) Restoration of brain energy metabolism as the function of sleep. *Prog. Neurobiol.* 45:347–360.
- Bernard C, Wheal HV (1994) Model of local connectivity patterns in CA3 and CA1 areas of the hippocampus. *Hippocampus* 4(5):497–529.
- Blum KI, Abbott LF (1995) Functional significance of long-term potentiation in recurrent networks. *Intl. J. Neural Syst. (Supp.)* 6:25–31.
- Blum KI, Abbott LF (1996) A model of spatial map formation in the hippocampus of the rat. *Neural Comp.* 8(1):85–93.
- Buhl EH, Halasy K, Somogyi P (1994) Diverse sources of hippocampal unitary inhibitory postsynaptic potentials and the number of synaptic release sites. *Nature* 368:823–828.
- Bunsey M, Eichenbaum H (1995) Selective damage to the hippocampal region blocks long-term retention of a natural and nonspatial stimulus-stimulus association. *Hippocampus* 5:546–556.
- Buzsaki G (1984a) Feed-forward inhibition in the hippocampal formation. *Prog. Neurobiol.* 22:131–153.
- Buzsaki G (1984b) Long-term changes of hippocampal sharp-waves following high-frequency afferent activation. *Brain Res.* 300:179–182.
- Buzsaki G (1986) Hippocampal sharp waves: Their origin and significance. *Brain Res.* 398:242–252.
- Buzsaki G (1989) Two-stage model of memory trace formation: A role for “noisy” brain states. *Neurosci.* 31(3):551–570.
- Buzsaki G (1996) The hippocampo-neocortical dialogue. *Cerebral Cortex* 6:81–92.
- Buzsaki G, Chrobak JJ (1995) Temporal structure in spatially organized neuronal ensembles: A role for interneuronal networks. *Curr. Opin. Biol.* 5:504–510.
- Buzsaki G, Horvath Z, Urioste R, Hetke J, Wise K (1992) High-frequency network oscillation in the hippocampus. *Science* 256:1025–1027.
- Chrobak JJ, Buzsaki G (1994) Selective activation of deep layer (V–VI) retrohippocampal cortical neurons during hippocampal sharp waves in the behaving rat. *J. Neurosci.* 14(10):6160–6170.
- Deacon TW, Eichenbaum H, Rosenberg P, Eckmann KW (1983) Afferent connections of the perirhinal cortex in the rat. *J. Comp. Neurol.* 220:168–190.
- Gulyas AI, Miles R, Sik A, Toth K, Tamamaki N, Freund TF (1993) Hippocampal pyramidal cells excite inhibitory neurons through a single release site. *Nature* 366:683–687.
- Gustafsson B, Wigstrom H (1986) Hippocampal long-lasting potentiation produced by pairing single volleys and brief conditioning tetani evoked in separate afferents. *J. Neurosci.* 6(6):1575–1582.
- Holmes WR, Levy WB (1990) Insights into associative long-term potentiation from computational models of NMDA receptor-mediated calcium influx and intracellular calcium concentration changes. *J. Neurophys.* 63(5):1148–1168.
- Holmes WR, Levy WB (1994) Temporal requirements for associative LTP in the dentate. Dependence on modeled r_m and r_i values. In: FH Eeckman, ed. *Computation in Neurons and Neural Systems*. Kluwer, Boston. pp. 299–304.
- Ishizuka N, Weber J, Amaral DG (1990) Organization of intrahippocampal projections originating from CA3 pyramidal cells in the rat. *J. Comp. Neurol.* 295:580–623.
- Jensen O, Lisman JE (1996) Hippocampal CA3 region predicts memory sequences: Accounting for the phase precession of place cells. *Learning and Memory* 3:279–287.
- Leonard BJ, McNaughton BL, Barnes CA (1987) Suppression of hippocampal plasticity during slow-wave sleep. *Brain Res.* 425:174–177.
- Levy WB (1989) A computational approach to hippocampal function. In: RD Hawkins, GH Bower, eds. *Computational Modeling of Learning in Simple Neural Systems*. Academic Press, Orlando, FL. pp. 243–305.
- Levy WB (1994) Unification of hippocampal function via computational considerations. In: INNS World Congress on Neural Networks, vol. 4. Erlbaum, NJ. pp. 661–666.
- Levy WB, Colbert CM, Desmond, NL (1995a). Another network model bites the dust: Entorhinal inputs are no more than weakly excitatory in the hippocampal CA1 region. *Hippocampus* 5:137–140.
- Levy WB, Sederberg PB (1997) A neural network model of hippocampally mediated trace conditioning. In: *IEEE International Congress on Neural Networks*, vol. 1, IEEE, NJ. pp. 372–376.
- Levy WB, Sederberg PB, August DA (1998) Sequence compression by a hippocampal model: A functional dissection. *Comput. and Neural Sys.*
- Levy WB, Steward O (1983) Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neurosci.* 8(4):791–797.
- Levy WB, Wu X (1996) The relationship of local context codes to sequence length memory capacity. *Network* 7:371–384.
- Levy WB, Wu X, Baxter RA (1995b) Unification of hippocampal function via computational/encoding considerations. In: *Proceedings of the Third Workshop: From Biology to High-Energy Physics*. Intl. J. Network Sys., 6 suppl.:71–80.
- Levy WB, Wu X, Tyrcha JM (1996) Solving the transverse patterning problem by learning context present: A special role for input codes. In: INNS World Congress on Neural Networks. Erlbaum, NJ. pp. 1305–1309.

- Lisman JE, Idiart MAP (1995) Storage of 7 ± 2 short-term memories in oscillatory subcycles. *Science* 267:1512–1515.
- Marr D (1971) Simple memory: A theory for archicortex. *Phil. Trans. Royal. Soc. Lond.* 262:23–81.
- McBain C, Dingledine R (1992) Dual-component miniature excitatory synaptic currents in rate hippocampal CA3 pyramidal neurons. *J. Neurophys.* 68(1):16–27.
- McClelland JL, McNaughton BL, O'Reilly RC (1994) Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. Technical Report PDP.CNS.94.1, Carnegie Mellon University.
- McGaugh JL, Petrinovich L (1959) The effect of strychnine sulfate on maze learning. *Am. J. Psychol.* 72:99–102.
- McNaughton BL, Barnes CA, Gerrard JL, Gothard K (1996) Deciphering the hippocampal polyglot: The hippocampus as a path integration system. *J. Exp. Biol.* 199:173–185.
- Miles R (1990) Synaptic excitation of inhibitory cells by single CA3 hippocampal pyramidal cells of the guinea-pig *in vitro*. *J. Physiol.* 428:61–77.
- Minai AA, Levy WB (1993a) The dynamics of sparse random networks. *Biol. Cybern.* 70:177–187.
- Minai AA, Levy WB (1993b) Sequence learning in a single trial. In: INNS World Congress on Neural Networks, vol. 2. Erlbaum, NJ. pp. 505–508.
- Minai AA, Levy WB (1993c) Setting the activity level in sparse random networks. *Neural Computation* 6:85–99.
- Muller RU, Kubie JL (1989) The firing of hippocampal place cells predicts the future position of freely moving rats. *J. Neurosci.* 9(12):4101–4110.
- O'Keefe J, Conway DH (1978) Hippocampal place units in the freely moving rat: Why they fire when they fire. *Exp. Brain Res.* 31:573–590.
- O'Keefe J, Reece ML (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3:317–330.
- Prepscius C, Levy WB (1994) Sequence prediction and cognitive mapping by a biologically plausible neural network. In: INNS World Congress on Neural Networks, vol. 4. Erlbaum, NJ. pp. 164–169.
- Quirk GJ, Muller RU, Kubie JL, Jr., Ranck JB (1992) The positional firing properties of medial entorhinal neurons: Description and comparison with hippocampal place cells. *J. Neurosci.* 12(5):1945–1963.
- Ranck JB (1973) Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. I. Behavioral correlates and firing repertoires. *Exp. Neurol.* 43:461–531.
- Ranck JB (1975) Behavioral correlates and firing repertoires of neurons in the dorsal hippocampal formation and septum of unrestrained rats. In: RL Isaacson, KH Pribram, eds. *The Hippocampus*, vol. 2. Plenum Press, New York. pp. 207–244.
- Redish AD, Touretzky DS (1997) Cognitive maps beyond the hippocampus. *Hippocampus* 7(1):15–35.
- Shen B, McNaughton BL (1996) Modeling the spontaneous reactivation of experience-specific hippocampal cell assemblies during sleep. *Hippocampus* 6:685–692.
- Skaggs WE, McNaughton BL (1995) The hippocampal theta rhythm and memory for temporal sequences of events. *Intl. J. Neural Sys.* 6:101–105.
- Skaggs WE, McNaughton BL (1996) Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science* 271:1870–1873.
- Skaggs WE, McNaughton BL, Wilson MA, Barnes CA (1996) Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 6:149–172.
- Smith AC, Levy WB (1998) Controlling activity fluctuations in large, sparsely connected random networks. *Neural Comput.* submitted.
- Spruston N, Jonas P, Sakmann B (1995) Dendritic glutamate receptor channels in rat hippocampal CA3 and CA1 pyramidal neurons. *J. Physiol.* 482(2):325–352.
- Squire LR, Zola-Morgan S (1991) The medial temporal lobe memory system. *Science* 253:1380–1386.
- Staley KJ, Mody I (1992) Shunting of excitatory input to dentate gyrus granule cells by a depolarizing GABA-A receptor-mediated postsynaptic conductance. *J. Neurophys.* 68(1):197–212.
- Steward O, Scoville S (1976) Cells of origin of entorhinal cortical afferents to the hippocampus and fascia dentata of the rat. *J. Comp. Neurol.* 169:347–370.
- Suzuki SS, Smith GK (1987) Spontaneous EEG spikes in the normal hippocampus. I. Behavioral correlates, laminar profiles and bilateral synchrony. *Electroenceph. Clin. Neurophys.* 67(4):348–359.
- Thompson LT, Best PJ (1989) Place cells and silent cells in the hippocampus of freely-behaving rats. *J. Neurosci.* 9(7):2382–2390.
- Treves A, Rolls ET (1992) Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus* 2(2):189–200.
- Tsodyks MV, Skaggs WE, Sejnowski TJ, McNaughton BL (1996) Population dynamics and theta rhythm phase precession of hippocampal place cell firing: A spiking neuron model. *Hippocampus* 6:271–280.
- Wallenstein GV, Hasselmo ME (1997) GABAergic modulation of hippocampal population activity: Sequence learning, place field development, and phase precession effect. *J. Neurophys.* 78(1), 1993, 408.
- Whishaw IQ, Jarrard LE (1996) Evidence for extra-hippocampal involvement in place learning and hippocampal involvement in path integration. *Hippocampus* 6(5):513–524.
- Wiener SI, Paul CA, Eichenbaum H (1989) Spatial and behavioral correlates of hippocampal neuronal activity. *J. Neurosci.* 9(8):2737–2763.
- Wilson MA, McNaughton BL (1994) Reactivation of hippocampal ensemble memories during sleep. *Science* 265:676–679.
- Witter MP (1993) Organization of the entorhinal-hippocampal system: A review of current anatomical data. *Hippocampus* 3:33–44.
- Wu XB, Levy WB (1995) Controlling performance by controlling activity levels in a model of hippocampal region CA3. I. Overcoming the effect of noise by adjusting network excitability parameters. In: INNS World Congress on Neural Networks. vol. 1. Erlbaum, NJ. pp. 577–581.
- Wu X, Levy WB (1996) Goal finding in a simple, biologically inspired neural network. In: INNS World Congress on Neural Networks. Erlbaum, NJ. pp. 1279–1282.
- Ylinen A, Bragin A, Nadasdy Z, Jando G, Szabo I, Sik A, Buzsaki G (1995) Sharp wave-associated high-frequency oscillation (200 Hz) in the intact hippocampus: Network and intracellular mechanisms. *J. Neurosci.* 15(1):30–46.