T-maze Training of a Recurrent CA3 Model Reveals the Necessity of Novelty-Based Modulation of LTP in Hippocampal Region CA3

J. D. Monaco and W. B Levy
Dept. of Neurological Surgery
University of Virginia
P.O. Box 800420
Charlottesville, VA 22908-0420

Abstract-**The rat hippocampus has been shown to mediate a large set of spatial navigation tasks such as the simple T-maze. We investigated the performance of a minimal computational, but biologically based, model of CA3 on this task. For successful performance, the model needs to generate and maintain neuronal codes for each of the two arms of the T-maze. Moreover, each code must be distinctively recalled in a goal-dependent manner. The development of such neuronal codes is aided by the appearance of repetitively firing recurrent neurons – known as local context units, analogous to hippocampal place cells – which promote spatiotemporal association within the T-maze training sequences. The number, longevity, and connectivity of local context units exclusively coding for each arm of the maze grow with training. Although with too much training, the coding for one arm uncontrollably dominates over the other code, and goal appropriate choice-behavior is lost. That is, successful network codes can easily deteriorate with overtraining. The amount of training that produces this deterioration in performance depends on other network parameters. Rather than a failure of the model, we believe these results tell us something important about the biology of the hippocampal system. That is, this result provides support for the hypothesis of a hippocampal afferent system which down-regulates LTP once a task has been successfully learned. Modulatory systems (e.g., dopaminergic, generally D1/D5r) exist which are candidates for this functional role.**

I. Introduction

The T-maze task has a long history as an important tool of behavioral psychology [1-5]. Tolman studied rats running complex T-mazes and systematized the theory of cognitive maps [1,6]. Others have used it to investigate behaviors in rats ranging from odor discrimination [4] to height aversion on an elevated maze [7]. It is also one of the simplest spatial navigation paradigms, and as such is a good minimal case for studying goal-dependent sequence learning and recall in hippocampal memory formation. Positive reinforcement tasks provide food or a more pleasant odor at one of the arms of the T-maze more often than the other [1-5], and the rats learn quickly where to go. That is, the animals have both a goal in mind and a learned spatial map allowing them to get there when being tested. In this paper, we apply an established computational model [8-11] of the CA3 region of the rat hippocampus to goal finding in a T-maze scenario.

Many recent studies have investigated the modulatory effect of novelty on hippocampal plasticity and the strength of memory formation. Even brief exposure to novelty for neonatal rats was shown to elicit a lasting enhancement of hippocampal long-term potentiation (LTP) [12]. Physiologically, novel environments were shown to increase the rate of phosphorylation of the protein CREB in pyramidal cells [13]; this process seems to lead to the generation of new dendritic spines in female rats [14]. Increases in postsynaptic spine density might serve to enhance hippocampal excitability as well as LTP. Such observations suggest that stronger hippocampal memories are formed in novel situations than familiar ones. Indeed, exposure to a novel environment just prior to a single trial avoidance task has a significant enhancing effect on task recall [15]. The hippocampal formation receives projections, primarily at D1/D5 receptors, that possibly modulate long-term synaptic modification [16]. It then projects directly to nucleus accumbens (N.Acc) from the ventral subiculum [17]; this projection is activated within the hippocampal formation by novel stimuli and increases extracellular dopamine in the N.Acc [18].

The implication of hippocampal involvement in a novelty-reward pathway is clear with regard to the learning of various cognitive tasks. Novel stimuli, such as those at the beginning of a training session, both reward the animal and enhance the hippocampal LTP necessary to learn the task. A reasonable hypothesis then emerges regarding continued training on the same task: there will be a point at which more training is unnecessary and may even be detrimental. This is the point at which the animal becomes bored with the exercise, as it is no longer stimulating its dopamine system, and the heightened plasticity of the hippocampus is down-regulated. That is, learning is effectively turned off once the task is learned. The demonstration of a simple hippocampal-mediated task whose solution requires LTP down-regulation would provide strong support for this hypothesis. As described above, the necessary novelty-reward mechanisms are in place [12-18].

In the present report, we find that overtraining of this CA3 model on the T-maze task leads to increased similarity between the recurrent neuronal firing that codes for each of the two arm subsequences. Such similarity destroys previously successful learning. That is, we found a large class of biologically reasonable networks which are successful T-maze learners after a given number of training trials, but then fail at some point if training is continued. Since rats quickly learn T-mazes [1-6], the model predicts that the hippocampal aspect of this learning be down-regulated as learning proceeds.

Overview of the Model

The CA3 model is a sparsely connected, asymmetric, recurrent neural network composed of McCulloch-Pitts units [8]. As described in [9], synaptic modification is based on a local Hebbian rule with a time-spanning associative capability. The total network activity comes from both externally driven input units and recurrent processing units, with the majority of network activity resulting from the latter. As such, the model processes spatiotemporal sequences of entorhinal cortex/dentate gyrus input.

Two sequence learning tasks that the model successfully solves are subsequence disambiguation and goal finding [9,19]. They are complementary in the sense that they are trained identically yet tested differently. In these tasks, the network is trained on two sequences which share a middle subsequence (see Fig. 2b in [9]). In testing for disambiguation, the proper sequence must be recalled when prompted with the beginning of either sequence. In testing goal finding, a goal code is present which represents the end of one sequence, and recall for both sequences must be completely dependent on the goal code. Behaviorally, a rat learns to associate

0-7803-7898-9/03/$17.00 ©2003 IEEE

the end of one of the two sequences with a reward, such as food. During testing, the animal is motivated by a goal, such as hunger, and finds the path to the spatial location that it associated with food during training. In this form of the goal finding task, a novel path to the goal must be found. That is, for success, a goal code must cause a simulated network to recall its respective goal pattern.

The T-maze task is similar to the above goal finding task, except that there are only two subsequences. Each of the two training sequences are composed of a shared "stem" subsequence followed by one of two nonshared (i.e., orthogonal) "arm" subsequences. These are analogous to the stem and arm pathways on a physical T-maze, where the final patterns of each of the two arms are the goal boxes. Simulations, to be successful, must demonstrate flexible, goal-code-dependent recall. Successive patterns of recurrent activity called local context codes are characteristic of our CA3 model's ability to resolve the inherent ambiguity of this task [20]. Local context neurons are recurrent neurons that identify a subsequence of a larger sequence, usually by firing repetitively in response to its particular input subsequence [9]. Accordingly, these neurons are analogous to hippocampal place cells (for review, see [21]). By firing continuously over more than one pattern in the sequence, they allow time-spanning associativity to flexibly associate temporally nonadjacent input patterns [9]. This repetitive firing of the model allows us to talk about attractors in the state space of neuronal codes [9]. Indeed, the neuronally encoded goals at the end of each T-maze arm are two fixed-point attractors learned by the network. The task then tests the network's ability to induce one or the other attractor.

II. Model and Methods

A. The Network Model

The hippocampal CA3 model is a sparsely and randomly connected recurrent network. The probability of a recurrent interconnection is 10%. The neurons are simple McCulloch-Pitts binary units. Synaptic weights are modified using a temporally asymmetric rule of association [22]. This asymmetry allows the recurrent network model to form context codes, which are critical to its problem-solving capabilities [8,9,20]. The Hebbian-like learning rule is

$$W_{ij}(t+1) = W_{ij}(t) + \mu Z_j(t)\left(\overline{Z}_i(t-1) - W_{ij}(t)\right), \quad (1)$$

where W_{ij} is the weight of the synapse from neuron i to neuron j, μ is the learning rate constant, Z_j is a binary indicator of the firing state of neuron j, and $\overline{Z}_i(t)$ is the neuronal signal decay of neuron i at time-step t such that

$$\overline{Z}_i(t) = \begin{cases} 1, & \text{if } Z_i(t) = 1 \\ \alpha \cdot \overline{Z}_i(t-1), & otherwise \end{cases}, \quad (2)$$

where α determines the off-rate time constant of the NMDA receptor. For the data in this report, we used $\alpha = 0.4$ and a learning rate constant of $\mu = 0.5$. Also, all neuronal firing was determined using a *k-winners-take-all* competitive paradigm. All externally driven neurons for a given timestep t are indicated by the binary vector $x(t)$, and $Z_i(t) = 1$ whenever $x_i(t) = 1$. The internal excitation for neuron j is

$$y_j(t) = \sum_{i=1}^{n} c_{ij} \cdot W_{ij}(t-1) \cdot Z_i(t-1), \quad (3)$$

where c_{ij} is a binary indicator of a synaptic connection from neuron i to neuron j. For recurrently excited neuron j, $y_j(t)$ determines that $Z_i(t) = 1$ if it is among the k largest neural excitations. Otherwise, $Z_i(t) = 0$.

The value k is the number of neurons that need to be active in order to maintain the predetermined network activity level a, which includes externally driven activity. For a network of size n, k is the largest integer less than na. If the lowest excitation value of the *k-winners* is shared among several neurons, not all of which are allowed to fire, then a random subset of these neurons is fired.

B. Input Sequences, Training, and Testing

There are two training sequences for the T-maze task, one each for the "left path" and "right path" of the maze. They share the same stem subsequence. The stem is 6 patterns long, while each of the arms is 4 patterns. The number of neurons in each input pattern in a sequence is the same, and is determined by the network size (n), total network activity level (a), and external fraction of activity (m_e). The variable a, as a fraction of the network size, determines how many neurons, both recurrent and externally driven, are active per timestep. The variable m_e specifies the number of neurons assigned to any given external input pattern as a fraction of the number of active neurons. For all data in this report, network size is constant at $n = 4096$, but a and m_e are independently variable.

Input firing patterns within subsequences are slowly shifting. That is, temporally adjacent patterns are spatially adjacent in that they share one third of their neurons. However, the three input subsequences that constitute the T-maze (i.e., stem, "left side" arm, and "right side" arm) are mutually orthogonal. Also, each of the 10 patterns in a training sequence are repeated ("stuttered") for 3 timesteps [10]. So, each input sequence is 30 timesteps long, with the arm subsequences beginning at $t = 19$.

A single training trial consists of the presentation of both training sequences to the network. The network is initialized for every sequence presentation with a random firing vector at the network activity level. Synaptic modification is turned off during the testing trials. Testing consists of two sequence trials in which the stem input subsequence is presented to the network. This is analogous to a rat physically moving though the T-maze up to the point where a decision must be made. From the beginning of the sequence and up to the final goal pattern in the first trial, 25% of the external neurons representing the goal pattern of the "left side" arm subsequence are also turned on. This is the goal code [9]. A normalized cosine of the activity of the external input neurons determines whether or not the network recalls the "left side" goal. The same process is done for the "right side" arm subsequence. If both the "left" and "right" goals are appropriately recalled, for at least 8 out of 10 pairs of random goal codes, then this network learned the T-maze.

III. Code Similarity and T-maze Performance

Successful learning of the T-maze task is mediated by the development during training of learned firing patterns, called local context codes [9]. The number of local context (LC) neurons in any particular simulation for a given input sequence depends on many different factors, the most general being the proportion of recurrent to externally driven activity [19]. The firing patterns of LC neurons allow the network to flexibly learn the two paths leading away from the end of the T-maze stem. However, T-maze performance is not linearly correlated with either the average length of context codes or the number of LC neurons (data not shown).

When we discuss the "codes" developed during a simulation, we are referring to the sequences of recurrent neuronal firing that occur during the presentation of the two input sequences within a given training trial. If we restrict the discussion to the "codes" for the arm subsequences, then we are only interested in the sequences of recurrent firing in the range of timesteps from $t = 19$ to $t = 30$. So,

the codes recalled during a training trial always consist of both "left side" sequence codes and "right side" sequence codes, regardless of which subsequences are being discussed. Given this, it is always possible to refer to the degree of similarity between the "left" and "right" training codes, which we term the "between-sequence similarity". Of course, this similarity is timestep dependent, since it is possible for the "left" codes to be very similar to the "right" codes during one range of timesteps of a trial and orthogonal to them during another range.

A. Arm Subsequence Similarity

This CA3 model develops an end of sequence attractor [9]. In the language of attractors, a successful simulation has developed two noisy fixed-point goal attractors. When tested, perturbations in the direction of one or the other goal attractor are able to influence the otherwise free recall of the network. The recall trajectory, through the state space of neuronal firing vectors, then falls into the basin of attraction of the appropriate goal attractor. This is what is meant by a goal code's ability to "induce" a goal attractor. The network training codes for the two goal patterns must be nearly orthogonal since any similarity only decreases the efficacy of goal codes. That is, as the linear dependence of the attractor points increases, so does the probability that a random goal code will push the recall trajectory toward both goal attractors simultaneously. Successful simulations are those that have resolved such ambiguity.

So, successful T-maze learning requires very low similarity in the recurrent codes for each goal pattern. However, the solution requires flexibility in the associations between recurrent LC neurons and external neurons representing the patterns of the arms. That is, there must be a set of recurrent neurons which fire repeatedly over two or more nongoal patterns and which innervate, directly or indirectly, external neurons from both arm subsequences. These LC neurons must themselves be activated during both sequences in a training or testing trial. This can only happen if they are sufficiently innervated by externals in both arm subsequences or are activated by the end of the stem subsequence. There must be a set of such recurrent context neurons to develop the flexible goal associations necessary to the solution of this task. That is, substantial similarity between the recurrent codes at the beginning of the arm subsequences is necessary for good T-maze performance.

Thus, there is a tension between the two requirements: the recurrent codes that are goal attractors must be orthogonal, but there must be a substantial number of common LC neurons that fire at the beginning of both arm subsequences. For a simulation to be successful, then, there must be a progression in the between-sequence similarity of the arm subsequences from substantial similarity, starting around timestep $t = 19$, to orthogonality, occurring before timestep $t = 30$.

B. A Useful Measure of Between-Sequence Similarity

We define similarity in terms of a normalized cosine between two recurrent firing vectors. Let $z_L(t)$ and $z_R(t)$ be the recurrent network states at timestep t for the "left" and "right" training sequences, respectively, of a particular training trial. The timestep-dependent between-sequence similarity, $s(t)$, is

$$s(t) = \frac{z_L(t) \bullet z_R(t)^T}{na}, \forall t \in \{1,...,30\}. \quad (4)$$

A value of $s(t) = 1$ indicates that the "left" and "right" recurrent network states are identical at timestep t; between-sequence orthogonality is indicated by $s(t) = 0$. As discussed below and in the statement of our results, the timestep t at which the half-maximal similarity occurs will be important. First, define the maximal similarity between the "left side" and "right side" recurrent training codes,

$$m = \max_{t \in \{1,...,30\}} [s(t)]. \quad (5)$$

Then, define the timestep nearest the end of the sequence that produces half of that value,

$$B_{SIM} = \max\left[t: s(t) \geq \frac{m}{2}\right]. \quad (6)$$

Formally, we call this timestep the "similarity boundary". It measures the last time during the sequence, from the beginning to the end, when the recurrent codes for the "left side" and "right side" training sequences are mostly similar to each other.

According to the account of T-maze performance presented in III.A, the between-sequence similarity of the arm subsequences is critical. The maximal similarity m does not vary much and always occurs during the stem codes. That is, the stem codes are a trivial case during training since the network receives the same external inputs for the stem subsequence in both sequences. There is some variability in the similarity of the stem codes, due to the random processes in the network (i.e., competitive firing and sequence initialization), but m will always be very close to 1. So, the similarity boundary B_{SIM} is based on a the maximal similarity m of a given simulation as a normalizing factor.

However, for successful performance to be possible, the arm codes must diverge; that is, the arm codes must tend toward orthogonality as t approaches the end of the sequence. This requirement, coupled with the necessity of nonzero similarity at the beginning of the arm subsequence, implies that between-sequence similarity must decrease to 0 as the timestep increases to 30. Thus, for any given similarity value s, which is less than the maximum, we can find a timestep t_s such that all timesteps up to and including t_s yield between-sequence similarity greater than or equal to s, and all later timesteps yield similarity values less than s. This timestep can be considered a "boundary" dividing the sequence based on a chosen similarity value. Letting $s = m/2$, then, we can see that $t_s = B_{SIM}$, where m and B_{SIM} are defined in (5) and (6), respectively. So, B_{SIM} is a similarity boundary, a timestep that partitions an ordered sequence based on between-sequence similarity. Certainly, there are other possible measures of subsequence similarity. However, this particular definition allows a necessary, yet insufficient, condition for successful T-maze performance to be inferred, in the form of a critical range for B_{SIM} at the end of training.

C. Success and Failure Modes

We define a specific set of identically parameterized simulations to be successful if at least 80% of randomly seeded simulations successfully solve the given task. Every such set in this report is comprised of 15 simulations. Graphs of a typical T-maze success are shown in Figs. 1a-d. Note that in the two testing trials, Figs. 1b and 1c, the appropriate external arm subsequence neurons, up to the end of the sequence, are activated in each case. One of the cosine similarity diagrams used for network decoding is depicted in Fig. 1d. It is apparent that the first 3 patterns of each arm subsequence are active regardless of which goal code is present. However, the goal codes are able to influence recall during testing and induce the goal attractor for each of the arm subsequences.

There are two failure modes, each arising from opposing difficulties in solving the T-maze. The first, call it a Type I failure, is evident in parameterizations with a low number of active recurrents during training and testing. This reduces the number of possible LC

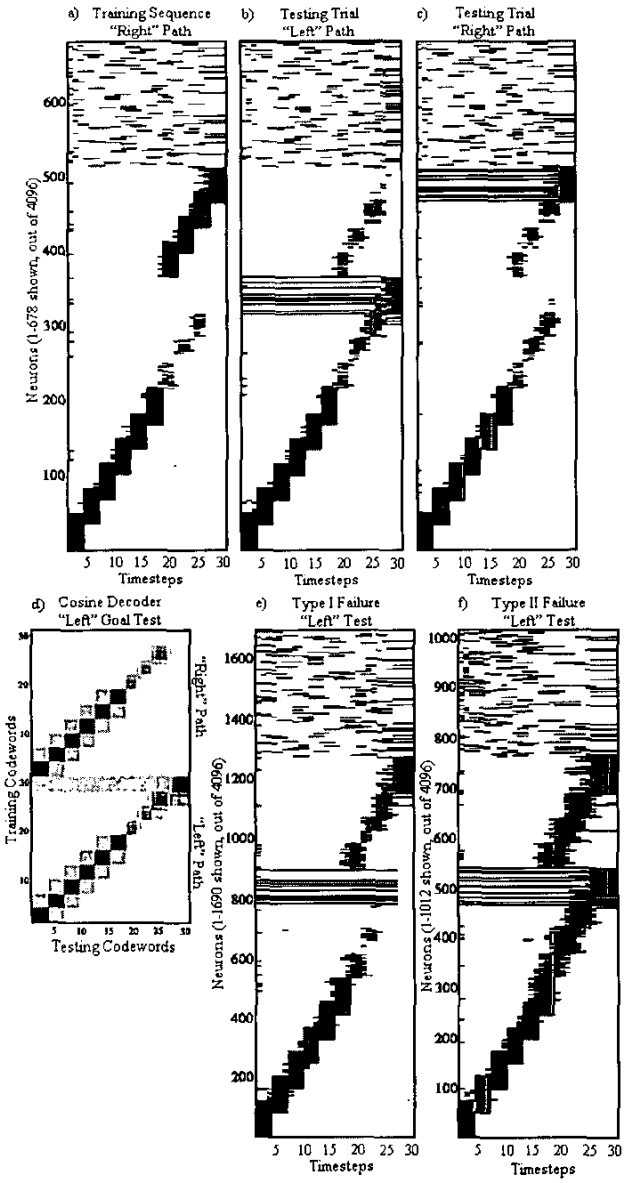

Fig. 1. Neuronal firing diagrams and an example cosine similarity matrix. Each of the firing diagrams shows all external neurons and a small subset of the recurrent neurons in the network. A "right path" training sequence for a typical successful network (m_e = 0.15, a = 0.08) is shown in (a). Note the concurrent "left side" arm firing that ceases around t = 27. The complete testing trial for this network is shown in (b) and (c). The horizontal lines are the fractional goal codes activated for the duration of the testing trial, except during the final pattern. The external neurons for both arm subsequences fire in both cases except for the appropriate goal pattern. The cosine similarities, as defined in (4), between the codewords recalled during the "left" testing sequence in (b) and the "left" and "right" training codewords are shown in the matrix (d). (In (d), a black cell indicates code identity, a white cell indicates orthogonality, and gray cells are intermediate.) This type of similarity matrix is used to decode network testing recall. Notice that the testing codewords at the end of the sequence are strongly similar to the final codewords of the "left" training sequence and orthogonal to those of the "right" training sequence. This is decoded as being a "left" choice. Finally, the neuronal firing diagrams for "left" testing sequences are shown for example Type I and II failure modes in (c) and (f), respectively.

neurons. Of those LC neurons, only a few, if any, will innervate the external neurons of both arm subsequences. As discussed above, this directly affects the ability of a network to have goal-code-dependent recall of both goals. Thus, a Type I failure mode occurs when a simul-ation predicts a single goal to either of the two test input sequences. For example, the simulation chooses the "right side" when the "right side" is the correct answer and when the "left side" is the correct answer. Fig. 1e contains a neur-onal firing diagram showing a Type I failure, in which this is the case.

All other failed simulations are Type II failures. A Type II failure results from parameterizations which activate a large number of recurrent neurons per timestep. A firing diagram of a typical Type II failure is shown in Fig. 1f. Due to the development of numerous LC neurons, training only serves to increase the between-sequence similarity of the recurrent codes for both arm subsequences, including the two goal patterns. The end result is that a single goal attractor is created at the end of the sequence, making the induction of separate goals impossible. A Type II failure mode occurs when a simulation predicts both goals simultaneously. In Fig. 1f, a single goal code is present, yet both external goal patterns are recalled.

IV. Results

We examined the T-maze performance of a set of 15 network connectivities across broad ranges of both total network activity (a) and the external fraction of activity (m_e). We will show that robust T-maze performance necessarily depends on learned recurrent codes that correlate with having a similarity boundary within a critical range. Finally, we show that the value of the similarity boundary is training dependent. It tends to move towards the end of the sequence as training continues. As a result, overtraining destroys a previously learned solution.

A. Network Parameterization and Task Performance

We assessed the T-maze performance of a large set of network parameterizations, in which all simulations were trained for 40 trials. The results are shown as a contour plot in Fig. 2. The external fraction of activity, m_e, was tested from 0.100 to 0.300 in increments of 0.025. This is a broad, but reasonable, range of m_e for this CA3 model [19]. Total network activity was then tested from 0.06 to 0.13 in increments of 0.01 for each m_e. Similarly, this is a broad, yet biologically reasonable, range for the activity level in a 4096-neuron simulation [24]. There are no successful parameterizations with m_e = 0.1, and only the 7% activity level was successful with m_e = 0.125.

Similarly, only the 8% activity level produced successful simulations with m_e = 0.15. Lower activities result in Type I failures while at higher activities all failures were Type II. At 10% activity and up to 13%, none of the simulations learned the task. Now, consider parameterizations with m_e = 0.175, 0.200, and 0.225. These values of external activity are closer to those that lead to robust performance on the original goal finding task [19]. No simulations were able to learn at 6% activity, but all 15 performed successfully at 9% and 10% activity levels (except that only 11/15 of those with m_e = 0.175 at 10% activity learned the task). Again, the low activity failures are of Type I, and higher activity levels produce Type II failures. There is another region of successful parameterizations at higher external fractions of activity. Simulations with m_e = 0.250, 0.275, and 0.300 were successful at both 11% and 12% activity levels. Each simulation at 8% and lower activity levels was a Type I failure, and each simulation at 13% activity was a Type II failure.

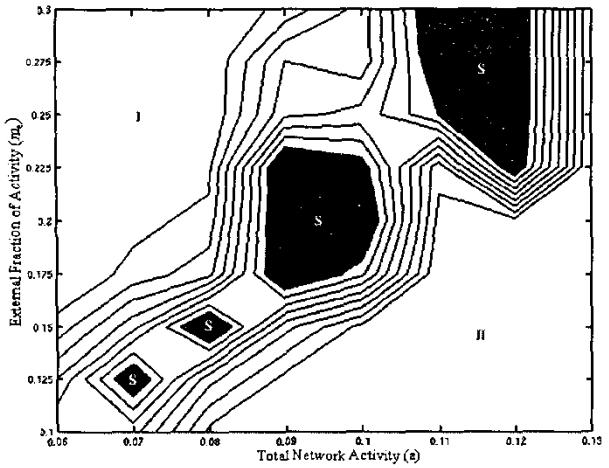

Fig. 2. T-maze success across both the total network activity (a) and the external fraction of activity (m_e). The two-dimensional parameter grid defined by m_e = {0.100, 0.125, ..., 0.300} and a = {0.60, 0.70, ..., 0.13} was tested for task performance after 40 training trials. A contour plot of this performance, as a fraction of successful simulations out of the 15 tested for each parameterization, is shown. The 80% contour is shaded in, representing the parameter values that lead to successful performance during testing. It should be noted that there are no data points between the apparently disjoint regions of success (i.e., those shaded in and marked "S"); that is, the evident division of the success region is an artifact of the interpolation necessary of a contour plot. The regions characterized by Type I or Type II failures are marked "I" and "II", respectively. There were 260 successful network simulations in the total 1,080 simulations of this data set.

B. *Critical Range of the Similarity Boundary*

Each of the 260 successful simulations in this data set of 1,080 total simulations has a similarity boundary in the range from t = 21 to t = 30. Fig. 3a shows the conditional probability of the success of a simulation given its similarity boundary at the end of 40 training trials. Each bar represents the fraction of simulations that learned successfully out of all simulations with a particular B_{SIM}. Task success correlates strongly with a similarity boundary between timesteps 24 and 26, with success most likely at t = 25. Specifically, 95.3% of simulations with a similarity boundary of B_{SIM} = 25 succesfully learned the T-maze. Simulations having B_{SIM} = 24 and 26 have similarly high percentages of successes: 85.7% and 77.2%, respectively. Outside of that range, successful T-maze performance becomes very unlikely and much less robust. Specifically, no simulations having B_{SIM} = 19 or 20 solved the task. Also, simulations with similarity boundaries between timesteps 21 and 23 have less than a 40% probability of performing the task. Lastly, of the 388 simulations with codes such that the goals are similar, having B_{SIM} = 30, only one was successful. This correlation of performance with a small range of the similarity boundary timestep is especially significant considering that it emerged from a diverse set of simulations. With a narrow critical range, it becomes important to see how the similarity boundary can change over the course of training.

C. *Similarity Boundary Dependence on Training*

The similarity between the codes of the training sequences is related to the ability of these networks to perform the T-maze task. As measured by the similarity boundary B_{SIM}, this dynamic aspect of the networks can change greatly with training. An example of a network simulation, with 8% network activity and an external activity of m_e = 0.15, is shown in Fig. 3b. Both B_{SIM} and the performance of the network are plotted across 120 training trials. Performance is measured as the fraction out of 15 goal code pairs that recall their respective goals. At trial 0, and up to trial 20, the similarity boundary comes before timestep 24 and testing results in the Type I failure mode. Between trials 30 and 40, the similarity boundary is between timesteps 24 and 25, and 100% performance is achieved at both of those trials. After trial 50, B_{SIM} increases to timestep 30 and there is no successful performance from that point on. All of these failures are of Type II. Thus, there is a very small window in which this particular simulation learned the task. Continued training led to the deterioration of the encoding of the task that made successful T-maze performance possible.

A comparison of this data set to a similar set, where each simulation was trained for 65 trials, reveals that only a very small

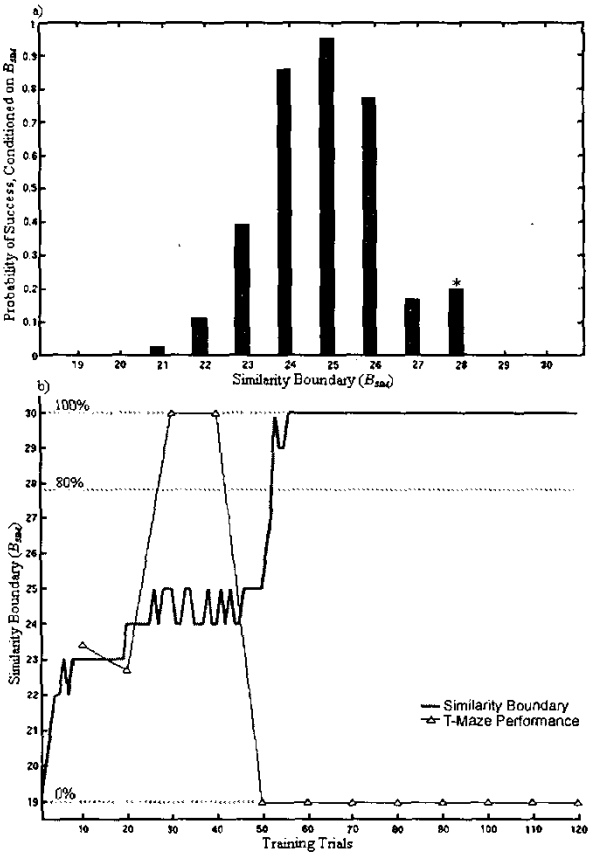

Fig. 3. T-maze performance depends critically upon the timestep at which the recurrent codes for the two arms of the maze cease to be mostly similar. This timestep, measured by the similarity boundary B_{SIM}, changes with the amount of training. Thus, the simulations can be overtrained such that the similarity boundary moves to the end of the sequence, destroying learned performance. A graph of the conditional probability of the success of a simulation, given the B_{SIM} of that simulation after training, is shown in (a). Thus, each bar in (a) represents the fraction of simulations, from the data set in Fig. 2, having a given B_{SIM} which are successful. (Only 5 simulations in the data set had B_{SIM} = 28, so the corresponding bar in (a) is due to only a single successful simulation.) A typical example of overtraining a simulation is shown in (b). In this figure, the changing value of B_{SIM} is shown, across training, along with the task performance of the simulation (measured every 10 training trials). At trials 10 and 20, this simulation is characterized by theType I failure mode. The simulation is only successful for trials 30 and 40, for which B_{SIM} = 24 or 25. Additional training only serves to increase the similarity boundary to the point that task performance results in the Type II failure mode.

region (i.e., around the parameterization with $m_e = 0.225$, $a = 0.10$) that was successful at 40 training trials remained successful (data not shown). That is, a mere 15 additional trials of training effectively destroyed all of the successful learning evident after just 40 trials.

V. Discussion

The T-maze problem is a special learning problem for a number of reasons. It is the simplest task requiring the induction of attractors, as in goal finding and other tasks (e.g., [19]). It is also a fundamental unit of many spatial learning and navigation tasks [2-5]. Tolman used a complex array of T-mazes to systematize the theory of cognitive maps [1,6], which laid the foundation for hippocampal place cell research [21]. In this report, we showed that a biologically reasonable, computational model of the rat CA3 can learn the T-maze over a wide range of parameterizations. Further, the task solution is critically dependent upon the similarity of the codes for the two arms of the maze. The neuronal codes for the stem of the T-maze will be identical, and local context codes facilitate the sharing of neurons at the beginning of the arm subsequences. However, this similarity must disappear at some point before the goal. In terms of the number of training trials, there is a narrow window when this between-sequence similarity transition is appropriate such that a simulation shows good learned performance.

We found that the point of this within trial transition in the sequence varies in a training dependent manner. This point comes increasingly later with more training, making successful performance more unlikely. This training dependent change in between-sequence similarity is mediated by local context codes, the very same coding feature which allows the network to solve the problem in the first place. Thus, there is a quantitative subtlety as to what constitutes a good encoding.

By documenting a large class of biologically reasonable parameterizations that learn the task well, but that can also be easily overtrained, we conjecture the necessity for a system capable of turning off hippocampal LTP once good learning has taken place.

There are plausible systems in place that could function in this role. For instance, the rat hippocampus receives dopaminergic projections from the ventral tegmental dopaminergic system [16], and projects to nucleus accumbens [17,18]. It is also known to function in the discrimination of novel experiences [13]. Additionally, novel stimuli have been shown to have modulatory effects on hippocampal LTP [12,15]. Therefore, the conjunction of the novelty and dopamine systems is at least a candidate for providing the hippocampal shut-off function necessitated by our study of the T-maze problem. That is, a rat that begins training on an unfamiliar task will be doing so with potentially stronger LTP, enabling it to learn reliably and fast. Training continues to the point at which the rat learns the task. Additional training serves to decrease the novelty of the situation, thus down-regulating LTP. This prevents the coding problem that destroyed, with overtraining, the learned solution developed by our networks. A rat that is familiar with the learned task, yet forced to continue training, will maintain its original solution.

This hypothesis seems reasonable, but certainly other possible systems exist which could serve this function. The primary result here is that the existence of such a system appears to be necessary to the maintenance of learned tasks in the rat hippocampus.

Acknowledgment

This work was supported by the John A. Harrison III Undergraduate Research Award of the University of Virginia Faculty Senate to J.D.M., and NIH MH48161, MH63855 and RR15205 to W.B.L.

References

[1] E.C. Tolman, B.F. Ritchie, and D. Kalish, "Studies in spatial learning. I. Orientation and the short-cut," *J. Exp. Psychol.: General*, vol. 36, pp. 13-24, 1946.
[2] R.A. Littman, "Latent learning in a T-maze after two degrees of training," *J. Comp. Physiol. Psychol.*, vol. 43, pp. 135-147, 1950.
[3] J.R. Ison and H. Kniaz, "T-maze performance as a function of consummatory activity," *Psych. Rep.*, vol. 12, pp. 107-110, 1963.
[4] J. Mendelson and S.L. Chorover, "Lateral hypothalamic stimulation in satiated rats: T-maze learning for food," *Science*, vol. 149, pp. 559-561, 1965.
[5] P.F. Southall and C.J. Long, "Odor stimuli, training procedures, and performance in a T-maze," *Psychonomic Science*, vol. 24, pp. 4-6, 1971.
[6] E.C. Tolman, "Cognitive maps in mice and men," *Psychol. Rev.*, vol. 55, pp. 189-208, 1948.
[7] H. Zangrossi Jr., et al., "Serotonergic regulation of inhibitory avoidance and one-way escape in the rat elevated T-maze," *Neurosci. Biobehav.*, vol. 25, pp. 637-645, 2001.
[8] W.B. Levy, "A computational approach to hippocampal function," in *Computational Models of Learning in Simple Neural Systems*, R.D Hawkins and G.H. Bower, Eds. New York: Academic, 1989, pp. 243-305.
[9] W.B. Levy, "A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks," *Hippocampus*, vol. 6, pp. 579-590, 1996.
[10] X.B. Wu, J. Tyrcha, and W.B. Levy, "A neural network solution to the transverse patterning problem depends on repetition of the input code," *Neurocomputing*, vol. 26, pp. 601-607, 1999.
[11] X.B. Wu and W.B. Levy, "Simulating the transverse non-patterning problem," *Neurocomputing*, vol. 44, pp. 1029-1034, 2002.
[12] H.E. Viola, et al., "Phosphorylated cAMP response element binding protein as a molecular marker of memory processing in rat hippocampus: effect of novelty," *J. Neurosci.*, vol. 20, art. no. RC112, 2000.
[13] A.C. Tang and B. Zou, "Neonatal exposure to novelty enhances long-term potentiation in CA1 of the rat hippocampus," *Hippocampus*, vol. 12, pp. 398-404, 2002.
[14] M. Segal and D.D. Murphy, "Estradiol induces formation of dendritic spines in hippocampal neurons: Functional correlates," *Horm. Behav.*, vol. 40, pp. 156-159, 2001.
[15] L.A. Izquierdo, et al., "Novelty enhances retrieval: molecular mechanisms involved in rat hippocampus," *Eur. J. Neurosci.*, vol. 13, pp. 1464-1467, 2001.
[16] S.L. Erickson, S.R. Sesack, and D.A. Lewis, "Dopamine innervation of monkey entorhinal cortex: postsynaptic targets of tyrosine hydroxylase-immunoreactive terminals," *Synapse*, vol. 36, pp. 47-56, 2000.
[17] C.D. Blaha, C.R. Yang, S.B. Floresco, A.M. Barr, and A.G. Phillips, "Stimulation of ventral subiculum of the hippocampus evokes glutamate-receptor mediated changes in dopamine efflux in the rat nucleus accumbens," *Eur. J. Neurosci.*, vol. 9, pp. 902-911, 1997.
[18] M. Legault and R.A. Wise, "Novelty-evoked elevations of nucleus accumbens dopamine: dependence on impulse flow from the ventral subiculum and glutamatergic neurotransmission in the ventral tegmental area," *Eur. J. Neurosci.*, vol. 13, pp. 819-828, 2001.
[19] S. Polyn, X.B. Wu, and W.B. Levy, "Entorhinal/dentate excitation of CA3: a critical variable in hippocampal models," *Neurocomputing*, vol. 32, pp. 493-499, 2000.
[20] X.B. Wu, R.A. Baxter, and W.B. Levy, "Context codes and the effect of noisy learning on a simplified hippocampal CA3 model," *Biol. Cybern.*, vol. 74, pp. 159-165, 1996.
[21] P.J. Best, A.M. White, and A.A. Minai, "Spatial processing in the brain: the activity of hippocampal place cells," *Annu. Rev. Neurosci.*, vol. 24, pp. 459-486, 2001.
[22] W.B. Levy and O. Steward, "Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus," *Neuroscience*, vol. 8, pp. 791-797, 1983.
[23] W.B. Levy and X.B. Wu, "The relationship of local context codes to sequence length memory capacity," *Network-Comp. Neural*, vol. 7, pp. 371-384, 1996.
[24] A.A. Minai and W.B. Levy, "Setting the activity level in sparse random networks," *Neural Comp.*, vol. 6, pp. 85-99, 1994.