INTERNATIONAL COMPUTER SCIENCE INSTITUTE

1947 Center St. • Suite 600 • Berkeley, California 94704-1198 • (510) 643-9153 • FAX (510) 643-7684

A Biological Grounding of Recruitment Learning and Vicinal Algorithms

Lokendra Shastri*

TR-99-009

April 1999

Abstract

Biological neural networks are capable of gradual learning based on observing a large number of exemplars over time as well as rapidly memorizing specific events as a result of a single exposure. The primary focus of research in connectionist modeling has been on gradual learning, but some researchers have also attempted the computational modeling of rapid (one-shot) learning within a framework described variably as *recruitment learning* and *vicinal algorithms*. While general arguments for the neural plausibility of recruitment learning and vicinal algorithms based on notions of neural plasticity have been presented in the past, a specific neural correlate of such learning has not been proposed. Here it is shown that recruitment learning and vicinal algorithms can be firmly grounded in the biological phenomena of long-term potentiation (LTP) and long-term depression (LTD). Toward this end, a computational abstraction of LTP and LTD is presented, and an "algorithm" for the recruitment of binding-detector cells is described and evaluated using biologically realistic data. It is shown that binding-detector cells of distinct bindings exhibit low levels of cross-talk even when the bindings overlap. In the proposed grounding, the specification of a vicinal algorithm amounts to specifying an appropriate network architecture and suitable parameter values for the induction of LTP and LTD.

*This work was partially funded by ONR grant N00014-93-1-1149 and NSF grant SBR-9720398. Thanks to Jeff Bilmes, Nancy Chang, Shawn Chang, Jerry Feldman, Steven Greenberg, Srini Narayanan, and Carter Wendelken. *E-mail:* shastri@icsi.berkeley.edu (http://www.icsi.berkeley.edu/~shastri)

1 Introduction

Biological neural networks are capable of slow gradual learning as well as rapid one-shot memorization. The former involves an exposure to a large number of exemplars and leads to the acquisition of perceptual-motor skills, category formation, language skills, and certain types of semantic knowledge. In contrast, one-shot memorization can result from a single exposure to an example and underlies, among other things, the acquisition of "episodic memories" of everyday events, and memories of faces.

The primary focus of research in connectionist and neural network models has been on slow gradual learning, but some researchers have also attempted the computational modeling of rapid one-shot learning within a framework described variably as *recruitment learning* (Feldman, 1982; Shastri, 1988; Diederich, 1989) and *vicinal algorithms* (Valiant, 1994). In simple terms recruitment learning can be described as follows: Learning occurs within a network of randomly connected nodes. Recruited nodes are those nodes in the network that have acquired a distinct "meaning" (or functionality) by virtue of their *strong* interconnections to other recruited nodes and/or other sensorimotor (i.e., input/output) nodes. Nodes that are not yet recruited can be viewed as "free" nodes. Such nodes are connected via weak links to a large number of free, recruited, and/or sensorimotor nodes. These free nodes form a primordial network from which suitably connected nodes may be recruited for representing new items. For example, a novel concept y which can be expressed as a conjunct of existing concepts x_1 and x_2 can be memorized by (i) identifying free nodes that receive links from nodes representing x_1 as well as nodes representing x_2 and (ii) "recruiting" one or more such free nodes by strengthening the weights of links incident on such nodes from x_1 and x_2 nodes.

Feldman (1982) showed that conjunctive concepts can be recruited with a high probability if one makes suitable assumptions about network connectivity. He presented a probabilistic analysis of recruitment learning based on the degree of connectivity and the number of intermediate layers in random interconnection networks. Shastri (1988) extended the notion of recruitment learning to relational concepts. He treated a concept as a collection of attribute-value bindings and suggested a two-stage memorization process. In the first stage, "binder" nodes are recruited for each attribute-value binding in a concept. In the second stage, these binder nodes are joined together by the recruitment of another conjunctive node. Diederich (1989) showed how this form of structured recruitment learning can be used to learn new concepts expressed as modifications of existing concepts. Valiant (1994) proposed a formal "neuroidal model" and described several algorithms for the recruitment learning of conjunctive and relational concepts. He also presented a quantitative analysis of these algorithms using plausible assumptions about connectivity in the neocortex. Valiant referred to these algorithms as "vicinal algorithms."

While general arguments in support of the neural plausibility of recruitment learning and vicinal algorithms have been presented in the past (see Feldman, 1982;

Valiant, 1994), a specific neural correlate of such learning has not been proposed. In this paper it is shown that recruitment learning can be firmly grounded in the biological phenomena of *long-term potentiation* (LTP) and *long-term depression* (LTD) that involve rapid, long-lasting, and highly specific changes in synaptic strength. Toward this end, a computational abstraction of LTP and LTD is proposed, and an "algorithm" for the recruitment of binding-detector cells is described and evaluated using biologically realistic data about region sizes and cell connectivity. In the proposed grounding, the specification of a vicinal algorithm amounts to choosing a suitable network architecture and a set of appropriate parameter values for the induction of LTP and LTD.

The rest of the paper is organized as follows: Section 2 briefly reviews the phenomena of LTP and LTD. Section 3 describes a computational abstraction of cells, synapses, LTP, and LTD. Section 4 describes how a transient pattern of activity can lead to the recruitment of binding-detector cells as a result of LTP (and optionally, LTD) within quasi-random network structures. Finally, section 6 presents some concluding remarks.

2 Long-term potentiation and depression

Long-term potentiation (LTP) refers to a long-term increase in synaptic strength[1] resulting from the pairing of presynaptic activity with postsynaptic depolarization, and has emerged as a promising cellular mechanism underlying activity-dependent learning in the brain (Bliss & Lomo, 1973; Bliss & Collingridge, 1993; Lynch & Ambros-Ingerson, 1993; Derrick & Martinez, 1996). LTP involves the unusual receptor NMDA which is activated by the neurotransmitter glutamate, but only if the postsynaptic membrane is already depolarized. Once the NMDA receptor is activated, calcium ions flood into the postsynaptic cell and lead to a complex series of biochemical changes that result in the induction of LTP.[2]

The conditions required for the activation of the NMDA receptor — presynaptic activity in the presence of postsynaptic depolarization — can be brought about by a

[1] A synapse is the site of communication between two cells. Typically, a synapse is formed when an axonal (output) fiber emanating from a "presynaptic" cell makes contact with the dendritic tree (input structure) of a "postsynaptic" cell.

A synapse can be excitatory or inhibitory. The arrival of activity at an excitatory synapse from its presynaptic cell leads to a depolarization of the local membrane potential of its postsynaptic cell and makes the postsynaptic cell more prone to firing. In contrast, the arrival of activity at an inhibitory synapse leads to a hyperpolarization of the local membrane potential of the postsynaptic cell and makes the postsynaptic cell less prone to firing. The strength of an excitatory (or inhibitory) synapse determines the degree of depolarization (or hyperpolarization) that will result from a given presynaptic activity. The greater the synaptic strength, the greater the depolarization (hyperpolarization). For more on cell and synapses refer to (Kandel, Schwartz, and Jessell, 1991).

[2] While the site of LTP *induction* is postsynaptic, the site of its *expression* continues to be a matter of controversy. Also, not all cases of LTP are mediated by NMDA receptors (see Zalutsky & Nicoll, 1990).

high frequency burst of activity arriving at a synapse or by multiple converging inputs arriving at a cell in close temporal proximity. The long-term increase in the efficacy of a synapse resulting from brief but high-frequency activity at the synapse is referred to as *homosynaptic* LTP. The long-term increase in the efficacy of synapses resulting from convergent activity arriving at multiple synapses sharing the same postsynaptic cell is called *associative* LTP.

LTP possesses several properties that make it suitable for rapid memory formation. It is induced very rapidly — within a few seconds, and is fully present within 20-30 seconds. Once stable, it can persist for a long time. Finally, it is synapse-specific, and hence, can express specific associations and correlations.

In addition to potentiation, synapses can also undergo long-term depression (LTD) (Linden, 1994; Artola & Singer, 1993). A synapse receiving no presynaptic activity can undergo heterosynaptic LTD if other synapses of the same postsynaptic cell receive high frequency presynaptic activity. Finally, prolonged low frequency stimulation of a synapse can lead to its homosynaptic LTD.[3]

3 Computational modeling of LTP and LTD

3.1 Computational abstraction of cell behavior

The computational abstraction of cells and synapses proposed below was guided by two considerations. The first consideration was to use an abstraction rich enough to capture temporal aspects critical for modeling LTP and LTD and to explain how LTP and LTD can lead to the recruitment of a variety of functional cells and circuits. The second consideration was to keep the model *discrete* and minimal so as to facilitate its analysis and computer simulation.

A cell is modeled as a highly idealized integrate-and-fire neuron (e.g., see Maass & Ruf, 1999). The spatio-temporal integration of activity arriving at the synapses of a postsynaptic cell is modeled as follows:

The *postsynaptic potential* at time t, resulting from presynaptic activity, $a_i(t_0)$, occurring at synapse s_i at time t_0, is given by:

$$psp_i(t\,|a_i(t_0)) = \begin{cases} a_i(t_0) * w_i(t_0) & t_0 \leq t < (t_0 + \omega_{int}) \\ \\ 0 & \text{otherwise} \end{cases}$$

where $w_i(t_0)$ is the *weight* (or the strength) of synapse s_i at time t_0, and ω_{int} is the *window of temporal integration* which denotes the maximum amount by which two incident activities may lead/lag and still be summated by the postsynaptic cell. Consequently, $psp_i(t)$, the postsynaptic potential at s_i at time t, is given by:

[3]There is also some evidence that a synapse may undergo associative LTD upon receiving presynaptic activity that is out of phase with strong rhythmic activity converging on other synapses of the postsynaptic cell (Stanton & Sejnowski, 1989).

$$psp_i(t) = \sum_{(0 \leq \tau < \omega_{int})} a_i(t-\tau) * w_i(t-\tau)$$

Note that the postsynaptic potential resulting from presynaptic activity is being modeled as a square-pulse of duration ω_{int}. Since actual *psp*s in biological networks have gradual onsets and gradual decays, the square-pulse approximation is a strong idealization. However, this idealization does not render the proposed abstraction biologically implausible for the following reason: It is known that models employing square-pulse *psp* functions are computationally *weaker* than models employing piecewise linear *psp* functions (Maass & Ruf, 1999). Since *psp* functions with graded onsets and decays can be approximated by piecewise linear *psp* functions, it follows that a model employing square-pulse *psp* functions is computationally *weaker* than a model employing *psp* functions with graded onsets and decays. Consequently, any solution developed using synapses with square-pulse *psp* functions can be realized using synapses with biologically plausible graded *psp* functions. Thus while the use of simple square-pulse *psp*s in the abstraction facilitates the analysis and simulation of a network model using the proposed abstraction, it does not detract from its biological plausibility.[4]

A cell's *potential* at time t resulting from the combined effect of presynaptic activity at its synapses is given by $pot(t)$, where

$$pot(t) = \sum_i psp_i(t)$$

the sum being taken over all synapses of the cell. We will use "potential" and "weighted sum of activity" interchangeably.

A cell may have two *response modes* (i) a normal spike response and (ii) a burst response consisting of a short-term generation of spikes at a high frequency.[5] A cell's response is governed by two *response thresholds* — one pertaining to the spike response and the other to the burst response — and one *refractory period*. The two response thresholds are referred to as $\theta_{sr}(t)$ (for spike response) and $\theta_{br}(t)$ (for burst-response). The refractory period is referred to as ω_{ref}. The response of a cell, $O(t)$, is characterized as follows:

[4]Of course, ease of simulation and analysis *per se* are not a virtue. For it to be interesting, the proposed abstraction must be capable of supporting a range of useful recruitment and vicinal learning algorithms. We will see one such application in this paper. Several others are described in (Shastri, 1999a; 1999c).

[5]The production of a burst response by a cell is not an uncommon phenomenon. In particular, pyramidal cells in the neocortex as well as the hippocampal formation may produce a burst response.

$$O(t) = \begin{cases} O_{br} & pot(t) \geq \theta_{br}(t) \\ O_{sr} & \theta_{sr}(t) \leq pot(t) < \theta_{br}(t)) \\ 0 & \text{otherwise} \end{cases}$$

In other words, a potential at, or above, the bursting-threshold leads to bursting, and a potential below the bursting-threshold, but at, or above, the spiking threshold leads to normal spiking. A potential below the spiking threshold leads to no response. After a cell responds at time t, the resulting spike (or burst) travels unattenuated and arrives at synapses downstream from the cell at time $t + d$, where d is the propagation delay (this activity constitutes presynaptic activity with respect to downstream synapses).

After a cell produces a response, it enters a *refractory* state for a duration ω_{ref}. During this interval, the cell does not produce a response irrespective of any inputs it might receive. This is modeled by assuming that once a cell responds at time t_0, its response thresholds become $+\infty$ during the interval $t_0 < t \leq t_0 + \omega_{ref}$. Subsequently, the response thresholds revert to their resting (or steady-state) values of Θ_{sr} and Θ_{br}, respectively.[6] That is,

$$\theta_{sr}(t) = \begin{cases} +\infty & \text{if the cell responded during the interval } [t - \omega_{ref}, t - 1] \\ \Theta_{sr} & \text{otherwise} \end{cases}$$

and

$$\theta_{br}(t) = \begin{cases} +\infty & \text{if the cell responded during the interval } [t - \omega_{ref}, t - 1] \\ \Theta_{br} & \text{otherwise} \end{cases}$$

Different types of cells may have different values of ω_{int}, ω_{ref}, Θ_{sr}, and Θ_{br}. A cell that can *only* produce a spike response can be modeled by setting Θ_{br} to ∞.

3.2 Modeling of Synaptic strengths

A *projection* refers to the set of links emanating from cells in a source region and impinging on cells in a target region. It is assumed that all the synapses formed by a projection are of the same *type*.

A synapse can be in any one of following three states: *naive*, *potentiated*, or *depressed*. The state of a synapse signifies its strength (weight). For a given synaptic type, the weights of all synapses in a given state are distributed within a restricted band, with the weight bands associated with each state being disjoint. The weight bands associated with synaptic states may differ from one synaptic type to another.

[6] The use of a discrete threshold function is also an idealization. As in the case of postsynaptic potential, while this idealization greatly simplifies simulation and analysis, it does not detract from the model's biological plausibility.

3.3 Computational modeling of LTP

LTP is governed by the following parameters:

- *potentiation threshold* θ_p,
- *weight increment* Δw_{ltp},
- *repetition factor* κ,
- *maximum inter-activity interval* τ_{iai}, and
- *probability of LTP*, ζ_{ltp}.

Consider a set of neighboring[7] synapses $s_1, \ldots, s_n$ on the same postsynaptic cell. Convergent presynaptic activity at $s_1, \ldots, s_n$ can lead to associative LTP of naive s_i's and increase their weights by Δw_{ltp} if the following conditions hold:

1. $\sum_{1 \leq i \leq n} psp_i(t) \geq \theta_p$

   The above entails that synchronous presynaptic activity arriving at $s_1, \ldots, s_n$ causes sufficient depolarization of the postsynaptic cell's membrane potential.

2. Such synchronous presynaptic activity recurs (repeats) $\geq \kappa$ times.

3. The interval between two *successive* arrivals of presynaptic activity at a synapse during the above repetition is $\leq \tau_{iai}$ time units.

The parameter ζ_{ltp} specifies the probability that a naive synapse will undergo LTP if conditions 1—3 are met. Here ζ_{ltp} provides a simple *computational* mechanism for controlling the occurrence of LTP in a region. Among other things, the value of ζ_{ltp} may be varied to model neuromodulatory effects arising as a result of a system's emotional and motivational states.

Homosynaptic LTP is modeled in an analogous manner. Repeated bursts of high-frequency activity arriving at a synapse i, can lead to its homosynaptic LTP if (i) the postsynaptic potential, psp_i, resulting from each burst is $\geq \theta_p$, (ii) the bursts repeat $\geq \kappa$ times, and (iii) the interval between two successive bursts is $\leq \tau_{iai}$ time units.

3.4 Computational modeling of LTD

Heterosynaptic LTD is also modeled using five parameters. These are:

- potentiation threshold θ_p,

[7] In the simplest case, all synapses associated with the same postsynaptic cell may be assumed to be neighboring synapses.

- weight decrement Δw_{ltd},
- repetition factor κ,
- maximum inter-activity interval τ_{iai}, and
- probability of LTD, ζ_{ltd}.

When naive or potentiated synapses of a postsynaptic cell receive convergent presynaptic activity, neighboring inactive naive synapses of the postsynaptic cell undergo heterosynaptic LTD and their weights decrease by Δw_{ltd}.

As in the case of LTP, θ_p dictates the minimum weighted sum of synchronous activity that neighboring synapses of the postsynaptic cell must receive, and κ specifies the number of times such presynaptic activity must recur in order to induce heterosynaptic LTD of naive inactive synapses. Also as before, τ_{iai} specifies the maximum permissible gap between the successive arrival of presynaptic activity.

The parameter ζ_{ltd} specifies the probability that an inactive naive synapse will undergo LTD when the above conditions are met. A value of $\zeta_{ltd} = 0$ means that there is no heterosynaptic LTD and a value of $\zeta_{ltd} = 1$ means that the occurrence of LTP can lead to the heterosynaptic LTD of all the neighboring inactive naive synapses of the postsynaptic cell.

A variant of the abstraction proposed above uses a synaptic weight modification rule similar in spirit to the BCM rule (Bienenstock, Cooper, & Munro, 1982). In this variant, θ_p of a cell is variable and increases upon the potentiation of its synapses. Hence, the potentiation of a certain number of synapses can raise the potentiation threshold of the postsynaptic cell to a sufficiently high level and thereafter, make the potentiation of its remaining naive synapses very unlikely. In computational terms, assuming a steep increase in the potentiation threshold (θ_p) of a postsynaptic cell after the potentiation of a small number of its synapses is analogous to assuming $\zeta_{ltd} \approx 1$ (i.e., widespread heterosynaptic LTD), whereby the LTP of a small number of synapses can lead to the LTD of all remaining naive synapses.

3.5 Emergence of cells and circuits responsive to specific functionalities

LTP and LTD can transform random networks into structures consisting of cells tuned to specific functionalities. Typically, a cell receives a large number of inputs (afferents), and hence, can potentially participate in a large number of functional circuits. If, however, the weights of selected synapses on the cell increase via LTP (and, optionally, the weights of other synapses decrease via LTD) the cell can become more selective and participate in a limited number of functional circuits. Thus LTP and LTD provide a promising neural mechanism for the formation of functional structures within random networks via recruitment learning.

In the following section we illustrate how transient activity propagating through neural circuits can automatically lead to the recruitment of functional cells via LTP (and optionally, LTD).

4 Recruitment of binding-detector cells

Our ability to remember events in our daily life demonstrates our capacity to rapidly acquire new memories. Typically, such memories record who did what to whom where and when, or describe states of affairs wherein multiple entities occur in particular configurations. This form of memory is often referred to as episodic memory (Tulving, 1978), and there is a broad consensus that the hippocampal formation (HF) and neighboring areas in the medial temporal lobes serve a critical role in its formation (O'Keefe & Nadel, 1978; Squire, 1992; Cohen & Eichenbaum, 1993; Treves & Rolls, 1994).

The persistent encoding of an event must be capable of encoding role-entity *bindings*. Consider the event *described* by "John gave Mary a book in the library on Tuesday". This event cannot be encoded by simply forming a conjunctive association between "John", "Mary", "a book", "Library", "Tuesday" and "give" since such an encoding would be indistinguishable from that of the event described by "Mary gave John a book in the library on Tuesday". In order to make the necessary distinctions, the encoding of an event should specify the *bindings* between the *entities* participating in the event and the *roles* they play in the event. For example, the encoding of the event in question should specify the following *role*-`entity` bindings: ($\langle$*giver*=`John`$\rangle$, $\langle$*recipient*=`Mary`$\rangle$, $\langle$*give-object*=`a-Book`$\rangle$, $\langle$*temporal-location*=`Tuesday`$\rangle$, $\langle$*location*=`Library`$\rangle$).

As explained in (Shastri, 1999c), it is possible to evoke a fleshed out representation of an event by "retrieving" the bindings pertaining to the event and activating the web of semantic and procedural knowledge with these bindings. Thus cortical circuits encoding generic "knowledge" about actions such as *give* and entities such as *persons*, *books*, *libraries*, and *Tuesday* can recreate the necessary gestalt and details about the event "John gave Mary a book on Tuesday in the library" upon being activated with the above bindings. This view is supported by work on "reflexive reasoning" (Shastri & Ajjanagadde, 1993; Shastri, 1999b) and "executing schemas" (Bailey, Chang, Feldman & Narayanan, 1998).

In view of the above, the recruitment of binding-detectors is expected to be a critical step in the memorization of episodic memory. The following describes how such binding-detector cells can arise spontaneously and rapidly within a biologically motivated network structure as a result of LTP (and optionally, LTD).

4.1 A structure for the encoding of binding-detector cells

A structure for the rapid formation of cells responsive to binding matches consists of three regions: ROLE, ENTITY, and BIND (see Figure 1(a)). Regions ROLE and ENTITY are assumed to have 1 million primary (excitatory) cells each, while region BIND is assumed to have 15 million cells. Regions ROLE and ENTITY have dense projections to region BIND, with each cell in ROLE and ENTITY regions making 17,000 connections with cells in region BIND. In other words, the projective field (PF) size[8] of the ROLE to BIND projection as well as the ENTITY to BIND projection is 17,000. It is assumed that these projections are uniformly distributed over BIND.

Each role and entity is encoded by a small ensemble of cells in the ROLE and ENTITY regions, respectively. Cells in role and entity ensembles are also assumed to be distributed uniformly within regions ROLE and ENTITY, respectively. Note that role ensembles may overlap, and so may entity ensembles.

4.2 A neural correlate of the structure for encoding binding-detectors

There is a direct correspondence between the model structure described above and the interaction between the entorhinal cortex (EC) which is a region in the medial temporal lobe, and the dentate gyrus (DG) which is a component of the HF. The ROLE and ENTITY regions correspond to subregions of the EC and the BIND region corresponds to the DG. The projections from high-level cortical areas to ROLE and ENTITY regions correspond to the well known cortical projections to EC (Van Hoesen, 1982; Insausti, Amaral & Cowan, 1987; Suzuki & Amaral, 1994). The dense and diffuse projections from ROLE and ENTITY to BIND correspond to the dense and diffuse projections from EC to DG (Amaral, Ishizuka & Claiborne, 1990; Amaral & Witter, 1995). Moreover, the projective field and region sizes shown in Figure 1(a) are based on anatomical findings presented in (Amaral & Witter, 1995; West, 1990).

4.3 The transient representation of role-entity bindings

It is assumed that the bindings constituting an event are expressed as a transient pattern of rhythmic activity over distributed high-level cortical circuits (HLCCs). These HLCCs project to cells in ENTITY and ROLE regions and, in turn, induce transient patterns of rhythmic activity within these regions. Figure 2 is an idealized depiction of the transient activity induced in ENTITY and ROLE regions by HLCCs to convey the relational instance RI: $(\langle r_1 = f_1 \rangle, \langle r_2 = f_2 \rangle)$. Here r_1 and r_2 are roles, and f_1 and f_2 are entities bound to r_1 and r_2, respectively. Each spike in the illustration signifies the synchronous firing of a cell ensemble. It is shown that cells in the r_1 and f_1 ensembles are firing in synchrony, and so are cells in the r_2 and f_2 ensembles. The

[8]The set of cells in the target region that receive links from a cell c in the source region is referred to as the projective field (PF) of c. The PF *size* of c refers to the number of synapses formed by c with cells in the target region.

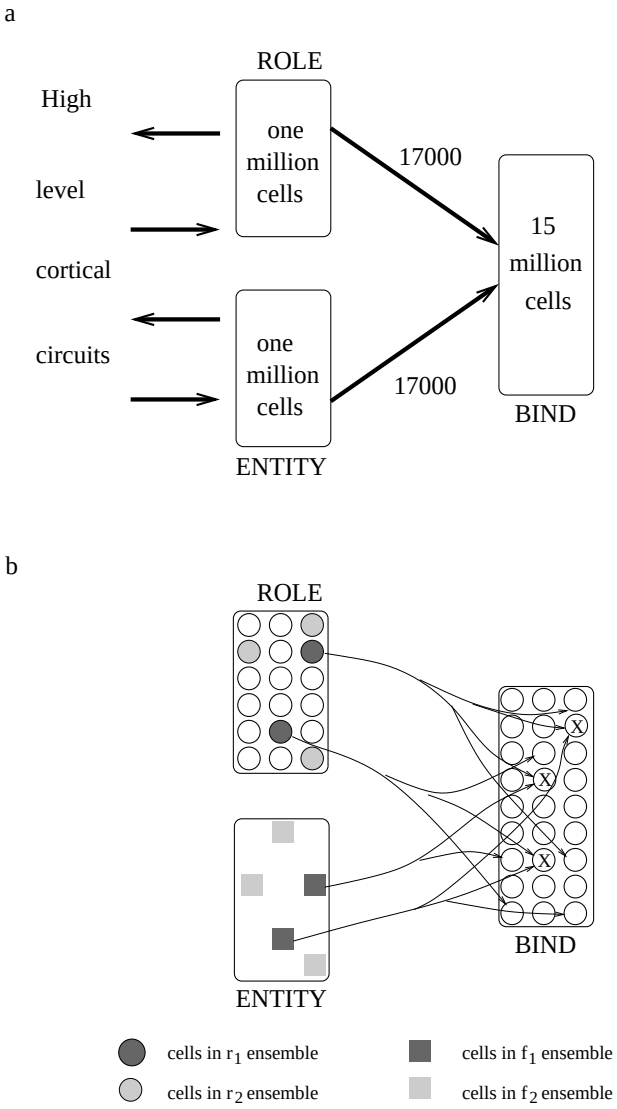

Figure 1: (a) A structure for the formation of binding-detector cells. Arcs indicate projections and the number on an arc indicates the projective field size. These projections are assumed to be uniformly distributed over the BIND region. Each role and entity is encoded by a small ensemble of cells in the ROLE and ENTITY regions, respectively. Cells in role and entity ensembles are also assumed to be distributed uniformly within regions ROLE and ENTITY, respectively. Binding-detector cells are recruited in region BIND. It is assumed that the ROLE and ENTITY regions lie in the entorhinal cortex and the region BIND corresponds to the dentate gyrus (a part of the hippocampus). The projective field and region sizes are based on (Amaral & Witter, 1995; West, 1990). (b) A schematic depiction of the ensembles of roles r_1 and r_2 and entities f_1 and f_2. Only links from cells in r_1 and f_1 ensembles to cells in BIND are shown. Cells marked with an "X" are candidates for recruitment as *binder* cells for the binding $\langle r_1 = f_1 \rangle$.

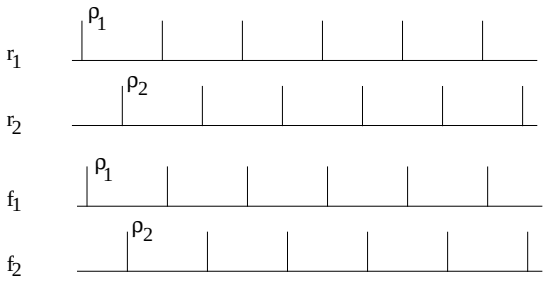

Figure 2: The transient encoding of a relational instance RI given by the bindings: $(\langle r_1 = f_1 \rangle, \langle r_2 = f_2 \rangle)$. Here r_1 and r_2 are roles, and f_1 and f_2 are entities bound to r_1 and r_2, respectively, in RI. Each spike in the illustration signifies the synchronous firing of a cell ensemble. Cells in the r_1 and f_1 ensembles fire in synchrony and so do cells in the r_2 and f_2 ensembles. The firing of cells in the r_1 and f_1 ensembles, however, is desynchronized with the firing of cells in the r_2 and f_2 ensembles. This desynchronization is assumed to be $\geq \omega_{int}$ time units. Moreover, the period of firing, π, is assumed to be $\leq \tau_{iai}$ time units. The dynamic encoding of RI can be viewed as a periodic pattern consisting of two *phases*: ρ_1 and ρ_2 (the *order* in which these phases appear has no significance).

firing of cells in the r_1 and f_1 ensembles, however, is desynchronized with the firing of cells in the r_2 and f_2 ensembles. This desynchronization is assumed to be $\geq \omega_{int}$ time units. Note that the dynamic encoding of RI can be viewed as a periodic pattern consisting of two *phases*: ρ_1 and ρ_2. Here ρ_1 and ρ_2 are mere labels and the *ordering* of phases has no significance.

In effect, a role-entity binding is expressed by the synchronous firing of the cell ensembles associated with the bound role and entity (von der Malsburg, 1986; Ajjanagadde & Shastri, 1991; Singer & Gray, 1995; Lisman & Idiart, 1995). In general, the transient encoding of a relational instance with n distinct entities participating as role-fillers involves n interleaved quasi-periodic activities having a period π. It is assumed that $\pi \leq \tau_{iai}$ time units. Such a spatio-temporal encoding enables multiple role-entity bindings to be expressed and propagated concurrently without cross-talk (Shastri & Ajjanagadde, 1993).

The following section explains how such a transient encoding of a relational instance may be transformed rapidly into persistent circuits for detecting bindings.

4.4 Recruitment of *binder* cells for memorizing role-entity bindings

BIND contains two kinds of cells: principal cells[9] and Type-1 inhibitory interneurons.[10] Each principal cell receives afferents from a number of cells in ROLE and ENTITY

[9]These correspond to granule cells in the dentate gyrus.

[10]The synapses formed by inhibitory interneurons on other cells have negative weights. The model makes the plausible assumption that such synapses cannot undergo LTP and LTD.

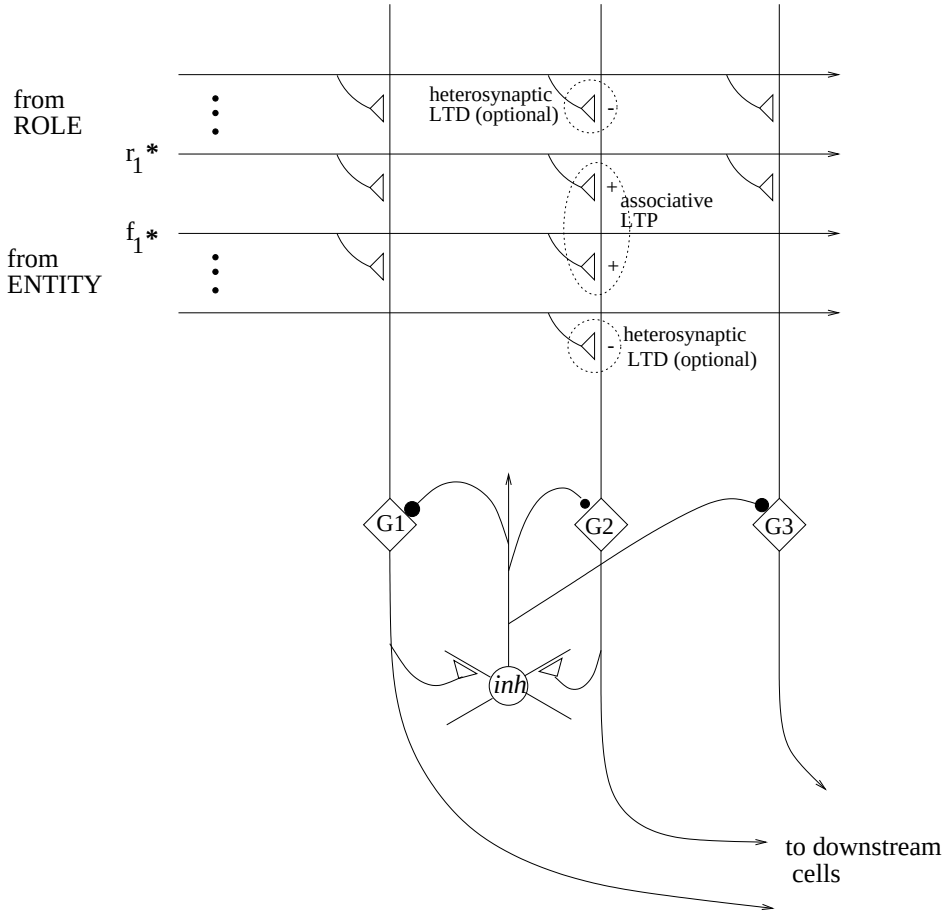

Figure 3: The BIND region consists of principal cells and inhibitory interneurons (Type-1). In the illustration, G1—G3 are principal cells and inh is a Type-1 interneuron. Afferents (incoming links) labeled r_1* and f_1* are from cells in the ensembles for role r_1 and entity f_1, respectively. Since G1 and G2 receive synchronous activity along afferents from r_1 and f_1 cells, they are *candidates* for becoming binding-detector cells for the binding $\langle r_1 = f_1 \rangle$. It is assumed that the inhibition from inh prevents the LTP of G1's synapses, and only G2 becomes a binding-detector cell for $\langle r_1 = f_1 \rangle$. Filled blobs denote inhibitory synapses, and the size of a filled blob is meant to convey the strength of the (inhibitory) synapse. See text for more details.

regions and makes synaptic contacts on a number of interneurons. The interneurons in turn make contacts on a number of principal cells, thereby forming inhibitory circuits within BIND (see Figure 3). The significance of inhibitory interneurons will be explained later.

The potentiation threshold, θ_p, of principal cells is sufficiently high, and hence, LTP of a synapse occurs only if multiple synapses of a postsynaptic cell receive coincident presynaptic activity.[11] Moreover, the response threshold, Θ_{sr}, of principal cells is such that a cell does not fire unless it receives impulses at multiple potentiated synapses. A set of values for θ_p, Θ_{sr}, synaptic weights, and other parameters of LTP and LTD are given below[12]:

$\Theta_{sr} = 1700; \quad \Theta_{br} = \infty; \quad \theta_p = 890;$

naive weight band = 100-110;

$\Delta_{ltp} = 100; \quad \Delta_{ltd} = 50;$

$\zeta_{ltp} = 1; \quad \zeta_{ltd} = 0;$

$\kappa = 5; \quad \omega_{int} = 2; \quad \omega_{ref} = 2.$

The choice of τ_{iai} is governed by ω_{int} and the number of role-entity bindings in an event. Thus any $\tau_{iai} \geq \omega_{int} * n$, where n is the number of bindings in the event, is appropriate.

The transient encoding of the relational instance RI shown in Figure 2 leads to the following events in BIND (refer to Figure 3). The synchronous firing of cells in the r_1 and f_1 ensembles (henceforth, r_1 and f_1 cells) leads to the associative LTP of active synapses of principal cells receiving sufficient afferents from r_1 and f_1 cells. At the same time, depending on the value of ζ_{ltd}, some of the inactive naive synapses of these principal cells may undergo heterogeneous LTD. The LTP of synapses formed by afferents arriving from r_1 and f_1 cells makes these principal cells behave as binding detector cells for the binding $\langle r_1 = f_1 \rangle$ and we will refer to such cells as $binder(\langle r_1 = f_1 \rangle)$ cells.[13]

The claim that $binder(\langle r_1 = f_1 \rangle)$ cells behave as binding-detector cells for $\langle r_1 = f_1 \rangle$ is substantiated quantitatively in Section 5, but it is not difficult to see why these

[11] Here and elsewhere in the paper, "coincidence" is defined with reference to ω_{int}, the window of temporal integration. See Section 3.

[12] LTD does not play a critical role in the recruitment of binding-detector cells described here. It might, however, play an important role in the recruitment of other functional circuits.

[13] To be precise, a cell is deemed to be recruited as a $binder(\langle r_1 = f_1 \rangle)$ cell if during the memorization of $\langle r_1 = f_1 \rangle$, the cell's synapses undergo LTP *and* the cell fires. The firing of the cell at the time of its recruitment is crucial if the cell is to become part of functional circuits lying downstream from the cell.

cells will behave in the desired manner. Note that a $binder(\langle r_1 = f_1 \rangle)$ cell will fire in response to the synchronous firing of r_1 and f_1 cells since the connectivity between r_1 and f_1 cells and a $binder(\langle r_1 = f_1 \rangle)$ cell required for the latter's recruitment during the memorization of $\langle r_1 = f_1 \rangle$ also suffices for the latter's firing during the retrieval of $\langle r_1 = f_1 \rangle$. At the same time, since Θ_{sr} is quite high (1700), a $binder(\langle r_1 = f_1 \rangle)$ cell is unlikely to fire as a result of stray impulses arriving at its synapses. Moreover, since a BIND cell requires concurrent activity at numerous (≥ 9) potentiated synapses in order to fire, it is unlikely that any $binder(\langle r_1 = f_1 \rangle)$ cell receives sufficient activity along potentiated links from r_1 cells alone or f_1 cells alone. Hence only the coincident arrival of impulses at potentiated synapses from r_1 *and* f_1 cells is likely to satisfy Θ_{sr} and cause a $binder(\langle r_1 = f_1 \rangle)$ cell to fire. Thus a $binder(\langle r_1 = f_1 \rangle)$ cell fires whenever r_1 cells in ROLE fire in synchrony with f_1 cells in ENTITY, and hence, behaves as a binding detector cell for the role-entity binding $\langle r_1 = f_1 \rangle$.

Similar LTP and LTD events occur at the synapses of principal cells that receive coincident activity along afferents from r_2 cells in ROLE and f_2 cells in ENTITY, and lead to their recruitment as $binder(\langle r_2 = f_2 \rangle)$ cells. A $binder(\langle r_2 = f_2 \rangle)$ cell fires whenever r_2 cells in ROLE and f_2 cells in ENTITY fire in synchrony and behaves as a binding detector cell for the role-entity binding $\langle r_2 = f_2 \rangle$. In general, numerous principal cells in BIND are recruited as *binder* binding detector cells for each binding.

4.5 Encoding and Response Times

The time required for the recruitment of *binder* cells is given by $\kappa * \tau_{iai}$. If we assume that the rhythmic activity encoding dynamic bindings corresponds to γ band activity (ca. 40 Hz) we get $\tau_{iai} \simeq 25$ msec. Assuming a plausible value of κ to be ~ 5 suggests that *binder* cells can be recruited in about 125 msec.[14] The time required for *binder* cells to respond to a retrieval cue is at most τ_{iai}. Thus both the recruitment and response times of the proposed model are consistent with the requirements of rapid (one-shot) memorization and recognition.

4.6 Potential problems in the formation of *binder* cells

The process by which *binder* cells are formed is susceptible to several problems. First, in order to form $binder(\langle r_i = f_j \rangle)$ cells there should exist cells that receive afferents from both r_i and f_j cells. Given the random connectivity between the regions ROLE and ENTITY and BIND this cannot be *guaranteed.*

[14] If instead, we assume that the rhythmic activity encoding dynamic bindings corresponds to θ band activity (ca. 8 Hz) then $\tau_{iai} \simeq 125$ msec. This suggest that *binder* cells can be recruited in about 525 msec. However, as discussed in (Shastri & Ajjanagadde, 1993; Shastri, 1999c) I believe that dynamic bindings of conceptual roles and entities are expressed as γ band activity, and the set of bindings pertaining to an event repeat over several (γ) cycles. It is this "block" of repetition that constitutes one period of θ band activity (cf. Lisman & Idiart, 1995; Luck & Vogel, 1997).

Second, there may exist cells that receive sufficient activity ($\geq \theta_p$) along afferents from r_i cells alone. Upon recruitment, such a cell could behave spuriously and produce false-positive responses since it will fire in response to the firing of r_i cells alone, even if there is no coincident activity of f_j cells. Similarly, cells receiving sufficient links from f_j cells alone could also produce false-positive responses.

Third, the same cell may get recruited as a *binder* cell for multiple bindings. This could also lead to spurious activity. Consider a cell that gets recruited as a *binder* cell for two bindings $\langle r_i = f_k \rangle$ and $\langle r_j = f_l \rangle$. This cell will fire in response to subsequent inputs containing these two bindings as well as the bindings: $\langle r_i = f_l \rangle$, and $\langle r_j = f_k \rangle$. Consequently, other cells connected downstream to this cell could receive false-positive binding-match signals in certain circumstances.

As stated below, the probability of not finding cells for recruitment as *binder* cells as well as the probabilities that *binder* cells will respond in a false-positive manner can be shown to be very small using biologically motivated values of various system parameters. The problem of too many cells becoming recruited for a binding turns out not to be very serious in the case under consideration, and hence, is not discussed here. In general, however, this problem can arise and can be alleviated partially by inhibitory feedback and feedforward local circuits formed by principal cells and Type-1 inhibitory interneurons. These inhibitory circuits act as soft-WTA and only allow synapses of a limited number of cells to undergo LTP (cf. Marr, 1971; McNaughton & Morris, 1987).

5 Quantitative Results

The following quantities have been calculated analytically[15] using the region and projective field sizes described in Section 4.1, the cell, synapse, and LTP parameters described in Section 4.4, and by assuming that each role and entity ensemble contains 600 cells.

1. P_{fail}, the probability that for a given binding *no* cells will be found in BIND (DG) for recruitment as binding detector cells is less than $< 10^{-18}$.

2. The expected number of cells in BIND (DG) that will receive appropriate connections and will be candidates for recruitment for a binding is 195.0173.

3. The expected number of *binder* cells of various bindings that will fire in response to a retrieval *cue* containing the three bindings $\langle r_1 = f_1 \rangle, \langle r_2 = f_2 \rangle$, and $\langle r_3 = f_3 \rangle$ (see Table 1). As shown in Table 1, a vast majority of *binder* cells of any given binding respond correctly to the cue. Even when the potential for cross-talk is maximal (e.g., in the case of $\langle r_1 = f_2 \rangle$ *binder* cells) less than 3.3% of the *binder* cells produce a false-positive response in any given time.

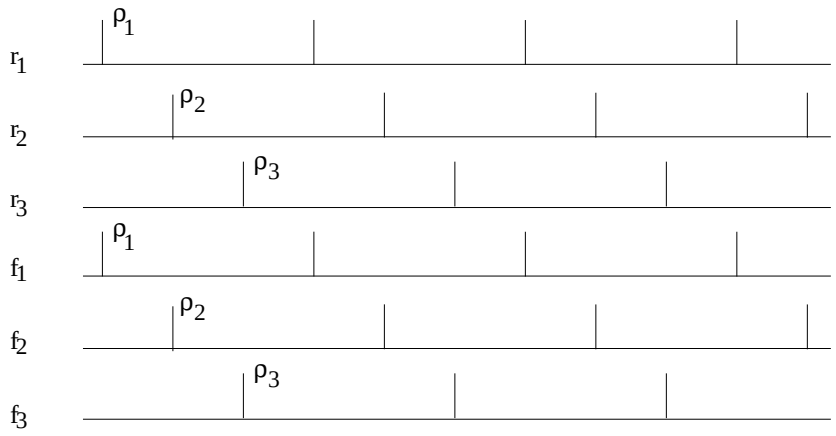

Figure 4: The transient encoding of a retrieval cue consisting of the bindings: $\langle r_1 = f_1 \rangle, \langle r_2 = f_2 \rangle$, and $\langle r_3 = f_3 \rangle$. Each spike in the illustration signifies the synchronous firing of a cell ensemble. Bindings are expressed as the synchronous firing of appropriate role and entity ensembles. Thus the cue is expressed as a rhythmic pattern of activity consisting of three interleaved periodic spike trains, one for each binding. This activity can be viewed as a periodic pattern consisting of three *phases*: ρ_1, ρ_2, and ρ_3.

binder cell (*binder* type)	Response					
	Phase ρ_1		Phase ρ_2		Phase ρ_3	
	OFF	ON	OFF	ON	OFF	ON
$\langle r_1 = f_1 \rangle$ ($\langle r_{eq} = f_{eq} \rangle$)	0.0000	**195.0173**	**195.0147**	0.0026	**195.0147**	0.0026
$\langle r_1 = f_2 \rangle$ ($\langle r_{inc} = f_{inc} \rangle$)	**186.6288**	6.3885	**188.6288**	6.3885	**195.0147**	0.0026
$\langle r_1 = f_{10} \rangle$ ($\langle r_{inc} = f_{exc} \rangle$)	**186.6288**	6.3885	**195.0147**	0.0026	**195.0147**	0.0026
$\langle r_{10} = f_1 \rangle$ ($\langle r_{exc} = f_{inc} \rangle$)	**186.6288**	6.3885	**195.0147**	0.0026	**195.0147**	0.0026
$\langle r_{10} = f_{10} \rangle$ ($\langle r_{exc} = f_{exc} \rangle$)	**195.0147**	0.0026	**195.0147**	0.0026	**195.0147**	0.0026

Table 1: The table describes how *binder* cells of all the memorized bindings respond to a retrieval cue containing the three bindings $\langle r_1 = f_1 \rangle$, $\langle r_2 = f_2 \rangle$, and $\langle r_3 = f_3 \rangle$. The expected number of *binder* cells responding **correctly** are marked in **boldface**. Note that a vast majority of *binder* cells for any given binding respond correctly to the retrieval cue. Each row in the table describes the response of *binder* cells of a specific binding. This response typifies the response of a whole class of *binder* cells. Thus the table completely characterizes the response of *binder* cells. This characterization holds irrespective of the number of bindings memorized. See text for details.

The retrieval cue is expressed as a rhythmic pattern of activity consisting of three interleaved periodic spike trains, one for each distinct role-entity binding (see Figure 4). This activity can be viewed as a periodic pattern consisting of three *phases*.

Each row in the table describes the response of *binder* cells of a specific binding. This response, however, typifies the response of a whole class of *binder* cells. Thus the data in Table 1 specifies the response of *binder* cells of *all* the bindings *memorized* in BIND prior to the posing of the retrieval cue.

The response of $\langle r_1 = f_1 \rangle$ *binder* cells typifies the response of *binder* cells of all bindings mentioned in the retrieval cue (this is indicated by the row label $\langle r_{eq} = f_{eq} \rangle$ which refers to all bindings of the form $\langle r_i = f_i \rangle$, $1 \leq i \leq 3$). For example, the response of $\langle r_2 = f_2 \rangle$ *binder* cells will be analogous to that of $\langle r_1 = f_1 \rangle$ *binder* cells, except that while the latter will respond maximally in phase ρ_1, the former will responds maximally in phase ρ_2.

Similarly, the response of $\langle r_1 = f_2 \rangle$ *binder* cells typifies the response of *binder* cells of any binding of the form $\langle r_i = f_j \rangle$ such that both r_i and f_j occur in the cue, but in distinct bindings (i.e., $1 \leq i, j \leq 3$, and $i \neq j$). This is indicated by the row label $\langle r_{inc} = f_{inc} \rangle$. For example, the response of $\langle r_3 = f_1 \rangle$ *binder* cells will be analogous to that of $\langle r_1 = f_2 \rangle$ *binder* cells, except that while the latter will produce more spurious responses in phases ρ_1 and ρ_2, the former will do so in phases ρ_1 and ρ_3.

Furthermore, the response of $\langle r_1 = f_{10} \rangle$ *binder* cells typifies the response of *binder* cells of any binding of the form $\langle r_i = f_j \rangle$ such that r_i occurs in the retrieval cue, but f_j does not (this is indicated by row label $\langle r_{inc} = f_{exc} \rangle$). For example, the response of $\langle r_2 = f_{15} \rangle$ *binder* cells will be analogous to that of $\langle r_1 = f_{10} \rangle$ *binder* cells, except that while the latter will produce more spurious responses in phase ρ_1, the former will do so in phase ρ_2.

Similarly, the response of $\langle r_{10} = f_1 \rangle$ *binder* cells typifies the response of *binder* cells of any binding of the form $\langle r_i = f_j \rangle$ such that f_j occurs in the retrieval cue, but r_i does not (row label $\langle r_{exc} = f_{inc} \rangle$).

Finally, the response of $\langle r_{10} = f_{10} \rangle$ *binder* cells typifies the response of *binder* cells of any binding involving roles and fillers that do not appear in the retrieval cue (row label $\langle r_{exc} = f_{exc} \rangle$).

Thus Table 1 completely characterizes the response of *binder* cells. Note that this characterization holds irrespective of the number of bindings memorized in BIND. In particular, it holds even if *all possible* bindings involving roles and entities encoded in ROLE and ENTITY have been memorized.

Note that an increase in the size of the retrieval cue (as measured by the number of bindings) would not degrade the quality of the response produced by *binder* cells. Consider a cue containing four bindings. The table describing the response of *binder* cells to this larger cue will simply have an additional column for phase ρ_4, whose entries will be identical to that of the column for phase ρ_3 in Table 1.

[15] Appendix 1 outlines the bases of these calculations.

Since each binding is redundantly encoded by multiple cells, and since these cells are physically dispersed in the BIND region, the probability that a diffuse cell loss will destroy many *binder* cells for any given binding remains extremely small. In particular, a diffuse loss of $x\%$ of the cells in region BIND will lead to an expected loss of only $x\%$ of the 195.0173 (expected) *binder* cells for a given binding. Thus the memorization of binding-detectors is robust with respect to diffuse cell loss.

The quantities in Table 1 are based on $\zeta_{ltd} = 0$ (i.e., no LTD). This condition results in a maximal sharing of *binder* cells among different bindings, and hence, these results provide a measure of the system's performance under conditions of maximal cross-talk. A non-zero value of ζ_{ltd} would reduce cross-talk, but it would also lead to a gradual reduction in the number of cells available for recruitment as more and more bindings are memorized.

A word on the size of role and entity ensembles: As pointed out in Section 4.2, the region and projective field sizes were chosen based on anatomical data. Plausible parameter values for synapses, cells and the induction of LTP in BIND were chosen to reflect the constraint that in the dentate gyrus, convergent activity at multiple synapses is required for the induction of LTP and also for the firing of principal (granule) cells. However, there did not exist sufficient empirical data to constrain the size of ROLE and ENTITY ensembles — *a priori*, the size could be anywhere from a few (~ 10) to several hundred thousand, or more. Various sizes were considered and a size of 600 was chosen because it resulted in satisfactory values of P_{fail}, the expected number of *binder* cells per binding, and the level of cross-talk.

6 Conclusion

A grounding of recruitment learning and vicinal algorithms in the biological phenomena of LTP and LTD has been described. A realization and specification of a vicinal algorithm using LTP has been illustrated by showing how *binder* cells responsive to specific role-entity bindings can be memorized rapidly in response to a transient pattern of activity encoding the bindings. Using biologically plausible values for the number of cells in the ROLE, ENTITY and BIND regions, and the density of projections from the ROLE and ENTITY regions to BIND, it has been shown that the existence of suitable *binder* cells for encoding arbitrary role-entity bindings is practically certain. It has also been shown that the interference between *binder* cells for different bindings remains extremely low.

The encoding of binding detectors is just one step in the memorization of episodic memory. As argued in (Shastri, 1997; 1999c) a proper encoding of episodic memory also requires the recruitment of *binding-error detector* circuits, *binding-error integrator* cells, *relational match* circuits, and *binding-extractor* cells. Shastri (1999a and 1999c) discuss how the requisite cells and circuits can also be recruited rapidly via LTP and LTD within quasi-random networks whose architecture and circuitry resembles that of the hippocampal formation. The formation of these cells and circuits

exercise the full-range of features included in the abstraction of LTP and LTD (e.g., spike versus burst firing modes) and further illustrate how interesting vicinal algorithms can arise from a suitable choice of network architecture and parameters for the induction of LTP and LTD.

Appendix 1

The following outlines the basic approach used in the computation of failure probability, the expected number of *binder* cells recruited for a binding, and the expected number of *binder* cells for a binding that will fire spuriously in response to another binding presented in a retrieval cue.

Let S_r refer to the ensemble of role r, let $|S_r|$ denote the number of cells in S_r, and let $|PF_{\text{ROLE}\rightarrow\text{BIND}}|$ denote the PF size for the projection from ROLE to BIND (recall that the PF size denotes the number of synapses made in the target region by afferents emanating from a single cell in the source region). Then n_r, the total number of synapses made by cells in S_r with cells in BIND, is given by:

$$n_r = |PF_{\text{ROLE}\rightarrow\text{BIND}}| * |S_r|$$

Similarly, let S_f refer to the ensemble of entity f, let $|S_f|$ denote the number of cells in S_f, and let $|PF_{\text{ENTITY}\rightarrow\text{BIND}}|$ denote the PF size for the projection from ENTITY to BIND. Then n_f, the total number of synapses made by cells in S_f with cells in BIND, is given by:

$$n_f = |PF_{\text{ENTITY}\rightarrow\text{BIND}}| * |S_f|$$

Since the PFs of ROLE and ENTITY cells are distributed *uniformly* and *independently* over BIND, p_r, the probability that a given synapse formed by afferents emanating from S_r impinges on a particular cell in BIND, can be approximated as $p_r = \frac{1}{|\text{BIND}|}$, where $|\text{BIND}|$ denotes the number of cells in BIND. Similarly, p_f, the probability that a given synapse formed by afferents from S_f impinges on a particular cell in BIND, can be approximated as $p_r = \frac{1}{|\text{BIND}|}$.

Furthermore, $p(S_r = k)$, the probability that a given cell in BIND receives exactly k afferents from S_r, is given by the binomial distribution:

$$p(S_r = k) = \left(\begin{array}{c} n_r \\ k \end{array} \right) p_r^k (1 - p_r)^{n_r - k}$$

and the probability, $p(S_r \geq k)$, the probability that a given cell in BIND receives *at least* k afferents from S_r, is given by:

$$p(S_r \geq k) = \left(1 - \sum_{i=0}^{i<k} \left(\begin{array}{c} n_r \\ i \end{array} \right) p_r^i (1 - p_r)^{n_r - i} \right)$$

The probabilities $p(S_f = k)$ and $p(S_f \geq k)$ have analogous interpretations and may be computed in an analogous manner.

Since the projections from ENTITY and ROLE to BIND are independent, $p(S_r \geq k_r, S_f \geq k_f)$, the probability that a given cell in BIND receives at least k_r afferents from S_r *and* at least k_f afferents from S_f, is given by: $p(S_r \geq k_r, S_f \geq k_f) = p(S_r \geq k_r) * p(S_f \geq k_f)$. In general, it is possible to compute: $p(S_r \ o_r \ k_r, S_f \ o_f \ k_f) = p(S_r \ o_1 \ k_r) * p(S_f \ o_2 \ k_f)$, where o_1 and o_2 could be "=", "$\geq$", "$\leq$", etc.

Appropriate combinations of k_f and k_r values required for the recruitment of a cell in BIND are determined by θ_p and synaptic weights along the projections from ENTITY and ROLE. If all combinations of k_r and k_f values satisfying θ_p are identified and expressed as a set of mutually exclusive conditions, then p_{cand}, the probability that a given cell in BIND is a candidate for recruitment as a *binder* cell for $\langle r = f \rangle$, can be computed as follows:

$$p_{cand} = \sum_{(S_r \ o_{r_i} \ k_{r_i}, \ S_f \ o_{f_i} \ k_{f_i})} p(S_r \ o_{r_i} \ k_{r_i}) * p(S_f \ o_{f_i} \ k_{f_i})$$

where the sum is taken over mutually exclusive conditions that jointly cover the space of possibilities under which a cell in BIND is a candidate for recruitment.

If p_{cand} is known, P_{fail}, the probability that none of the cells in BIND are candidates for recruitment as *binder* cells for $\langle r = f \rangle$, equals:

$$P_{fail} = (1 - p_{cand})^{|\text{BIND}|}$$

and $\text{E}\langle cand \rangle$, the expected number of candidates for recruitment as *binder* cells for $\langle r = f \rangle$, equals:

$$\text{E}\langle cand \rangle = p_{cand} * |\text{BIND}|$$

The recruitment of BIND cells as *binder* cells for multiple bindings can lead to spurious responses. In particular, the number of $\langle r_i = f_j \rangle$ *binder* cells that will fire spuriously upon the presentation of the binding $\langle r_k = f_l \rangle$ in a retrieval cue is given by the number of BIND cells that were recruited for both, the binding $\langle r_i = f_j \rangle$ and the binding $\langle r_k = f_l \rangle$. Note that the two recruitment events are *not* necessarily independent (for example, consider $i = k$, or $j = l$). An additional complication arises because role ensembles can overlap in ROLE and so can entity ensembles in ENTITY. Thus probabilities such as $p(S_{r_1} = k_1)$ and $p(S_{r_2} = k_2)$ cannot be treated as being independent. The same holds for probabilities such as $p(S_{f_1} = k_1)$ and $p(S_{f_2} = k_2)$. Consequently, the joint probability that a cell is recruited for a given pair of bindings cannot be computed by simply multiplying the individual recruitment probabilities. Instead, the joint probability has to be computed by enumerating the set of mutually exclusive conditions that together cover the space of possibilities in which a BIND cell can get recruited for both the bindings. Once this set is enumerated, the expected number of cells recruited by both bindings can be computed using the techniques described above.

References

Ajjanagadde, V. & Shastri, L. (1991). Rules and Variables in Neural Nets. *Neural Computation* 3, 121—134.

Amaral, D.G., Ishizuka, N. & Claiborne, B. (1990) Neurons, numbers and the hippocampal network. In *Progress in Brain Research: Understanding the brain through the hippocampus* (eds Storm-Mathisen, J., Zimmer, J. & Ottersen, O.P.) 1–11, Elsevier Science, Amsterdam.

Amaral, D.G. & Witter, M.P. (1995) Hippocampal Formation. In *The Rat Nervous System*, 2nd edn (ed Paxinos, G.) 443–493, Academic Press, London.

Artola, A. & Singer, W. (1993) Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation. *Trends in Neuroscience* 16: 480–487.

Bailey, D., Chang, N, Feldman, Narayanan, S. (1998) Extending Embodied Lexical Development. In *Proc. Twentieth Annual Conference of the Cognitive Science Society*, Madison, WI. August, 1998.

Bienenstock, E.L., Cooper, L.N., & Munro, P.W. (1982). Theory for the development of neural selectivity, orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience, 2*: 32—48.

Bliss, T.V.P. & Lomo, T. (1973) Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology 232,* pp 331–356.

Bliss, T.V.P. & Collingridge, G.L. (1993) A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* 361, 31–39.

Cohen, N.J. & Eichenbaum, H. (1993) *Memory, Amnesia, and the Hippocampal System.* M.I.T. Press, Cambridge, Massachusetts.

Derrick, B.E. & Martinez Jr., J.L. (1996) Associative, bidirectional modifications at the mossy-fibre-CA3 synapse. *Nature* 381, 429-434.

Diederich, J. (1989) Instruction and high-level learning in connectionist networks. *Connection Science* 1, 161–180.

Feldman, J. A. (1982) Dynamic connections in neural networks. *Bio-Cybernetics* 46, 27–39.

Insausti, R., Amaral, D.G. & Cowan, W.M. (1987) The entorhinal cortex of the monkey: II Cortical afferents. *Journal of Comparative Neurology* 264, 356–395.

Kandel, E., Schwartz, J. & Jessell, T. (1991) (eds.) *Principles of Neural Science*, Third Edition. Elsevier, New York.

Linden, D.J. (1994) Long-term synaptic depression in the mamallian brain. *Neuron* 12, 457–472.

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

Luck, S. J. & Vogel, E. K., (1997) The capacity of visual working memory for features and conjunctions. *Nature* 390:279-281.

Lynch, G. & Ambros-Ingerson, J. (1993) *Memory's Highway: Brain rules for encoding and organizing information.* Center for the neurobiology of learning and memory, University of California, Irvine, CA.

Marr, D. (1971) Simple memory: a theory for archicortex. *Philosophical Transactions of the Royal Society* B 262, 23–81.

Maass, W. & Ruf, B. (1999) On Computation with Pulses. *Information and Computation*, 148, 202-218.

McNaughton, B.L. & Morris, R.G.M. (1987) Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neuroscience* 10, 408–415.

O'Keefe, J. & Nadel, L. (1978) *The hippocampus as a cognitive map.* Clarendon Press, Oxford.

Shastri, L. (1988) *Semantic Networks: An evidential formalization and its connectionist realization.* (see p. 181–191). Morgan Kaufmann, Los Altos/Pitman Publishing Company, London, 1988.

Shastri, L. (1997) A Model of Rapid Memory Formation in the Hippocampal System, In the *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*, 680–685, Stanford University, August 1997.

Shastri, L. (1999a) Recruitment of binding-match and binding-error detector circuits via long-term potentiation. *Neurocomputing* (In press).

Shastri, L. (1999b) Advances in SHRUTI — A neurally motivated model of relational knowledge representation and rapid inference using temporal synchrony. *Applied Intelligence* (In press).

Shastri, L. (1999c) From transient patterns to persistent structures: a computational model of rapid memory formation in the hippocampal system. In preparation.

Shastri, L. & Ajjanagadde V. (1993) From simple associations to systematic reasoning: A connectionist encoding of rules, variables and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences* 16, 417–494.

Singer, W. & Gray, C.M. (1995) Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience* 18, 555–586.

Squire, L.R. (1992) Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review* 99, 195–231.

Suzuki, W.A. & Amaral, D.G. (1994) Perirhinal and parahippocampal cortices of the macaque monkey: Cortical Afferents. *Journal of Comparative Neurology* 350, 497–533.

Treves, A & Rolls, E.T. (1994) Computational analysis of the role of the hippocampus in memory. *Hippocampus* 4, 374–391.

Tulving, E. (1983) *Elements of Episodic Memory.* Clarendon Press, Oxford.

Valiant, L. (1994). *Circuits of the mind.* New York : Oxford University Press.

Van Hoesen, G.W. (1982) The primate hippocampus gyrus: New insights regarding its cortical connections. *Trends in Neuroscience* 5, 345–350.

von der Malsburg, C. (1986) Am I thinking assemblies? In *Brain Theory* (eds Palm, G. & Aertsen, A.) 161–176. Springer-Verlag, Berlin.

West, M.J. (1990) Stereological studies of the hippocampus: a comparison of the hippocampal subdivisions of diverse species including hedgehogs, laboratory rodents, wild mice and men. In *Progress in Brain Research: Understanding the brain through the hippocampus* (eds Storm-Mathisen, J., Zimmer, J. & Ottersen, O.P.) 13–36, Elsevier Science, Amsterdam, 1990.

Zalutsky, R.A. & Nicoll, R.A. (1990) Comparison of two forms of long-term potentiation in single hippocampal neurons. *Science*, 248: 1619–1624.