A NEURAL NETWORK REALIZATION OF FUZZY ART

Gail A. Carpenter† , Stephen Grossberg‡ and David B. Rosen¶

Center for Adaptive Systems
and
Department of Cognitive and Neural Systems
Boston University
111 Cummington Street
Boston, Massachusetts 02215 USA

Technical Report CAS/CNS–TR–91–021

Boston, MA: Boston University

Key Words: fuzzy ART, adaptive resonance theory, neural network, fuzzy set theory, categorization, pattern recognition

† Supported in part by British Petroleum (89-A-1204), DARPA (AFOSR 90-0083), the National Science Foundation (NSF IRI 90-00530) and the Office of Naval Research (ONR N00014-91-J-4100).

‡ Supported in part by the Air Force Office of Scientific Research (AFOSR 90-0175), DARPA (AFOSR 90-0083) and the Office of Naval Research(ONR N00014-91-J-4100).

¶ Supported in part by DARPA (AFOSR 90-0083).

Acknowledgements: The authors wish to thank Diana J. Meyers for her valuable assistance in the preparation of the manuscript.

ABSTRACT

A neural network realization of the fuzzy Adaptive Resonance Theory (ART) algorithm is described. Fuzzy ART is capable of rapid stable learning of recognition categories in response to arbitrary sequences of analog or binary input patterns. Fuzzy ART incorporates computations from fuzzy set theory into the ART 1 neural network, which learns to categorize only binary input patterns, thus enabling the network to learn both analog and binary input patterns. In the neural network realization of fuzzy ART, signal transduction obeys a path capacity rule. Category choice is determined by a combination of bottom-up signals and learned category biases. Top-down signals impose upper bounds on feature node activations.

1. ART systems for unsupervised and supervised learning

Adaptive Resonance Theory, or ART, was introduced as a theory of human cognitive information processing (Grossberg, 1976, 1980). The theory has since led to an evolving series of real-time neural network models for unsupervised category learning and pattern recognition. These models self-organize stable recognition categories in response to arbitrary input sequences with either fast or slow learning. Model families include ART 1 (Carpenter and Grossberg, 1987a), which learns to categorize binary input patterns; ART 2 (Carpenter and Grossberg, 1987b), which learns to categorize either analog or binary input patterns; and ART 3 (Carpenter and Grossberg, 1990), which includes a medium-term memory that enables the network to carry out parallel search of distributed recognition codes. Like ART 2, fuzzy ART (Carpenter, Grossberg, and Rosen, 1991) learns to categorize either analog or binary inputs. However, the fuzzy ART measure of pattern similarity is the city-block (L_1) metric, rather than the euclidean (L_2) metric of ART 2. In fact, fuzzy ART generalizes ART 1 in the sense that it reduces to ART 1 when all inputs are binary. However, the neural network that realizes ART 1 does not naturally extend to a network realization of fuzzy ART. The system introduced herein does realize fuzzy ART as a self-organizing neural network that uses only local computations. By extension, the network is also a new realization of ART 1.

Fuzzy ART is an example of how computations from fuzzy set theory can be incorporated naturally into ART systems. For example, the intersection ($\cap$) operator that describes ART 1 dynamics is replaced by the AND operator ($\wedge$) of fuzzy set theory (Zadeh, 1965) in the choice, search, and learning laws of ART 1 (Figure 1). Noteworthy is the close relationship between the computation that defines fuzzy subsethood (Kosko, 1986) and the computations that define category choice and matching in ART 1. Replacing operation $\cap$ by operation $\wedge$ leads to a more powerful version of ART 1.

In fuzzy ART, learning converges because all adaptive weights are monotone nonincreasing. This useful stability property could lead to the unattractive property of category proliferation as too many adaptive weights converge to zero. A preprocessing step, called *complement coding*, uses on-cell and off-cell responses to prevent category proliferation (Carpenter, Grossberg, and Rosen, 1991). Complement coding concatenates an input vector $\mathbf{a}$ with its complement. The input to fuzzy ART then becomes $\mathbf{I} = (\mathbf{a}, \mathbf{a}^c)$. This process normalizes input vectors while preserving the amplitudes of individual feature activations. Without complement coding, an ART category memory encodes the degree to which critical features are consistently present in the training exemplars of that category. With complement coding, both the degree of absence and the degree of presence of features are represented by the category weight vector. The network described in Section 4 realizes fuzzy ART with or without complement coding.

ART modules have recently been used to construct network hierarchies for supervised learning. In particular, an architecture called ARTMAP rapidly self-organizes categorical mappings between m-dimensional input vectors and n-dimensional output vectors (Carpenter, Grossberg, and Reynolds, 1991). ARTMAP's internal control mechanisms create stable recognition categories of optimal size by maximizing predictive generalization while minimizing predictive error in an on-line setting. The first ARTMAP used ART 1 modules to learn mappings between binary input and binary output vectors. For supervised learning of analog maps, fuzzy ART replaces ART 1 to form the fuzzy ARTMAP architecture (Carpen-

ART 1 (BINARY)	FUZZY ART (ANALOG)
CATEGORY CHOICE	
$T_j = \dfrac{\lvert I \cap w_j \rvert}{\alpha + \lvert w_j \rvert}$	$T_j = \dfrac{\lvert I \wedge w_j \rvert}{\alpha + \lvert w_j \rvert}$
MATCH CRITERION	
$\dfrac{\lvert I \cap w \rvert}{\lvert I \rvert} \geq \rho$	$\dfrac{\lvert I \wedge w \rvert}{\lvert I \rvert} \geq \rho$
FAST LEARNING	
$w_j^{(new)} = I \cap w_j^{(old)}$	$w_j^{(new)} = I \wedge w_j^{(old)}$
$\cap$ = logical AND intersection	$\wedge$ = fuzzy AND minimum

Figure 1. Analogy between ART 1 and fuzzy ART. In ART 1, $\mathbf{w}_j$ denotes, for category j, the index set of top-down LTM traces that exceed a prescribed positive threshold value (Carpenter, Grossberg, and Rosen, 1991).

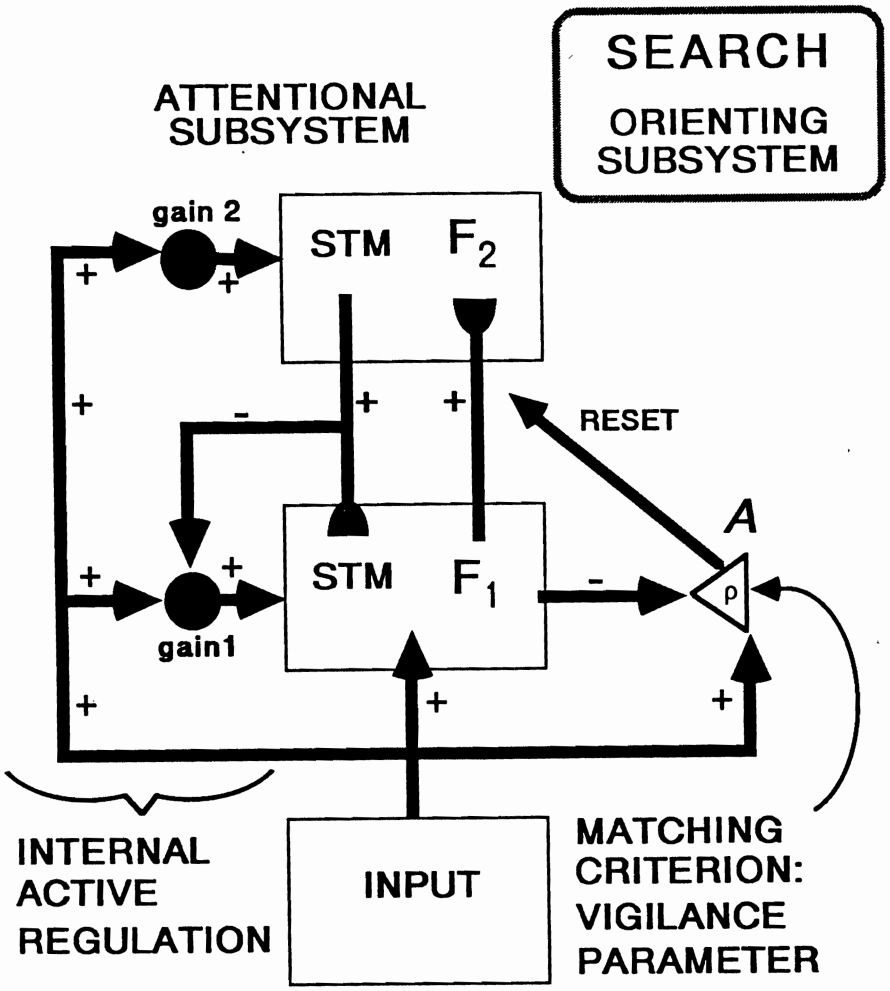

Figure 2. Typical ART 1 neural network (Carpenter and Grossberg, 1987a).

ter, Grossberg, Markuzon, Reynolds, and Rosen, 1991). This system, which learns stable categorical mappings between analog or binary input and output vectors, has performed successfully on benchmark problems that compare ARTMAP performance with machine learning, genetic algorithms, and other neural networks.

The main properties of ART system dynamics will now be outlined (Section 2) followed by a summary of the fuzzy ART algorithm (Section 3). Section 4 includes the specification of a neural network realization of fuzzy ART.

2. ART and fuzzy ART

Fuzzy ART incorporates the basic features of all ART systems, notably, pattern matching between bottom-up input and top-down learned prototype vectors. This matching process leads either to a resonant state that focuses attention and triggers stable prototype learning or to a self-regulating parallel memory search. If the search ends by selecting an established category, then the category's prototype may be refined to incorporate new information in the input pattern. If the search ends by selecting a previously untrained node, then learning of a new category takes place.

Figure 2 illustrates a typical ART 1 model, and Figure 3 illustrates an ART search cycle. As shown in Figure 3a, an input vector **I** registers itself as a pattern **X** of activity across level F_1. The F_1 output vector **S** is then transmitted through the multiple converging and diverging adaptive filter pathways emanating from F_1. This transmission event multiplies the vector **S** by a matrix of adaptive weights, or long term memory (LTM) traces, to generate a new input vector **T** to level F_2. The internal competitive dynamics of F_2 contrast-enhance vector **T**. A compressed activity vector **Y** is thereby generated across F_2. In ART 1, the competition is tuned so that the F_2 node that receives the maximal $F_1 \rightarrow F_2$ input is selected. Only one component of **Y** is nonzero after this choice takes place. Activation of such a winner-take-all node defines the category, or symbol, of the input pattern **I**. Such a category represents all the inputs **I** that maximally activate the corresponding node.

Activation of an F_2 node may be interpreted as "making a hypothesis" about an input **I**. When **Y** is activated, it generates a signal vector **U** that is sent top-down through the second adaptive filter. After multiplications by the adaptive weight matrix of the top-down filter, a net vector **V** inputs to F_1 (Figure 3b). Vector **V** plays the role of a learned top-down expectation. Activation of **V** by **Y** may be interpreted as "testing the hypothesis" **Y**, or "reading out the category prototype" **V**. The ART 1 network is designed to match the "expected prototype" **V** of the category against the active input pattern, or exemplar, **I**.

This matching process may change the F_1 activity pattern **X** by suppressing activation of all the feature detectors in **I** that are not confirmed by **V**. The resultant pattern **X*** encodes the pattern of features to which the network "pays attention". If the expectation **V** is close enough to the input **I**, then a state of *resonance* occurs as the attentional focus takes hold. The resonant state persists long enough for learning to occur; hence the term *adaptive resonance* theory. ART 1 learns prototypes, rather than exemplars, because the attended feature vector **X***, rather than the input **I** itself, is learned.

The criterion of an acceptable match is defined by a dimensionless parameter called *vigilance* (ρ, Figure 2). Vigilance calibrates how close the input exemplar **I** must be to the top-down prototype **V** in order for resonance to occur. Because vigilance can vary across learning trials, recognition categories capable of encoding widely differing degrees of generalization, can be learned by a single ART system. Low vigilance leads to broad generalization and abstract prototypes that represent fewer input exemplars. In the limit of very high vigilance, prototype learning reduces to exemplar learning. Thus a single ART system may be used, say, to recognize abstract categories of faces and dogs, as well as individual faces and dogs.

If the combination of top-down expectation **V** and the bottom-up input **I** is too novel, or unexpected, to satisfy the vigilance criterion, then a bout of hypothesis testing or memory search, is triggered. Search leads to selection of a better recognition code, symbol, category, or hypothesis to represent input **I** at level F_2. An *orienting subsystem A* mediates the search process. The orienting subsystem interacts with the attentional subsystem, as in Figures 3c and 3d, to enable the attentional subsystem to learn about novel inputs without risking unselective forgetting of its previous knowledge. ART 3 specifies a search mechanism that incorporates medium-term memory (MTM) into the adaptive filter.

The search process prevents associations from forming between **Y** and **X*** if **X*** is too different from **I** to satisfy the vigilance criterion. The search process resets **Y** before such an

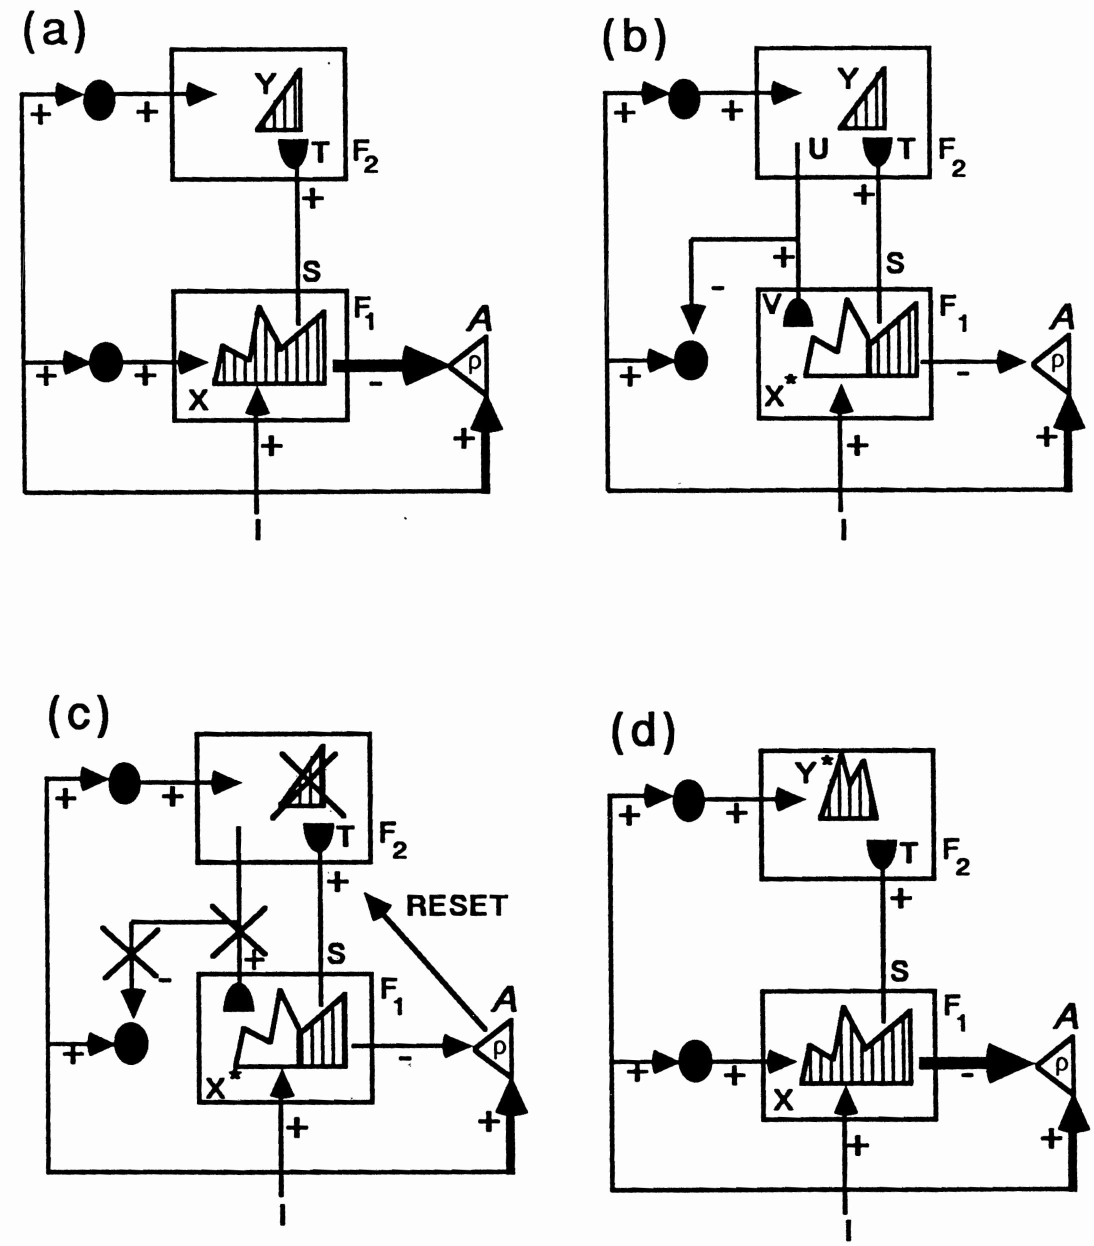

Figure 3. ART search for an F_2 code. (a) The input pattern **I** generates the specific STM activity pattern **X** at F_1 as it nonspecifically activates the orienting subsystem A. Pattern **X** both inhibits A and generates the output signal pattern **S**. Signal pattern **S** is transformed into the input pattern **T**, which activates the STM pattern **Y** across F_2. (b) Pattern **Y** generates the top-down signal pattern **U** which is transformed into the prototype pattern **V**. If **V** mismatches **I** at F_1, then a new STM activity pattern **X*** is generated at F_1. The reduction in total STM activity which occurs when **X** is transformed into **X*** causes a decrease in the total inhibition from F_1 to A. (c) If the matching criterion fails to be met for a given vigilance ρ, A releases a nonspecific signal which resets the STM pattern **Y** at F_2. (d) After **Y** is inhibited, its top-down prototype signal is eliminated, and **X** can be reinstated at F_1. ART 3 specifies how MTM in the adaptive filter leaves enduring traces of the prior reset that allow **X** to activate a different STM pattern **Y*** at F_2. Search continues until the matching criterion is satisfied.

association can form. A familiar category may be selected by the search if its prototype is similar enough to the input **I** to satisfy the vigilance criterion. The prototype may then be refined in light of new information carried by **I**. If **I** is too different from any of the previously learned prototypes, then an uncommitted F_2 node is selected and learning of a new category is initiated.

A network parameter, called the *choice parameter* (α, Figure 1), controls how deeply the search proceeds before an uncommitted node is chosen. As learning of a particular category self-stabilizes, all inputs coded by that category access it directly in a one-pass fashion, and search is automatically disengaged. The category selected is, then, the one whose prototype provides the globally best match to the input pattern. Learning can proceed on-line, and in a stable fashion, with familiar inputs directly activating their categories, while novel inputs continue to trigger adaptive searches for better categories, until the network's memory capacity if fully utilized.

3. Summary of the fuzzy ART algorithm

Fuzzy ART field activity vectors: Like all ART systems, fuzzy ART includes a field F_0 of nodes that represent a current input vector; a field F_2 that represents the active code, or category; and a field F_1 that combines bottom-up input from F_0 with top-down input from F_2 in an ART matching process. The F_0 output vector is denoted $\mathbf{I} = (I_1, \ldots I_i, \ldots I_M)$, with each component I_i in the interval [0,1]. The F_1 output vector is denoted $\mathbf{x} = (x_1, \ldots x_i, \ldots x_M)$ and the F_2 output vector is denoted $\mathbf{y} = (y_1, \ldots y_j, \ldots y_N)$. With complement coding (Section 1), **I** is the output of a preprocessing stage that normalizes **I**.

Weight vector: Associated with each F_2 category node $j (j = 1, \ldots, N)$ is a vector $\mathbf{w}_j \equiv (w_{j1}, \ldots w_{ji}, \ldots w_{jM})$ of adaptive weights, or long-term memory (LTM) traces. Initially

$$w_{j1}(0) = \ldots = w_{jM}(0) = 1, \tag{1}$$

and each category is said to be *uncommitted*. After a category is selected for coding it becomes *committed*. As shown below, each LTM trace w_{ji} is monotone nonincreasing through time and hence converges to a limit. In the fuzzy ART algorithm, the weight vector $\mathbf{w}_j$ subsumes both the bottom-up and top-down weight vectors of ART 1 (Figure 2). The fuzzy ART neural network includes both bottom-up weight vectors $\mathbf{w}_j^{BU}$ and top-down weight vectors $\mathbf{w}_j^{TD}$, in order to realize fuzzy ART using only local computations (Figure 4). However, $\mathbf{w}_j^{BU} = \mathbf{w}_j^{TD} = \mathbf{w}_j$ at all times.

Parameters: Fuzzy ART dynamics are determined by a choice parameter $\alpha > 0$; a learning parameter $\beta \in [0,1]$; and a vigilance parameter $\rho \in [0,1]$.

Category choice: For each input **I** and F_2 node j, the *choice function* T_j is defined by

$$T_j(\mathbf{I}) = \frac{|\mathbf{I} \wedge \mathbf{w}_j|}{\alpha + |\mathbf{w}_j|} \tag{2}$$

where the fuzzy AND (Zadeh, 1965) operator $\wedge$ is defined by

$$(\mathbf{p} \wedge \mathbf{q})_i \equiv \min\,(p_i, q_i) \tag{3}$$

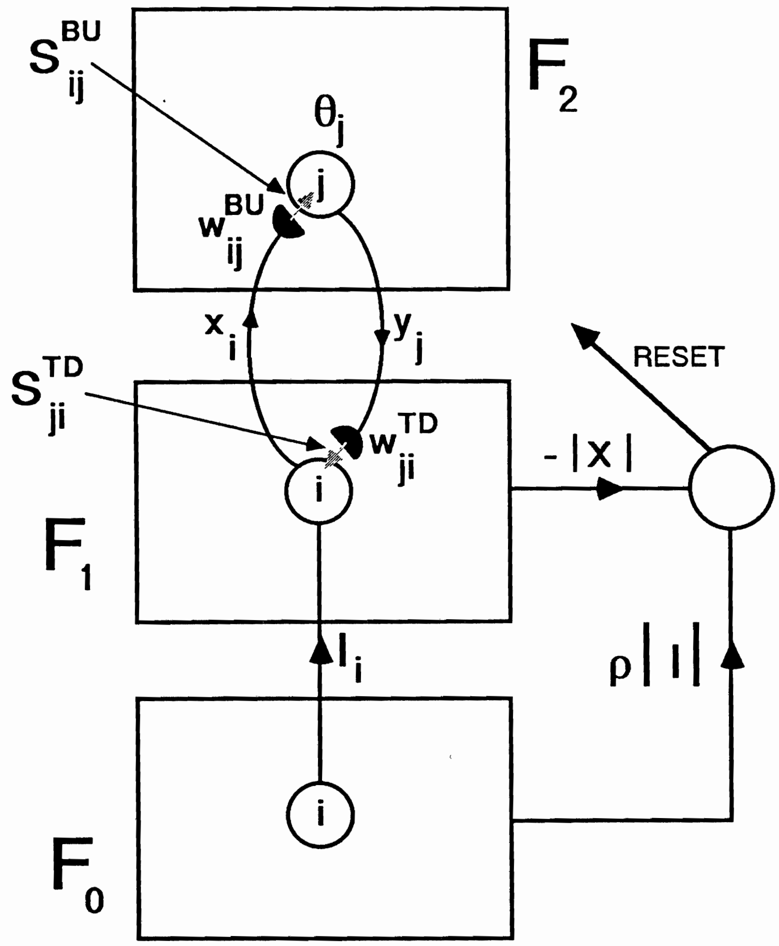

Figure 4. A neural network realization of fuzzy ART. Each F_2 category node j has a trainable strength, or bias, θ_j that can grow when the j^{th} node is active.

for any M-dimensional vectors **p** and **q**; and where the city-block, or L_1, norm $|\cdot|$ is defined by

$$|\mathbf{p}| \equiv \sum_{i=1}^{M} |p_i|. \tag{4}$$

For notational simplicity, $T_j(\mathbf{I})$ in (2) is often written as T_j when the input **I** is fixed.

The system is said to make a *category choice* when at most one F_2 node can become active at a given time. The category choice is indexed by J, where

$$T_J = \max\{T_j : j = 1 \ldots N\}. \tag{5}$$

If more than one T_j is maximal, the category j with the smallest index is chosen. In particular, nodes become committed in order $j = 1, 2, 3, \ldots$. When the J^{th} category is chosen $y_J = 1$; and $y_j = 0$ for $j \neq J$. In a system that makes a choice, the F_1 output vector **x** obeys the equation

$$\mathbf{x} = \begin{cases} \mathbf{I} & \text{if } F_2 \text{ is inactive} \\ \mathbf{I} \wedge \mathbf{w}_J & \text{if the } J^{th}\ F_2 \text{ node is chosen.} \end{cases} \tag{6}$$

Resonance or reset: *Resonance* occurs if the *match function* $|\mathbf{x}|/|\mathbf{I}|$ of the chosen category meets the vigilance criterion; that is, if:

$$\frac{|\mathbf{x}|}{|\mathbf{I}|} \geq \rho. \tag{7}$$

By (6), when the J^{th} category is chosen, resonance occurs if:

$$|\mathbf{x}| = |\mathbf{I} \wedge \mathbf{w}_J| \geq \rho|\mathbf{I}|. \tag{8}$$

Learning then ensues, as defined below. *Mismatch reset* occurs if:

$$\frac{|\mathbf{x}|}{|\mathbf{I}|} < \rho \tag{9}$$

(Figure 4). Thus, when the J^{th} category is chosen, reset occurs if:

$$|\mathbf{x}| = |\mathbf{I} \wedge \mathbf{w}_J| < \rho|\mathbf{I}|. \tag{10}$$

Thereafter the value of the choice function $T_J(\mathbf{I})$ is set to 0 for the duration of the presentation of input $\mathbf{I}$, in order to prevent the persistent selection of the same category during search. A new index J is then chosen, by (5). The search process continues until the chosen J satisfies the matching criterion (7). By (1), (6), and (7), search always ends if J is an uncommitted node.

Learning: Once search ends, the weight vector $\mathbf{w}_J$ is updated according to the equation:

$$\mathbf{w}_J^{(new)} = \beta(\mathbf{I} \wedge \mathbf{w}_J^{(old)}) + (1-\beta)\mathbf{w}_J^{(old)}. \tag{11}$$

Fast learning corresponds to setting the learning parameter β equal to 1.

4. A neural network realization of fuzzy ART

The computations described in Section 3 are derived via a direct translation from the binary, set-theoretic description of ART 1 to the analog, fuzzy set-theoretic description of fuzzy ART (Figure 1). However the neural network that realizes ART 1 (Carpenter and Grossberg, 1987a) does not have such a direct extension to a network realization of fuzzy ART. A different neural network that both reduces to ART 1 in the binary case and performs the computations of algorithmic fuzzy ART in the analog case will now be described.

Activity output vectors: As in Section 3, let $\mathbf{I} = (I_1, \ldots I_i, \ldots I_M)$ denote the output vector of an ART field F_0, with $0 \leq I_i \leq 1$; let $\mathbf{x} = (x_1, \ldots x_i, \ldots x_M)$ denote the output vector of a field F_1; and let $\mathbf{y} = (y_1, \ldots y_j, \ldots y_N)$ denote the output vector of a field F_2 (Figure 4).

$\mathbf{F_2}$ choice: The field F_2 is assumed to be a competitive network designed to make a choice (winner-take-all). That is, at most one node ($j = J$) can be active at a time, and

$$y_j = \begin{cases} 1 & \text{if } F_2 \text{ is active and } j = J \\ 0 & \text{if otherwise.} \end{cases} \tag{12}$$

LTM weights: An ART system characteristically includes both bottom-up and top-down adaptive filters (Figure 2) that play dual roles in the architecture's neural computation. The description of fuzzy ART in Section 3 included a single weight vector $\mathbf{w}_j$ for each F_2 category index $j = 1, \ldots, N$. In order to specify a system that uses only local computations, the neural network realization of fuzzy ART includes both bottom-up ($F_1 \to F_2$) weight vectors $\mathbf{w}_j^{BU} \equiv (w_{1j}^{BU}, \ldots, w_{ij}^{BU}, \ldots w_{Mj}^{BU})$ and top-down ($F_2 \to F_1$) weight vectors $\mathbf{w}_j^{TD} \equiv (w_{j1}^{TD}, \ldots, w_{ji}^{TD}, \ldots w_{jM}^{TD})$. The differential equations that determine $\mathbf{w}_j^{BU}$ and $\mathbf{w}_j^{TD}$ imply that, at all times, the size of the bottom-up weight w_{ij}^{BU} in the path from the i^{th} F_1 node to the j^{th} F_2 node equals the size of the top-down weight w_{ji}^{TD} from the j^{th} F_2 node to the i^{th} F_1 node. Equations (13)-(14) describe the dynamics of these adaptive weights during learning:

$$\frac{d\mathbf{w}_j^{BU}}{dt} = y_j[\mathbf{x} - \mathbf{w}_j^{BU}] \tag{13}$$

$$\frac{d\mathbf{w}_j^{TD}}{dt} = y_j[\mathbf{x} - \mathbf{w}_j^{TD}]. \tag{14}$$

Initial values are given by:

$$\mathbf{w}_{ij}^{BU}(0) = \mathbf{w}_{ji}^{TD}(0) = 1 \tag{15}$$

for $i = 1, \ldots, M$ and $j = 1, \ldots, N$.

Signal transduction: In this network realization of fuzzy ART, a weight is interpreted as a path capacity, or upper bound on the maximum size signal that can be transmitted through the weight's corresponding path. That is, when a signal x traveling along a path with weight w transmits a net signal S to a target cell, S is determined by the equation:

Path capacity rule $$S = x \wedge w. \tag{16}$$

Thus the net top-down signal S_{ji}^{TD} from the j^{th} F_2 node to the i^{th} F_1 node obeys the equation:

$$S_{ji}^{TD} = y_j \wedge w_{ji}^{TD}; \tag{17}$$

and the net bottom-up signal S_{ij}^{BU} from the i^{th} F_1 node to the j^{th} F_2 node obeys the equation:

$$S_{ij}^{BU} = x_i \wedge w_{ij}^{BU}. \tag{18}$$

Note that the path capacity rule (16) specifies a signal transduction mechanism that is fundamentally different from that of equation:

Mass action rule $$S = xw, \tag{19}$$

widely used in neural network models. Equation (19) postulates that a path weight w acts multiplicatively, or by mass action, upon a signal x. The path capacity rule and the mass action rule are identical, however, if x is binary (and w is scaled so that $0 \leq w \leq 1$); or if

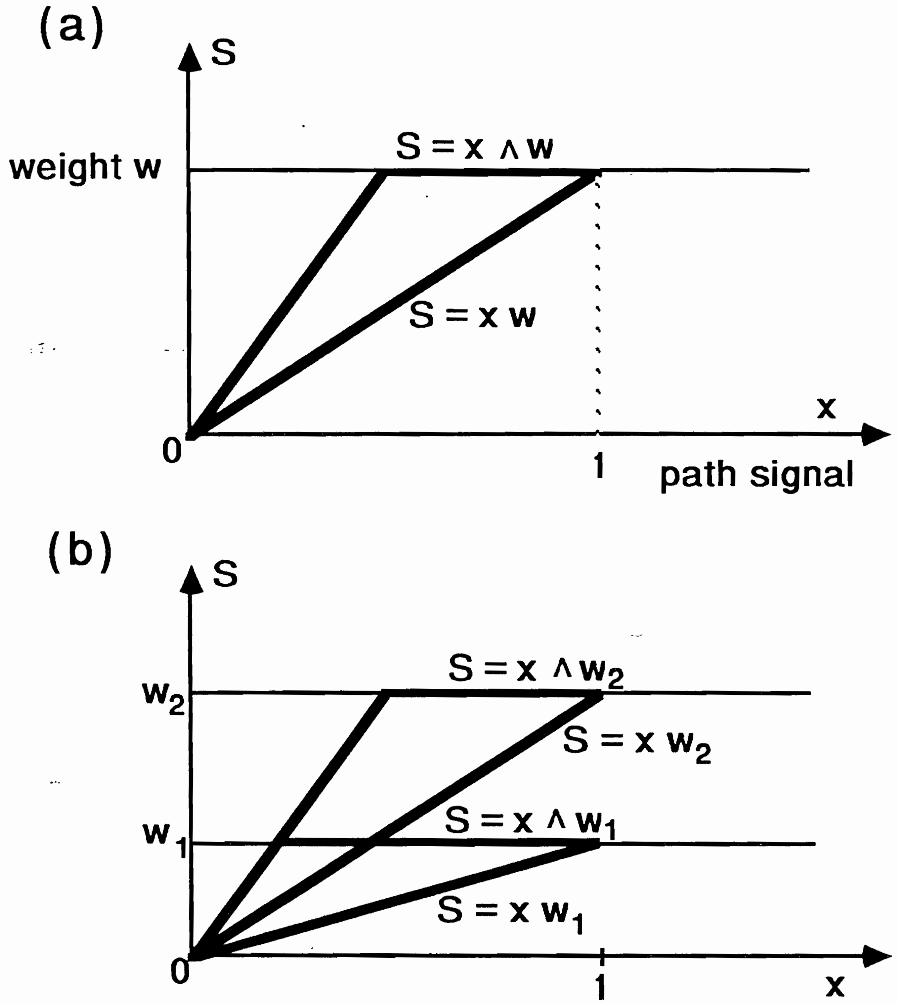

Figure 5. (a) The fuzzy ART path capacity rule for signal transduction implies that the net signal $S = x \wedge w$ transmitted to a target cell is distinct from the signal $S = xw$ postulated by a mass action rule, except when the signal path x is binary. (b) The path capacity rule (16) implies that small (low frequency) path signals x are transmitted identically through synapses with large weights (w_2) and small weights (w_1). In contrast, a mass action rule (19) implies that any signal x, large or small, is transmitted differently through any two paths with different weights.

either $x = 0$ or $w = 0$ (Figure 5a). When path signals are analog, the rules (17) and (19) imply both distinct physical mechanisms and distinct network computations (Figure 5b). In particular, the path capacity rule implies that, for a large weight w, the net signal S is proportional to the path signal x; but that, for a small weight the maximum net path signal is limited to the small upper bound (capacity) imposed by the weight.

The total signal σ_i^{TD} from F_2 to the i^{th} F_1 node is the sum:

$$\sigma_i^{TD} = \sum_j S_{ji}^{TD} = \sum_j y_j \wedge w_{ji}^{TD}, \tag{20}$$

by (17). When F_2 makes a choice, with some $y_J = 1$ and all other $y_j = 0$,

$$\sigma_i^{TD} = y_J \wedge w_{Ji}^{TD} = w_{Ji}^{TD}. \tag{21}$$

In this case, therefore, the $F_2 \to F_1$ signal vector $\boldsymbol{\sigma}^{TD} \equiv (\sigma_1^{TD}, \ldots, \sigma_M^{TD})$ reduces to the top-down weight vector of the J^{th} F_2 node:

$$\boldsymbol{\sigma}^{TD} = \mathbf{w}_J^{TD}. \tag{22}$$

$\mathbf{F_1}$ activation: The fuzzy ART $F_2 \to F_1$ signal $\boldsymbol{\sigma}^{TD}$ is assumed to have a net *inhibitory* effect on nodes in F_1. In fact, the F_1 output vector $\mathbf{x}$ obeys the equation:

$$\mathbf{x} = \begin{cases} \mathbf{I} & \text{if } F_2 \text{ is inactive} \\ \mathbf{I} \wedge \boldsymbol{\sigma}^{TD} & \text{if } F_2 \text{ in active.} \end{cases} \tag{23}$$

By (23), the total top-down signal σ_i^{TD} to the i^{th} F_1 node imposes an upper bound on the signal I_i that can be transferred through that node, without truncation. A small top-down signal σ_i^{TD} quenches the F_1 output x_i, whereas a large signal σ_i^{TD} permits I_i to pass through F_1 undiminished, with $x_i = I_i$. When F_2 makes a choice,

$$\mathbf{x} = \begin{cases} \mathbf{I} & \text{if } F_2 \text{ is active} \\ \mathbf{I} \wedge \mathbf{w}_J^{TD} & \text{if the } J^{th}\ F_2 \text{ node is active,} \end{cases} \tag{24}$$

by (21) and (23).

Mismatch and search: F_2 activation induces a search if total F_1 activity is thereby reduced to the point where the vigilance criterion fails to be met. This occurs when the net input to the orienting subsystem is positive:

$$\rho|\mathbf{I}| - |\mathbf{x}| > 0 \tag{25}$$

(Figure 4). An ART search process requires not only that STM at F_1 and F_2 then be reset, but also that a selective bias against recently active F_2 nodes endures, allowing a new category representation to be tested. ART 3 (Carpenter and Grossberg, 1990) can realize such a parallel search for systems in which F_2 either makes a choice or has distributed activity. ART 3 constructs a medium-term memory (MTM) that biases the $F_1 \to F_2$ adaptive filter against selection of recently active F_2 nodes. ART 3 MTM also provides a means whereby the input vector $\mathbf{I}$ may vary continuously. Then competition at F_2 holds category representations constant through small fluctuations. When the input drifts too far, reset occurs, automatically shifting the attentional focus to features of the new input.

Resonance and learning: Resonance occurs if the vigilance criterion is met; that is if:

$$\rho|\mathbf{I}| - |\mathbf{x}| \leq 0. \tag{26}$$

Then, (14) and (24) imply that, during learning,

$$\frac{dw_{ji}^{TD}}{dt} = y_j[(I_i \wedge w_{Ji}^{TD}) - w_{ji}^{TD}] = \begin{cases} 0 & \text{if } j \neq J \\ 0 & \text{if } j = J \text{ and } I_i \geq w_{Ji}^{TD} \\ I_i - w_{Ji}^{TD} & \text{if } j = J \text{ and } I_i < w_{Ji}^{TD}. \end{cases} \tag{27}$$

Note that (24) and (27) imply that w_{ji}^{TD} is monotone non-increasing and that $\mathbf{x}$ remains constant during learning. Also, when the F_2 node J becomes active at time $t = t_0$, $\mathbf{w}_J^{TD}(t)$ obeys equation (28) for as long as $\mathbf{I}$ remains constant:

$$\mathbf{w}_J^{TD}(t) = \beta(t)(\mathbf{I} \wedge \mathbf{w}_J^{TD}(t_0)) + (1 - \beta(t))\mathbf{w}_J^{TD}(t_0), \tag{28}$$

where $\beta(t_0) = 0$ and $\beta(t)$ approaches 1 exponentially as $t \to \infty$. Thus, $\mathbf{w}_J^{TD}(t)$, as well as its equal $\mathbf{w}_J^{BU}(t)$, obey the algebraic learning law (11).

Category choice: In order to complete the neural network interpretation of Fuzzy ART, one more hypothesis is needed to realize the $F_1 \to F_2$ choice function T_j defined by Equation (2). By (18), (22), and (23), the net signal S_{ij}^{BU} from the i^{th} F_1 node to the j^{th} F_2 node obeys the equation:

$$S_{ij}^{BU} = \begin{cases} I_i \wedge w_{ij}^{BU} & \text{if } F_2 \text{ is inactive} \\ I_i \wedge w_{Ji}^{TD} \wedge w_{ij}^{BU} & \text{if the } J^{th} \text{ } F_2 \text{ node is active.} \end{cases} \tag{29}$$

Note that since $w_{Ji}^{TD} = w_{iJ}^{BU}$, once category J is chosen, the J^{th} bottom-up signal vector $\mathbf{S}_J^{BU} \equiv (S_{1J}^{BU}, \ldots, S_{MJ}^{BU})$ remains equal to its original value, $\mathbf{I} \wedge \mathbf{w}_J^{BU}$. However, all other components $S_{ij}^{BU} (j \neq J)$ could become smaller than they were before node J became active. Stable category choice requires that the initial bottom-up category choice be confirmed after top-down signals alter the F_1 signal pattern $\mathbf{x}$, as in ART 1 (Carpenter and Grossberg, 1987a). This suggests that the choice function T_j might be realized by a computation of the form

$$T_j(\mathbf{I}) = |\mathbf{S}_j^{BU}| \ \theta_j. \tag{30}$$

Each term θ_j is independent of both F_1 and F_2 STM activity and thus remains the same before and after F_2 becomes active. The size $|\mathbf{S}_J^{BU}|$ of the new $F_1 \to F_2$ signal to the J^{th} node is also unaffected by the reduction of F_1 activity following F_2 activation; but the net signal $|\mathbf{S}_j^{BU}|$ to any other F_2 node may have become smaller, by (29). Thus any choice function $T_j(\mathbf{I})$ defined by an equation of the form (30) realizes the principle that top-down read-out should confirm the original bottom-up choice.

Term θ_j in (30) may be interpreted as an adaptive node weight, or *bias*, that might be realized as the size or strength or number of receptors of the j^{th} F_2 node. To build a neural network whose computations are equivalent to fuzzy ART, θ_J should approach $(\alpha + |\mathbf{w}_J^{BU}|)^{-1}$ during learning, by (2), (29), and (30). This relationship is satisfied if θ_j obeys the equation:

$$\frac{d\theta_j}{dt} = y_j[(1 - \alpha\theta_j) - \theta_j|\mathbf{S}_j^{BU}|]. \tag{31}$$

In other words, when the J^{th} F_2 node is active, θ_J grows toward a maximum level (α^{-1}); but this growth is countered by the sum of all signals to that node. When the J^{th} node is active,

$$\mathbf{S}_J^{BU} = \mathbf{I} \wedge \mathbf{w}_J^{BU} = \mathbf{I} \wedge \mathbf{w}_J^{TD} = \mathbf{x}. \tag{32}$$

Thus during learning:

$$\frac{d\theta_J}{dt} = y_J[(1 - \alpha\theta_J) - \theta_J|\mathbf{x}|] = y_J(\alpha + |\mathbf{x}|)\left[\frac{1}{\alpha + |\mathbf{x}|} - \theta_J\right]. \tag{33}$$

Therefore as $\mathbf{w}_J^{BU} \to \mathbf{x}$ (13),

$$\theta_J \to \frac{1}{\alpha + |\mathbf{w}_J^{BU}|}. \tag{34}$$

Initial values of $\theta_1, \ldots, \theta_N$ are given by:

$$\theta_j(0) = \frac{1}{\alpha + |\mathbf{w}_j^{BU}(0)|} = \frac{1}{\alpha + M}. \tag{35}$$

Since $|\mathbf{w}_J^{BU}|$ decreases during learning, (34) and (35) imply that the node bias θ_J increases monotonically during learning. In the fast-learn limit,

$$T_J(\mathbf{I}) = |\mathbf{S}_J^{BU}|\theta_J \rightarrow \frac{|\mathbf{I} \wedge \mathbf{w}_J^{BU}|}{\alpha + |\mathbf{w}_J^{BU}|} \tag{36}$$

when the J^{th} F_2 node is active. Thus (30)-(35) implement the choice function $T_j(\mathbf{I})$ in (2).

5. Conclusion

This neural network realization of fuzzy ART indicates the type computations that can interpret fuzzy set theory in the neural context. The construction illustrates how a formal algorithm such as fuzzy ART can suggest neural network hypotheses, such as the path capacity rule, and conversely. Other examples are likely to develop as concepts from fuzzy logic and neural network theories continue to enrich one another.

References

Carpenter, G.A. and 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. and Grossberg, S. (1987b). ART 2. Stable self-organization of pattern recognition codes for analog input patterns. *Applied Optics*, **26**, 4919-4930.

Carpenter, G.A. and Grossberg, S. (1990). ART 3. Hierarchical search using chemical transmitters in self-organizing pattern recognition architectures. *Neural Networks*, **3**, 129-152.

Carpenter, G.A., Grossberg, S., Markuzon, N., Reynolds, J.H. and Rosen, D.B. (1991). Fuzzy ARTMAP: A neural network architecture for incremental supervised learning of analog multidimensional maps. *IEEE Transactions on Neural Networks*, **3**, in press. Technical Report CAS/CNS-TR-91-016, Boston, MA: Boston University.

Carpenter, G.A., Grossberg, S. and Reynolds, J.H. (1991). ARTMAP: Supervised real-time learning and classification of nonstationary data by a self-organizing neural network. *Neural Networks*, **4**, 565-588.

Carpenter, G.A., Grossberg, S. and Rosen, D.B. (1991). Fuzzy ART: Fast stable learning and categorization of analog patterns by an adaptive resonance system. *Neural Networks*, **4**, 759-771.

Grossberg, S. (1976). Adaptive pattern classification and universal recoding, II: Feedback, expectation, olfaction, and illusions. *Biological Cybernetics*, **23**, 187-202.

Grossberg, S. (1980). How does a brain build a cognitive code? *Psychological Review*, **1**, 1-51.

Kosko, B. (1986). Fuzzy entropy and conditioning. *Information Sciences*, **40**, 165-174.

Zadeh, L. (1965). Fuzzy sets. *Information Control*, **8**, 338-353. .