

Goal Finding in a Simple, Biologically Inspired Neural Network

Xiangbao Wu & William B Levy
Department of Neurological Surgery
University of Virginia Health Sciences Center
Charlottesville, Virginia 22908, USA
xw3f@virginia.edu & wbl@virginia.edu

Abstract. Recurrent networks are often used as content addressable memories and as an explanatory hypothesis for human cognitive pattern completion (Anderson et al. 1984, Grossberg 1976). In addition, sparsely connected recurrent networks trained for dynamic representations can successfully perform sequence completion under ambiguous conditions (Minai et al. 1994, Levy et al. 1995, Wu et al. 1996). If we take the view that problem solving is often a sequence prediction problem, then we find that these sequence learning networks can, without search, hypothesize novel paths to goals that are only partially specified. Such results demonstrate the possibility that network dynamics contribute to creative problem solving.

Introduction. The hippocampus is critical for animals as they learn and find their way around in the world. The mental basis of this ability is called cognitive mapping, (Tolman 1932), and it is fundamental to current theories of hippocampal function (e.g. O'Keefe & Nadel 1978). One of the obvious reasons to have a cognitive map is to plan a strategy to get from one's present location to another place (a goal) where there is something desirable. For instance, when a rat is thirsty, it will need to go where water is and when it is hungry, it needs to go where food is. Thus, the desire an animal has in its mind (e.g. food or water) should affect the planned sequence that will carry it through its cognitive map of the world. Indeed, the spatial location of the goal that the animal might hope to head toward must be affected by the desires of the organism itself.

Here we present a model of the hippocampus (Levy 1989; Minai & Levy 1993b) that can hypothesize a novel path to a goal and that can produce an appropriate path to a goal without search and without actually going down that path.

The Network. The network, inspired by a small piece of hippocampal region CA3, has a sparse recurrent connectivity (10%) between its 512 excitatory neurons and a paucity (one) of inhibitory neurons that maintain control over net activity levels in a crude, imprecise way (Minai & Levy 1993a). The simultaneously updated McCulloch-Pitts neurons transmit $z_i(t) \in \{0,1\}$ every increment of time, t . The inhibition is of the shunting form although there is an equivalent form in terms of a subtractive inhibition (Minai & Levy 1993a). The excitation y_j of CA3 neuron j is given by:

$$y_j(t) = \frac{\sum_i w_{ij} c_{ij} z_i(t-1)}{\sum_i w_{ij} c_{ij} z_i(t-1) + K_I \sum_i x_i(t) + K_R \sum_i z_i(t-1)}$$

and its output of neuron j is

$$z_j(t) = \begin{cases} 1 & \text{if } y_j(t) \geq \theta \text{ or if } x_j(t) = 1; \\ 0 & \text{otherwise} \end{cases}$$

where each neuron j receives a single external input $x_j(t) \in \{0,1\}$. Average activity levels — controlled by the inhibition constants K_R , K_I , and threshold θ — are set to just under 15%.

Synaptic modification proceeds adaptively via a self-supervised, local modified Hebbian process of the form $w_{ij}(t+1) = w_{ij}(t) + 0.01 \cdot z_j(t) (z_i(t-1) - w_{ij}(t))$. The time delay, similar to a proposal of Amari (1972), is qualitatively consistent with physiological observations (Levy & Steward 1983). The use of this postsynaptically controlled equation, however, is not critical to the qualitative performance of this model (see Minai & Levy

1993b, Minai et al. 1994).

Learning occurs by repetitively presenting the input 300 times. Additional learning trials have no effect on performance. After each presentation of a sequence, all neurons are randomly set to zero or one. When there are two different sequences, a random process determines which sequence to present as an input. During testing, the feedback inhibition, K_R , is raised.

So far, there is nothing too far out of the ordinary here. But let us now define the inputs and outputs as purely dynamic. Thus, instead of defining stable points or transient attractors (Hopfield 1984, Tank & Hopfield 1987) as code words which the network should learn, let us view each successive state of the network as its representation of the world (Amari 1972) or, in the absence of input, its prediction on the world. In addition, specifying input activities (8 out of 512) much lower than network activities (ca. 40 out of the other 504 neurons) leads to a largely random recoder because the inputs directly command so little synaptic modification. The possibility for a broad range of self-determined codings is purposely enhanced by the sparse random connectivity. It is this sparsity and the random placement of synapses that lead to an essentially asymmetric connectivity. As a result of this asymmetry, the initial network dynamic is, for all practical purposes, unpredictable (Minai & Levy 1993a), which implies that recoding by the self-modifying network will appear to be a random process.

Basic testing consists of sequence completion. The network is randomly excited, and the first pattern of a sequence is turned on for one time step. The network states are then allowed to evolve on their own. From our point of view, we are asking the network to perform a sequence prediction problem.

The cosine function for comparing two vectors is used to understand the output of the network during testing. Specifically, after learning, the network states that evolve in response to the test input are compared with the states produced by a complete sequence of input patterns. By definition, the largest value at each time step is the decoded network state at that moment. Although other techniques (e.g., maximum likelihood) would produce superior decoding, such methods would obscure our object of our study, i.e., what the CA3-like recoding network is actually doing.

Results. Here we look at an analog of a useful skill in problem solving — the ability to combine separately learned inferences. In this problem the network learns two deterministic sequences of 12 patterns (see Fig. 1). Each sequence contains three orthogonal segments. Within each segment, the successive external inputs share 7 of 8 active neurons. The two sequences share a common subsequence $[\alpha, \beta, \gamma]$ that is three patterns long. We initially studied this problem of shared subsequences (Minai et al 1994 and again in Levy et al. 1995, Wu et al. 1996) to prove that the network, as a model of the hippocampus, can learn and can use context. Thus, we start by studying the disambiguation problem. As before, the network successfully disambiguates two related sequences based only on context. After learning, when the network is given pattern O, it produces a sequence culminating in pattern Z. Similarly (see Fig. 2a), giving the network pattern A produces a sequence culminating in pattern L. Once again, we emphasize that there is no explicit spectrum of delays (e.g. Kleinfeld 1986) nor is there any microscopic time spanning properties greater than one step so that the solution derives from the code created by the network itself (Levy 1989).

In addition to solving sequence prediction under ambiguous conditions, we now want to show goal finding by the network. This is a compositional problem involving a recombination of subsequences. Therefore, in addition to giving the network an initial starting point as an input, we also give it a partial description of a goal that requires a novel path. Thus, the input is initially pattern A as before, but now two neurons of pattern Z are continuously activated — as if the model has some partial idea about what the goal looks like. In this case the network hypothesizes the path $[A, B, C, \alpha, \beta, \gamma, U, V, W, X, Y, Z]$ which reaches the appropriate goal pattern. Fig. 2b helps show how the network succeeds. Thus, the model decomposed the two learned sequences and recomposed the pieces as appropriate to solve the problem.

To avoid unrepresentative results, here we have present only robust results where robust is defined as replicable in 4 out of 5 randomly constructed networks. There are four test cases, i.e. given pattern A, given pattern O, given pattern A and two neurons of pattern Z, and given pattern O and two neurons of pattern L. Out of the 5 networks we tried, the problems were correctly solved between 80-90% of the time with an average of 85%.




Fig. 1 Pictorial representation of the environment. There are two sequences of 12 input patterns to learn. Each sequence contains three orthogonal segments. The two sequences share a subsequence of three patterns.

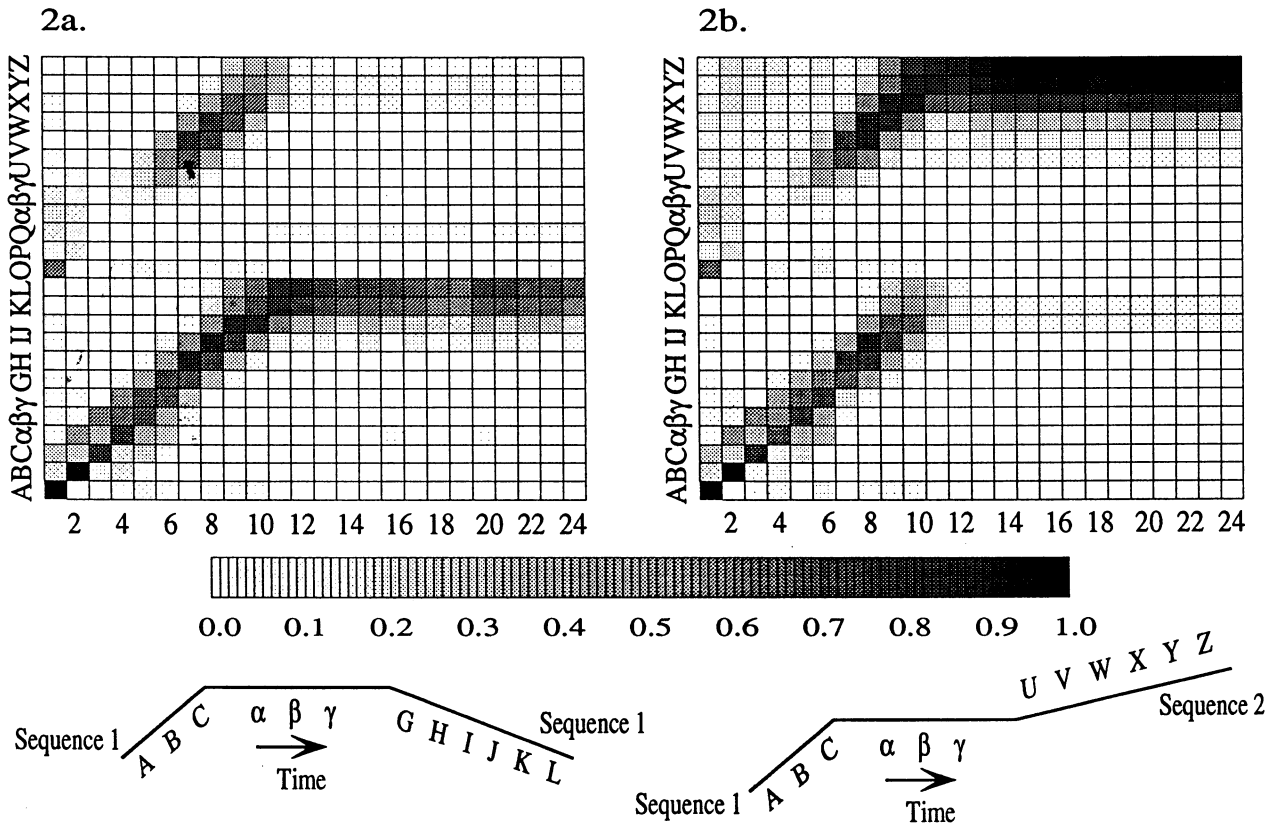


Fig. 2 Sequence prediction after learning the two partially overlapping sequences illustrated in Fig. 1. These similarity matrices use the cosine function. They compare the network states generated by each full pattern sequence after learning (ordinate) with the network states generated over time during testing (abscissa). Decode the output during testing by finding the darkest square in each column; the letter on the ordinate is the decoded answer. 2a. Sequence completion in the disambiguation problem is made difficult by the shared subsequence $[\alpha\beta\gamma]$. Here we show the similarity values (used for decoding) when the network is transiently given pattern A. Appropriately enough for the learning and for the starting point, the states go to pattern L, a pattern essentially orthogonal to the representation of the other learned goal pattern, Pattern Z. 2b. When the same network is given the same transient input but two neurons of goal Z are also turned on, the network produces a sequence of representations leading to this partially specified goal. To create this path, the network must produce a novel sequence that appropriately combines its knowledge of the two separately learned sequences.

Discussion. Because the network solves the context-dependent disambiguation problem, representations of α through γ must differ depending on where the network starts its sequence of representations. Even so, there is enough similarity between each set of the α - γ states for the network to follow either path depending on the relative strength of the two attractors at the end of each sequence. In addition to the network's immediate representation of γ , the winning attractor is a function of experience and of external biasing.

Hopfield's research (Hopfield 1982, 1984) has inspired a great amount of analytical work on the properties of recurrent networks. As a result, we now know that even networks with a time spanning Hebb rule can form, at the very least, transient attractors (e.g. Sompolinsky & Kanter 1986, Wu & Liljenström 1994, Liljenström & Wu 1995). However, using the dynamics of such networks to solve sequence prediction problems with novel goals seems a new use of such networks.

In sum, the hippocampal model studied here is capable of goal finding, albeit a highly simplified form of this problem. As such, these results further validate this model as a simplified, functional paradigm of the hippocampus.

Acknowledgments. This work was supported by NIH MH48161, MH00622, and EPRI RP8030-08 and Pittsburgh Supercomputing Center Grant #1 P41 RR06009 from the NIH National Center for Research Resources to WBL, and by the Department of Neurosurgery, Dr. John A. Jane, Chairman. The authors thank Dr. Robert A. Baxter for his feedback during the course of the research and Mr. A. Amarasingham for his help on the programming.

References.

- Abbott, L. F. & Blum, K. I. (1995) *Cerebral Cortex*, in press.
- Amari, S-I (1972) *IEEE Trans. on Computers* C-21:1197-1206.
- Andersson, J. A., Silverstein, J. W., Ritz, S. A., & Jones, R. S. (1977) *Psych. Rev.* 84, 413-451.
- Grossberg, S. (1976) *Biol. Cybern.* 23, 187-202.
- Hopfield, J. J. (1982) *Proc. Natl. Acad. Sci. USA* 79, 2554-2558.
- Hopfield, J. J. (1984) *Proc. Natl. Acad. Sci. USA* 81, 3088-3092.
- Levy, W. B (1989) in *Computational Models of Learning in Simple Neural Systems*, eds. Hawkins, R. D. & Bower, G. H. (Academic Press, New York), pp. 243-305.
- Levy, W. B (1994) *INNS World Congress on Neural Networks IV*-661-666.
- Levy, W. B & Steward, O. (1983) *Neurosci.* 8, 791-797.
- Levy, W. B & Wu, X. B. (1995) *INNS World Congress on Neural Networks I*-582-586.
- Levy, W. B & Wu, X. B. (1996) *Network: Compt. Neural Syst.*, 7, 371-384.
- Levy, W. B, Wu, X. B., & Baxter, R. A. (1995) *Intl. J. Neural Syst. Supplement*, 71-80.
- Liljenström, H. & Wu, X. B. (1995) *Intl. J. Neural Syst.* 6, 19-29.
- Minai, A. A., Barrows, G., & Levy, W. B (1994) *World Congress on Neural Networks IV*-176-181
- Minai, A. A. & Levy, W. B (1993a) In: *Advances in Neural Information Processing Systems*, eds. Giles, C. L., Hanson, S. J., & Cowen, J. D. (Morgan Kaufmann, San Mateo, CA), pp. 556-563
- Minai, A. A. & Levy W. B (1993b) *INNS World Congress on Neural Networks II*-505-508
- Newell, A. & Simon H. (1972) *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.
- O'Keefe, J. & Nadel, L. (1978) *The Hippocampus as a Cognitive Map*. (Clarendon Press:Oxford).
- Perkins, D. N. (1988) in *The Nature of Creativity*, ed. Sternberg, R. J. (Cambridge University Press, New York), pp. 362-385.
- Sompolinsky, H. & Kanter, I. (1986) *Phys. Rev. Ltrs.* 57, 2861-2864.
- Tank, D. W. & Hopfield, J. J. (1987) *Proc. Natl. Acad. Sci. USA* 84, 1896-1900.
- Tolman, E. C. (1948) *Psychol. Rev.* 55, 189-208.
- Kleinfeld, D. (1986) *Proc. Natl. Acad. Sci. USA* 83, 9469-9473.
- Wu, X. B., Baxter, R. A., & Levy, W. B (1996) *Biol. Cybern.* 74, 159-165.
- Wu, X. B. & Levy, W. B (1995) *INNS World Congress on Neural Networks I*-577-581.
- Wu, X. B. & Liljenström, H. (1994) *Network: Comput. Neural Syst.* 5, 47-60.