

ORIGINAL CONTRIBUTION

ART 3: Hierarchical Search Using Chemical Transmitters in Self-Organizing Pattern Recognition Architectures

GAIL A. CARPENTER AND STEPHEN GROSSBERG

Center for Adaptive Systems, Boston University

(Received 29 June 1989; revised and accepted 21 August 1989)

Abstract—A model to implement parallel search of compressed or distributed pattern recognition codes in a neural network hierarchy is introduced. The search process functions well with either fast learning or slow learning, and can robustly cope with sequences of asynchronous input patterns in real-time. The search process emerges when computational properties of the chemical synapse, such as transmitter accumulation, release, inactivation, and modulation, are embedded within an Adaptive Resonance Theory architecture called ART 3. Formal analogs of ions such as Na^+ and Ca^{2+} control nonlinear feedback interactions that enable presynaptic transmitter dynamics to model the postsynaptic short-term memory representation of a pattern recognition code. Reinforcement feedback can modulate the search process by altering the ART 3 vigilance parameter or directly engaging the search mechanism. The search process is a form of hypothesis testing capable of discovering appropriate representations of a nonstationary input environment.

Keywords—Neural network, Pattern recognition, Adaptive Resonance Theory, Hypothesis testing, Search, Transmitter, Modulator, Synapse, Reinforcement, Competition.

1. INTRODUCTION: DISTRIBUTED SEARCH OF ART NETWORK HIERARCHIES

This article incorporates a model of the chemical synapse into a new Adaptive Resonance Theory (ART) neural network architecture called ART 3. ART 3 system dynamics model a simple, robust mechanism for parallel search of a learned pattern recognition code. This search mechanism was designed to implement the computational needs of ART systems embedded in network hierarchies, where there can, in general, be either fast or slow learning and distributed or compressed code representations. The search mechanism incorporates a code reset property that serves at least three distinct functions: to correct erroneous category choices, to learn from reinforcement feedback, and to respond to changing

input patterns. The three types of reset are illustrated, by computer simulation, for both maximally compressed and partially compressed pattern recognition codes (Sections 20–26).

Let us first review the main elements of ART. ART architectures are neural networks that carry out stable self-organization of recognition codes for arbitrary sequences of input patterns. ART first emerged from an analysis of the instabilities inherent in feedforward adaptive coding structures (Grossberg, 1976a). More recent work has led to the development of two classes of ART neural network architectures, specified as systems of differential equations. The first class, ART 1, self-organizes recognition categories for arbitrary sequences of binary input patterns (Carpenter & Grossberg, 1987a). A second class, ART 2, does the same for either binary or analog inputs (Carpenter & Grossberg, 1987b).

Both ART 1 and ART 2 use a maximally compressed, or choice, pattern recognition code. Such a code is a limiting case of the partially compressed recognition codes that are typically used in explanations by ART of biological data (Grossberg, 1982a, 1987a, 1978b). Partially compressed recognition codes have been mathematically analysed in models for competitive learning, also called self-organizing feature maps, which are incorporated into ART models as part of their bottom-up dynamics (Grossberg, 1976a,

This research was supported in part by the Air Force Office of Scientific Research (AFOSR F49620-86-C-0037 and AFOSR F49620-87-C-0018), the Army Research Office (ARO DAAL03-88-K-0088), and the National Science Foundation (NSF DMS-86-11959 and IRI-87-16960). We thank Diana Meyers, Cynthia Suchta, and Carol Yanakakis for their valuable assistance in the preparation of the manuscript.

Requests for reprints should be sent to Gail A. Carpenter, Center for Adaptive Systems, Boston University, 111 Cummington Street, Boston, MA 02215.

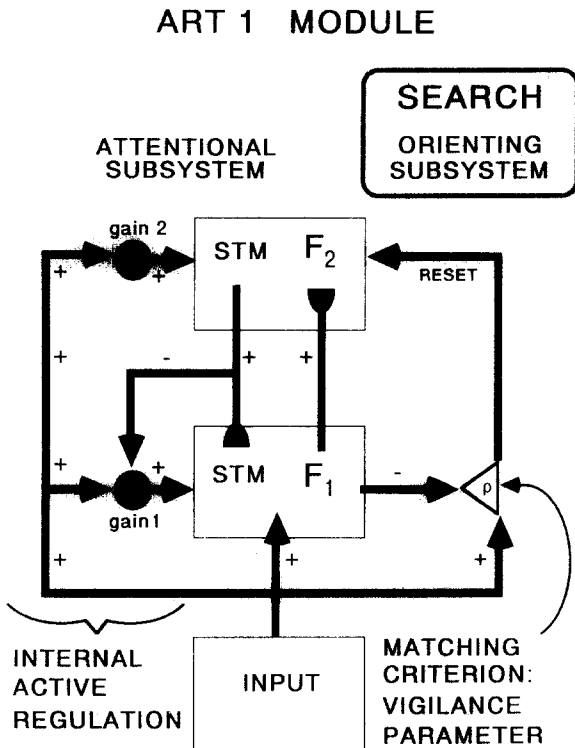


FIGURE 1. Typical ART 1 neural network module (Carpenter & Grossberg, 1987a).

1982a; Kohonen, 1984). Maximally compressed codes were used in ART 1 and ART 2 to enable a rigorous analysis to be made of how the bottom-up and top-down dynamics of ART systems can be joined together in a real-time self-organizing system capable of learning a stable pattern recognition code in response to an arbitrary sequence of input patterns. These results provide a computational foundation for designing ART systems capable of stably learning partially compressed recognition codes. The present results contribute to such a design.

The main elements of a typical ART 1 module are illustrated in Figure 1. F_1 and F_2 are fields of network nodes. An input is initially represented as a pattern of activity across the nodes, or feature detectors, of field F_1 . The pattern of activity across F_2 corresponds to the category representation. Because patterns of activity in both fields may persist after input offset yet may also be quickly inhibited, these patterns are called short-term memory (STM) representations. The two fields, linked both bottom-up and top-down by adaptive filters, constitute the Attentional Subsystem. Because the connection weights defining the adaptive filters may be modified by inputs and may persist for very long times after input offset, these connection weights are called long-term memory (LTM) variables.

An auxiliary Orienting Subsystem becomes active

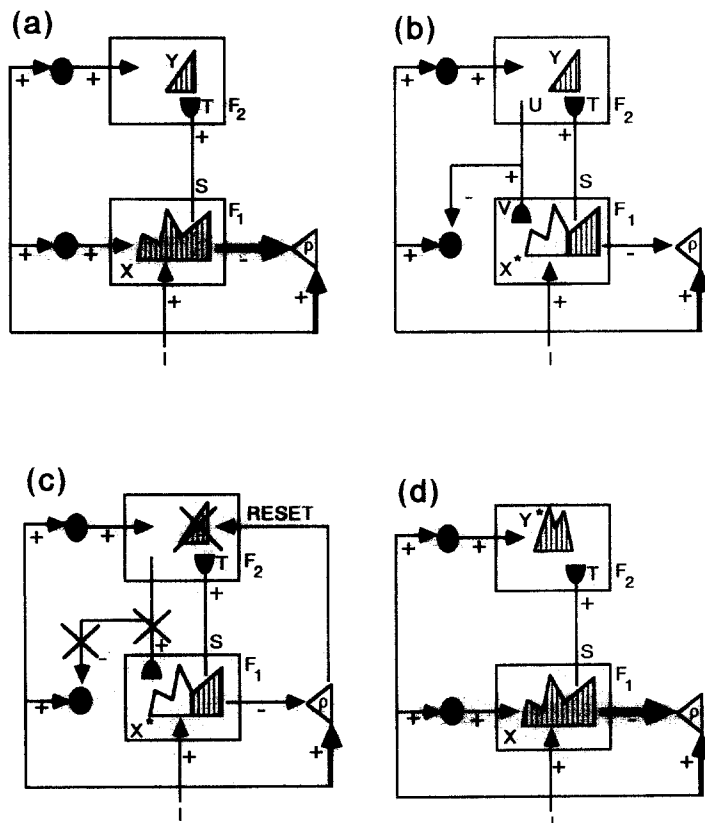


FIGURE 2. ART search cycle (Carpenter & Grossberg, 1987a).

during search. This search process is the subject of the present article.

2. AN ART SEARCH CYCLE

Figure 2 illustrates a typical ART search cycle. An input pattern \mathbf{I} registers itself as a pattern \mathbf{X} of activity across F_1 (Figure 2a). The F_1 output signal vector \mathbf{S} is then transmitted through the multiple converging and diverging weighted adaptive filter pathways emanating from F_1 , sending a net input signal vector \mathbf{T} to F_2 . The internal competitive dynamics of F_2 contrast-enhance \mathbf{T} . The F_2 activity vector \mathbf{Y} therefore registers a compressed representation of the filtered $F_1 \rightarrow F_2$ input and corresponds to a category representation for the input active at F_1 . Vector \mathbf{Y} generates a signal vector \mathbf{U} that is sent top-down through the second adaptive filter, giving rise to a net top-down signal vector \mathbf{V} to F_1 (Figure 2b). F_1 now receives two input vectors, \mathbf{I} and \mathbf{V} . An ART system is designed to carry out a matching process whereby the original activity pattern \mathbf{X} due to input pattern \mathbf{I} may be modified by the *template pattern* \mathbf{V} that is associated with the current active category. If \mathbf{I} and \mathbf{V} are not sufficiently similar according to a matching criterion established by a dimensionless *vigilance parameter* ρ , a reset signal quickly and enduringly shuts off the active category representation (Figure 2c), allowing a new category to become active. Search ensues (Figure 2d) until either an adequate match is made or a new category is established.

In earlier treatments (e.g., Carpenter & Grossberg, 1987a), we proposed that the enduring shut-off of erroneous category representations by a non-specific reset signal could occur at F_2 if F_2 were organized as a gated dipole field, whose dynamics depend on depletable transmitter gates. Though the new search process does not here use a gated dipole field, it does retain and extend the core idea that transmitter dynamics can enable a robust search process when appropriately embedded in an ART system.

3. ART 2: THREE-LAYER COMPETITIVE FIELDS

Figure 3 shows the principal elements of a typical ART 2 module. It shares many characteristics of the ART 1 module, having both an input representation field F_1 and a category representation field F_2 , as well as Attentional and Orienting Subsystems. Figure 3 also illustrates one of the main differences between the examples of ART 1 and ART 2 modules so far explicitly developed; namely, the ART 2 examples all have three processing layers within the F_1 field. These three processing layers allow the ART 2 system to stably categorize sequences of analog

ART 2 MODULE




FIGURE 3. Typical ART 2 neural network module, with three-layer F_1 field (Carpenter & Grossberg, 1987b). Large filled circles are gain control nuclei that nonspecifically inhibit target nodes in proportion to the Euclidean norm of activity in their source fields, as in eqn (33).

input patterns that can, in general, be arbitrarily close to one another. Unlike in models such as back propagation, this category learning process is stable even in the fast learning situation, in which the LTM variables are allowed to go to equilibrium on each learning trial. In Figure 3, one F_1 layer reads in the bottom-up input, one layer reads in the top-down filtered input from F_2 , and a middle layer matches patterns from the top and bottom layers before sending a composite pattern back through the F_1 feedback loop. Both F_1 and F_2 are shunting competitive networks that contrast-enhance and normalize their activation patterns (Grossberg, 1982a).

4. ART BIDIRECTIONAL HIERARCHIES AND HOMOLOGY OF FIELDS

In applications, ART modules are often embedded in larger architectures that are hierarchically organized. Figure 4 shows an example of one such hierarchy, a self-organizing model of the perception and production of speech (Cohen, Grossberg, &

SPEECH

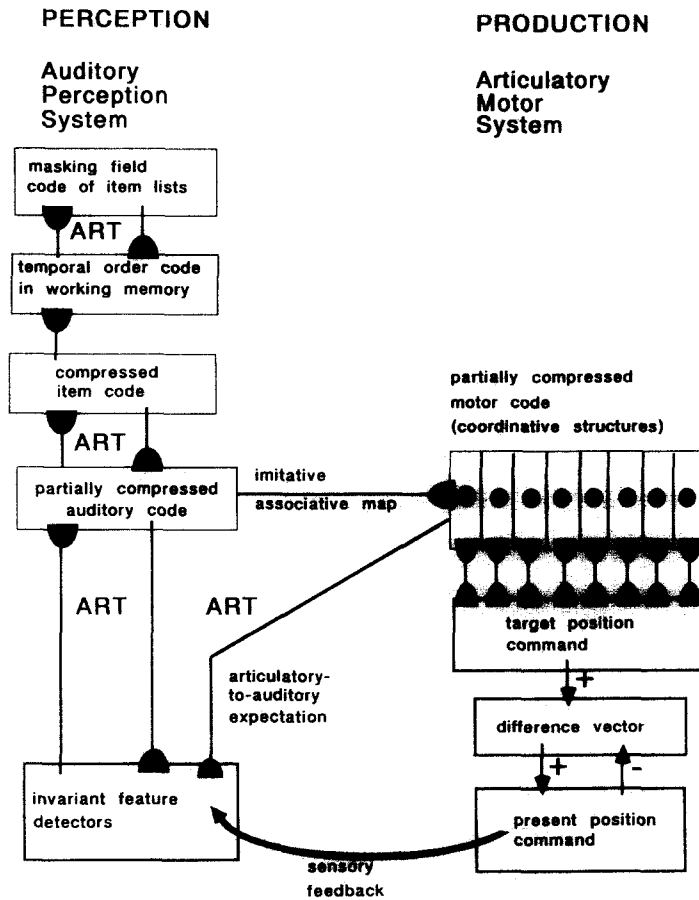


FIGURE 4. Neural network model of speech production and perception (Cohen, Grossberg, & Stork, 1988).

Stork, 1988). In Figure 4, several copies of an ART module are cascaded upward, with partially compressed codes at each level. Top-down ART filters both within the perception system and from the production system to the perception system serve to stabilize the evolving codes as they are learned. We will now consider how an ART 2 module can be adapted for use in such a hierarchy.

When an ART module is embedded in a network hierarchy, it is no longer possible to make a sharp distinction between the characteristics of the input representation field F_1 and the category representation field F_2 . For example, within the auditory perception system of Figure 4, the partially compressed auditory code acts both as the category representation field for the invariant feature field and as the input field for the compressed item code field. For them to serve both functions, the basic structures of all the network fields in a hierarchical ART system should be homologous in so far as possible (Figure 5). This constraint is satisfied if all fields of the hierarchy are endowed with the F_1 structure of an ART 2 module (Figure 3). Such a design is sufficient for the F_2 field as well as the F_1 field because the principal property required of a category representation field,

ART BIDIRECTIONAL HIERARCHY

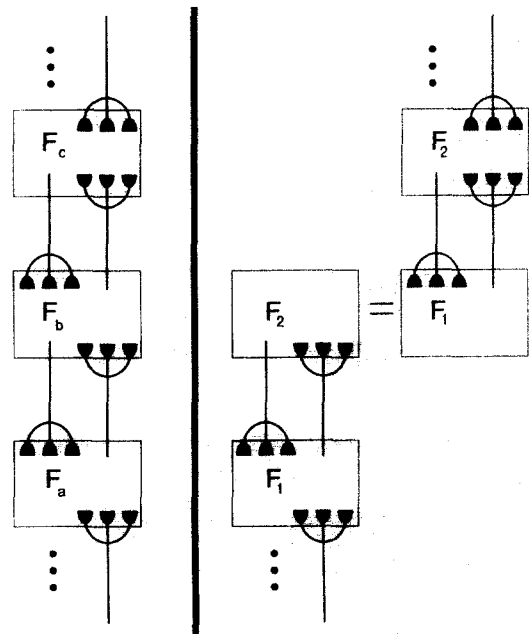


FIGURE 5. Homology of fields F_a, F_b, F_c, \dots in an ART bidirectional hierarchy.

namely that input patterns be contrast-enhanced and normalized, is a property of the three-layer F_1 structure. The system shown in Figure 5 is called an *ART bidirectional hierarchy*, with each field homologous to all other fields and linked to contiguous fields by both bottom-up and top-down adaptive filters.

5. ART CASCADE

For the ART hierarchy shown in Figure 5, activity changes at any level can ramify throughout all lower and higher levels. It is sometimes desirable to buffer activity patterns at lower levels against changes at higher levels. This can be accomplished by inserting a bottom-up pathway between each two-field ART module. Figure 6 illustrates a sequence of modules $A, B, C \dots$ forming an *ART cascade*. The "category representation" field F_{2A} acts as the input field for the next field F_{1B} . As in an ART 2 module (Figure 3), connections from the input field F_{2A} to the first field F_{1B} of the next module are nonadaptive and unidirectional. Connections between F_{1B} and F_{2B} are adaptive and bidirectional. This scheme repeats itself throughout the hierarchy. Activity changes due to a reset event at a lower level can be felt at higher levels

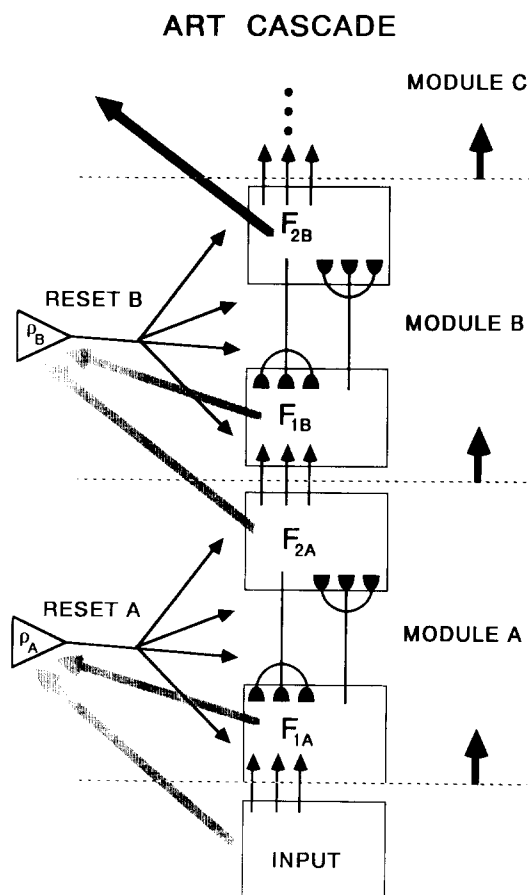


FIGURE 6. An ART cascade. Nonadaptive connections terminate in arrowheads. Adaptive connections terminate in semicircles.

via an ascending cascade of reset events. In particular, reset at the lowest input level can lead to a cascade of input reset events up the entire hierarchy.

6. SEARCH IN AN ART HIERARCHY

We now consider the problem of implementing parallel search among the distributed codes of a hierarchical ART system. Assume that a top-down/bottom-up mismatch has occurred somewhere in the system. How can a reset signal search the hierarchy in such a way that an appropriate new category is selected? The search scheme for ART 1 and ART 2 modules incorporates an asymmetry in the design of levels F_1 and F_2 that is inappropriate for ART hierarchies whose fields are homologous. The ART 3 search mechanism described below eliminates that asymmetry.

A key observation is that a reset signal can act upon an ART hierarchy *between* its fields $F_a, F_b, F_c \dots$ (Figure 7). Locating the site of action of the reset signal between the fields allows each individual field to carry out its pattern processing function without introducing processing biases directly into a field's internal feedback loops.

The new ART search mechanism has a number of useful properties. It: (a) works well for mismatch, reinforcement, or input reset; (b) is simple; (c) is homologous to physiological processes; (d) fits naturally into network hierarchies with distributed codes and slow or fast learning; (e) is robust in that it does

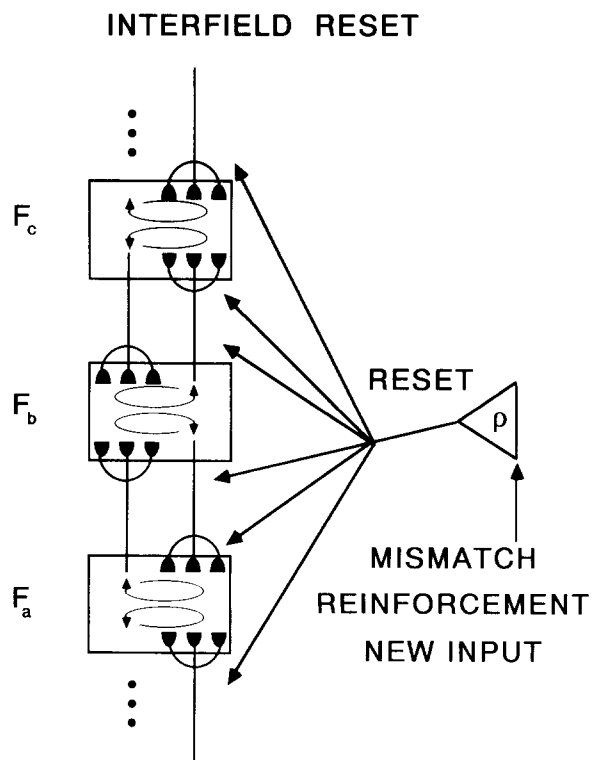


FIGURE 7. Interfield reset in an ART bidirectional hierarchy.

not require precise parameter choices, timing, or analysis of classes of inputs; (f) requires no new anatomy, such as new wiring or nodes, beyond what is already present in the ART 2 architecture; (g) brings new computational power to the ART systems; and (h) although derived for the ART system can be used to search other neural network architectures as well.

7. A NEW ROLE FOR CHEMICAL TRANSMITTERS IN ART SEARCH

The computational requirements of the ART search process can be fulfilled by formal properties of neurotransmitters (Figure 8), if these properties are appropriately embedded in the total architecture model. The main properties used are illustrated in Figure 9, which is taken from Ito (1984). In particular, the ART 3 search equations incorporate the dynamics of production and release of a chemical transmitter substance, the inactivation of transmitter at postsynaptic binding sites, and the modulation of these processes via a nonspecific control signal. The net effect of these transmitter processes is to alter the ionic permeability at the postsynaptic membrane site, thus effecting excitation or inhibition of the postsynaptic cell.

The notation to describe these transmitter properties is summarized in Figure 10 for a synapse between the i th presynaptic node and the j th postsynaptic node. The presynaptic signal, or action potential, S_i arrives at a synapse whose adaptive weight, or long-term memory trace, is denoted z_{ij} . The variable z_{ij} is identified with the maximum amount of available transmitter. When the transmitter at this synapse is fully accumulated, the amount of transmitter u_{ij} available for release is equal to z_{ij} . When a signal S_i arrives, transmitter is typically released. The vari-




FIGURE 8. The ART search model specifies rate equations for transmitter production, release, and inactivation.




FIGURE 9. Schematic diagram showing electrical, ionic, and chemical events in a dendritic spine synapse. Open arrows indicate steps from production of neurotransmitter substance (p) to storage (s) or release (r) to reaction with subsynaptic receptors (c), leading to change of ionic permeability of subsynaptic membrane (i) or to removal to extracellular space (m), enzymatic destruction (d), or uptake by presynaptic terminal (u). t, action of trophic substance. Note. From *The Cerebellum and Neural Control* (p. 52) by M. Ito, 1984, New York: Raven Press. Copyright 1984 by Raven Press. Reprinted by permission.

able v_{ij} denotes the amount of transmitter released into the extracellular space, a fraction of which is assumed to be bound at the postsynaptic cell surface and the remainder rendered ineffective in the extracellular space. Finally, x_j denotes the activity, or membrane potential, of the postsynaptic cell.




FIGURE 10. Notation for the ART chemical synapse.

8. EQUATIONS FOR TRANSMITTER PRODUCTION, RELEASE, AND INACTIVATION

The search mechanism works well if it possesses a few basic properties. These properties can be realized using one of several closely related sets of equations, with corresponding differences in biophysical interpretation. An illustrative system of equations is described below.

Equations (1)–(3) govern the dynamics of the variables z_{ij} , u_{ij} , v_{ij} , and x_j at the ij th pathway and j th node of an ART 3 system.

Presynaptic Transmitter

$$\frac{du_{ij}}{dt} = (z_{ij} - u_{ij}) - u_{ij}[\text{release rate}]. \quad (1)$$

Bound Transmitter

$$\begin{aligned} \frac{dv_{ij}}{dt} &= -v_{ij} + u_{ij}[\text{release rate}] - v_{ij}[\text{inactivation rate}] \\ &= -v_{ij} + u_{ij}[\text{release rate}] - v_{ij}[\text{reset signal}]. \end{aligned} \quad (2)$$

Postsynaptic Activation

$$\begin{aligned} \varepsilon \frac{dx_j}{dt} &= -x_j + (A - x_j)[\text{excitatory inputs}] \\ &\quad - (B + x_j)[\text{inhibitory inputs}] \\ &= -x_j + (A - x_j) \\ &\quad \times \left[\sum_i v_{ij} + \{\text{intrafield feedback}\} \right] \\ &\quad - (B + x_j)[\text{reset signal}]. \end{aligned} \quad (3)$$

Equation (1) says that presynaptic transmitter is produced and/or mobilized until the amount u_{ij} of transmitter available for release reaches the maximum level z_{ij} . The adaptive weight z_{ij} itself changes on the slower time scale of learning, but remains essentially constant on the time scale of a single reset event. Available presynaptic transmitter u_{ij} is released at a rate that is specified below.

A fraction of presynaptic transmitter becomes postsynaptic bound transmitter after being released. For simplicity, we ignore the fraction of released transmitter that is inactivated in the extracellular space. Equation (2) says that the bound transmitter is inactivated by the reset signal.

Equation (3) for the postsynaptic activity x_j is a shunting membrane equation such that excitatory inputs drive x_j up toward a maximum depolarized level equal to A ; inhibitory inputs drive x_j down toward a minimum hyperpolarized level equal to $-B$; and activity passively decays to a resting level equal to 0 in the absence of inputs. The net effect of bound trans-

mitter at all synapses converging on the j th node is assumed to be excitatory, via the term.

$$\sum_i v_{ij}. \quad (4)$$

Internal feedback from within the target field (Figure 3) is excitatory, while the nonspecific reset signal is inhibitory. Parameter ε is small, corresponding to the assumption that activation dynamics are fast relative to the transmitter accumulation rate, equal to 1 in eqn (1).

The ART 3 system can be simplified for purposes of simulation. Suppose that $\varepsilon \ll 1$ in (3); the reset signals in (2) and (3) are either 0 or $\gg 1$; and net intrafield feedback is excitatory. Then eqns (1), (5), and (6) below approximate the main properties of ART 3 system dynamics.

Simplified ART 3 Equations

$$\frac{du_{ij}}{dt} = (z_{ij} - u_{ij}) - u_{ij}[\text{release rate}] \quad (1)$$

$$\begin{cases} \frac{dv_{ij}}{dt} = -v_{ij} + u_{ij}[\text{release rate}] & \text{if reset} = 0 \\ v_{ij}(t) = 0 & \text{if reset} \gg 1. \end{cases} \quad (5)$$

$$\begin{cases} x_j(t) = \left\{ \sum_i v_{ij} + [\text{intrafield feedback}] \right\} & \text{if reset} = 0 \\ 0 & \text{if reset} \gg 1. \end{cases} \quad (6)$$

9. ALTERNATIVE ART 3 SYSTEMS

In eqns (2) and (3), the reset signal acts in two ways, by inactivating bound transmitter and directly inhibiting the postsynaptic membrane. Alternatively, the reset signal may accomplish both these goals in a single process if all excitatory inputs in (3) are realized using chemical transmitters. Letting w_j denote the net excitatory transmitter reaching the j th target cell via intrafield feedback, an illustrative system of this type is given by eqns (1), (2), (7), and (8) below.

Presynaptic Transmitter

$$\frac{du_{ij}}{dt} = (z_{ij} - u_{ij}) - u_{ij}[\text{release rate}]. \quad (1)$$

Bound Transmitter

$$\frac{dv_{ij}}{dt} = -v_{ij} + u_{ij}[\text{release rate}] - v_{ij}[\text{reset signal}]. \quad (2)$$

$$\begin{aligned} \frac{dw_j}{dt} &= -w_j + [\text{intrafield feedback}] \\ &\quad - w_j[\text{reset signal}]. \end{aligned} \quad (7)$$

Postsynaptic Activation

$$\epsilon \frac{dx_j}{dt} = -x_j + (A - x_j) \left(\sum_i v_{ij} + w_i \right). \quad (8)$$

The reset signal now acts as a chemical modulator that inactivates the membrane channels at which transmitter is bound. It thus appears in eqns (2) and (7), but not in eqn (8) for postsynaptic activation.

When the reset signal can be only 0 or ≥ 1 , the simplified system in Section 8 approximates both versions of the ART 3 system. However, if the reset signal can vary continuously in size, eqns (2), (7), and (8) can preserve relative transmitter quantities from all input sources. Thus, this system is a better model for the intermediate cases than eqns (2) and (3).

An additional inhibitory term in the postsynaptic activation eqn (8) helps to suppress transmitter release, as illustrated in Section 25.

10. TRANSMITTER RELEASE RATE

To further specify the ART search model, we now characterize the transmitter release and inactivation rates in eqns (1) and (2). Then we trace the dynamics of the system at key time intervals during the presentation of a fixed input pattern (Figure 11). We first observe system dynamics during a brief time interval after the input turns on ($t = 0^+$), when the signal S_i first arrives at the synapse. We next consider the effect of subsequent internal feedback signals from within the target field, following contrast-enhancement of the inputs. We observe how the ART 3 model responds to a reset signal by implementing a rapid and enduring inhibition of erroneously selected pattern features. Then we analyze how the ART 3 model responds if the input pattern changes.

We will begin with the *ART Search Hypothesis 1*: Presynaptic transmitter u_{ij} is released at a rate jointly proportional to the presynaptic signal S_i and a function $f(x_j)$ of the postsynaptic activity. That is, in eqns (1), (2), and (5),

$$\text{release rate} = S_i f(x_j). \quad (9)$$

SYSTEM DYNAMICS




FIGURE 11. The system is designed to carry out necessary computations at critical junctures of the search process.

The function $f(x_j)$ in eqn (9) has the qualitative properties illustrated in Figure 12. In particular, $f(x_j)$ is assumed to have a positive value when x_j is at its 0 resting level, so that transmitter u_{ij} can be released when the signal S_i arrives at the synapse. If $f(0)$ were equal to 0, no excitatory signal could reach a postsynaptic node at rest, even if a large presynaptic signal S_i were sent to that node. The function $f(x_j)$ is also assumed to equal 0 when x_j is significantly hyperpolarized, but to rise steeply when x_j is near 0. In the simulations, $f(x_j)$ is linear above a small negative threshold.

The form factor $S_i f(x_j)$ is a familiar one in the neuroscience and neural network literatures. In particular, such a product is often used to model associative learning, where it links the rate of learning in the ij th pathway to the presynaptic signal S_i and the postsynaptic activity x_j . Associative learning occurs, however, on a time scale that is much slower than the time scale of transmitter release. On the fast time scale of transmitter release, the form factor $S_i f(x_j)$ may be compared to interactions between voltages and ions. In Figure 9, for example, note the dependence of the presynaptic signal on the Na^+ ion; the postsynaptic signal on the Ca^{2+} ion; and transmitter release on the *joint* fluxes of these two ions. The ART Search Hypothesis 1 thus formalizes a known type of synergetic relationship between presynaptic and postsynaptic processes in effecting transmitter release. Moreover, the rate of transmitter release is typically a function of the concentration of Ca^{2+} in the extracellular space, and this function has qualitative properties similar to the function $f(x_j)$ shown in Figure 12 (Kandel & Schwartz, 1981, p. 84; Kuffler, Nicholls, & Martin, 1984, p. 244).

11. SYSTEM DYNAMICS AT INPUT ONSET: AN APPROXIMATELY LINEAR FILTER

Some implications of the ART Search Hypothesis 1 will now be summarized. Assume that at time $t = 0$ transmitter u_{ij} has accumulated to its maximal level z_{ij} and that activity x_j and bound transmitter v_{ij} equal

ART SEARCH HYPOTHESIS 1




FIGURE 12. The ART Search Hypothesis 1 specifies the transmitter release rate.

0. Consider a time interval $t = 0^+$ immediately after a signal S_i arrives at the synapse. During this brief initial interval, the ART equations approximate the linear filter dynamics typical of many neural network models. In particular, eqns (2) and (9) imply that the amount of bound transmitter is determined by equation

$$\frac{dv_{ij}}{dt} = -v_{ij} + u_{ij}S_i f(x_j) - v_{ij}[\text{inactivation rate}]. \quad (10)$$

Thus, at times $t = 0^-$,

$$\frac{dv_{ij}}{dt} \approx z_{ij}S_i f(0) \quad (11)$$

and so

$$v_{ij}(t) \approx K(t)S_i z_{ij} \quad \text{for times } t = 0^-. \quad (12)$$

Because eqn (12) holds at all the synapses adjacent to cell j , eqn (6) implies that

$$x_j(t) \approx \sum_i K(t)S_i z_{ij} = K(t)\mathbf{S} \cdot \mathbf{z}_j \quad \text{for times } t = 0^+. \quad (13)$$

Here \mathbf{S} denotes the vector $(S_1 \dots S_n)$, \mathbf{z}_j denotes the vector $(z_{1j} \dots z_{nj})$, and $i = 1 \dots n$. Thus, in the initial moments after a signal arrives at the synapse, the small amplitude activity x_j at the postsynaptic cell grows in proportion to the dot product of the incoming signal vector \mathbf{S} times the adaptive weight vector \mathbf{z}_j .


12. SYSTEM DYNAMICS AFTER INTRAFIELD FEEDBACK: AMPLIFICATION OF TRANSMITTER RELEASE BY POSTSYNAPTIC POTENTIAL

In the next time interval, the intrafield feedback signal contrast-enhances the initial signal pattern (13) via eqn (6) and amplifies the total activity across field F_c in Figure 13a. Figure 13b shows typical contrast-enhanced activity profiles: partial compression of the initial signal pattern; or maximal compression, or choice, where only one postsynaptic node remains active due to the strong competition within the field F_c .

In all, the model behaves initially like a linear filter. The resulting pattern of activity across postsynaptic cells is contrast-enhanced, as required in the ART 2 model as well as in the many other neural network models that incorporate competitive learning (Grossberg, 1988). For many neural network systems, this combination of computational properties is all that is needed. These models implicitly assume that intracellular transmitter u_{ij} is always accumulated up to its target level z_{ij} and that postsynaptic activity x_j does not alter the rate of transmitter release:

$$u_{ij} \approx z_{ij} \quad \text{and} \quad v_{ij} \approx z_{ij}S_i. \quad (14)$$

$$(a) \quad x_j \approx K(t)\mathbf{S} \cdot \mathbf{z}_j \quad (t = 0^+)$$



$$(b) \quad \text{FEEDBACK CONTRAST-ENHANCES } x_j$$




FIGURE 13. (a) If transmitter is fully accumulated at $t = 0$, low-amplitude postsynaptic STM activity x_j is initially proportional to the dot product of the signal vector \mathbf{S} and the weight vector \mathbf{z}_j . Fields are labeled F_b and F_c for consistency with the ART 3 system in Figure 21. (b) Intrafield feedback rapidly contrast-enhances the initial STM activity pattern. Large-amplitude activity is then concentrated at one or more nodes.

If the linear filtering properties implied by (14) work well for many purposes, why complicate the system by adding additional hypotheses? Even a new hypothesis that makes a neural network more realistic physiologically needs to be justified functionally or it will obscure essential system dynamics. Why, then, add two additional nonlinearities to the portion of a neural network system responsible for transmitting signals from one location to another? The following discussion suggests how nonlinearities of synaptic transmission and neuromodulation can, when embedded in an ART circuit, help to correct coding errors by triggering a parallel search, allow the system to respond adaptively to reinforcement, and rapidly reset itself to changing input patterns.

In eqn (10), term

$$u_{ij}S_i f(x_j) \quad (15)$$

for the amount of transmitter released per unit time implies that the original incoming weighted signal $z_{ij}S_i$ is distorted both by depletion of the presynaptic transmitter u_{ij} and by the activity level x_j of the postsynaptic cell. If these two nonlinearities are significant, the net signal in the ij th pathway depends jointly on the maximal weighted signal $z_{ij}S_i$; the prior activity in the pathway, as reflected in the amount of depletion of the transmitter u_{ij} ; and the immediate context in which the signal is sent, as reflected in the

target cell activity x_j . In particular, once activity in a postsynaptic cell becomes large, this activity dominates the transmitter release rate, via the term $f(x_j)$ in (15). In other words, although linear filtering properties initially determine the small-amplitude activity pattern of the target field F_c , once intrafield feedback amplifies and contrast-enhances the postsynaptic activity x_j (Figure 13b) it plays a major role in determining the amount of released transmitter v_{ij} (Figure 14). In particular, the postsynaptic activity pattern across the field F_c that represents the recognition code (Figure 13b) is imparted to the pattern of released transmitter (Figure 14), which then also represents the recognition code, rather than the initial filtered pattern $S \cdot z_j$.

13. SYSTEM DYNAMICS DURING RESET: INACTIVATION OF BOUND TRANSMITTER CHANNELS

The dynamics of transmitter release implied by the ART Search Hypothesis 1 can be used to implement the reset process, by postulating the *ART Search Hypothesis 2*: The nonspecific reset signal quickly inactivates postsynaptic membrane channels at which transmitter is bound (Figure 15). The reset signal in eqns (5) and (6) may be interpreted as assignment of a large value to the inactivation rate in a manner analogous to the action of a neuromodulator (Figure 9). Inhibition of postsynaptic nodes breaks the strong intrafield feedback loops that implement ART 2 and ART 3 matching and contrast-enhancement (eqn (3) or (6)).

Let us now examine system dynamics following transmitter inactivation. The pattern of released transmitter can be viewed as a representation of the postsynaptic recognition code. The arrival of a reset signal implies that some part of the system has judged





FIGURE 14. The ART Search Hypothesis 1 implies that large amounts of transmitter (v_{ij}) are released only adjacent to postsynaptic nodes with large-amplitude activity (x_j). Competition within the postsynaptic field therefore transforms the initial low-amplitude distributed pattern of released and bound transmitter into a large-amplitude contrast-enhanced pattern.



$$\frac{dv_{ij}}{dt} = -v_{ij} + u_{ij} \left[\text{RELEASE RATE} \right] - v_{ij} \left[\text{INACTIVATION RATE} \right]$$




FIGURE 15. The ART Search Hypothesis 2 specifies a high rate of inactivation of bound transmitter following a reset signal. Postsynaptic action of the nonspecific reset signal is similar to that of a neuromodulator.

this code to be erroneous, according to some criterion. The ART Search Hypothesis 1 implies that the largest concentrations of bound extracellular transmitter are adjacent to the nodes which most actively represent this erroneous code. The ART Search Hypothesis 2 therefore implies that the reset process selectively removes transmitter from pathways leading to the erroneous representation.

After the reset wave has acted, the system is biased against activation of the same nodes, or features, in the next time interval: Whereas the transmitter signal pattern $S \cdot u_j$ originally sent to target nodes at times $t = 0^+$ was proportional to $S \cdot z_j$, as in eqn (12), the transmitter signal pattern $S \cdot u_j$ after the reset event




FIGURE 16. Following reset, the system is selectively biased against pathways that had previously released large quantities of transmitter. After a mismatch reset, therefore, the adaptive filter delivers a smaller signal to the previous category representation, the one that generated the reset signal.

is no longer proportional to $\mathbf{S} \cdot \mathbf{z}_j$. Instead, it is selectively biased against those features that were previously active (Figure 16). The new signal pattern $\mathbf{S} \cdot \mathbf{u}_j$ will lead to selection of another contrast-enhanced representation, which may or may not then be reset. This search process continues until an acceptable match is found, possibly through the selection of a previously inactive representation.

14. PARAMETRIC ROBUSTNESS OF THE SEARCH PROCESS

This search process is relatively easy to implement, requiring no new nodes or pathways beyond those already present in ART 2 modules. It is also robust, since it does not require tricky timing or calibration. How the process copes with a typical slow learning situation is illustrated in Figure 17. With slow learning, an input can select and begin to train a new category so that the adaptive weights correspond to a perfect pattern match during learning. However, the input may not be on long enough for the adaptive weights to become very large. That input may later activate a different category node whose weights are




FIGURE 17. An erroneous category representation with large weights (z_j) may become active before another representation that makes a good pattern match with the input but which has small weights. One or more mismatch reset events can decrease the functional value (u_j) of the larger weights, allowing the “correct” category to become active.

large but whose vector of adaptive weights forms a poorer match than the original, smaller weights.

Figure 17a shows such a typical filtered signal pattern $\mathbf{S} \cdot \mathbf{z}_j$. During the initial processing interval ($t = 0^+$) the transmitted signal $\mathbf{S} \cdot \mathbf{u}_j$ and the post-synaptic activity x_j are proportional to $\mathbf{S} \cdot \mathbf{z}_j$. Suppose that the weights z_{ij} in pathways leading to the J th node are large, but that the vector pattern \mathbf{z}_j is not an adequate match for the signal pattern \mathbf{S} according to the vigilance criterion. Also suppose that dynamics in the target field F_c lead to a “choice” following competitive contrast-enhancement (Figure 17b) and that the chosen node J represents a category. Large amounts of transmitter will thus be released from synapses adjacent to node J , but not from synapses adjacent to other nodes. The reset signal will then selectively inactivate transmitter at postsynaptic sites adjacent to the J th node. Following such a reset wave, the new signal pattern $\mathbf{S} \cdot \mathbf{u}_j$ will be biased against the J th node relative to the original pattern. However, it could happen that the time interval prior to the reset signal is so brief that only a small fraction of available transmitter is released. Then $\mathbf{S} \cdot \mathbf{u}_j$ could still be large relative to a “correct” $\mathbf{S} \cdot \mathbf{u}_j$ after reset occurs (Figure 17c). If this were to occur, the J th node would simply be chosen again, then reset again, leading to an accumulating bias against that choice in the next time interval. This process could continue until enough transmitter v_{ij} is inactivated to allow another node, with smaller weights z_{ij} but a better pattern match, to win the competition. Simulations of such a reset sequence are illustrated in Figures 23–26.

15. SUMMARY OF SYSTEM DYNAMICS DURING A MISMATCH-RESET CYCLE

Figure 18 summarizes system dynamics of the ART search model during a single input presentation. Initially, the transmitted signal pattern $\mathbf{S} \cdot \mathbf{u}_j$, as well

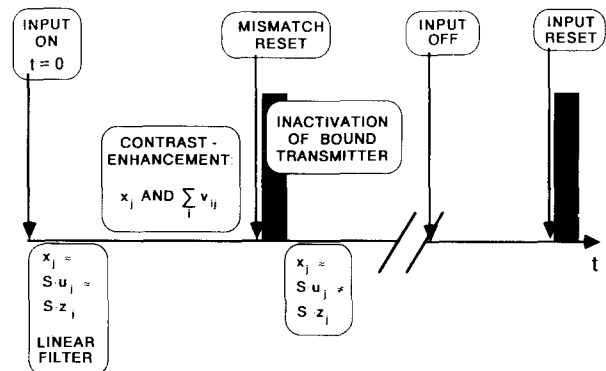


FIGURE 18. ART Search Hypotheses 1 and 2 implement computations to carry out search in an ART system. Input reset employs the same mechanisms as mismatch reset, initiating search when the input pattern changes significantly.

as the postsynaptic activity x_j , are proportional to the weighted signal pattern $\mathbf{S} \cdot \mathbf{z}_j$ of the linear filter. The postsynaptic activity pattern is then contrast-enhanced due to the internal competitive dynamics of the target field. The ART Search Hypothesis 1 implies that the transmitter release rate is greatly amplified in proportion to the level of postsynaptic activity. A subsequent reset signal selectively inactivates transmitter in those pathways that caused an error. Following the reset wave, the new signal $\mathbf{S} \cdot \mathbf{u}_j$ is no longer proportional to $\mathbf{S} \cdot \mathbf{z}_j$ but is, rather, biased against the previously active representation. A series of reset events ensue until an adequate match or a new category is found. Learning occurs on a time scale that is long relative to that of the search process.

16. AUTOMATIC STM RESET BY REAL-TIME INPUT SEQUENCES

The ART 3 architecture serves other functions as well as implementing the mismatch–reset–search cycle. In particular, it allows an ART system to dispense with additional processes to reset STM at onset or offset of an input pattern. The representation of input patterns as a sequence, $\mathbf{I}_1, \mathbf{I}_2, \mathbf{I}_3, \dots$, corresponds to the assumption that each input is constant for a fixed time interval. In practice, an input vector $\mathbf{I}(t)$ may vary continuously through time. The input need never be constant over an interval, and there may be no temporal marker to signal offset or onset of “an input pattern” per se. Furthermore, feedback loops within a field or between two fields can maintain large amplitude activity even when $\mathbf{I}(t) = 0$. Adaptive resonance develops only when activity patterns across fields are amplified by such feedback loops and remain stable for a sufficiently long time to enable adaptive weight changes to occur (Grossberg, 1976b, 1982a). In particular, no reset waves are triggered during a resonant event.

The ART reset system functionally defines the onset of a “new” input as a time when the orienting subsystem emits a reset wave. This occurs, for example, in the ART 2 module (Figure 3) when the angle between the vectors $\mathbf{I}(t)$ and $\mathbf{p}(t)$ becomes so large that the norm of $\mathbf{r}(t)$ falls below the vigilance level $\rho(t)$, thereby triggering a search for a new category representation. This is called an *input reset* event, to distinguish it from a *mismatch reset* event, which occurs while the bottom-up input remains nearly constant over a time interval but mismatches the top-down expectation that it has elicited from level F_2 (Figure 2).

This property obviates the need to mechanistically define the processing of input onset or offset. The ART Search Hypothesis 3, which postulates restoration of a resting state between successive inputs (Carpenter & Grossberg, 1989), is thus not needed.

Presynaptic transmitter may not be fully accumulated following an input reset event, just as it is not fully accumulated following a mismatch reset event. For both types of reset, the orienting subsystem judges the active code to be incorrect, at the present level of vigilance, and the system continues to search until it finds an acceptable representation.

17. REINFORCEMENT FEEDBACK

The mechanisms described thus far for STM reset are part of the recognition learning circuit of ART 3. Recognition learning is, however, only one of several processes whereby an intelligent system can learn a correct solution to a problem. We have called recognition, reinforcement, and recall the “3 R’s” of neural network learning (Carpenter & Grossberg, 1988).

Reinforcement, notably reward and punishment, provides additional information in the form of environmental feedback based on the success or failure of actions triggered by a recognition event. Reward and punishment calibrate whether the action has or has not satisfied internal needs, which in the biological case include hunger, thirst, sex, and pain reduction, but may in machine applications include a wide variety of internal cost functions.

Reinforcement can shift attention to focus upon those recognition codes whose activation promises to satisfy internal needs based on past experience. A model to describe this aspect of reinforcement learning was described in Grossberg (1982a, 1982b, 1984; reprinted in Grossberg, 1987a) and was supported by computer simulations in Grossberg and Levine (1987; reprinted in Grossberg, 1988). An attention shift due to reinforcement can also alter the structure and learning of recognition codes by amplifying (or suppressing) the STM activations, and hence the adjacent adaptive weights, of feature detectors that are active during positive (or negative) reinforcement.

A reset wave may also be used to modify the pattern of STM activation in response to reinforcement. For example, both green and yellow bananas may be recognized as part of a single recognition category until reinforcement signals, contingent upon eating the bananas, differentiate them into separate categories. Within ART 3, such a reinforcement signal can alter the course of recognition learning by causing a reset event. The reset event may override a bias in either the learned path weights (Figure 19) or in the input strengths (Figure 20) that could otherwise prevent a correct classification from being learned. For example, both green and yellow bananas may initially be coded in the same recognition category because features that code object shape (e.g., pathway A in Figures 19 and 20) prevent features

REINFORCEMENT RESET CAN OVERRIDE PATH WEIGHT BIAS




FIGURE 19. A system whose weights are biased toward feature A over feature B over feature C. (a) Competition amplifies the weight bias in STM, leading to enhanced transmitter release of the selected feature A. (b) Transmitter inactivation following reinforcement reset allows feature B to become active in STM.

REINFORCEMENT RESET CAN OVERRIDE INPUT BIAS




FIGURE 20. A system whose input signals are biased towards A over B over C. (a) Competition amplifies the input bias in STM, leading to enhanced transmitter release of the selected feature A. (b) Transmitter inactivation following reinforcement reset allows feature B to become active in STM.

that code object color (e.g., pathway B in Figures 19 and 20) from being processed in STM. Reset waves triggered by reinforcement feedback can progressively weaken the STM activities of these shape features until both shape and color features can simultaneously be processed, and thereby incorporated into different recognition codes for green bananas and yellow bananas.

In technical applications, such a reset wave can be implemented as a direct signal from an internal representation of a punishing event. The effect of the reset wave is to modify the spatial pattern of STM activation whose read-out into an overt action led to the punishing event. The adaptive weights, or LTM traces, that input to these STM activations are then indirectly altered by an amount that reflects the new STM activation pattern. Such a reinforcement scheme differs from the competitive learning scheme described by Kohonen (1984, p. 200), in which reinforcement acts directly, and by an equal amount, on all adaptive weights that lead to an incorrect classification.

Reinforcement may also act by changing the level of vigilance (Carpenter & Grossberg, 1987a, 1987b). For example, if a punishing event increases the vigilance parameter, then mismatches that were tolerated before will lead to a search for another recognition code. Such a code can help to distinguish pattern differences that were previously considered too small to be significant. Such a role for reinforcement is illustrated by computer simulations in Figures 25–28.

All three types of reaction to reinforcement feedback may be useful in applications. The change in vigilance alters the overall sensitivity of the system to pattern differences. The shift in attention and the reset of active features can help to overcome prior coding biases that may be maladaptive in novel contexts.

18. NOTATION FOR HIERARCHIES

Table 1 and Figure 21 illustrate notation suitable for an ART hierarchy with any number of fields F_a, F_b, F_c, \dots . This notation can also be adapted for related neural networks and algorithmic computer simulation.

Each STM variable is indexed by its field, layer, and node number. Within a layer, x denotes the activity of a node receiving inputs from other layers, while y denotes the (normalized) activity of a node that sends signals to other layers. For example, x_i^{a2} denotes activity at the i th input node in layer 2 of field $F_a (i = 1 \dots n_a)$ and y_i^{a2} denotes activity of the corresponding output node. Parameters are also indexed by field (p_1^a, p_2^a, \dots), as are signal functions (g^a). Variable r_i^b denotes activity of the i th reset node

TABLE 1
Notation for ART 3 Hierarchy

$F_{\text{field}} = F_a$	STM field <i>a</i>
$i = i_a = 1 \dots n_a$	node index, field <i>a</i>
$L = 1, 2, 3$	index, 3 layers of an STM field
x_i^{aL}	STM activity, input node <i>i</i> , layer <i>L</i> , field <i>a</i>
y_i^{aL}	STM activity, output node <i>i</i> , layer <i>L</i> , field <i>a</i>
$g^a(y_i^{aL}) \equiv S_i^{aL}$	signal function, field <i>a</i>
p_k^a	parameter, field <i>a</i> , $k = 1, 2, \dots$
r_i^b	STM activity, reset node <i>i</i> , field <i>b</i>
ρ^b	vigilance parameter, field <i>b</i>
z_{ij}^{bc}	LTM trace, pathway from node <i>i</i> (field <i>b</i>) to node <i>j</i> (field <i>c</i>)
u_{ij}^{bc}	intracellular transmitter, pathway from node <i>i</i> (field <i>b</i>) to node <i>j</i> (field <i>c</i>)
v_{ij}^{bc}	released transmitter, pathway from node <i>i</i> (field <i>b</i>) to node <i>j</i> (field <i>c</i>)

of field F_b , and ρ^b is the corresponding vigilance parameter.

Variable z denotes an adaptive weight or LTM trace. For example, z_{ij}^{bc} is the weight in the bottom-up pathway from the i th node of field F_b to the j th node of field F_c . Variables u_{ij}^{bc} and v_{ij}^{bc} denote the

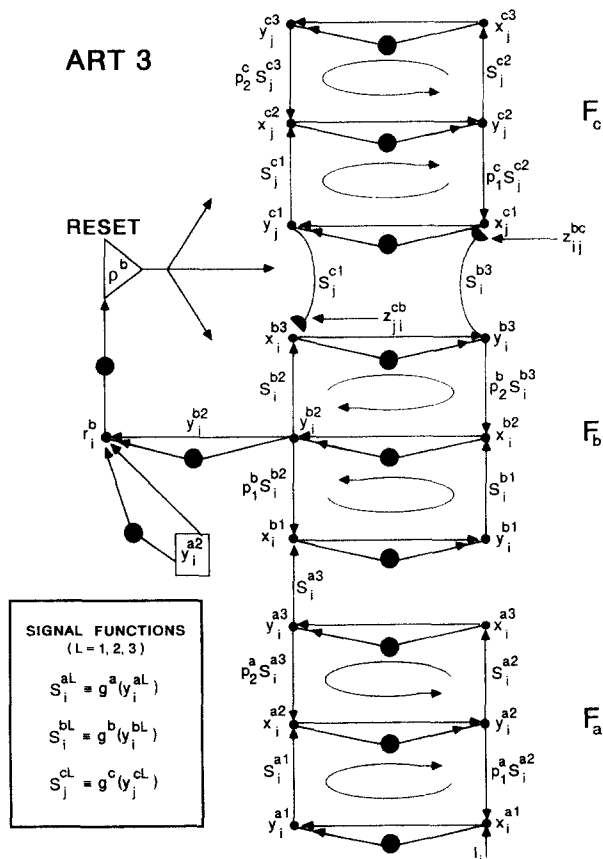


FIGURE 21. ART 3 simulation neural network. Indices $i = 1 \dots n_a = n_b$ and $j = 1 \dots n_c$. The reset signal acts at all layers 1 and 3 (Section 8).

corresponding presynaptic and bound transmitter quantities, respectively. Variables for the top-down pathways are z_{ij}^{cb} , u_{ij}^{cb} , and v_{ij}^{cb} .

Complete simulation equations are specified in Section 26.

19. TRADE-OFF BETWEEN WEIGHT SIZE AND PATTERN MATCH

The simulations in Sections 20–24 illustrate the dynamics of search in the ART 3 system shown in Figure 21. The simulation time scale is assumed to be short relative to the time scale of learning, so all adaptive weights z_{ij}^{bc} and z_{ij}^{cb} are held constant. The weights are chosen, however, to illustrate a problem that can arise with slow learning or in any other situation in which weight vectors are not normalized at all times. Namely, a category whose weight vector only partially matches the input vector may become active because its weights are large. This can prevent initial selection of another category whose weight vector matches the input vector but whose weight magnitudes are small due, say, to a brief prior learning interval.

The search process allows the ART 3 system to reject an initial selection with large weights and partial pattern match, and then to activate a category with smaller weights and a better pattern match. As in ART 2, when weights are very small (nodes $j = 6, 7, \dots$, Figure 22) the ART system tolerates poor pattern matches to allow new categories to become established. During learning, the weights can become larger. The larger the weights, the more sensitive the ART system is to pattern mismatch (Carpenter & Grossberg, 1987b).

Figure 22 illustrates the trade-off between weight size and pattern match in the system used in the simulations. In Figures 22a and 22b, vector \mathbf{S} illustrates the STM pattern stored in F_a and sent from F_a to F_b when an input vector \mathbf{I} is held constant. The S_i values were obtained by presenting to F_a an input function \mathbf{I} with I_i a linearly decreasing function of i . Vector \mathbf{S} is also stored in F_b , as long as F_c remains inactive. Initially, \mathbf{S} is the signal vector in the bottom-up pathways from F_b to F_c . In Figure 22a, $S_1 > S_2 > \dots > S_5$; for $i = 6, 7 \dots 15 (= n_a = n_b)$, S_i is small. Each vector $\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3$, and \mathbf{z}_4 , plotted in columns within the square region of Figure 22a, partially matches the signal vector \mathbf{S} . These weights are significantly larger than the weights of vector \mathbf{z}_5 . However, \mathbf{z}_5 is a perfect match to \mathbf{S} in the sense that the angle between the two vectors is 0:

$$\cos(\mathbf{S}, \mathbf{z}_5) = 1. \tag{16}$$

The relationship

$$\mathbf{S} \cdot \mathbf{z}_i = \|\mathbf{S}\| \|\mathbf{z}_i\| \cos(\mathbf{S}, \mathbf{z}_i) \tag{17}$$

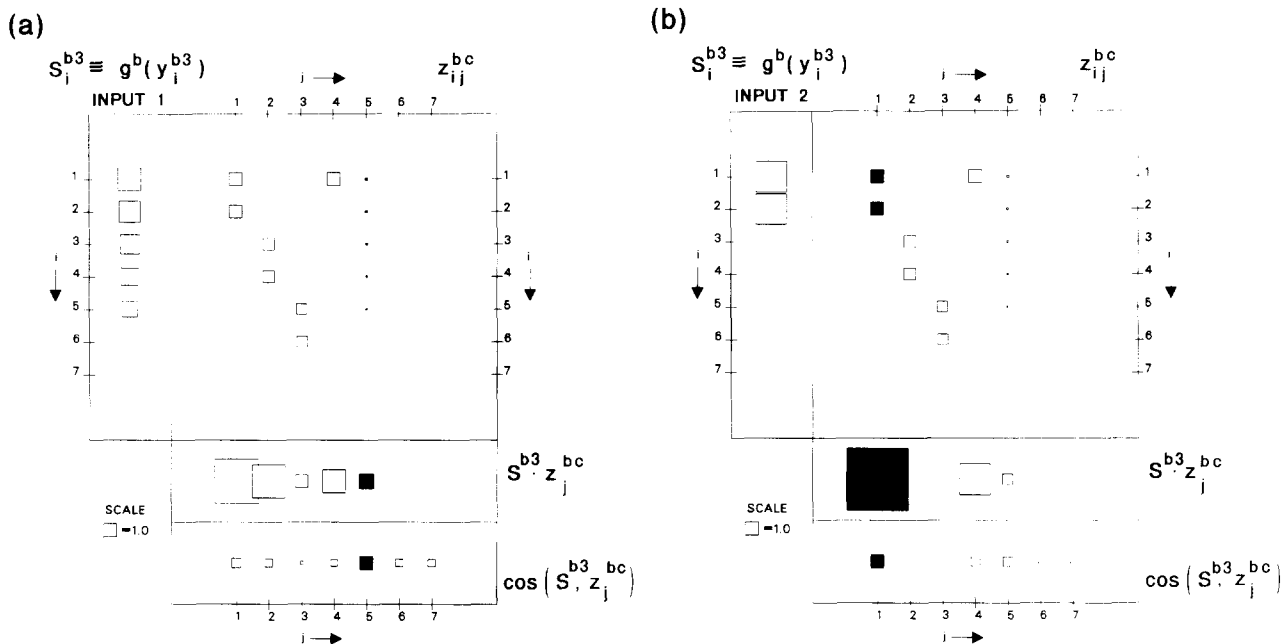


FIGURE 22. The length of a side of the square centered at position i or j or (i,j) gives the value of a variable with the corresponding index. Shown are quantities S_i^{b3} , z_j^{bc} , $S_i^{b3} \cdot z_j^{bc}$, and $\cos(S_i^{b3}, z_j^{bc})$. (a) Vector S^{b3} is the signal response to Input 1 in the simulations. Vector z_j^{bc} (filled squares) is parallel to S^{b3} , but $|z_j^{bc}|$ is small. Thus $S^{b3} \cdot z_j^{bc}$ is smaller than $S^{b3} \cdot z_j^{bc}$ for $j = 1, 2, \text{ and } 4$, despite the fact that $\cos(S^{b3}, z_j^{bc})$ is maximal. (b) Vector S^{b3} is the signal response to Input 2 in the simulations. Vector z_j^{bc} (filled squares) is parallel to S^{b3} .

implies a trade-off between weight size, as measured by the length $\|z_j\|$ of z_j , and pattern match, as measured by the angle between S and z_j . If the initial signal from F_b to F_c is proportional to $S \cdot z_j$, as in (13), then the matched node ($j = 5$) may receive a net signal that is smaller than signals to other nodes. In fact, in Figure 22a,

$$S \cdot z_1 > S \cdot z_2 > S \cdot z_4 > S \cdot z_5 > \dots \quad (18)$$

Figure 22b shows a signal vector S that is parallel to the weight vector z_1 .

20. ART 3 SIMULATIONS: MISMATCH RESET AND INPUT RESET OF STM CHOICES

The computer simulations summarized in Figures 23–26 use the inputs described in Figure 22 to illustrate the search process in an ART 3 system. In these simulations, the F_c competition parameters were chosen to make a choice; hence, only the node receiving the largest filtered input from F_b is stored in STM. The signal function of F_c caused the STM field to make a choice. In Figure 27, a different signal function at F_c , similar to the one used in F_a and F_b , illustrates how the search process reorganizes a distributed recognition code. The simulations show how, with high vigilance, the ART search process rapidly causes a series of mismatch resets that alter the transmitter vectors u_1, u_2, \dots until $S \cdot u_5$ becomes maximal. Once node $j = 5$ becomes active in STM it

amplifies transmitter release. Since the pattern match is perfect, no further reset occurs while Input 1 (Figure 22a) remains on. Input reset is illustrated following an abrupt or gradual switch to Input 2 (Figure 22b).

Each simulation figure illustrates three system variables as they evolve through time. The time axis (t) runs from the top to the bottom of the square. A vector pattern, indexed by i or j , is plotted horizontally at each fixed time. Within each square, the value of a variable at each time is represented by the length of a side of a square centered at that point. In each figure, part (a) plots y_i^{b1} , the normalized STM variables at layer 1 of field F_c . Part (b) plots $\sum_i v_{ij}^{bc}$, the total amount of transmitter released, bottom-up, in paths from all F_b nodes to the j th F_c node. Part (c) plots $\sum_i v_{ji}^{cb}$, the total amount of transmitter released, top-down, in paths from all F_c nodes to the i th F_b node. The ART Search Hypothesis 1 implies that the net bottom-up transmitter pattern in part (b) reflects the STM pattern of F_c in part (a); and that the net top-down transmitter pattern in part (c) reflects the STM pattern of F_b .

In Figure 23, the vigilance parameter is high and fixed at the value

$$\rho = .98. \quad (19)$$

For $0 \leq t < .8$, the input (Figure 22a) is constant. The high vigilance level induces a sequence of mismatch resets, alternating among the category nodes $j = 1, 2, \text{ and } 4$ (Figure 23a), each of which receives

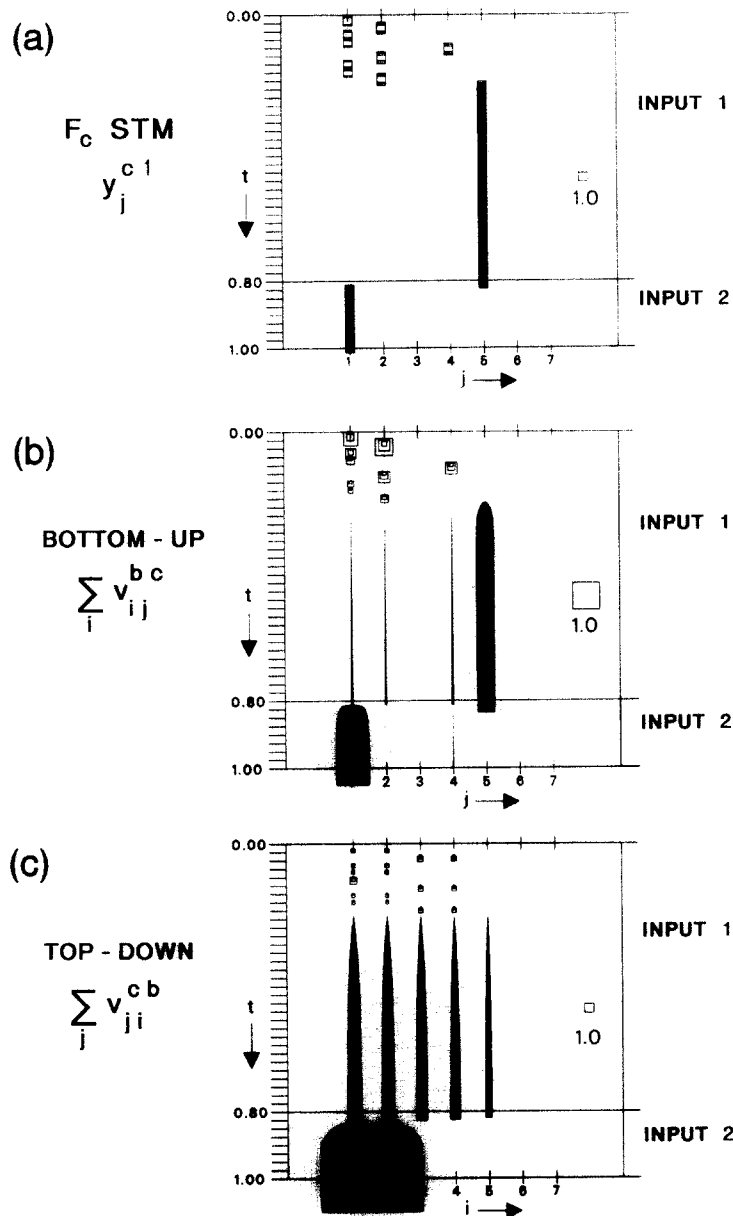


FIGURE 23. ART 3 simulation with $\rho = .98$. A series of 9 mismatch resets lead to activation of the matched category ($j = 5$) at $t = .215$. Input 1 switches to Input 2 at $t = .8$, causing an input reset and activation of a new category representation ($j = 1$).

an initial input larger than the input to node $j = 5$ (Figure 22a). At $t = .215$, the F_c node $j = 5$ is selected by the search process (Figure 23a). It remains active until $t = .8$. Then, the input from F_a is changed to a new pattern (Figure 22b). The mismatch between the new STM pattern at F_a and the old reverberating STM pattern at F_b leads to an input reset (Figures 18 and 23). The ART Search Hypothesis 2 implies that bound transmitter is inactivated and the STM feedback loops in F_b and F_c are thereby inhibited. The new input pattern immediately activates its category node $j = 1$, despite some previous depletion at that node (Figure 23a).

Large quantities of transmitter are released and bound only after STM resonance is established. In Figure 23b, large quantities of bottom-up transmitter

are released at the F_c node $j = 5$ in the time interval $.215 < t < .8$, and at node $j = 1$ in the time interval $.8 < t < 1$. In Figure 23c, the pattern of top-down bound transmitter reflects the resonating matched STM pattern at F_b due to Input 1 at times $.215 < t < .8$ and due to Input 2 at times $.8 < t < 1$.

21. SEARCH TIME INVARIANCE AT DIFFERENT VIGILANCE VALUES

Figure 24 shows the dynamics of the same system as in Figure 23 but at the lower vigilance value

$$\rho = .94. \quad (20)$$

The F_c node $j = 5$ becomes active slightly sooner ($t = .19$, Figure 24a) than it does in Figure 23a,

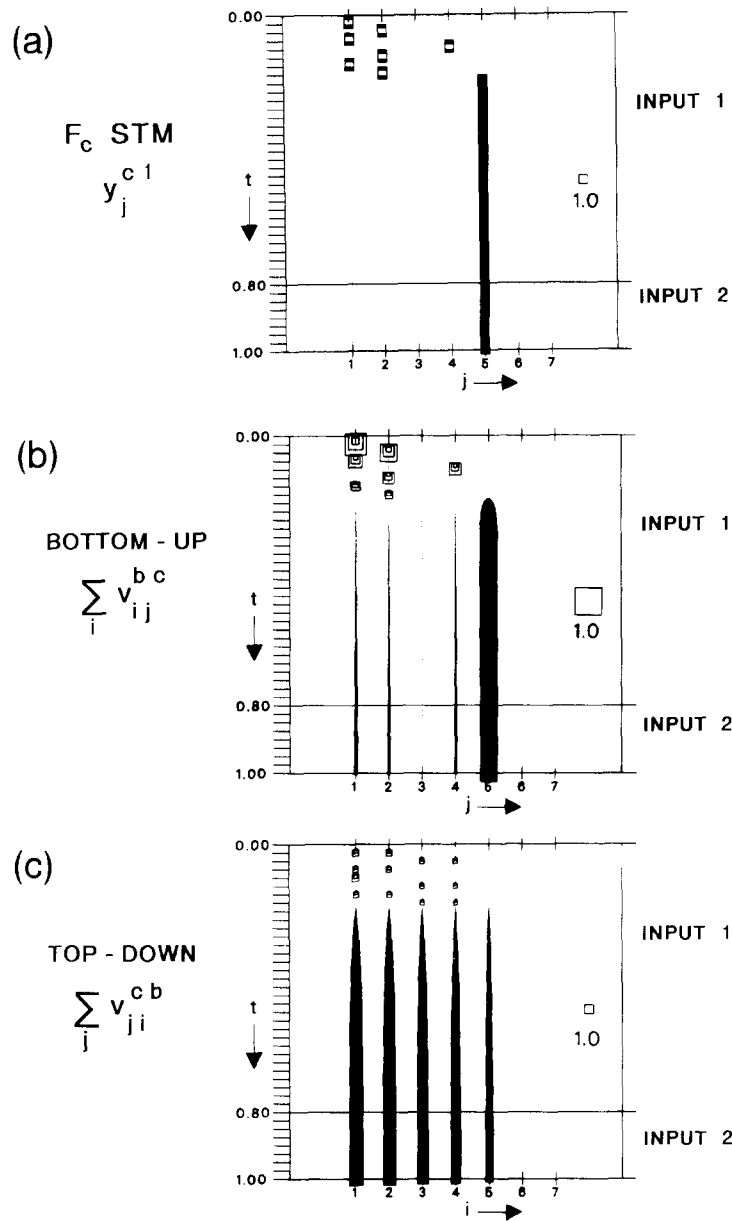


FIGURE 24. ART 3 simulation with $\rho = .94$. A series of seven mismatch resets lead to activation of the matched category ($j = 5$) at $t = .19$ Input 1 switches to Input 2 at $t = .8$, but no input reset occurs, and node $j = 5$ remains active, due to the lower vigilance level than in Figure 23.

where $\rho = .98$. At a lower vigilance, more transmitter needs to be released before the system reacts to a mismatch so that each "erroneous" category node is active for a longer time interval than at higher vigilance. When $\rho = .98$ (Figure 23b), node $j = 1$ is searched five times. When $\rho = .94$ (Figure 24b), node $j = 1$ is searched only three times, but more transmitter is released during each activation/reset cycle than at comparable points in Figure 23b. Inactivation of this extra released transmitter approximately balances the longer times to reset. Hence, the total search time remains approximately constant over a wide range of vigilance parameters. In the present instance, the nonlinearities of transmitter release terminate the search slightly sooner at lower vigilance.

Figure 24a illustrates another effect of lower vigilance: the system's ability to tolerate larger mismatches without causing a reset. When the input changes at $t = .8$, the mismatch between the input pattern at F_a and the resonating pattern at F_b is not great enough to cause an input reset. Despite bottom-up input only to nodes $i = 1, 2$, the strong resonating pattern at nodes $i = 1, \dots, 5$ maintains itself in STM at F_b (Figure 24c).

22. REINFORCEMENT RESET

In Figure 25, vigilance is initially set at value

$$\rho = .9, \tag{21}$$

in the time interval $0 < t < .1$. At this low vigilance

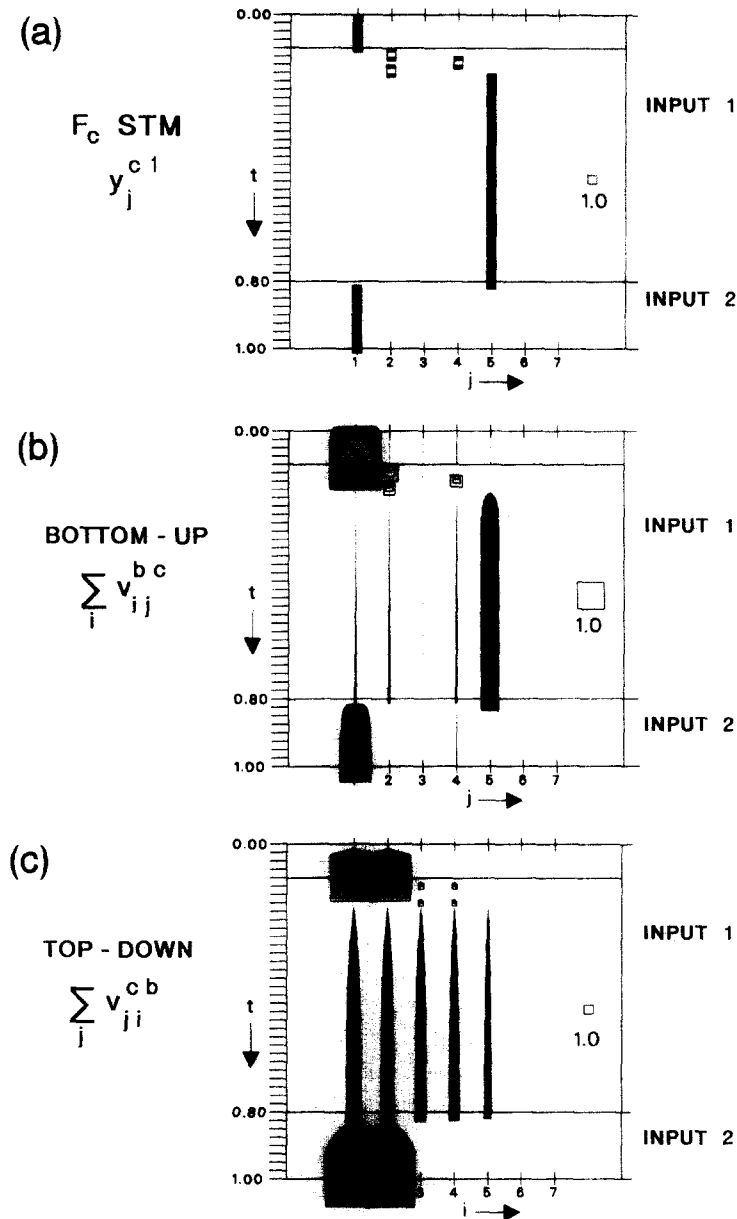


FIGURE 25. ART 3 simulation with $\rho = .9$ ($0 \leq t < .1$) and $\rho = .98$ ($.1 < t \leq 1$). At low vigilance, activation of node $j = 1$ leads to resonance. When vigilance is suddenly increased due, say, to reinforcement feedback, a series of 4 mismatch resets lead to activation of the matched category ($j = 5$) at $t = .19$. As in Figure 23, switching to input 2 at $t = .8$ causes an input reset and activation of a new category representation ($j = 1$).

level, the STM pattern of F_b does not experience a mismatch reset series. Node $j = 1$ is chosen and resonance immediately ensues (Figure 25a), as is also reflected in the amplification of transmitter release (Figure 25b). The simulation illustrates a case where this choice of category leads to external consequences, including reinforcement (Section 17), that feed back to the ART 3 module. This reinforcement teaching signal is assumed to cause vigilance to increase to the value

$$\rho = .98 \tag{22}$$

for times $t \geq .1$. This change triggers a search that ends at node $j = 5$, at time $t = .19$. Note that, as in Figure 24, enhanced depletion of transmitter at

$j = 1$ shortens the total search time. In Figure 23, where ρ also equals .98, the search interval has length .215; in Figure 25, the search interval has length .09, and the system never again activates node $j = 1$ during search.

23. INPUT HYSTERESIS SIMULATION

The simulation illustrated in Figure 26 is nearly the same as in Figure 25, with $\rho = .9$ for $0 \leq t < .1$ and $\rho = .98$ for $t > .1$. However, at $t = .8$, Input 1 starts to be slowly deformed into Input 2, rather than being suddenly switched, as in Figure 25. The $F_a \rightarrow F_b$ input vector becomes a convex combination of Input 1 and Input 2 that starts as Input 1 ($t \leq .8$) and is linearly

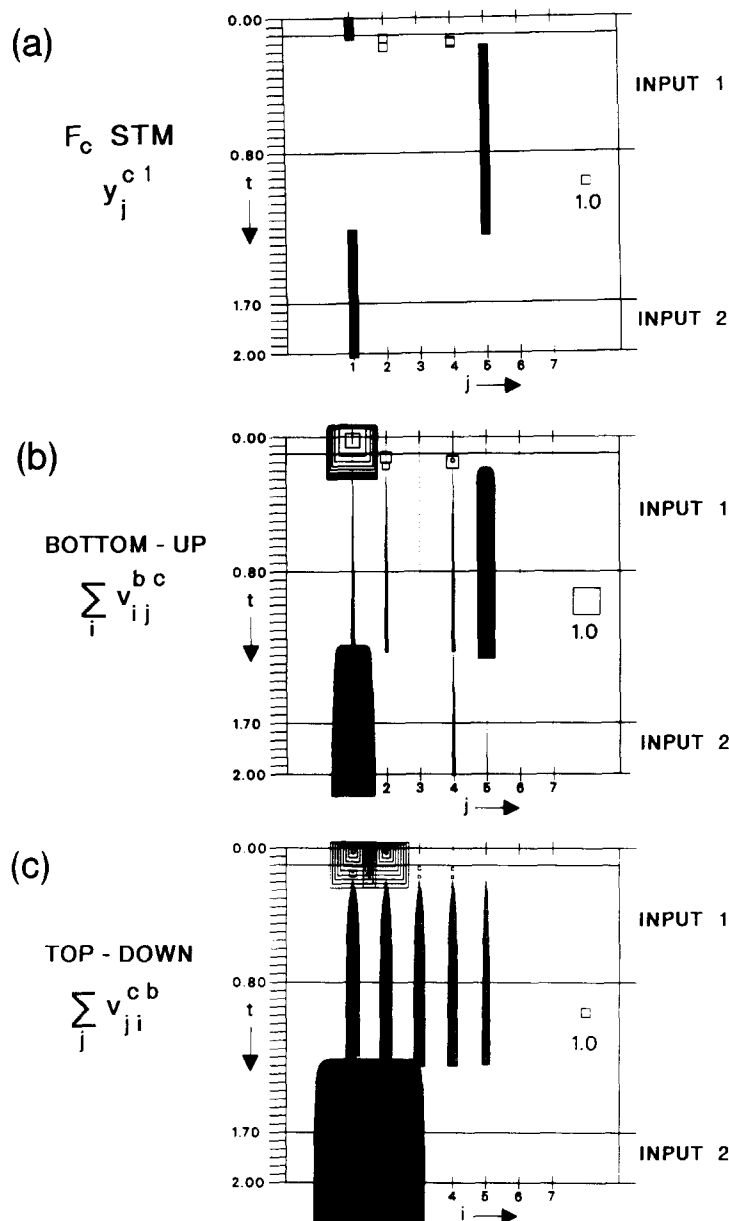


FIGURE 26. ART 3 simulation with $\rho = .9$ ($0 \leq t < .1$) and $\rho = .98$ ($1 < t \leq 2$). Input 1 is presented for $0 \leq t < .8$. For $.8 < t < 1.7$ the input to F_b is a convex combination of Input 1 and Input 2. Then Input 2 is presented for $1.7 < t \leq 2$. At $t = 1.28$ an input reset causes the STM choice to switch from node $j = 5$, which matches Input 1, to node $j = 1$, which matches Input 2.

shifted to Input 2 ($t \geq 1.7$). Despite the gradually shifting input, node $j = 5$ remains active until $t = 1.28$. Then an input reset immediately leads to activation of node $j = 1$, whose weight vector matches Input 2. Competition in the category representation field F_c causes a history-dependent choice of one category or the other, not a convex combination of the two.

24. DISTRIBUTED CODE SIMULATION

Issues of learning and code interpretation are subtle and complex when a code is distributed. However, the ART 3 search mechanism translates immediately

into this context. The simulation in Figure 27 illustrates how search operates on a distributed code. The only difference between the ART 3 system used for these simulations and the one used for Figures 23–26 is in the signal function at F_c . In Figures 23–26, a choice is always made at field F_c . The signal function for Figure 26 is, like that at F_a and F_b , piecewise linear: 0 below a threshold, linear above. With its fairly high threshold, this signal function compresses the input pattern; but the compression is not so extreme so as to lead inevitably to choice in STM.

Distributed code STM activity is shown in Figure 27a. At a given time more than one active node may represent a category ($2.6 < t < 7$), or one node may be chosen ($7.7 < t \leq 9$).

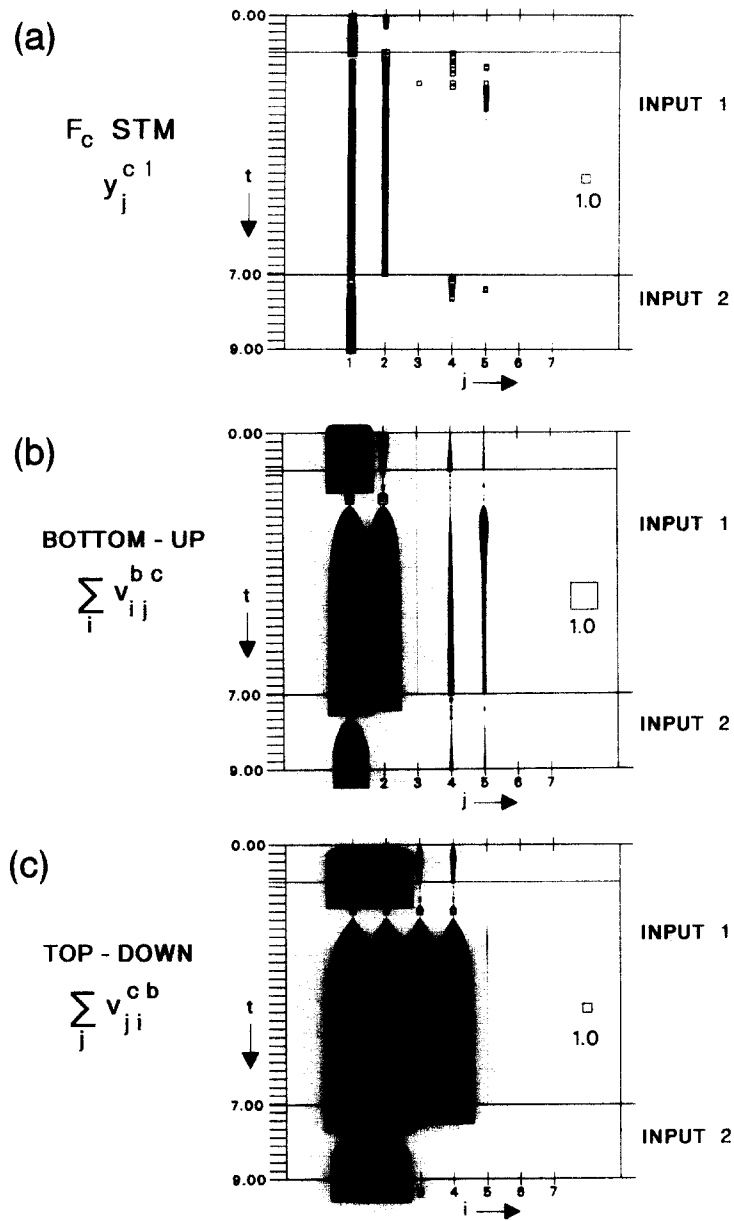


FIGURE 27. ART 3 simulation of a distributed code. Parameter $\rho = .9$ for $0 \leq t < 1$, and $\rho = .98$ for $1 < t \leq 9$. Input 1 is presented for $0 \leq t < 7$ and Input 2 is presented for $7 < t \leq 9$. At resonance, the single node $j = 1$ is active for $t < 1$ and $t > 7.7$. For $2.6 < t < 7$, simultaneous activity at two nodes ($j = 1$ and $j = 2$) represents the category. Top-down weights z_j^c are large for $j = 1$ and $i = 1$ and 2 ; and for $j = 2$ and $i = 3$ and 4 . Together the top-down signals (c) match enough of the bottom-up input pattern to satisfy the vigilance criterion.

25. ALTERNATIVE ART 3 MODEL SIMULATION

ART 3 systems satisfy the small number of design constraints described above. In addition, ART 3 satisfies the ART 2 stability constraints (Carpenter & Grossberg, 1987b). For example, top-down signals need to be an order of magnitude larger than bottom-up signals, all other things being equal, as illustrated below by (24) and parameters p_1 and p_2 in Table 4 and eqns (31) and (34). At least some of the STM fields need to be competitive networks. However, many versions of the ART systems exist within these

boundaries. A simulation of one such system is illustrated in Figure 28, which duplicates the conditions on ρ and input patterns of Figure 25. However, the system that generated Figure 28 uses a different version of the ART 3 STM field F_c than the one described in Section 26. In particular, in the STM equation (3), $B > 0$. STM nodes can thus be hyperpolarized, so that $x_j < 0$, by intrafield inhibitory inputs. The transmitter release function $f(x_j)$ (eqn (9)) equals 0 when x_j is sufficiently hyperpolarized. The system of Figure 28 thus has the property that transmitter release can be terminated at nodes that become inactive during the STM competition. Since

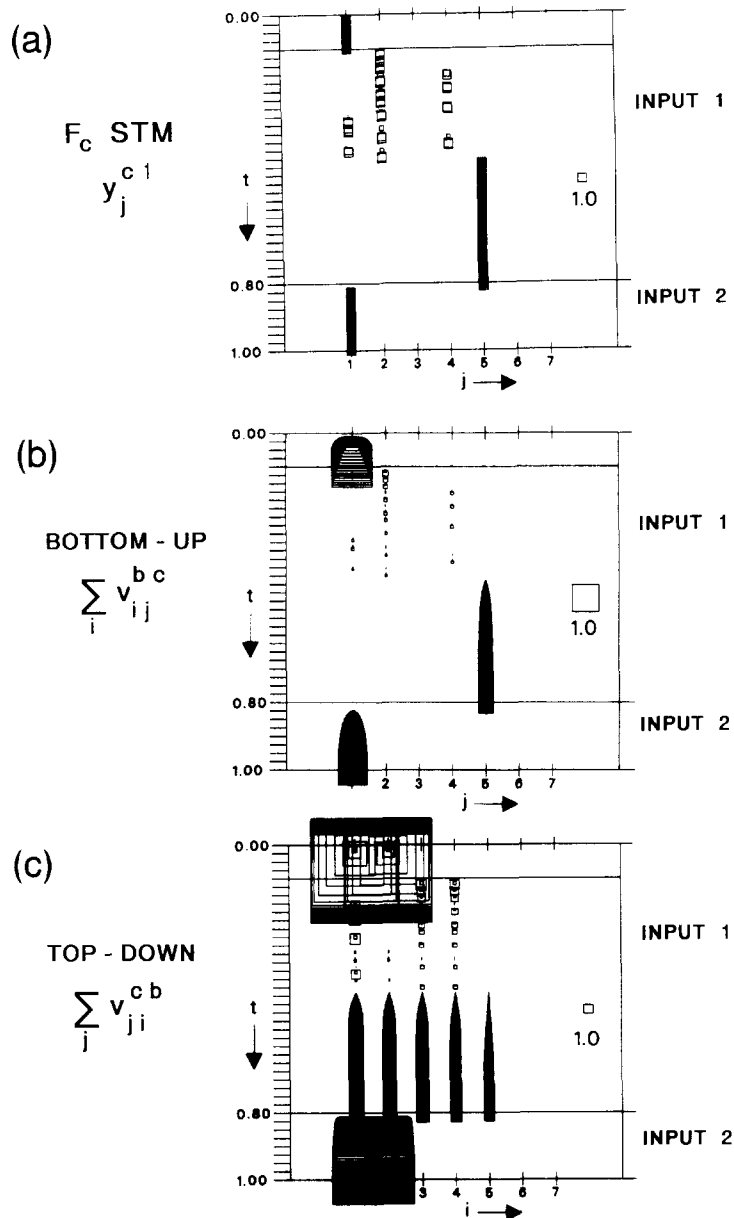


FIGURE 28. An alternative ART 3 model, which allows hyperpolarization of x_j^{c1} , gives results similar to those illustrated in Figure 25. As in Figure 25, $\rho = .9$ for $0 \leq t < .1$, $\rho = .98$ for $t > .1$, and Input 1 switches to Input 2 at $t = .8$. Category node $j = 5$ becomes active at $t = .44$, but immediately switches to node $j = 5$ at $t = .8$, when Input 2 is presented.

$f(0)$ needs to be positive to allow transmitter release to begin (Figure 12), low-level transmitter release by nodes without significant STM activity is unavoidable if nodes cannot be hyperpolarized. Figure 28 shows that a competitive STM field with hyperpolarization gives search and resonance results similar to those of the other simulations.

Similarly, considerable variations in parameters also give similar results.

26. SIMULATION EQUATIONS

Simulation equations are described in an algorithmic form to indicate the steps followed in the computer program that generated Figures 23–27.

Time Scale

The simulation time scale is fixed by setting the rate of transmitter accumulation equal to 1. The intrafield STM rate is assumed to be significantly faster and the LTM rate significantly slower. Accordingly, STM equations are iterated several times each time step and LTM weights are held constant. The simulation time step is

$$\Delta t = .005. \tag{23}$$

Integration Method

Transmitter variables u and v are integrated by first-order approximation (Euler's method). The IMSL

Gear package gives essentially identical solutions but requires more computer time.

LTM Weights

The bottom-up LTM weights z_{ij}^{bc} illustrated in Figure 22 are specified in Table 2. At "uncommitted" nodes ($j \geq 6$) $z_{ij}^{bc} \equiv 0.001$. Top-down LTM weights z_{ji}^{bc} are constant multiples of corresponding z_{ij}^{bc} weights:

$$z_{ji}^{cb} = 10 \cdot z_{ij}^{bc}. \tag{24}$$

This choice of LTM weights approximates a typical state of an ART system undergoing slow learning. Weights do not necessarily reach equilibrium on each presentation, but while the J th F_c node is active,

$$z_{ji}^{cb} \longrightarrow x_i^{b3} \tag{25}$$

and

$$z_{ij}^{bc} \longrightarrow S_i^{b3}. \tag{26}$$

Given the parameters specified below, as STM and LTM variables approach equilibrium,

$$x_i^{b3} \cong 10 \cdot S_i^{b3}. \tag{27}$$

Equations (25)–(27) imply that eqn (24) is a good approximation of a typical weight distribution.

Initial Values

Initially,

$$u_{ij}^{bc}(0) = z_{ij}^{bc} \tag{28}$$

and

$$u_{ji}^{cb}(0) = z_{ji}^{cb}. \tag{29}$$

All other initial values are 0.

Input Values

The F_b input values (S_i^{b3}) are specified in Table 3. All simulations start with Input 1. Several of the simulations switch to Input 2 either with a jump or

TABLE 3
 $F_a \rightarrow F_b$ Input Values (S_i^{b3})

i	Input 1	Input 2
1	1.76	2.36
2	1.62	2.36
3	1.48	0.0
4	1.34	0.0
5	1.20	0.0
6	0.0	0.0
7	0.0	0.0
...

gradually. Input 1 values are obtained by presenting a linear, decreasing function I_i to F_a . Input 2 values are obtained by setting $I_1 = I_2 = 1$ and $I_i = 0$ ($i \geq 3$).

Implicit in this formulation is the assumption that a changing input vector \mathbf{I} can register itself at F_a . This requires that STM at F_a be frequently "reset." Otherwise, new values of I_i may go unnoticed, due to strong feedback within F_a . Feedback within F_b allows the STM to maintain resonance even with fluctuating amplitudes at F_a .

STM Equations

Except during reset, equations used to generate the STM values for Figures 23–27 are similar to the ART 2 equations (Carpenter & Grossberg, 1987b). Dynamics of the fields F_a , F_b , and F_c are homologous, as shown in Figure 21. Steady-state variables for the field F_b , when the reset signal equals 0, are given by eqns (31)–(36). Similar equations hold for fields F_a and F_c .

Layer 1, input variable

$$\epsilon \frac{dx_i^{b1}}{dt} = -x_i^{b1} + S_i^{b3} + p_1^b S_i^{b2}. \tag{30}$$

In steady state,

$$x_i^{b1} \cong S_i^{b3} + p_1^b S_i^{b2}. \tag{31}$$

Table 4 specifies parameter p_1^b , p_2^b , . . . values and the signal function

$$g^b(y_i^{bL}) \cong S_i^{bL} \tag{32}$$

for layers $L = 1, 2, 3$. Equation (31) is similar to the simplified STM eqn (6), with x_i^{b1} equal to the sum of an interfield input (S_i^{b3}) and an intrafield input ($p_1^b S_i^{b2}$).

Layer 1, output variable

$$y_i^{b1} \cong \frac{x_i^{b1}}{p_3^b + \|x^{b1}\|}. \tag{33}$$

TABLE 2
LTM Weights z_{ij}^{bc}

						$j \longrightarrow$	
1	2	3	4	5	6		
1.0	0.0	0.0	1.0	0.176	0.0001	1	$i \downarrow$
1.0	0.0	0.0	0.0	0.162	0.0001	2	
0.0	0.9	0.0	0.0	0.148	0.0001	3	
0.0	0.9	0.0	0.0	0.134	0.0001	4	
0.0	0.0	0.8	0.0	0.120	0.0001	5	
0.0	0.0	0.8	0.0	0.0	0.0001	6	
0.0	0.0	0.0	0.0	0.0	0.0001	7	
...	
...	

$i = 1 \dots n_a = n_b = 15$
 $j = 1 \dots n_c = 20$

TABLE 4

Parameters		
$p_1^a = p_1^b = p_1^c = 10.0$		
$p_2^a = p_2^b = p_2^c = 10.0$		
$p_3^a = p_3^b = p_3^c = 0.0001$		
$p_4^c = 0.9$		
$p_5^b = p_5^c = 0.1$		
$p_6^b = p_6^c = 1.0$		
Signal Functions g^a, g^b, g^c		
F_a, F_b Distributed	F_c Choice	F_c Distributed
$p_7^a = p_7^b = 0.0$	$p_7^c = 1/\sqrt{n_c}$	$p_7^c = 0.0$
$p_8^a = p_8^b = 0.3$	$p_8^c = 0.2$	$p_8^c = 0.4$
Distributed		Choice
$g(w) = \begin{cases} 0 & \text{if } w \leq p_7 + p_8 \\ \left(\frac{w - p_7}{p_8}\right) & \text{if } w > p_7 + p_8 \end{cases}$		$g(w) = \begin{cases} 0 & \text{if } w \leq p_7 \\ \left(\frac{w - p_7}{p_8}\right)^2 & \text{if } w > p_7 \end{cases}$

Layer 2, input variable

$$x_i^{b2} \cong S_i^{b1} + p_2^b S_i^{b3}. \quad (34)$$

Layer 2, output variable

$$y_i^{b2} \cong \frac{x_i^{b2}}{p_3^b + \|\mathbf{x}^{b2}\|}. \quad (35)$$

Layer 3, input variable

$$x_i^{b3} \cong S_i^{b2} + p_4^c \sum_j v_j^{cb}. \quad (36)$$

Layer 3, output variable

$$y_i^{b3} \cong \frac{x_i^{b3}}{p_5^c + \|\mathbf{x}^{b3}\|}. \quad (37)$$

Normalization of the output variables in eqns (33), (35), and (37) accomplishes two goals. First, since the nonlinear signal function g^b in eqn (32) has a fixed threshold, normalization is needed to achieve orderly pattern transformations under variable processing loads. This goal could have been reached with other norms, such as the L^1 norm ($\|\mathbf{x}\| \equiv \sum_i x_i$). The second goal of normalization is to allow the patterns to have direct access to category representations, without search, after the code has stabilized (Carpenter & Grossberg, 1987a, 1987b). Equations (13) and (17) together tie the Euclidean norm to direct access in the present model. If direct access is not needed, or if another measure of similarity of vectors is used, the Euclidean norm may be replaced by L^1 or another norm.

Transmitter Equations

When the reset signal equals 0, levels of presynaptic and bound transmitter are governed by equations of the form (1) and (5), as follows.

Presynaptic transmitter, $F_b \rightarrow F_c$

$$\frac{du_{ii}^{bc}}{dt} = (z_{ii}^{bc} - u_{ii}^{bc}) - u_{ii}^{bc} p_5^c (x_i^{c1} + p_6^c) S_i^{b3}. \quad (38)$$

Bound transmitter, $F_b \rightarrow F_c$

$$\frac{dv_{ii}^{bc}}{dt} = -v_{ii}^{bc} + u_{ii}^{bc} p_5^c (x_i^{c1} + p_6^c) S_i^{b3}. \quad (39)$$

Presynaptic transmitter, $F_c \rightarrow F_b$

$$\frac{du_{ii}^{cb}}{dt} = (z_{ii}^{cb} - u_{ii}^{cb}) - u_{ii}^{cb} p_5^b (x_i^{b3} + p_6^b) S_i^{c1}. \quad (40)$$

Bound transmitter, $F_c \rightarrow F_b$

$$\frac{dv_{ii}^{cb}}{dt} = -v_{ii}^{cb} + u_{ii}^{cb} p_5^b (x_i^{b3} + p_6^b) S_i^{c1}. \quad (41)$$

Note that eqns (38) and (39) imply that

$$u_{ii}^{bc} + v_{ii}^{bc} \longrightarrow z_{ii}^{bc} \quad (42)$$

and eqns (40) and (41) imply that

$$u_{ii}^{cb} + v_{ii}^{cb} \longrightarrow z_{ii}^{cb} \quad (43)$$

Reset Equations

Reset occurs when patterns active at F_a and F_b fail to match according to the criterion set by the vigilance parameter. In Figure 21,

$$r_i^b \cong \frac{y_i^{a2} + y_i^{b2}}{p_i^a + \|y^{a2}\| + \|y^{b2}\|}. \quad (44)$$

Reset occurs if

$$\|r^b\| < \rho^b, \quad (45)$$

where

$$0 < \rho^b < 1. \quad (46)$$

As in eqns (5) and (6), the effect of a large reset signal is approximated by setting input variables x_i^{b1} , x_i^{b3} , x_j^{c1} , x_j^{c3} and bound transmitter variables v_{ij}^{bc} , v_{ji}^{cb} equal to 0.

Iteration Steps

Steps 1–7 outline the iteration scheme in the computer program used to generate the simulations.

- Step 1. $t \rightarrow t + \Delta t$.
- Step 2. Set ρ and S_i^{a3} values.
- Step 3. Compute r_i^b and check for reset.
- Step 4. Iterate STM equations F_b , F_c five times, setting variables to 0 at reset.
- Step 5. Iterate transmitter eqns (38)–(41).
- Step 6. Compute sums $\sum_i v_{ij}^{bc}$ and $\sum_j v_{ji}^{cb}$.
- Step 7. Return to Step 1.

27. CONCLUSION

In conclusion, we have seen that a functional analysis of parallel search within a hierarchical ART architecture can exploit processes taking place at the chemical synapse as a rich source of robust designs with natural realizations. Conversely, such a neural network analysis embeds model synapses into a processing context that can help to give functional and behavioral meaning to mechanisms defined at the intracellular, biophysical, and biochemical levels.

REFERENCES

- Carpenter, G. A., & Grossberg, S. (1987a). A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing*, **37**, 54–115.
- Carpenter, G. A., & Grossberg, S. (1987b). ART 2: Self-organization of stable category recognition codes for analog input patterns. *Applied Optics*, **26**, 4919–4930.
- Carpenter, G. A., & Grossberg, S. (1988). The ART of adaptive pattern recognition by a self-organizing neural network. *Computer*, **21**, 77–88.
- Carpenter, G. A., & Grossberg, S. (1989). Search mechanisms for Adaptive Resonance Theory (ART) architectures. *Proceedings of the International Joint Conference on Neural Networks* (pp. 1 201–205). Washington, DC: IEEE.
- Cohen, M. A., Grossberg, S., & Stork, D. (1988). Speech perception and production by a self-organizing neural network. In Y. C. Lee (Ed.), *Evolution, learning, cognition, and advanced architectures* (pp. 217–231). Hong Kong: World Scientific Publishers.
- Grossberg, S. (1976a). Adaptive pattern classification and universal recoding, I: Parallel development and coding of neural feature detectors. *Biological Cybernetics*, **23**, 121–134.
- Grossberg, S. (1976b). Adaptive pattern classification and universal recoding, II: Feedback, expectation, olfaction, and illusions. *Biological Cybernetics*, **23**, 187–202.
- Grossberg, S. (1982a). *Studies of mind and brain: Neural principles of learning, perception, development, cognition, and motor control*. Boston: Reidel Press.
- Grossberg, S. (1982b). Processing of expected and unexpected events during conditioning and attention: A psychophysiological theory. *Psychological Review*, **89**, 529–572.
- Grossberg, S. (1984). Some psychophysiological and pharmacological correlates of a developmental, cognitive, and motivational theory. In R. Karrer, J. Cohen, & P. Tueting (Eds.), *Brain and information: Event-related potentials* (pp. 58–151). New York: New York Academy of Sciences.
- Grossberg, S. (Ed.) (1987a). *The adaptive brain, I: Cognition, learning, reinforcement, and rhythm*. Amsterdam: North-Holland.
- Grossberg, S. (Ed.) (1987b). *The adaptive brain, II: Vision, speech, language, and motor control*. Amsterdam: North-Holland.
- Grossberg, S. (Ed.) (1988). *Neural networks and natural intelligence*. Cambridge, MA: MIT Press.
- Grossberg, S., & Levine, D. S. (1987). Neural dynamics of attentionally modulated Pavlovian conditioning: Blocking, inter-stimulus interval, and secondary reinforcement. *Applied Optics*, **26**, 5015–5030.
- Ito, M. (1984). *The cerebellum and neural control*. New York: Raven Press.
- Kandel, E. R., & Schwartz, J. H. (1981). *Principles of neural science*. New York: Elsevier/North-Holland.
- Kohonen, T. (1984). *Self-organization and associative memory*. New York: Springer-Verlag.
- Kuffler, S. W., Nicholls, J. G., & Martin, A. R. (1984). *From neuron to brain* (2nd ed.). Sunderland, MA: Sinauer Associates.