Tracing Recurrent Activity in Cognitive Elements (TRACE):

A Model of Temporal Dynamics in a Cell Assembly

Stephen Kaplan
Department of Electrical Engineering and Computer Science, and Department of Psychology, The University of Michigan

Martin Sonntag and Eric Chown
Department of Electrical Engineering and Computer Science, The University of Michigan

Running head: ACTIVITY IN COGNITIVE ELEMENTS

Address reprint requests to:
Professor Stephen Kaplan
Psychological Laboratories
Mason Hall
The University of Michigan
Ann Arbor, MI 48109-1027

phone: 313-764-0426
e-mail: Stephen_Kaplan@um.cc.umich.edu

Table of Contents

Hebb's introduction of the cell assembly concept marks the beginning of modern connectionism, yet its implications remain largely unexplored and its potential unexploited. Lately, however, promising efforts have been made to utilize recurrent connections, suggesting the timeliness of a reexamination of the cell assembly as a key element in a cognitive connectionism. Our approach emphasizes the psychological functions of activity in a cell assembly. This provides an opportunity to explore the dynamic behavior of the cell assembly considered as a continuous system, an important topic that we feel has not been given sufficient attention. A step-by-step analysis leads to an identification of characteristic temporal patterns and of necessary control systems. Each step of this analysis leads to a corresponding building block in a set of emerging equations. A series of experiments is then described that explore the implications of the theoretically derived equations in terms of the time course of activity generated by a simulation under different conditions. Finally the model is evaluated in terms of whether the various constraints deemed appropriate can be met, whether the resulting solution is robust, and whether the solution promises sufficient utility and generality.

Keywords: Neural networks, connectionism, cell assembly, neural fatigue, perception, simulation, peak activation, consolidation

1. Introduction

In 1949, Hebb presented a cautious yet revolutionary theory to account for the psychology and the physiology of perception as it was understood at the time. Central to Hebb's theory was the cell assembly[1] construct. Despite the venerable age and admitted incompleteness of cell assembly theory, this construct provides a cognitive foundation for connectionism not possessed by other models. In a cognitive connectionism the cell assembly serves as a symbol-like unit of thought. Unlike other symbol systems (and most connectionist systems) the cell assembly is grounded in an explicit neural mechanism.

Taken a step further, the cell assembly provides automatic mechanisms for learning and structure generation.

Despite the potential power of the cell assembly concept, the theory is, as Hebb himself emphasized repeatedly, incomplete. Further theoretical development is called for to harness the potential power of this deceptively simple concept. Over the years since the *Organization of Behavior* was written, it has became increasingly clear that several new theoretical components are needed. It has also become apparent that the resulting system of interacting constructs would necessarily be a reformulation rather than a simple extension of Hebb's original model.

The purpose of this paper is to present such a reformulation. The cell assembly provides the cognitive system with flexibility far beyond the simple activation of concepts. Instead of viewing the assembly as simply active or latent we see the activation of the assembly as coming in a series of phases. Each phase of activity serves a different purpose, giving the theory the power and flexibility to handle a wide range of psychological data.

This time course of activity in the cell assembly is due to the introduction of new control mechanisms to the basic model. Milner (1957) added one control mechanism, inhibition, critical to the theory. We postulate two additional control mechanisms, fatigue and short term connection strength, that not only are important functionally in controlling activity, but also serve the role of breaking the time course of activity into distinct phases.

After examining the cell assembly concept from physiological and psychological perspectives, we focus on what we see as a key conceptual issue, namely the potential roles of persisting activity in the cell assembly. We then turn our attention to the development of the model. TRACE is a mathematical model and computer simulation of a single Hebbian cell assembly. In keeping with our conceptual focus we have decided upon what might be called a "molar strategy." The cell assembly is represented in our simulation not as a network, but as a complex temporal process.

We then report on a series of simulation experiments in which the model is found to substantially agree with several theoretical predictions. The newly introduced factors not only serve their required functions as control mechanisms, but provide the cell assembly construct with additional information processing capabilities.

2. The Cell Assembly: Neurophysiological and Psychological Perspectives

A cell assembly is a highly interconnected group of neurons that form a reverberatory circuit which can sustain activity. This reverberatory activity, through the action of the learning rule, is the mechanism by which cell assemblies are formed and strengthened. These two processes are mutually reinforcing: the learning rule operates through the action of reverberatory activity, and the operation of the learning rule provides for enhanced reverberatory activity in the future. Storing information in memory is achieved by the strengthening of appropriate cell assemblies; retrieval from memory is the reactivation of these same assemblies.

Hebb assumed that long term memory is realized by changes in synaptic strength and that this synaptic strengthening is in turn determined by contiguity and repetition of neural activity. Hebb (1949) expressed this assumption as an explicit, qualitative rule:

> "Whenever an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cell's firing B, is increased." (p. 62)

Unlike previous theories invoking the cell assembly mechanism for dynamic memory storage, which have been refuted by empirical evidence (Hilgard & Marquis, 1940), Hebb's cell assembly resolves the apparent contradiction between dynamic and structural theories of memory. Dynamic memory stores information for a relatively short period of time. It also serves as a mechanism for laying down long term memory, which is structural. It is, in other words, a mechanism to implement the consolidation process, first proposed by Müller and Pilzecker (1900). The dynamic activity creates the structural changes which the data indicate are necessary for the durable character of long term memory.

2.1. Neurophysiological Evidence for the Cell Assembly

Much more is known today about cortical structure and function than was available to Hebb in 1949. None of the new information contradicts the cell assembly hypothesis, and some of it provides, if not conclusive empirical support, then at least strong encouragement for the notion.

The assumption that long term memory is implemented by an increase in synaptic efficacy is now widely accepted. An effect which operates similarly to Hebb's contiguity learning rule, known as long term potentiation (LTP) has been demonstrated and studied extensively in the hippocampus and to some extent in the cortex as well (Lynch, 1986). A likely candidate for the mechanism of this change has been found in the action of NMDA receptors, which seem to play a role in use-dependent synaptic strengthening during development and learning (Cotman & Iverson, 1987).

Detailed studies of the source and destination of axonal projections between cortical processing areas in primates have revealed that whenever an area receives input from another, it usually has a reciprocal influence on that area through feedback projections (Van Essen, 1985). Such recurrent connections are required for the self-excitatory loops which support developing cell assemblies.

A major shortcoming in Hebb's theory, which he recognized, is the need for some negative influence to prevent a rapid spread of neural activity through the strong positive synapses. Hebb cautiously omitted the inhibitory synapses necessary to correct this problem because of a lack of empirical evidence for them. Milner (1957), with Hebb's encouragement, added inhibition to the theory. Today several types of inhibitory interneurons are known to be abundant throughout the cortex (Shepherd, 1988), returning Hebb's proposal to a solid neurophysiological foundation.

Few research efforts have been undertaken to identify cell assemblies per se. There is a strong tendency to study the detailed molecular workings of neurons and synapses before proceeding to explore networks of neurons. In addition, there are technical difficulties facing any search for cell assemblies. While the components of cell assemblies near sensory cortex may be localized, the neurons composing other cell assemblies may be scattered throughout the cortex and possibly subcortical regions as well (Mishkin and Appenzeller, 1987). Thus it is far from clear where to look for the individual cells that comprise a single assembly.

Single cell recordings have identified extremely selective neurons in the inferior temporal cortex of monkeys (reviewed in (Perrett *et al.*, 1987)). Mishkin and Appenzeller (1987) suggest that these neurons are the extreme end of a processing pathway along which cells respond to more of the visual field and more complex combinations of features until a

complete representation of an object is achieved. This suggests the possibility that recordings are being made from a single cell of an extensive cell assembly.

Measuring the concurrent activity of multiple cells constitutes yet another challenge to any search for cell assemblies. Simultaneous in vivo multi-cell recording, once limited to monitoring less than fifty neurons (Gerstein *et al.*, 1983), can now monitor a few hundred neurons, but this number is still short of the potentially thousands of neurons which might comprise a cell assembly. However, information revealed by extracellular measurements, although it represents an average of many neurons' firings, may be useful in making inferences about the underlying neural activity.

Early evidence for self-excitation within recurrently connected groups of cortical neurons was found by Burns (1951). He isolated an area of a cat's cerebral cortex from all connecting neural tissue while maintaining its blood supply. He then measured changes in electrical potential on the surface of the isolated cortex in response to brief point stimulations. Because the response was a long lasting, asynchronous, repetitive pattern which spread across the entire cortical area before ending abruptly after 2 - 4*s*, he concluded that it resulted from excitation travelling around closed axodendritic loops.

Reports of widely separated neurons in cat visual cortex which fire in synchrony when responding to features from the same object (Gray *et al.*, 1989) suggest that groups of neurons can act in concert and might be a part of the object recognition process.

Work done by Walter Freeman's group (Freeman, 1991) provides the most convincing evidence to date that neurons act in concert rather than individually. They believe that this cooperative activity can be explained by appealing to the existence of cell assemblies. They have studied the processing of scents in rabbits by monitoring EEG signals in an array of locations across the olfactory bulb and olfactory cortex. The EEG reading at each point in the array provides a time series of average activity values for a localized group of neurons rather than a single neuron. Taken together, the EEG readings yield a picture of the collective activity within the regions of study over time. The data suggest that unique individual scents are represented in the olfactory bulb by a unique spatial pattern of activity amplitudes across the array. Thus for a given scent, some locations are highly active while others are not; for a different scent a different set of locations are highly active. These patterns appear in a "burst" in response to olfactory input,

during the period between inhalation and exhalation, and last less than a second. Such a concerted, unitary response is very much what might be expected if cell assemblies were responsible for the activity.

2.2. Psychological Functions of the Cell Assembly

During the period when Milner was modifying Hebb's theory, a set of computer simulations were conducted to investigate whether cell assemblies would form under the conditions Hebb assumed (Rochester *et al.*, 1956). This work is one of the earliest examples of modern connectionism. Currently, connectionist interest is focused on pattern classification tasks, following the direction pioneered by other early work (Rosenblatt, 1962). The simulation of cell assembly like aggregations has attracted little attention since the Rochester et al. study. This is unfortunate, since we believe the cell assembly construct could provide the critical link between pattern classification networks and networks which model higher cognitive processes.

Modern connectionist studies of higher cognitive processes (e.g. Rumelhart and McClelland (1986)) exhibit a peculiar discordance between the representational power assigned to a single unit and the neurophysiological basis for the unit. The units of the network are either assigned, if they are input or output units, or come to represent, if they are hidden units, specific semantic features. In other words, the unit represents a feature at a very high level of neural processing. Yet the function of a unit is loosely based on the computations that might be performed by a single, highly simplified neuron. The evidence suggests that, on the contrary, an extensive set of neurons contributes to the recognition of high level features. An even greater disparity is evident in the propositional models of classical artificial intelligence (AI), which dispense with any attempt to provide a grounding in neural mechanisms and begin their analysis at the level of the symbol.

It seems there may be a question which is not addressed in existing models of higher cognition, be they classical or connectionist. What exactly are the units of which thought is composed? They are probably not single neurons, yet the units in connectionist models function like simplified neurons[2]. They may be symbols, but classical AI so far has revealed nothing of their origins. Hofstadter (1985, chap. 26) refers to this problem when

he criticizes the passive symbols employed in AI models. Harnad (1990) has called this the symbol grounding problem.

Hebb's cell assembly is ideally suited to play the role of the unit of thought. Its development and function can be ascribed to explicit neural mechanisms. It is composed of an extensive set of neurons, and so it has properties beyond those of a single neuron. Once formed, a cell assembly acts as a unit. Its strong recurrent connections cause the entire assembly to become active if enough of its constituent neurons are activated, giving it the ability to "fill in" features missing from its input. This provides a mechanism for what Bruner (1957) referred to as "Going beyond the information given" and is closely akin to Gibson's (1979) concept of "affordances." Since a cell assembly is built by perceptual learning it represents an object of experience. This means it is a symbol which is firmly grounded in the environment and represents something meaningful in that environment. Margolis (1987) presents a useful analysis of why perception and pattern recognition are central to cognitive theory.

We suggest that the unit used in connectionist models of higher cognition ought to to be modelled after the cell assembly rather than the single neuron. The problems and benefits of using the cell assembly as an "active symbol" of something in external reality which can be combined with and influence other symbols in a connectionist cognitive system are presented in detail in Kaplan *et al.* (1990).

2.3. The Roles of Activity in Cognition

There is, then, a sound basis for considering the cell assembly or recurrent circuit to be a central construct in a connectionist model of mind. A key aspect of the Hebbian synthesis is the insight that the amount of activity among the elements making up the cell assembly determines the role the assembly plays. To oversimplify, an inactive assembly represents latent structure, while an active assembly is a cognitive element participating in ongoing cognition. In simplest terms, activity signifies presence as opposed to possibility. This distinction is a crucial one in human cognition, since the domain of "possibility" is very large while what can be considered "present" is severely limited. The human cognitive system, in other words, is characterized by a vast capacity for inactive or latent structure.

The portions of this structure that may be active concurrently is orders of magnitude smaller.

This limited capacity for what can be thought of and/or perceived at any one time may seem at a first glance to be an unfortunate bottleneck in an otherwise powerful system. Further consideration, however, suggests a quite different conclusion. Unlike many other animals, humans have the capacity to consider multiple alternatives, even possibilities that are hypothetical or contrary to fact. Such a system could all too readily become lost in thought. Indeed, early steps toward the development of a cognitively oriented learning theory by Tolman (1932) received exactly this criticism (Guthrie, 1935). The constraint of a limited capacity for activity means that the number of alternatives considered at one time will be modest, thus reducing the gulf between cognition and action.

The role of activity in distinguishing presence as opposed to possibility does not, however, exhaust the potential of this concept. Activity is a powerful and flexible code, capable of playing roles more differentiated and more varied than this simple binary distinction would suggest. Four roles are of particular interest in the present context:

1. Activation and perception. A complex organism presumably has a great many tightly knit associative structures, or, in the language of our framework, latent cell assemblies. Only a few of these can be actively experienced at any one time. Activity provides an appropriate means for distinguishing these few from the rest. In our model perception occurs when an environmental event leads to a high level of activity in a corresponding cell assembly.

2. Peak activation and dominance. Connectionist models tend to view perception as a competitive process. Those potential percepts receiving the greatest support, both external and internal, are presumed most likely to inhibit alternative interpretations of the same stimulus configuration. In this way a relatively small number of cell assemblies will ultimately become substantially more active than any others, and thus dominate. This same competitive process presumably applies to thought as well; the crucial differences are the distance of the cell assembly in question from the sensory interface and the degree to which its activation was due to the activity of other cell assemblies rather than to input from the environment.

The state of dominance characteristic of cell assemblies undergoing peak activation has important functional significance, since these highly active cell assemblies are the most likely to transmit information to other cell assemblies and/or to generate behavioral output. This mechanism provides an orderly way in which a small subset of the total stored information in the system can control the output of the system at a given time.

3. Activation and primary memory. James (1892) described primary memory as memory whose retrieval required no effort. Primary memory immediately follows perception. We propose that a lower level of activation than that required for perception serves as a code for a "just perceived" state. In this state the retrieval potential is high, but the stimulus in question is not experienced as being present.

4. Activation and consolidation. Considerable evidence points to a brief post perceptual process that is crucial for information to be stored in long term memory (Weingartner and Parker, 1984). The intensity and duration of this process impact whether information is stored, and how strongly it is stored. This post perceptual activity, called the consolidation process, can be viewed as an internal form of repetition that allows for a learning process considerably more efficient than one completely dependent on external repetition. While the function is different from primary memory, the same intermediate level of activation is assumed to provide the mechanism for both.

To fulfill these four roles then, the proposed model should include at least three distinct levels of activity. The lowest level corresponds to the latent, or inactive state. The highest level represents peak activation and perception; a middle level embodies primary memory and consolidation.

Each of these levels might be expected to have a distinctive time course. These characteristic temporal patterns, in turn, must also reflect the properties of the cell assembly as a whole. The most basic of these properties is that patterns must occur in a tightly knit structure. The danger in such a configuration is that activity has the potential of spreading exponentially. Such a tendency toward positive feedback necessarily requires a negative feedback control.

Activity would be expected to display initial positive feedback effects, followed by the influence of negative feedback controls. Furthermore, the precise form of this function might be expected to be susceptible to influence by a number of key parameters beyond the

duration and amplitude of the stimulus. In this way, what Hebb (1963) has called the "semi-autonomous process" has a semi-autonomous time course that is by no means identical to the time course of the stimulus that initiated it.

3. Building the TRACE Model

Since the cell assembly has the potential to play such a pivotal role in a theory of cognition, it seems promising to explore this construct more fully. At the time of the Rochester et al. (1956) studies, computer simulation was the method of choice. In vivo neuron recording was in its infancy and formal analytic methods could be applied only to very simple networks. Little has happened since then to favor another approach. While single cell recording technology is quite sophisticated, and extracellular recording techniques have proved useful for monitoring aggregate neuronal activity, many problems stand in the way of using multi-cell recording techniques to investigate cell assemblies (Gerstein *et al.*, 1989). Formal methods have been extended to the analysis of multi-layer networks, but only in certain restricted cases, e.g. fully symmetric networks (Hopfield, 1982) and back propagation networks (Rumelhart *et al.*, 1986). Thus computer simulation is still the preferred method. Computer simulation makes it possible to study phenomena too complex to analyze completely and too elusive to capture in the real world (Wolfram, 1984).

Since most of the important questions about the anatomy and physiology of cell assemblies remain unanswered, how do we go about constructing a model of one? Ecologists have faced a similar problem in their study of population dynamics. The behaviors of individuals are highly variable and yet their combined effect is often predictable. Ecologists have employed dynamic systems models which capture the essential interactions between a selected set of aggregate state variables. We pursue a similar course in modelling the global dynamic behavior of a cell assembly without specifying the detailed interactions between each of its component neurons.

The construction of a molar model affords us an opportunity to gather and organize the many constraints applicable to cell assembly functioning. These constraints can in the future guide construction of a network model as well. Furthermore, the molar model can specify quantitative behavior which future network models can strive to duplicate.

3.1 The Factors Underlying Activity: A Conceptual Analysis

In order to simulate the workings of this complex system it is necessary to identify the key factors at work. These become state variables in a mathematical description of how these factors influence themselves and each other. The resulting description takes the form of a system of simultaneous equations. The focus of these equations is, of course, activity. This construct is designated P, in memory of perseveration, a term for persisting neural activity used by Müller and Pilzecker in 1900. P stands for the level of activity in a single cell assembly.

The variables that influence activity can be grouped into four primary conceptual domains: Input, Prior experience, Competition, and Control. Let us examine each of these in turn.

3.1.1. Input.

An input pulse in this simulation is designated I. Input provides the initial impetus for activity, but in a well learned cell assembly the many strong interconnections assure the spread of activity after that activity has reached some critical level. This allows a subset of features to trigger the perception of an object, making available to the organism more information about the object than was available in the stimulus configuration.

3.1.2. Prior experience

The connection strength between elements in a cell assembly reflects the corresponding prior experience of the organism. This latent long term structure is expressed as L in the simulation. Long term structure brings with it a potential hazard. Consider a stimulus configuration for which the individual possesses only a weak cell assembly, or in the extreme case, none at all. In such a situation the activity generated by sensory input from the stimulus is likely to be drawn off into other, better learned cell assemblies. This tendency of a cell assembly to function as an attractor, while adaptive under many circumstances, thus constitutes a potential drawback as well. In the extreme case, such a tendency would dictate that only preexisting cell assemblies could be activated and hence new ones could never be learned.

Consider a bird-watcher-in-training who starts out with firm knowledge of only two birds, the robin and the blue jay. Then our inexperienced protagonist sees a kingfisher. This new stimulus clearly fits the blue jay category far better than the robin category. It is not an excellent fit, but if there is not some way for temporary structure to arise rapidly, the pattern of activity generated by the stimulus will not be supported. Instead, the overlap with the preexisting category, plus the positive feedback capacity of the underlying structure, will lead to a shift in the center of gravity of the resulting pattern. The kingfisher, in other words, would be categorized as a blue jay every time.

In a cell assembly that is not yet well learned, some factor is required that could substitute, albeit partially and temporarily, for the stable connections that will come to define the cell assembly and make positive feedback possible. A temporary connection change would be sufficient to permit the temporary operation of new cell assemblies. A similar construct has been proposed by Peterson (1966) and by Hinton and Plaut (1987). This hypothetical short term connection strength (STCS) is designated S.

3.1.3. Competition

The context for the cell assembly we are simulating is a larger neurocognitive system. This larger system necessarily has, as we have seen, a limited capacity as far as the number of active cell assemblies is concerned. The concrete expression of this limited capacity is in terms of competition from other active cell assemblies. The mechanism for this competition is the activity of inhibitory neurons such as those described by Milner (1957). These neurons are assumed to be sensitive to nearby activity, and to respond by generating nonspecific inhibitory activity in the region. This inhibitory activity serves to protect the very cell assemblies that engendered it, making it difficult for competing assemblies to become active. Thus, as long as P is large, there will be a correspondingly high level of regional inhibition, which will tend to restrict competition. Conversely, as P declines, there will be less inhibitory protection and a greater likelihood of competing activity. This competing activity, in turn, generates inhibition, which will tend to further undermine P.

3.1.4. Control

A system so well endowed with positive feedback requires comparably powerful controls. Paralleling other biological systems that tend to expend resources and build up deficits, it seems appropriate to employ a fatigue factor (designated F) for this purpose. Such a component would make it impossible for the activity of a given cell assembly to persist beyond reasonable bounds, thus providing an effective control on duration of activity. A highly active cell assembly is in a position of competitive advantage. It is in a dominant role with respect to the system as a whole, and it generates inhibition that constrains the activity of potentially competing cell assemblies. It is for this reason that fatigue plays such a vital role. Without it a dominant percept or thought could remain dominant indefinitely.

The upper bound on the amount of activity in a cell assembly is provided by the fact that it is comprised of a finite number of neurons. Another control on the amount of activity is the stochastic tendency of units to drop out over time. This is assumed to occur most often at high levels of P, since the neurons brought into play then are the most tenuously connected and the most likely to drop out.

3.1.5. Summarizing Remarks

Together these various considerations suggest a general form for the expected activity function. It begins to rise gradually until the short term connection factor builds to the point of aiding the process, and until there are enough active units to help stimulate other units, producing a snowballing effect. There is then a rapid rise as more and more elements are caught up in the positive feedback. After a certain point, however, there are fewer and fewer units left in the cell assembly that are not already activated, and the rise gradually flattens out. By this point some of the initial elements would be beginning to show the effect of fatigue, and a gradual decline in activity begins. The activity in the cell assembly is protected from interference by the inhibition it generates. However, as the activity declines, the inhibition lessens, and there is an ever increasing likelihood that competing activity will arise. This newly active competition, itself now generating inhibitory protection, will finally terminate the initial activity.

3.2. Some Empirical Landmarks

While these psychological and theoretical considerations provide important constraints in developing the model, further input is needed to bridge the gap between these rather broad considerations and the specification required for a mathematical description. This additional input is largely empirical. Its primary contribution is to help structure the temporal characteristics of key theoretical constructs.

3.2.1. The differentiated activity code: Perception and primary memory

Not only is activity a central theme in this model; it is also an important code. The highest levels of activity indicate perception, while a somewhat lower level of activity denotes primary memory. Thus, in order to determine the general form one might expect of the curve of activity over time, it is necessary to assign approximate durations to each of these codes.

The duration of perception was a central concern to Hebb in *The Organization of Behavior*; his comment on the value he chose is instructive:

> What I have in mind in emphasizing half a second or so as the duration of a reverberatory activity, is the observed duration of a single content in perception (Pillsbury, 1913; Boring, 1933). Attention wanders, and the best estimate one can make of the duration of a single "conscious content" is of this time order. (Hebb, 1949, p. 74)

What are experienced as percepts of longer duration might then be the expression of what Hebb (1963) called "a serially ordered activity of mediating processes or cell assemblies" (p. 26). James (1892) took a similar position. He pointed out that the focus of sustained interest "is not an identical *object* in the psychological sense, but a succession of mutually related objects forming an identical *topic* only, upon which the attention is fixed" (p. 92).

To this estimate of 0.5 *s* for perception it is possible to add a bit of further specification based on short term visual storage data. The transition to positive feedback should come at approximately 20 *ms* and the beginning of perception at approximately 50 *ms* after the onset of stimulation.

The duration of primary memory can be approximated in two different ways. One approach is to assume that it is equatable with the memory span, as that construct is usually measured. The traditional measurement procedure involves presenting digits at a rate of one per second. Thus, since the memory span has been determined to encompass 5 ± 2 items (Mandler, 1975), the approximate duration of primary memory can be estimated to be in the five *s* range.

A second approach to estimating the duration of primary memory depends upon the equation of primary memory with the period after the presentation of an item during which the information may be consolidating. Since consolidation and primary memory are both post-perceptual active processes, this would seem to be an appropriate assumption. Miller and Marlin (1984) have noted that the best controlled studies of the consolidation process also arrive at a maximum value of approximately five *s*.

3.2.2. The primary supporting variables: STCS and fatigue

Much of the rationale for our assumptions about the behavior of STCS and fatigue depend upon a line of research begun almost thirty years ago by Kleinsmith and Kaplan (1963, 1964). The findings of these and other studies using what we shall refer to as the "Kleinsmith paradigm" were so radical that they have not yet been successfully incorporated into the traditional theoretical frameworks in psychology. Since we feel that these findings can be readily understood in terms of a connectionist model, and since TRACE was greatly influenced by this line of research, it may be useful to describe this work in some detail.

The Kleinsmith paradigm emerged from what began as a straightforward investigation of the relationship between learning and arousal. The intention had been simply to measure differences in a subject's learning at high and low arousal levels. The learning task involved incidental learning of a paired associate list presented only a single time. Arousal was measured by item rather than by subject. Based on skin resistance reactions, each subject's items were divided into high and low arousal groupings. Although the initial expectation had been that high arousal material would be better recalled, it only required a few pretests to discover that quite the opposite was the case.

Initial puzzlement gave way to the insight that the intense neural activity resulting from high arousal might be producing a temporary fatigue effect that was leading to an

apparent failure of learning. Thus, if there were indeed a facilitating effect of arousal on learning, it would only be apparent at a later time, after the fatigue had dissipated. Based on this new set of expectations the experiment was redesigned with recall measured not only immediately after learning, but at various points thereafter (20 minutes, 45 minutes, one day, and one week in the first experiment in the series). To avoid contamination, a different group of subjects was tested at each recall interval.

For the low arousal items the recall was as would be expected. Recall of these items was relatively high soon after the learning trial, but declined steadily over time. The high arousal results were dramatically different, confirming the hypothesis that a fatigue effect was involved. Kleinsmith and Kaplan (1963) found that for items where arousal was high, the recall started out low and slowly improved. Thus subjects would be more likely to recall an item 20 *min* after seeing it than 2 *min* after seeing it. In addition to the surprising fact that the recall curve actually rose in the high arousal case these results were striking in that for a period of 10 to 20 *min* after the presentation of the material, recall was actually better for low arousal items than for high arousal items.

Because these results directly contradicted what any known model would have predicted, Kleinsmith and Kaplan (1964) repeated the experiment using nonsense syllables instead of words. This second experiment generated essentially identical results. A number of other attempts have been made to replicate these findings (Walker & Tarte, 1963; McLean, 1969; Butter, 1970). These studies have also obtained comparable results.

In order to make inferences about the time courses of STCS and fatigue based on these findings, it is necessary to assign some theoretical linkages. It is our assumption that the "low arousal" curves reflect some sort of short term memory, since recall for that material steadily declines over time. It cannot be a curve of primary memory, since that process appears to have a much briefer duration. Rather it seems to fit far better the STCS construct, posited to provide a basis for positive feedback in cases where the requisite long term structure, L, is not present. Based on the Kleinsmith data, STCS would be expected to persist for a considerable period, say on the order of 10 to 20 *min*.

The "high arousal" curves obtained using the Kleinsmith paradigm reflect what on theoretical grounds would be expected to be a fatigue effect, since the high arousal would cause intense activity with a concomitant rise in fatigue. It is evident that fatigue must not

build up too rapidly, lest it destroy the trace activity before perception occurs. On the other hand, it must decay reasonably rapidly, since the Kleinsmith results show recovery of previously inaccessible material by 5 or 6 *min* after presentation. At its peak, of course, it must be able to overcome both L and S combined, since these same studies show *below chance* performance where fatigue would be expected to be operative.

3.3. Quantitative Specifications

These various conceptual issues come together in a system of three simultaneous difference equations[3]. Of these, the equation for defining ΔP, the change in P over one time step, is by far the most complex and most significant. The discussion that follows focuses on the derivation of an equation for ΔP. Along the way we will introduce the other state variables of the system as they are needed. It may be useful to think of P(t) as the population of neural elements belonging to a given cell assembly that are active at any moment in time. In our simulation the change in this population is influenced by four parameters: (1) the population of elements active currently, P(t), (2) the fatigue level, F(t), (3) short term connection strength, S(t), and (4) the preexisting structure or long term connection factor, L(t). For convenience we designate a sensitivity function, V, which is a function of L, S, and F. It turns out to be conceptually useful to collect all the factors which affect sensitivity into one place. We will require that all variables are bounded between 0 and 1. Thus P(t) can be thought of as the proportion of elements (out of all possible elements in the cell assembly) that are active at time t.

The cell assembly is an internal structure that corresponds to some class of events in the environment. When the corresponding environmental event occurs, the cell assembly should become active. Eventually its activity should cease. Thus one would expect P to rise and fall over time in response to an input. It is possible to separate those factors contributing to the rise of activity from those contributing to its fall, such that

$$\Delta P = \Delta P_{rise} - \Delta P_{fall} \qquad (1)$$

Let us begin by specifying the rise portion of the process, ΔP_{rise}. A major limitation on the spread of activity is the fact that a cell assembly has a finite number of neural units; the more of these that are already active, the fewer that are left to become active. If the

proportion of active units is P(t), then the proportion of units left to stimulate would be (1 − P(t)). Another key component of the rise portion is V, the sensitivity variable. The higher the value of V, the weaker the stimulus necessary to activate the cell assembly. Combining these three factors yields

$$\Delta P_{rise} = P(t)\,(1 - P(t))\,V(t) \qquad (2)$$

which partly specifies the rise portion of ΔP. The three factors, P(t), (1 − P(t)), and V(t), have a noncompensatory relationship with each other. In other words, a deficit in one cannot be made up by an excess in another. Thus, for example, when P(t) = 1, there are no elements left to be activated, and ΔP_{rise} should be zero. It should not be possible to compensate for the zero value of (1 − P(t)). For this reason the relationship among these factors is multiplicative.

The sensitivity variable V represents a gathering together of all the factors that make a latent cell assembly more readily activated. This can be represented by[4]

$$V(t) = \frac{(L + S)\,(1 - F)}{\nu} \qquad (3)$$

Long term and short term connection strengths are positive factors in V, and fatigue is a negative factor. The effects of long term connection strength, L, and short term connection strength, S, on sensitivity are compensatory; either component can make up for the lack of the other. This is necessary given the important role of S in permitting learning to occur in a weak cell assembly, as discussed previously. The compensatory relationship is expressed by the fact that the relationship between L and S in equation (3) is additive. Fatigue produces a proportional decline in sensitivity; since its effect is negative it is represented by its complement, (1 − F), in equation (3). The divisor in equation (3) prevents V from exceeding unity, which might otherwise occur with S and L high and F low.

The equation generated so far for ΔP_{rise} incorporates positive feedback, the limited availability of units, and the role of the sensitivity factor. The remaining component of the rise portion of the equation is input, designated I. Input has the capacity to activate neural elements, just as does activity, P. These two factors are directly compensatory; a higher value of one can compensate for a lower value of the other. Their relationship thus must be additive. And here again, the number of available (unactivated) units places a limit upon the

possible influence of I. For this reason, the value I (1 − P) is added to P in Equation 1 to yield[5]

$$\Delta P_{rise} = (P + I\,(1 - P))\,(1 - P)\,V \qquad (4)$$

This completes the rise component of the equation for ΔP. With only this component, activity would grow rapidly given stimulation, as indeed it should. It would, however, grow gradually even without stimulation, and once large, would tend to level off and remain large. Thus, in addition to the growth or rise component, the decline or fall component is essential.

The decline component is composed of two factors, namely competition and loss. Just as the rise component benefits from a high sensitivity, V, the reverse is true for the fall component. Thus the decline component is not multiplied by V, but by its complement, (1 − V).

An important aspect of the decline component is the tendency of a cell assembly with a high level of activity to lose participating neurons as fatigue mounts. Some neurons are, as we have seen, less strongly connected to other neurons in the cell assembly, and these would be the first to drop out. Since this effect is most marked at very high levels of activity and tends to decrease as activity drops, it is expressed as P raised to a power, θ_l.

Another factor in the decline of activity is competition from other activity. If our model encompassed multiple cell assemblies this competition could be included explicitly. Since our model specifies only a single cell assembly, this competition must be approximated by a factor in the decline component. External competitive pressure on the cell assembly is greater when its own level of activity is low. This is because the protective inhibition, with which it holds down competitors, is less effective at low levels of activity. Thus this factor is expressed as the complement of P, i.e. (1 − P), to a power, θ_c. However in this form, the lower the activity, the greater the competition, until P disappears entirely. It is unlikely that spontaneous activity is completely eliminated under typical conditions. Thus, in order to reduce the likelihood of this occurring, this expression is moderated by P, such that the competitive effect when the activity is very low is itself low. Drawing together all the factors in the fall component ΔP_{fall}, we obtain

$$\Delta P_{fall} = (P^{\theta_l} + P (1 - P)^{\theta_c}) (1 - V) \quad (5)$$

It shoud be noted in equation (5) that these factors are compensatory, and furthermore, that they are multiplied by the complement of the sensitivity, (1 − V). In other words, the decline in V, in general due to increasing fatigue, enhances the decline component of the equation just as the increasing sensitivity enhances growth.

Substituting the rise component, ΔP_{rise}, and the fall component, ΔP_{fall}, into equation (1) yields the complete expression for ΔP, such that

$$\Delta P = (P + I (1 - P)) (1 - P) V - (P^{\theta_l} + P (1 - P)^{\theta_c}) (1 - V) \quad (6)$$

The functions for fatigue, ΔF, and for short term connection strength ΔS

$$\Delta F = \varphi_g P (1 - F)^2 - \varphi_d F \quad (7)$$

$$\Delta S = \sigma_g P (1 - S)^2 - \sigma_d S \quad (8)$$

are essentially similar: they both depend on the activity level, P; they are both limited by their maximum value of unity; they both decay gradually. Each incorporates constants which modify the rates of growth and decline. Adjustments of the constants is such as to make fatigue rise more slowly and decay more rapidly than short term connection strength, a decision based on logical and empirical considerations, to be discussed subsequently.

In the set of experiments presented here, we simplify by modelling a single cell assembly without learning, i.e.

$$\Delta L = 0 \quad (9)$$

In TRACE simple equations are used to model the aggregate effect of large numbers of processes. The input equation is no different. Further, the equation is not designed to model a particular input, but instead a prototypic input[6]. The input function

$$I(t) = \begin{cases} \alpha & \text{when } 0 < t \le \delta \\ 0.0 & \text{when } t > \delta \end{cases} \quad (10)$$

makes no assumptions about the nature of the input beyond the fact that it has amplitude α for some duration δ and then ceases.

4. Simulating the TRACE Model

In our discussion of TRACE we have focused on the constraints that have influenced the generation of the model. These constraints range from the neurophysiological to the cognitive. A series of experiments have been carried out to explore the implications and capabilities of the proposed model with regard to these constraints. These experiments vary in complexity and tone from the simple but critical verification of the model's ability to generate the desired time course of activity, to an examination of the model's potential to support distinctive modes of functioning.

Conceptually the experiments are divided into three categories (Table I). The first and most important focuses on the model's ability to generate the general form of the activity function. The second category consists of experiments two through five. These experiments deal with the key modulators of activity, fatigue, and STCS. While fatigue and STCS are supposed to play well specified roles in the model, it is not clear in a system as complex as TRACE, especially where timing is so critical, whether or not the

Table I. Summary of simulation experiments

Experiment	Purpose
1	Verification of model for the general case
2, 3, 4, 5	Examination of key modulators: Their effects and interactions
6, 7	Exploration of model's generality

interactions of such factors will have unforeseen consequences and whether the individual components will play their expected roles. The first two of these experiments examine the role that fatigue plays; they focus on the growth and decline parameters, respectively. The next two experiments view STCS in two contexts, one where the percept in question is already well learned and one where it is not. The final category of experiments, consisting of experiments six and seven, explores the potential of the model to be generalized to meet the special requirements of differing ends of the neurocognitive spectrum, specifically high level functioning on the one hand and near the sensory interface on the other.

The TRACE difference equations were simulated as shown in Table II. Since the reader has by now seen the derivation of these equations in some detail, we simply summarize the variables and parameters in Table II.

In each simulation run, all parameters except for those being manipulated were held constant at the indicated unmodified value. Initial conditions were always as given in Table II. In the experiments reported here we do not simulate the learning of a cell assembly. Thus within each experimental run, the value of L, long term connection strength, remains at its initial value, λ. The unmodified value of $\lambda = 0.5$ represents a well learned assembly. Each variable is bounded in the range [0,1] by the defining equations.

In interpreting the solution curves, we make a convenient assignment of 10 *ms* as the interval between time steps. Perception is considered to occur when the activity, P, rises above 0.55. This level is somewhat arbitrary in that the perceptual level of a cell assembly in this context cannot be directly determined. The proposed value, however, represents a useful rule of thumb based on extensive experience with the simulation.

Table II. The TRACE difference equation system as simulated

$$P(t+1) = P(t) + \Delta P \qquad \Delta P = (P + \bar{P}\, I)\, \bar{P}\, V - (P\theta_l + P\, \bar{P}\, \theta_c)\, \bar{V}$$

$$V = \frac{(L + S)\, \bar{F}}{\nu}$$

$$F(t+1) = F(t) + \Delta F \qquad \Delta F = \varphi_g P\, \bar{F}^2 - \varphi_d F$$

$$S(t+1) = S(t) + \Delta S \qquad \Delta S = \sigma_g P\, \bar{S}^2 - \sigma_d S$$

$$L(t+1) = L(t) + \Delta L \qquad \Delta L = 0.0$$

$$I(t) = \begin{cases} \alpha & \text{when } 0 < t \le \delta \\ 0.0 & \text{when } t > \delta \end{cases}$$

Variable	Interpretation	Initial condition
P(t)	perseverance, activity	P(0) = 0.01
F(t)	fatigue	F(0) = 0.01
S(t)	short term connection strength	S(0) = 0.01
L(t)	long term connection strength	L(0) = λ
I(t)	external input	I(0) = 0.0

Parameter	Interpretation	Unmodified value
θ_l	unit loss	5
θ_c	inhibitory competition	9
ν	normalization factor	1.5
φ_g	fatigue growth	0.14
φ_d	fatigue decline	0.0001
σ_g	short term connection strength growth	0.4
σ_d	short term connection strength decline	0.00015
λ	initial long term connection strength	0.5
α	input amplitude	1.0
δ	input duration	10

$\bar{X}$ denotes the quantity (1−X). The state variables P, F, S, and L are evaluated at time t in the equations for ΔP, ΔF, V, and ΔS; the input variable I is evaluated at t+1.

4.1. Experiment 1: The General Case

Let us begin with the basic curve of activity over time that is the focus of the proposed model.

The cell assembly being modelled is some distance from the sensory interface and therefore will not receive direct stimulation from the environment. This assumption is reflected by reducing the amplitude of input from its maximum possible value of 1.0. In Experiment 1 the amplitude of input, α, is set to 0.2. The duration for input, δ, is set at 10 time steps, or 100 *ms*.

The parameters manipulated in this experiment were short term growth rate, σ_g, short term decline rate, σ_d, fatigue growth rate, φ_g, and fatigue decline rate, φ_d. By setting these at 0.4, 0.00015, 0.14, and 0.0001 respectively not only do short term connection strength and fatigue match the theoretical expectations, but the activity trace does as well. These parameter values will serve as the unmodified settings for subsequent experiments.

The results of Experiment 1, shown in Figure 1, show that TRACE is indeed capable of modelling the general theoretical case. However, generating one curve is only a beginning in supporting TRACE's claim to be an interesting and useful cognitive model. TRACE was constructed conceptually with each component playing a well defined role. It is important that each of these components fill those roles as specified. Further, it is important that each of these additional components is free of unforeseen side effects. In addition the basic form of activity trace generated in Experiment 1 should be the rule rather than the exception. That is, the TRACE model should tend to generate activity curves essentially similar to the one in the basic model. That way the modest variations in parameters that will occur from individual to individual would not be expected to result in dramatic behavioral differences.

In this context the first experiment serves as only the first step in testing TRACE. The following experiments were designed to further test TRACE under a wide variety of conditions.

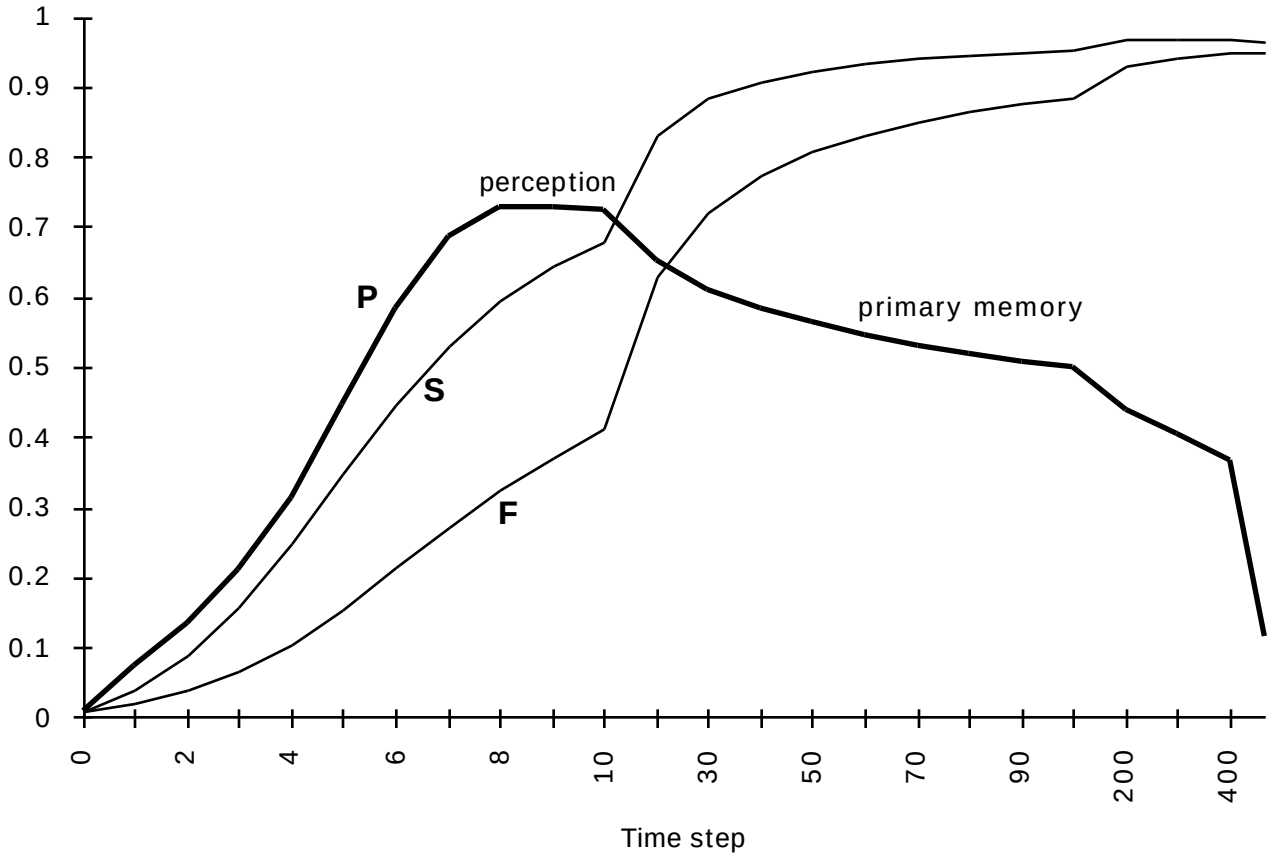

Figure 1. Results of Experiment 1: The basic time course of activity in a well learned cell assembly. Input occurs over the first 10 time steps (100 msec). The time scale is logarithmic.

4.2. Key Modulators: Fatigue

The Rochester *et al*. (1956) results suggest that activity based systems must have negative controls to hold the spread of activity in check. In most connectionist networks negative feedback is expressed in terms of inhibition. While the TRACE model uses inhibition as a method of selecting between competing percepts, an additional factor, fatigue, has been postulated that is not only complimentary to inhibition, but equally critical in holding activity in check. Competition between cell assemblies will mean that only the strongest circuits will become percepts. However, once an assembly becomes strongly

active an additional mechanism is necessary to ensure that it does not stay active indefinitely. The role of fatigue therefore should not be so much in preventing percepts from forming, but in providing a kind of shutoff mechanism once perception has occurred.

Table III. Effects of fatigue parameter changes

Parameter	Peak	Perception	Memory
	Experiment 2		
Fatigue growth, φ_g			
0.10	.750	75	873
0.11	.745	67	756
0.12	.739	60	548
0.13	.734	55	478
0.15	.724	48	388
0.16	.720	44	356
0.17	.715	41	331
0.20	.703	34	275
0.25	.685	27	218
0.40	.643	15	144
	Experiment 3		
Fatigue decline, φ_d			
0.00001	.729	52	293
0.00005	.729	52	333
0.00008	.729	52	381
0.00009	.729	52	402
0.00011	.729	52	457
0.00012	.729	52	498
0.00013	.729	52	559
0.00014	.729	52	668
0.00015	.729	52	∞

The left column identifies the parameter manipulated and the values tested. The right three columns give the values of three dependent variables: the height of the peak of the TRACE curve, the time steps until the onset of perception, and primary memory as measured by duration in time steps. (A time step equals 10 msec.)

4.2.1. Experiment 2: Effect of growth of fatigue

The depletion of resources within the cell assembly leads to the growth of fatigue. As fatigue grows the cell assembly loses its ability to sustain activity. Therefore the rate at which fatigue grows should impact how long the assembly can remain active. The next experiment was designed to test how changing the growth rate affected the durations of the perceptual and primary memory phases. The major goal of this experiment is to show that fatigue plays a critical role in shutting off the activity in a cell assembly. However, fatigue should also play a decisive role in the duration of perception while only having a minimal impact on its onset.

The mechanics of the experiment were very simple; vary the fatigue growth rate parameter, φ_g, to determine the nature of its impact on the activity trace. The results in Table 3 show the substantial effect of fatigue upon a cell assembly. By keeping the growth rate between 0.1 and 0.2 it is possible to vary the primary memory phase by as much as 6 *s* and the perceptual phase by as much as 0.4 *s*. The results also show that fatigue plays a lesser role during the onset of perception with activity varying only by 0.035 as the rate is moved between 0.1 and 0.2. Even with a rate as high as 0.4 the perceptual peak is still relatively high while perception and primary memory have dwindled considerably.

4.2.2. Experiment 3: Effect of decline of fatigue

The results of Experiment 2 were conclusive in showing the important role that fatigue plays in the cell assembly. They did not necessarily, however, tell the whole story. Fatigue is brought on by the depletion of resources. If the cell assembly is to recover from fatigue then those resources must be replenished. Just as the activity trace has a characteristic time course so does the buildup of fatigue. Fatigue is a consequence of activity and therefore should not have an effect on activity until later in the time course. Experiment 2 showed that by varying the rate at which fatigue builds up, it is possible to change the time that fatigue begins to take effect. Recovery from fatigue is a different issue, however. Just as activity must precede fatigue so must fatigue precede recovery from fatigue. It follows that changing the recovery rate of fatigue should only impact the activity trace late in its time course. Experiment 3 was designed to test this hypothesis. The object of the experiment was to determine at what point in the time course of activity recovery from fatigue becomes a factor.

The results of the experiment, shown in Table 3, are clear cut. Variations in the fatigue decline rate, φ_d, have a huge impact on the length of primary memory, but none at all on peak activity or the length of perception. If the recovery rate is fast enough, the primary memory phase can last indefinitely because the rates of depletion and recovery reach an equilibrium.

4.2.3. Reflection on robustness

Fatigue serves a definite function in shaping the activity trace. Once reverberation begins, a mechanism is needed to eventually shut it back down. Fatigue, even at widely varying parameter settings, serves this purpose. It dampens activity sufficiently to end perception fairly quickly and then later, in conjunction with inhibition, it shuts down primary memory. The parameters controlling fatigue can be used to vary these durations, and small changes in the parameters lead to similarly small changes in these durations.

Experiment 3 provides an interesting example of the tendency of the system to maintain stability. The effects of changing the recovery rate are asymmetrical around the base rate. A drop in the recovery rate will have a much smaller effect than a corresponding rise. In the context of a system consuming resources this is a reasonable reaction. It is unlikely that resources will ever suddenly become more readily available. However, it is easy to conceive of situations where resources are scarcer than usual. The results show that a drop in the incoming resources will have only minor effects.

4.3. Key Modulators: Short Term Connection Strength

The previous two experiments examined the critical role that fatigue plays in determining the length of the various phases of activity. Conceptually fatigue's counterpart is STCS. Fatigue serves to dampen activity, ultimately even to extinguish it, but only after activity becomes strong. STCS provides temporary positive support to permit activity to get started. It should be critical during the early rise of activity.

We have not called this modulator "short term memory" because of our conviction that there are in fact several types of short term memory, each with its own mechanism and characteristic range of durations. Primary memory and short term connection strength are,

for example, distinct in both respects; yet both are appropriately included under the broad heading, "short term memory." The special role of STCS is, as its name suggests, to provide a transient increment in connection strength. This is particularly critical in poorly learned or loosely connected assemblies where the capacity for reverberatory feedback is small. In these cases STCS can provide the necessary extra impetus to keep the assembly active. In well learned assemblies, however, where the capacity for feedback is large, the impact of STCS should be minimal.

4.3.1. Experiment 4: The impact of STCS on well learned cell assemblies

Basic TRACE simulates a well learned cell assembly. In such an assembly the impact of STCS should be small. The feedback in the assembly should be enough to drive perception on its own, and STCS should serve only to strengthen and speed the process. Once activity is at high levels the combined effect of feedback and fatigue should virtually nullify any impact of STCS.

Like fatigue, STCS is characterized by two parameters which determine the its rates of growth and decline. However, in this case growth and decline serve different purposes. As with the previous experiments, in this experiment we explore these purposes by modifying each parameter systematically while holding all the other parameters constant. The parameters varied were short term growth rate, σ_g, and short term decline rate, σ_d.

The results in Table IV show that STCS has the predicted effect for a well learned cell assembly. The major impact of short term growth rate comes during the rise of activity. However, the durations of perception and primary memory are virtually unchanged. Peak activity varies from .659 to .782, a small change considering the large variation of the parameter. The minimal impact of short term decline rate is even more striking. Large variations in the decline rate have virtually no impact on the system. It is only at a level more than 100 times the normal level that any noticeable, though still minor, change is seen.

Table IV. Experiment 4: effects of short-term connection strength parameter changes

Parameter	Peak	Perception	Memory

Short term growth, σ_g			
0.1	.659	43	413
0.3	.714	50	424
0.5	.743	51	428
1.0	.782	53	431
Short term decline, σ_d			
0.00005	.729	51	434
0.00020	.729	51	423
0.00050	.729	51	410
0.00200	.728	50	379
0.02000	.725	43	295

The left column identifies the parameter manipulated and the values tested. The right three columns give the values of three dependent variables: the height of the peak of the TRACE curve, the time steps until the onset of perception, and primary memory as measured by duration in time steps. (A time step equals 10 *ms*.)

4.3.2. Experiment 5: The role of STCS in poorly learned cell assemblies

The results of Experiment 4 indicate that STCS has a minimal effect on the activity curves under the conditions tested. Much as this outcome is theoretically reasonable, it does raise a concern as to whether STCS plays any role whatsoever. Experiment 5 was designed to explore this issue. In our theoretical framework, STCS should make its primary contribution when the cell assembly in question is not well learned. Thus in Experiment 5 the methods of Experiment 4 were repeated with this modification: the value of long term connection strength, L, normally constant at 0.5, was lowered to a constant 0.2. The model predicts that under these circumstances STCS will be essential to support enough activity to achieve and sustain perception. In addition the rate at which STCS increases should be directly proportional to how quickly perception occurs.

Table V. Experiment 5: effects of STCS parameter changes when L is low

Parameter	Peak	Perception	Memory
High long-term connection strength (λ = 0.5)			

Short term growth, σ_g			
0.0	.597	10	133
Low long-term connection strength ($\lambda = 0.2$)			
Short term growth, σ_g			
0.0	.061	–	–
0.1	.090	–	–
0.2	.131	–	–
0.3	.633	33	295
0.4	.650	36	295
0.5	.660	36	295
0.6	.672	38	295
1.0	.698	40	296
Short term decline, σ_d			
0.00005	.650	37	299
0.00007	.650	37	297
0.00009	.650	37	296
0.00015	.650	36	295
0.00017	.650	36	294
0.00020	.650	36	294
0.00050	.650	36	284

The left column identifies the parameter manipulated and the values tested. The right three columns give the values of three dependent variables: the height of the peak of the TRACE curve, the time steps until the onset of perception, and primary memory as measured by duration in time steps. (A time step equals 10 msec.) A dash (–) signifies failure to achieve a peak or primary memory phase.

Table V shows that STCS does indeed have the predicted effect on the buildup of activity. In conditions where the growth rate is low, activity cannot build quickly enough in conjunction with input to achieve perception. Once a critical level is reached then activity can be sustained for a reasonable time. Even at the highest levels, however, the activity curve will never reach the levels of the basic TRACE model. For an organism functioning in an uncertain world this makes good sense. One should rely on the learned before the new; thus perception of the new is achieved at a more tentative level.

Although not shown in the table, the experiment also established that the growth of STCS is related to the speed of perception. To take two examples, when short term growth,

σ_g, was set to 0.3 the onset of perception took 0.18 *s* (slow compared with the well learned case). At 0.6 this was speeded up to 0.1 *s*. It is worth noting here that in Experiment 3 manipulation of fatigue growth rate yielded only negligible effects on this onset rate.

Experiments 4 and 5 demonstrate that STCS plays considerably different roles depending upon how well learned the cell assembly is. When the circuit is not well learned, STCS is critical in providing the support to generate enough activity for perception. In well learned assemblies, however, the role of STCS is minor. It is worth noting here that the decline rate of STCS had little effect on activity in either case. This does not mean that it plays no significant role. This rate would be expected to influence later perceptions as well as the ease of retrieval from primary memory.

4.3.3. Reflection on robustness

The role of STCS is essentially limited to providing temporary connections that allow reverberation in new cell assemblies. Taken together, Experiments 4 and 5 show exactly this effect. Even at extreme parameter settings STCS has little effect on a well learned assembly. And in the case where the assembly is not well learned, STCS has remarkably little effect beyond getting reverberation started. Once reverberation has begun, the boost given to activity by STCS is minimized by the other factors present. The experiments show that the model is so robust with respect to STCS as to be almost completely insensitive to parameter settings.

4.4. Modelling Special Purpose Variants

Human neurocognitive functioning cannot be unitary. If one thinks of the cognitive system as being organized in a hierarchical fashion, then some assemblies will be relatively close to the sensory interface while others would be engaged in high level cognitive functions quite far from the direct influence of sensory input. It seems reasonable to expect assemblies at these levels to behave differently. For example, one might expect neurons in the sensory interface to be characterized by brief persistence and rapid recovery, since it is necessary for the sensory system to "clear" promptly in order to be ready for new input. In contrast, units further along in the neurocognitive hierarchy, i.e. farther from the sensory

interface, might be expected to show exactly the opposite pattern. The guidance of planful behavior, for example, would require high persistence and considerable resistance to fatigue.

At the same time it would be economical for the system to use the same basic mechanisms throughout despite the expected differences in behavior. The purpose of these final two experiments is to determine if these contrasting behaviors are within the capacity of the basic underlying set of mechanisms which TRACE is intended to model.

4.4.1. Experiment 6: High level functioning

Experiment 6 was designed to test the ability of TRACE to generate behavior associated with high level assemblies. Consider a higher level cognitive activity such as planning. The landmark treatment of this topic in *Plans and the Structure of Behavior* (Miller *et al*., 1960) made it clear that the circuitry involved in planning should be capable of activity that is persistent and highly resistant to fatigue. As Tolman (1932) pointed out in his classic analysis of purposive behavior, the plan must continue to guide behavior as the organism moves through the environment, encountering obstacles and impediments of various kinds along the way. The nature of the experiment intended to capture this essential property of high level assemblies was simple; the idea was to reduce fatigue with the expected result that activity should be more persistent. Resistance to fatigue implies that the available resources are used more efficiently. Therefore fatigue decline rate, φ_d, was raised, signifying that resources were replenished more quickly.

The experiment proved successful. By increasing fatigue decline, φ_d, from the base rate of 0.0001 it was possible to extend the length of the activity curve indefinitely. A fatigue decline rate equal to 0.00014 generates a persistent activity trace appropriate for the high level assembly being modelled. These results can be seen in Figure 2.

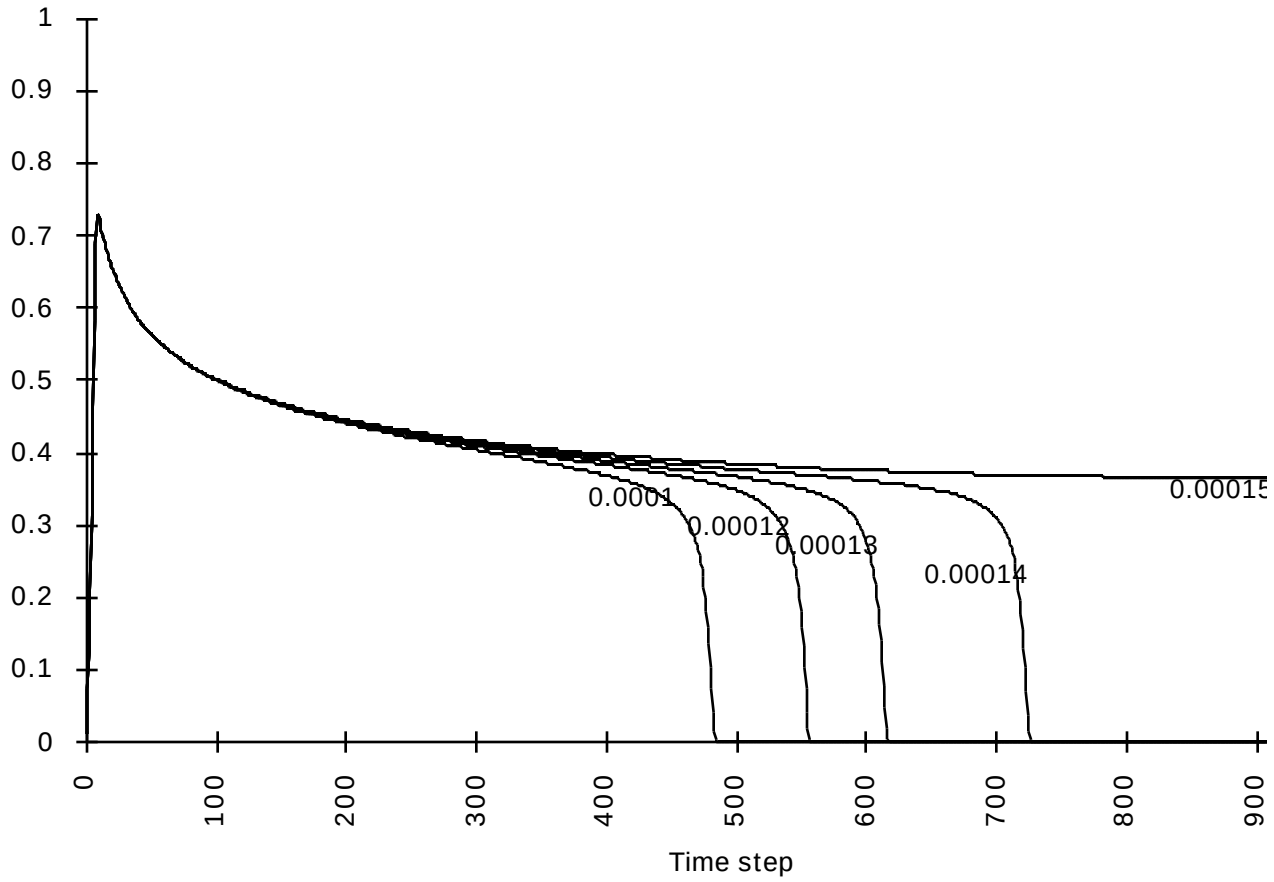

Figure 2. Results of Experiment 6: Activity traces generated by varying the decline rate of fatigue, φ_d. As φ_d is raised from the base rate of 0.0001 to a maximum of 0.00015, the duration of primary memory is extended. In each case the input occurs over the first 10 time steps (100 msec).

4.4.2. Experiment 7: Near the sensory interface

A more complex set of issues was raised by an experiment dealing with behavior at the lower end of the processing continuum, as might be exhibited at the sensory interface. The target behavior calls for relatively short activity traces with STCS and fatigue also clearing rapidly. At the perceptual interface there can be little or no long term structure. After all, the elements at this end of the hierarchy are shared components that temporarily

participate in the link between information from the environment and cognitive activities at a higher level. For this reason, the constant value of long term connection strength, L, was lowered to 0.2. In addition there is the issue of competition. As James (1892) points out, perception is of "the probable and the definite." Perception is characteristically experienced as clear and definitive despite uncertainties and ambiguities. The system is biased towards a single dominant percept as opposed to several alternate conceptions of what one might be looking at. For this reason the inhibitory competition parameter, θ_c, was reduced to 4, thereby raising the level of competition. In addition the amplitude of the input, α, was increased to 1.0 to reflect the fact that these elements are receiving direct stimulation.

In the absence of long term structure the role of STCS becomes more vital in providing the impetus for activity. With this in mind short term growth, σ_g, was set at 1.0 to maximize the onset rate. Even with short term connection strength at its maximum, without long term structure the sensitivity is particularly susceptible to fatigue. For this reason the growth rate of fatigue, φ_g, was set only at 0.1. Both fatigue and short term connection strength should recover rapidly since the system must be ready for another input. For this reason their decline rates, σ_d and φ_d, were set at high levels (0.2 and 0.01 respectively).

Using these parameters a trace is generated which peaks at an activity level of 0.729 with a perceptual phase of 0.25 *s* and a primary memory phase which last less than 0.75 *s*. Recovery for short term connection strength is virtually instantaneous while fatigue recovers in approximately 3 *s*. This curve is shown in Figure 3.

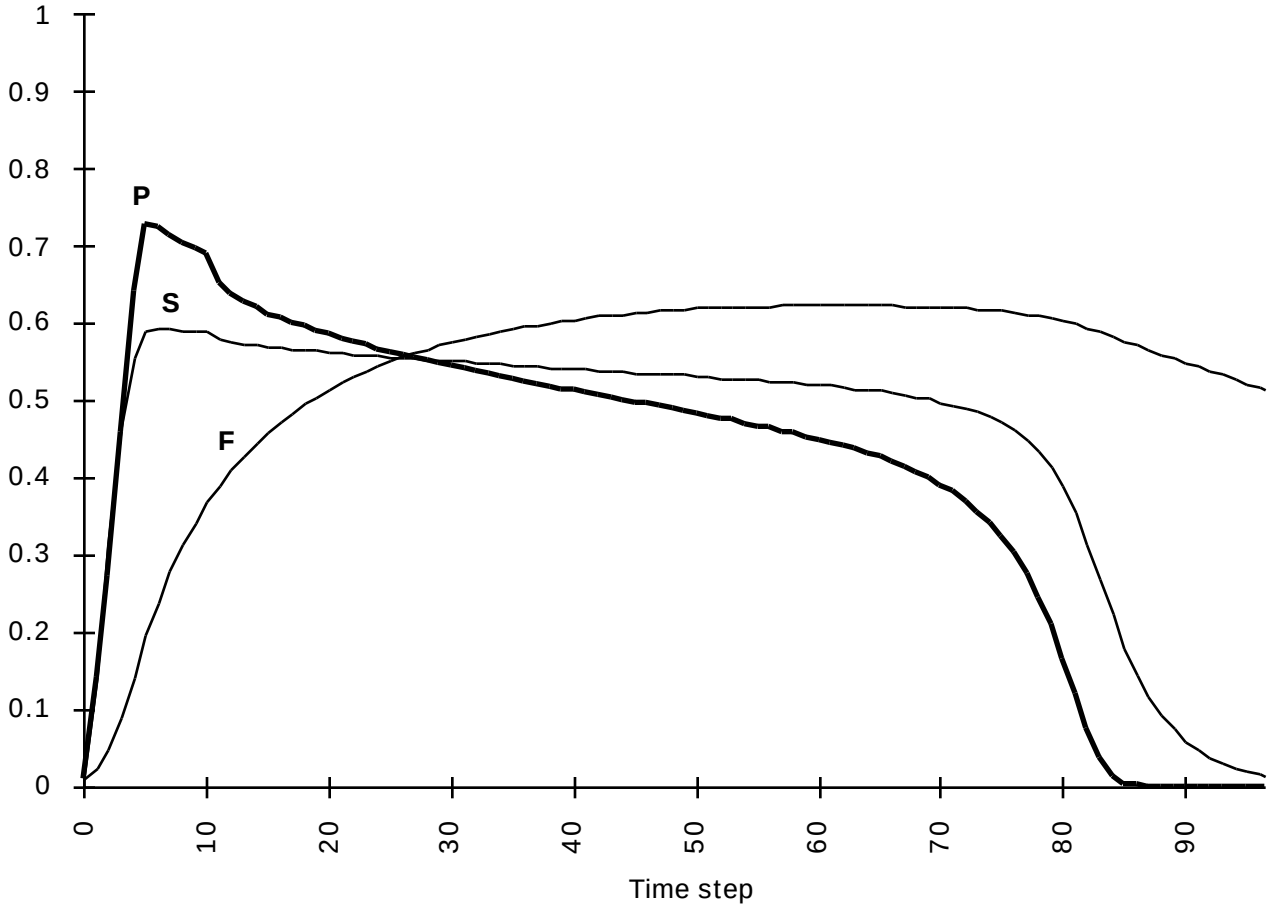

Figure 3. Results of Experiment 7: The time course of activity in a cell assembly near the sensory interface. Input occurs over the first 10 time steps (100 msec).

5. Conclusions

Hebb introduced the cell assembly construct as a theoretical bridge between neural mechanisms and psychological processes. We feel that there is still a need for such an intermediate construct in modern cognitive modelling efforts and that the cell assembly is ideally suited to serve in this capacity.

Perceptual experience guides the operation of neural learning mechanisms to form cell assemblies which represent environmental regularities and objects. We have argued that

the resultant cell assemblies provide a plausible neural implementation of grounded symbols. For this reason we propose the cell assembly as the "unit of thought" in any analysis of cognition. In order to achieve a closer correspondence between neural mechanisms and the analysis of cognitive tasks, we further propose that network units which stand for cell assemblies replace the simple neuron-like units often employed in connectionist models of higher cognition. The dynamic properties of this type of unit provide additional capabilities for a neurocognitive system. The time course of unit activity can serve to denote just perceived input, to denote primary memory contents, and to consolidate long term memory contents.

Thus we agree with Hebb in considering the cell assembly to be a basic and pervasive element in the way that the human brain processes information. As such it is surrounded by many constraints, both theoretical and empirical. This circumstance is central to the first question our research addresses: Is there a set of theoretically meaningful equations that meets these various constraints? Based on the results of our research, it seems clear that such a set of equations does indeed exist. We do not claim that TRACE represents a complete account of the behavior of a cell assembly. Instead we believe that we have captured some key constructs critical to understanding the nature of the cell assembly's functionality. As such the experiments in this paper do not represent exact, numerical solutions but instead are intended as examples of the kind of exploration possible within the context of TRACE.

The second question that arises in this context is the following: Does the set of equations we arrived at manage to meet these various constraints only within a highly constricted range of parameters? In other words, is this solution fragile, unstable and tenuous, or is it robust? Here the results are unequivocal; the simulation is well behaved across a broad parameter range. In the areas where it is somewhat more sensitive to parameter variation (particularly with respect to the fatigue construct), it is theoretically reasonable for this to be the case.

The fact that the model is well behaved should not be taken to suggest that it is, or that it should be, insensitive. The model is based on multiple, interacting parameters. In a small way this characteristic reflects the system it is intended to model, which is also a multi-parameter interactive system. Like the human mind, one would not expect the model to function properly at all possible parameter settings. Also like the human mind, however,

one would expect considerable robustness; it should be possible to function at least passably under a relatively wide range of settings. As we have seen, the model displays both of these characteristics.

Finally, it is appropriate to ask whether the solution we have arrived at has any generality. One way to address this issue is in terms of the range of content to which our model appears to be pertinent. From this perspective the model seems to be reasonably broad. It speaks to research in the perceptual domain, where it would seem particularly appropriate for addressing such issues as brief presentations, multiple presentations, and the results of presenting stimuli that are more or less well learned. It also speaks to fatigue and satiation effects, such as those studied by Pomerantz *et al.*, (1969).

The fact that the model incorporates the consolidation process makes it pertinent to issues of learning, memory, and reinforcement. Perhaps particularly salient here are mechanisms of unsupervised learning, and the related (and recently rejuvenated) study of the phenomenon of human consolidation, as demonstrated by subsequent studies following the Kleinsmith paradigm.

Furthermore, as we have attempted to demonstrate, parameter variations allow one to tailor the model to widely differing information processing requirements, including the sensory interface at one extreme, and high level cognitive processes such as planning on the other. One such application is the study of counterfactuals by French and Weaver (1987). Work is in progress on applying a network of TRACE elements to the problem of learning temporal sequences and building cognitive structure from them.

It may be that the model we have described can serve as a fairly general test bed for exploring connectionist solutions at a relatively molar level for a wide range of interesting psychological processes. Work is currently under way to model the dynamics of the learning process over time, with special attention to the rather remarkable differences in short vs. long term memory effects that are revealed by the use of the Kleinsmith procedure. By explicitly identifying consolidation with activity, TRACE provides the flexibility and automaticity necessary for a wide ranging learning theory.

The path we have taken to attack the cell assembly problem is far from traditional. Rather than building a network model we have built a simulation using difference equations. While this type of simulation may not have the kind of power in the long run that a network

model would have, it is an important piece in a very complicated puzzle. There are a host of problems facing the construction of a neuron level model including the number of neurons in an assembly, the connection patterns involved, the control of activity, and the maintenance of discreteness among cooperating cell assemblies. What the TRACE simulation does is to isolate some of these issues in order that they may be tackled in manageable steps.

By way of example, most connectionist models do not address temporal issues. In our paper we have repeatedly stressed the importance of timing in a cell assembly model. The TRACE simulation provides an environment where these issues can be examined closely. We have been able to study the interactions of factors such as fatigue and STCS in a fairly high level environment. An additional benefit is that it was not necessary to make assumptions about numbers of neurons or patterns of connections.

We see the construction of a network model as a necessary long term goal of our research. However, as with any complex problem, by reaching more manageable subgoals along the way our path should be smoother. We have already described some of the other work that is being done with TRACE. These are all part of our larger strategy, namely to carry out small, mutually constraining simulations at multiple levels. We see each of these efforts as both useful explorations in their own right, and as components of an eventual larger understanding. The more we understand about the dynamics of a cell assembly the easier the task of building a neuron level network model will become. And when such a simulation is ultimately carried out, its results can then be placed in the context of a space that has already been explored.

Acknowledgements

The TRACE model was first presented at the Annual Meeting of the Society for Mathematical Psychology at Ann Arbor, Michigan in August 1969. The authors would like to thank Roger Howard and Thomas Plum who participated in that initial effort. The analysis, testing, and description that led to the current article greatly benefitted from the suggestions of the members of SESAME, and in particular those of Mark Weaver, Ron Lesperance, Frances Kuo, Eric Ivancich and Robert French. We would also like to thank James Anderson and James Levenick for their helpful comments.

Notes

1. We have adopted the expression 'cell assembly' because of its common usage in descriptions of Hebb's work. Hebb's analysis of cognitive processes actually included three levels of organization: cell assemblies, phase sequences and phase cycles. Cell assemblies served as elements in the higher level network that consituted the phase sequence (and likewise for the relationship between the phase sequence and the phase cycle). Since Hebb considered the cell assembly to be the smallest unit of conscious thought, our simulation, which is oriented to perception, cognition, and memory for objects, would be more properly considered as at the phase sequence level. For the purposes of our model, however, these distinctions may not be of great importance. We assume that the basic functions characteristic of these networks are likely to be the same across the three different levels.
2. The embedding fields theory of Grossberg (1969, 1971), in which the unit activation function has positive beedback and decay, provides a unique exception to this generalization.
3. We present here a system of difference equations. In fact, solutions to the corresponding system of differential equations were obtained by several approximation methods. These solutions were quite close to those obtained by employing Euler's method with a step size of 1. So, to simplify the exposition and reduce computational requirements, we present and simulate the difference equations.
4. In the remainder of the exposition we replace the functional notation $X(t)$ with the simpler notation X. It should be understood that a variable X which appears in the right hand side of an equation is evaluated at time t.
5. I is an instantaneous function, not a state variable, so it is evaluated at time $t+1$ rather than at time t as are the state variables in these equations.
6. TRACE has been simulated with a vareity of input equations with negligible difference in results.

References

Boring, E. G. (1933). *The physical dimensions of consciousness*. New York: Century.

Bruner, J. S. (1957). On going beyond the information given. In *Contemporary approaches to cognition* Cambridge, MA: Harvard, pp. 41-69

Burns, B. D. (1951). Some properties of isolated cerebral cortex in the unanaesthetized cat. *Journal of Physiology,* **112**, 156-175.

Butter, M. J. (1970). Differential recall of paired associates as a function of arousal and concreteness-imagery levels. *Journal of Experimental Psychology,* **84**, 252-256.

Cotman, C. W., & Iversen, L. L. (1987). Excitatory amino acids in the brain - focus on NMDA receptors. *Trends in Neuroscience,* **10**(7), 263-265.

Freeman, W. J. (1991). The physiology of perception. *Scientific American*, (February), 78-85.

French, R., & Weaver, M. (1987). The role of categories in the generation of counterfactuals: a connectionist interpretation. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society* Hillsdale, NJ: Lawrence Erlbaum,.pp. 938-944

Gerstein, G. L., Bedenbaugh, P., & Aertsen, A. M. H. J. (1989). Neuronal assemblies. *IEEE Transactions on Biomedical Engineering,* **36**(1), 4-14.

Gerstein, G. L., Bloom, M. J., Espinosa, I. E., Evanczuk, S., & Turner, M. R. (1983). Design of a laboratory for multineuron studies. *IEEE Transactions on Systems, Man, and Cybernetics,* **SMC-13**(5), 668-676.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin.

Gray, C. M., König, P., Engel, A. K., & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature,* **338**, 334-337.

Grossberg, S. (1969). Embedding fields: A theory of learning with physiological implications. *Journal of Mathematical Psychology*, **6**, 209-239.

Grossberg, S. (1971). Embedding fields: Underlying philosophy, mathematics, and applications to psychology, physiology, and anatomy. *Journal of Cybernetics,* **1**(1), 28-50.

Guthrie, E. R. (1935). *The psychology of learning*. New York: Harper.

Harnad, S. (1990). The symbol grounding problem. *Physica D*, **42**, 335-346.

Hebb, D. O. (1949). *The Organization of Behavior*. New York: John Wiley.

Hebb, D. O. (1963). The semiautonomous process: It's nature and nurture. *American Psychologist*, **18**, 16-22.

Hilgard, E. R., & Marquis, D. G. (1940). *Conditioning and learning*. New York: Appleton-Century-Crofts.

Hinton, G. E., & Plaut, D. C. (1987). Using fast weights to deblur old memories. In *The Ninth Annual Conference of the Cognitive Science Society* (pp. 177-186). Hillsdale, NJ: Lawrence Erlbaum.

Hofstadter, D. R. (1985). *Metamagical themas*. New York: Basic Books.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational properties. *Proceedings of the National Academy of Sciences*, **79**, 2554-2558.

James, W. (1892). *Psychology: The briefer course* (1961 ed.). New York: Harper & Row.

Kaplan, R., & Kaplan, S. (1969). The arousal-retention interval interaction revisited: The effects of some procedural changes. *Psychonomic Science*, **15**, 84-85.

Kaplan, S., Weaver, M., & French, R. (1990). Active symbols and internal models: Towards a cognitive connectionism. *AI & Society*, **4**, 51-71.

Kleinsmith, L. J., & Kaplan, S. (1963). Paired-associate learning as a function of arousal and interpolated interval. *Journal of Experimental Psychology*, **65**, 190-193.

Kleinsmith, L. J., & Kaplan, S. (1964). Interaction of arousal and recall interval in nonsense syllable paired-associate learning. *Journal of Experimental Psychology*, **67**, 124-126.

Lynch, G. (1986). *Synapses, circuits, and the beginnings of memory*. Cambridge, MA: MIT Press.

Mandler, G. (1975). Consciousness: respectable, useful, and probably necessary. In R. L. Solso (Ed.), *Information processing and cognitive psychology.* Hillsdale, NJ: Lawrence Erlbaum.

Margolis, H. (1987). *Patterns, thinking, and cognition: A theory of judgement*. Chicago: University of Chicago Press.

McLean, P. D. (1969). Induced arousal and time of recall as determinants of paired-associate recall. *British Journal of Psychology*, **60**, 57- 62.

Miller, G., Galanter, E., & Pribram, K. (1960). *Plans and the structure of behavior*. New York: Henry Holt.

Miller, R. R., & Marlin, N. A. (1984). The physiology and semantics of consolidation. In H. Weingartner, & E. S. Parker (Ed.), *Memory consolidation: Psychobiology of cognition* Hillsdale, NJ: Lawrence Erlbaum.

Milner, P. M. (1957). The cell assembly: Mark II. *Psychological Review*, **64**(4), 242-252.

Mishkin, M., & Appenzeller, T. (1987). The anatomy of memory. *Scientific American*, **256**(6), 80-89.

Müller, G. E., & Pilzecker, A. (1900). Lehre vom gedachtnis. *Z. Psychol., Suppl. No. 1*.

Perrett, D. I., Mistlin, A. J., & Chitty, A. J. (1987). Visual neurones responsive to faces. *Trends in Neuroscience*, **10**(9), 358-364.

Peterson, L. R. (1966). Short-term verbal memory and learning. *Psychological Review*, **73**(3), 193-207.

Pillsbury, W. B. (1913). "Fluctuations of attention" and the refractory period. *J. Phil. Psychol. Sci. Meth*., **10**, 181-185.

Pomerantz, J. R., Kaplan, S., & Kaplan, R. (1969). Satiation effects in the perception of single letters. *Perception and Psychophysics*, **6**, 129-132.

Rochester, N., Holland, J. H., Haibt, L. H., & Duda, W. L. (1956). Tests on a cell assembly theory of the action of the brain, using a large digital computer. *IRE Transactions on Information Theory*, **IT-2**, 80-93.

Rosenblatt, D. E. (1962). *Principles of neurodynamics*. New York: Spartan.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart, J. L. McClelland, & the. PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition: Volume 1: Foundations.* Cambridge, MA: MIT Press.

Rumelhart, D. E., & McClelland, J. L. (1986). On learning the past tenses of English verbs. In J. L. McClelland, D. E. Rumelhart, & the PDP Research Group (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition: Volume 2: Psychological and biological models.* Cambridge, MA: MIT Press.

Shepherd, G. M. (1988). A basic circuit of cortical organization. In M. S. Gazzaniga (Ed.), *Perspectives in memory research*. Cambridge, MA: MIT Press.

Tolman, E. C. (1932). *Purposive behavior in animals and men*. New York: Century.

Van Essen, D. C. (1985). Functional organization of primate visual cortex. In A. Peters & E. G. Jones (Eds.), *Cerebral cortex*. New York: Plenum.

Walker, E. L., & Tarte, R. D. (1963). Memory storage as a function of arousal and time with homogeneous and heterogeneous lists. *Journal of Verbal Learning and Verbal Behavior,* **2**, 113-119.

Weingartner, H., & Parker, E. S. (Eds.). (1984). *Memory consolidation: Psychobiology of cognition*. Hillsdale, NJ: Lawrence Erlbaum.

Wolfram, S. (1984). Computer software in science and mathematics. *Scientific American*, **251**(3), 188-203.

Footnotes

[1] We have adopted the expression "cell assembly" because of its common usage in descriptions of Hebb's work. Hebb's analysis of cognitive processes actually included three levels of organization: cell assemblies, phase sequences, and phase cycles. Cell assemblies served as elements in the higher level network that constituted the phase sequence (and likewise for the relationship between the phase sequence and the phase cycle). Since Hebb considered the cell assembly to be the smallest unit of conscious thought, our simulation, which is oriented to perception, cognition, and memory for objects, would be more properly considered as at the phase sequence level. For the purposes of our model, however, these

distinctions may not be of great importance. We assume that the basic functions characteristic of these networks are likely to be the same across the three different levels.

[2] The embedding fields theory of Grossberg (1969,1971), in which the unit activation function has positive feedback and decay, provides a unique exception to this generalization.

[3] We present here a system of difference equations. In fact, solutions to the corresponding system of differential equations were obtained by several approximation methods. These solutions were quite close to those obtained by employing Euler's method with a step size of 1. So, to simplify the exposition and reduce computational requirements, we present and simulate the difference equations.

[4] In the remainder of the exposition we replace the functional notation X(t) with the simpler notation X. It should be understood that a variable X which appears in the right hand side of an equation is evaluated at time t.

[5] I is an instantaneous function, not a state variable, so it is evaluated at time t+1 rather than at time t as are the state variables in these equations.

[6] TRACE has been simulated with a variety of input equations with negligible difference in results.