Large-scale Neural Systems for Vision and Cognition

Gail A. Carpenter

Department of Cognitive and Neural Systems
Boston University
677 Beacon Street
Boston, Massachusetts 02215 USA

Proceedings of the International Joint Conference on Neural Networks
June, 2009, Atlanta

Technical Report CAS/CNS TR-2009-001
Boston, MA: Boston University

Acknowledgements: This research was supported by the SyNAPSE program of the Defense Advanced Projects Research Agency (Hewlett-Packard Company, DARPA prime contract HR0011-09-3-0001 and HRL Laboratories LLC, subcontract #801881-BS under DARPA prime contract HR0011-09-C-0011) and by the Science of Learning Centers program of the National Science Foundation (NSF SBE-0354378).

Large-scale Neural Systems for Vision and Cognition

Gail A. Carpenter

***Abstract*—** Consideration of how people respond to the question *What is this?* has suggested new problem frontiers for pattern recognition and information fusion, as well as neural systems that embody the cognitive transformation of declarative information into relational knowledge. In contrast to traditional classification methods, which aim to find the single correct label for each exemplar (*This is a car*), the new approach discovers rules that embody coherent relationships among labels which would otherwise appear contradictory to a learning system (*This is a car, that is a vehicle, over there is a sedan*). This talk will describe how an individual who experiences exemplars in real time, with each exemplar trained on at most one category label, can autonomously discover a hierarchy of cognitive rules, thereby converting local information into global knowledge. Computational examples are based on the observation that sensors working at different times, locations, and spatial scales, and experts with different goals, languages, and situations, may produce apparently inconsistent image labels, which are reconciled by implicit underlying relationships that the network's learning process discovers. The ARTMAP information fusion system can, moreover, integrate multiple separate knowledge hierarchies, by fusing independent domains into a unified structure. In the process, the system discovers cross-domain rules, inferring multilevel relationships among groups of output classes, without any supervised labeling of these relationships. In order to self-organize its expert system, the ARTMAP information fusion network features distributed code representations which exploit the model's intrinsic capacity for one-to-many learning (*This is a car* and *a vehicle* and *a sedan*) as well as many-to-one learning (*Each of those vehicles is a car*). Fusion system software, testbed datasets, and articles are available from http://cns.bu.edu/techlab .

I. Introduction

THIS plenary talk will describe recently developed Adaptive Resonance Theory (ART) networks for rule discovery (Carpenter, Martens, & Ogas, 2005; Carpenter & Ravindran, 2008), as described in the Abstract. This paper provides an introduction to the computations and dynamics of the basic ART and ARTMAP networks.

II. ART and ARTMAP

ART neural networks model real-time prediction, search, learning, and recognition. ART networks serve both as models of human cognitive information processing (Grossberg, 1999, 2003; Carpenter, 1997) and as neural systems for technology transfer (Caudell et al., 1994; Lisboa, 2001; Parsons & Carpenter, 2003).

Design principles derived from scientific analyses and design constraints imposed by targeted applications have jointly guided the development of many variants of the basic networks, including fuzzy ARTMAP (Carpenter et al., 1992), ART-EMAP (Carpenter & Ross, 1995), ARTMAP-IC (Carpenter & Markuzon, 1998), and Gaussian ARTMAP (Williamson, 1998). A defining characteristic of various ARTMAP classes is the nature of the internal code representation. Early ARTMAP systems, including fuzzy ARTMAP, employ winner-take-all (WTA) coding, whereby each input activates a single category node during both training and testing. When a node is first activated during training, it is mapped to its designated output class.

Starting with ART-EMAP, subsequent systems have used distributed coding during testing, which typically improves predictive accuracy, while avoiding the computational problems inherent in the use of distributed code representations during training. In order address these problems, distributed ARTMAP (Carpenter, 1998; Carpenter, Milenova, & Noeske, 1998) introduces a new network configuration, in addition to new learning laws.

Comparative analysis of the performance of ARTMAP systems on a variety of benchmark problems has led to the identification of a *default ARTMAP* network (Carpenter, 2003), which features simplicity of design and robust performance in many application domains. Default ARTMAP employs winner-take-all coding during training and distributed coding during testing within a distributed ARTMAP network architecture. With winner-take-all coding during testing, default ARTMAP reduces to a version of fuzzy ARTMAP.

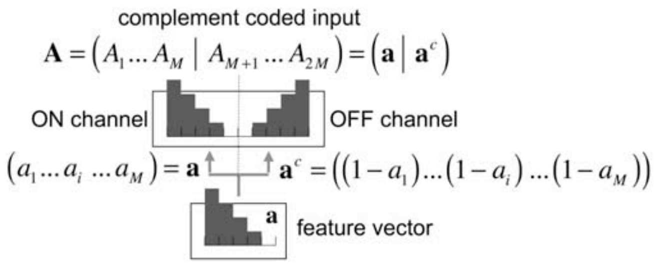

Figure 1. Complement coding transforms an *M*-dimensional feature vector **a** into a 2*M*-dimensional system input vector **A**. A complement-coded system input represents both the degree to which a feature *i* is present (a_i) and the degree to which that feature is absent $(1-a_i)$.

III. Complement Coding: Learning Both Absent and Present Features

ART and ARTMAP employ a preprocessing step called *complement coding* (Figure 1), which models the nervous system's ubiquitous use of the computational design known as *opponent processing* (Hurvich & Jameson, 1957). Balancing an entity against its opponent, as in agonist-antagonist muscle pairs, allows a system to act upon relative quantities, even as absolute magnitudes may vary unpredictably. In ART systems, complement coding (Carpenter, Grossberg, & Rosen, 1991) is analogous to retinal ON-cells and OFF-cells (Schiller, 1982). When the learning system is presented with a set of input features $\mathbf{a} \equiv (a_1 \ldots a_i \ldots a_M)$, complement coding doubles the number of input components, presenting to the network both the original feature vector and its complement.

Complement coding allows an ART system to encode within its critical feature patterns of memory features that are consistently *absent* on an equal basis with features that are consistently *present*. Features that are sometimes absent and sometimes present when a given category is learning become regarded as uninformative with respect to that category. Since its introduction, complement coding has been a standard element of ART and ARTMAP networks, where it plays multiple computational roles, including input normalization. However, this device is not particular to ART, and could, in principle, be used to preprocess the inputs to any type of system.

To implement complement coding, component activities a_i of a feature vector **a** are scaled so that $0 \le a_i \le 1$. For each feature *i*, the ON activity a_i determines the complementary OFF activity $(1-a_i)$. Both a_i and $(1-a_i)$ are represented in the 2*M*-dimensional system input vector $\mathbf{A} = (\mathbf{a} \mid \mathbf{a}^c)$ (Figure 1). Subsequent network computations then operate in this 2*M*-dimensional input space. In particular, learned weight vectors $\mathbf{w}_J$ are 2*M*-dimensional.

IV. ARTMAP Search and Match Tracking

The ART matching process triggers either learning or a parallel memory search (Figure 2). If search ends at an established code, the memory representation may either remain the same or incorporate new information from matched portions of the current input. While this dynamic applies to arbitrarily distributed activation patterns, the F_2 code will here be described as a single category node, in a winner-take all system.

Before ARTMAP makes a class prediction, the bottom-up input **A** is matched against the top-down learned expectation, or critical feature pattern, that is read out by the active node (Figure 2b). The matching criterion is set by a parameter ρ, called *vigilance*. Low vigilance permits the learning of abstract, prototype-like patterns, while high vigilance requires the learning of specific, exemplar-like patterns. When a new input arrives, vigilance equals a baseline level, $\bar{\rho}$. Baseline vigilance is set equal to zero by default, in order to maximize generalization. Vigilance rises only after the system has made a predictive error. The internal control process that determines how far ρ must rise in order to correct the error is called *match tracking* (Carpenter, Grossberg, & Reynolds, 1991). As vigilance rises, the network is required to pay more attention to how well top-down expectations match the current bottom-up input.

Match tracking (Figure 3) forces an ARTMAP system not only to reset its mistakes, but to learn from them. With match tracking and fast learning, each ARTMAP network passes the Next Input Test, which requires that, if a training input were re-presented immediately after a learning trial, it would directly activate the correct output class, with no predictive errors or search. Match tracking thus simultaneously implements the design goals of maximizing generalization and minimizing predictive error, without requiring the choice of a fixed matching criterion. ARTMAP memories thereby include both broad and specific pattern classes, with the latter typically formed as exceptions to the more general "rules" defined by the former. ARTMAP learning typically produces a wide variety of such mixtures, whose exact composition depends upon the order of training exemplar presentation.

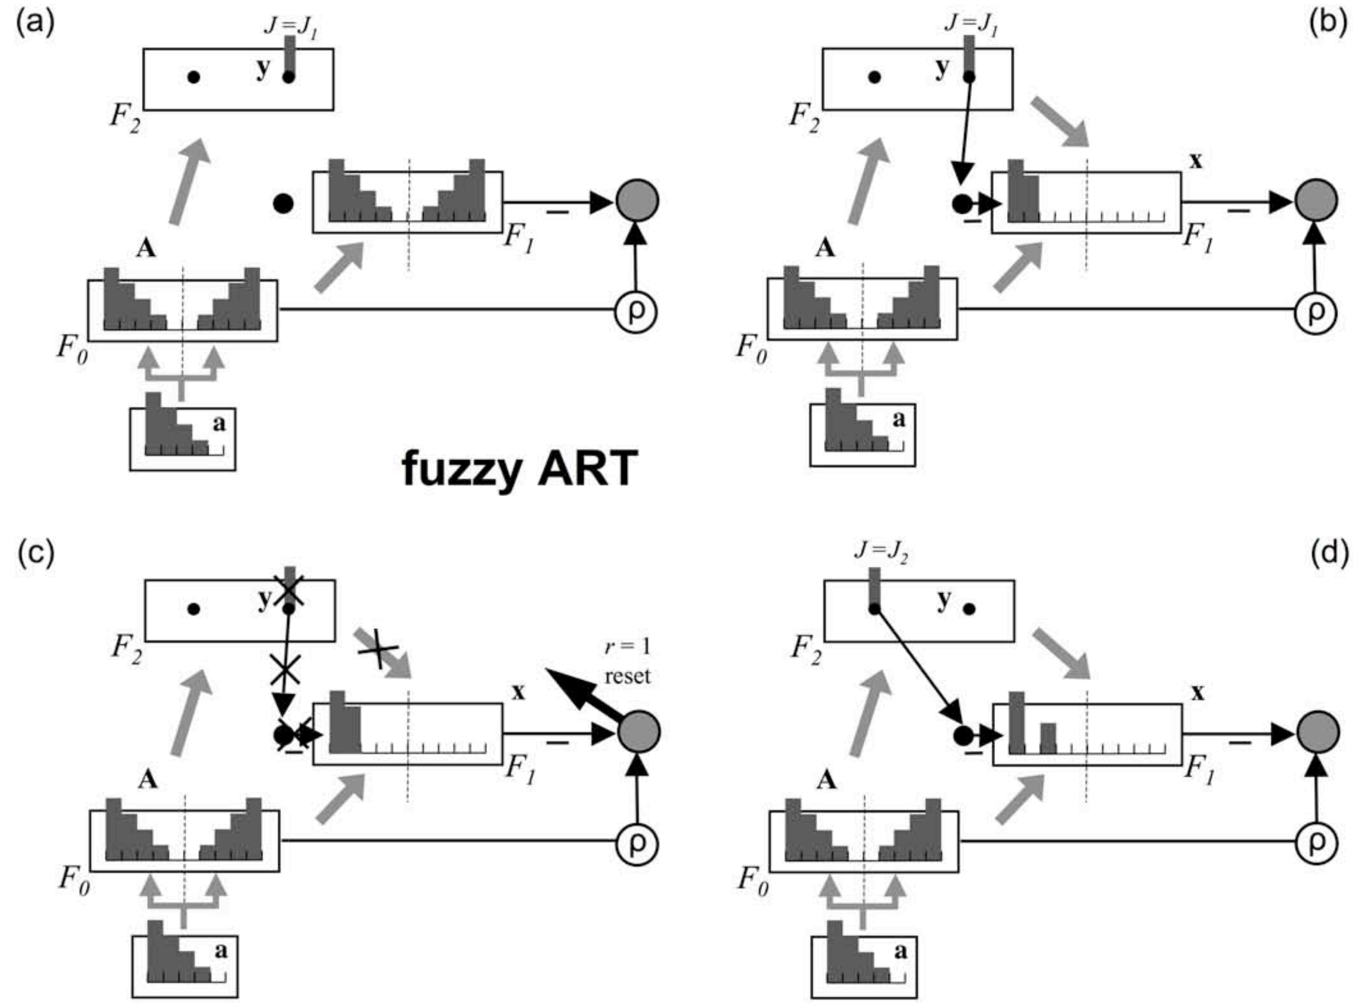

Figure 2. A fuzzy ART search cycle (Carpenter, Grossberg & Rosen, 1991), with a distributed ART network configuration (Carpenter, 1997). The ART 1 search cycle (Carpenter & Grossberg, 1987) is the same, but allows only binary inputs and did not originally feature complement coding. The match field F_1 represents the matched activation pattern $\mathbf{x} = \mathbf{A} \wedge \mathbf{w}_J$, where $\wedge$ denotes the component-wise minimum, or fuzzy intersection, between the bottom-up input **A** and the top-down expectation $\mathbf{w}_J$. If the matched pattern fails to meet the matching criterion, then the active code is reset at F_2, and the system searches for another code **y** that better represents the input. The match / mismatch decision in the ART orienting system. Each active feature in the input pattern **A** excites the orienting system with gain equal to the vigilance parameter ρ. Hence, with complement coding, the total excitatory input is $\rho\left|\mathbf{A}\right| = \rho\sum_{i=1}^{2M} A_i = \rho M$. Active cells in the matched pattern **x** inhibit the orienting system, leading to a total inhibitory input equal to $-\left|\mathbf{x}\right| = -\sum_{i=1}^{2M} x_i$. If $\rho\left|\mathbf{A}\right| - \left|\mathbf{x}\right| \le 0$, then the orienting system remains quiet, allowing resonance and learning to occur. If $\rho\left|\mathbf{A}\right| - \left|\mathbf{x}\right| > 0$, then the reset signal r=1, initiating search for a better matching code.

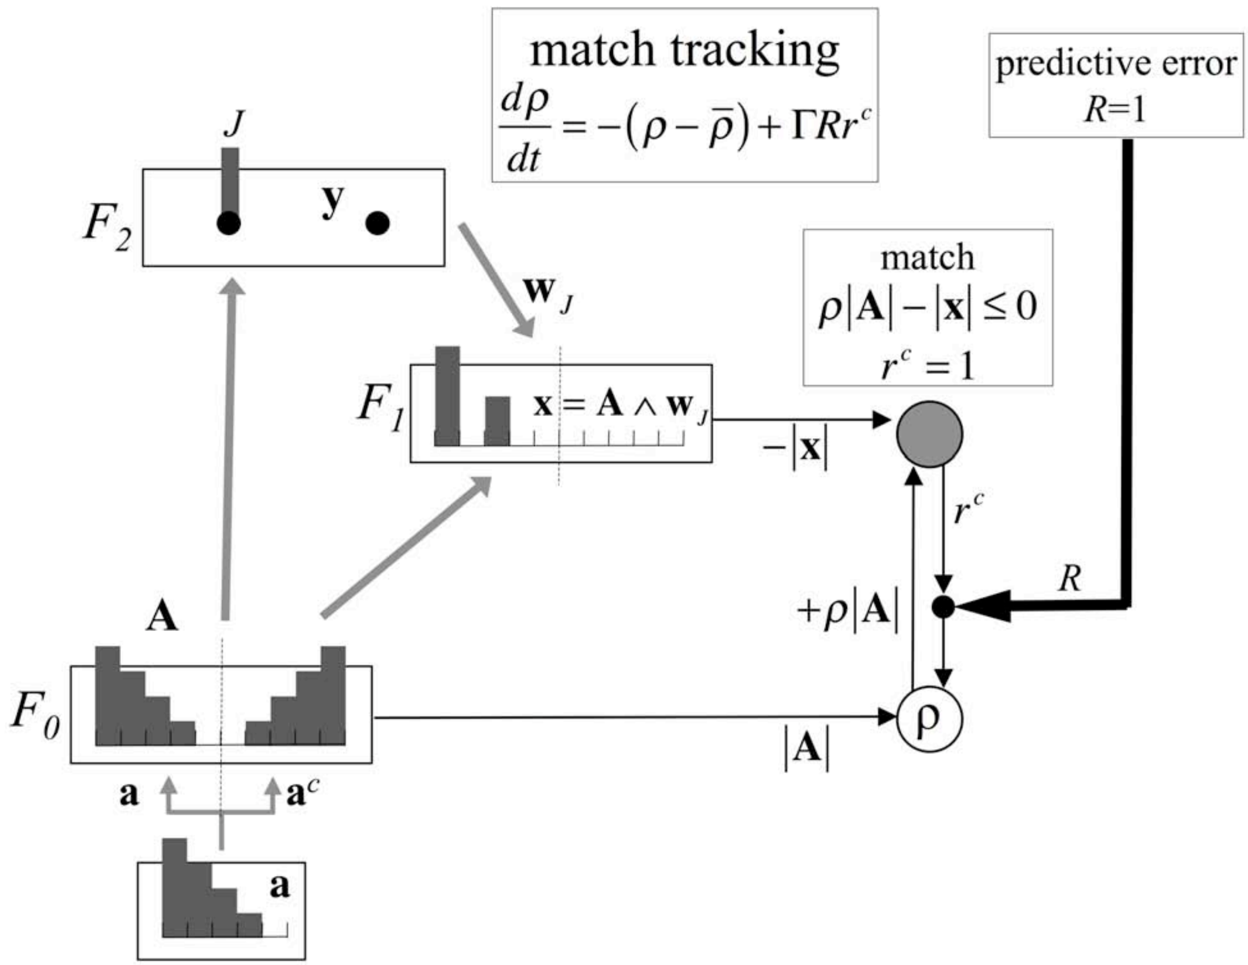

Figure 3. ARTMAP match tracking (Carpenter, Grossberg, & Reynolds, 1991). When an active node J meets the matching criterion $(\rho|\mathbf{A}| - |\mathbf{x}| \leq 0)$, the reset signal r=0 and the node makes an prediction. If the predicted output is incorrect, the feedback signal R=1. While $R = r^c = 1$, ρ increases rapidly. As soon as $\rho > \frac{|\mathbf{x}|}{|\mathbf{A}|}$, r switches to 1, which both halts the increase of ρ and resets the active F_2 node. From one chosen node to the next, ρ decays to slightly below $\frac{|\mathbf{x}|}{|\mathbf{A}|}$ (MT–: Carpenter & Markuzon, 1998). On the time scale of learning ρ returns to $\bar{\rho}$.

Unless they have already activated all their coding nodes, ARTMAP systems contain a reserve of nodes that have never been activated, with weights at their initial values. These *uncommitted* nodes compete with the previously active *committed* nodes, and an uncommitted node will be chosen over poorly matched committed nodes. An ARTMAP design constraint specifies that an active uncommitted node should not reset itself. Weights initially begin with $w_{iJ} = 1$. Thus, when the active node J is uncommitted, $\mathbf{x} = \mathbf{A} \wedge \mathbf{w}_J = \mathbf{A}$ at the match field. Then, $\rho|\mathbf{A}| - |\mathbf{x}| = \rho|\mathbf{A}| - |\mathbf{A}| = (\rho - 1)|\mathbf{A}|$. Thus $\rho|\mathbf{A}| - |\mathbf{x}| \leq 0$ and an uncommitted node does not trigger a reset, provided that $\rho \leq 1$.

V. ART GEOMETRY

ART long-term memories are visualized as hyper-rectangles, called *category boxes*. The weight vector $\mathbf{w}_J$ is interpreted geometrically as a box R_J whose ON-channel corner $\mathbf{u}_J$ and OFF-channel corner $\mathbf{v}_J$ are, in the format of the complement-coded input vector, defined by $\left(\mathbf{u}_J \mid \mathbf{v}_J^C\right) \equiv \mathbf{w}_J$ (Figure 4). For fuzzy ART with the choice-by-difference $F_0 \rightarrow F_2$ signal function T_J (Carpenter & Gjaja, 1994), an input **a** activates the node J of the closest category box R_J, according to the L_1 (city-block) metric. In case of a tie, as when **a** lies in more than one box, the node with the smallest R_J is chosen, where $|R_J|$ is defined as the sum of the edge lengths $\sum_{i=1}^{M}|v_{iJ}-u_{iJ}|$. The chosen node J will reset if $|R_J \oplus \mathbf{a}| > M(1-\rho)$, where $R_J \oplus \mathbf{a}$ is the smallest box enclosing both R_J and **a**. Otherwise, R_J expands toward $R_J \oplus \mathbf{a}$ during learning. With fast learning, $R_J^{new} = R_J^{old} \oplus \mathbf{a}$.

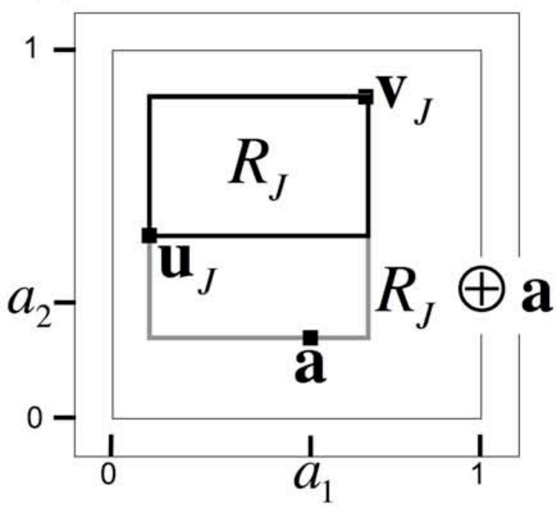

Figure 4. Fuzzy ART geometry. The weight of a category node J is represented in complement-coding form as $\mathbf{w}_J = \left(\mathbf{u}_J \mid \mathbf{v}_J^C\right)$, and the M-dimensional vectors $\mathbf{u}_J$ and $\mathbf{v}_J$ define the corners of the category box R_J. When M=2, the size of R_J equals its width plus its height. During learning, R_J expands toward $R_J \oplus \mathbf{a}$, defined as the smallest box enclosing both R_J and **a**. Node J will reset before learning if $|R_J \oplus \mathbf{a}| > M(1-\rho)$.

VI. BIASING AGAINST PREVIOUSLY ACTIVE CATEGORY NODES AND PREVIOUSLY ATTENDED FEATURES DURING ATTENTIVE MEMORY SEARCH

Activity **x** at the ART field F_1 continuously computes the match between the field's bottom-up and top-down input patterns. A reset signal r shuts off the active F_2 node J when **x** fails to meet the matching criterion determined by the value of the vigilance parameter ρ. Reset alone does not, however, trigger a search for a different F_2 node: unless the prior activation has left an enduring trace within the F_0-to-F_2 subsystem, the network will simply reactivate the same node as before. As modeled in ART 3 (Carpenter & Grossberg, 1990), biasing the bottom-up input to the coding field F_2 to favor previously inactive nodes implements search by allowing the network to activate a new node in response to a reset signal. The ART 3 search mechanism defines a medium-term memory (MTM) in the F_0-to-F_2 adaptive filter which biases the system against re-choosing a node that had just produced a reset. A presynaptic interpretation of this bias is transmitter depletion, or habituation (Figure 5).

Medium-term memory in all ART models allows the network to shift attention among learned categories at the coding field F_2 during search. The new biased ART network (Carpenter & Gaddam, 2009) introduces a second medium-term memory that shifts attention among input features, as well as categories, during search.

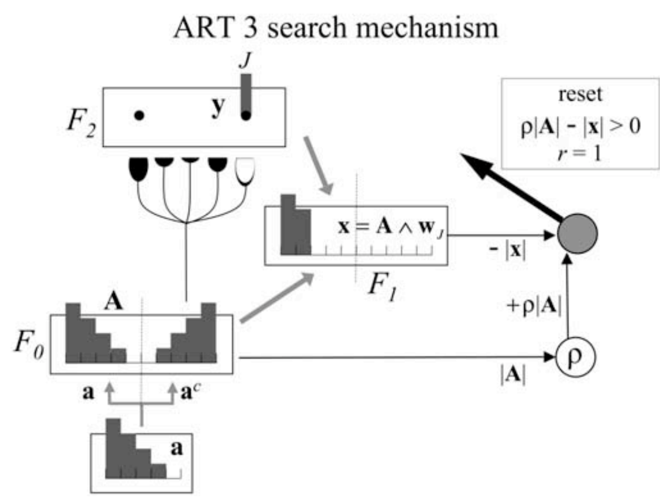

Figure 5. ART 3 search implements a medium-term memory within the F_0-to-F_2 pathways, which biases the system against choosing a category node that had just produced a reset.

REFERENCES

Carpenter G. A. (1997). Distributed learning, recognition, and prediction by ART and ARTMAP neural networks. *Neural Networks*, *10*, 1473–1494.

Carpenter, G. A. (2003). Default ARTMAP. *Proceedings of the international joint conference on neural networks (IJCNN'03)*, 1396–1401.

Carpenter, G.A., & Gaddam, S.C. (2009). Biased ART: A neural architecture that shifts attention toward previously disregarded features following an incorrect prediction. CAS/CNS Technical Report.

Carpenter, G.A., & Gjaja, M.N. (1994). Fuzzy ART choice functions. *Proceedings of the World Congress on Neural Networks (WCNN-94)*, Hillsdale, NJ: Lawrence Erlbaum Associates, **I**-713-722.

Carpenter G. A., & Grossberg S. (1987). 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. (1990). ART 3: Hierarchical search using chemical transmitters in self-organizing pattern recognition architectures. *Neural Networks*, *4*, 129-152.

Carpenter G. A., Grossberg S., Markuzon N., Reynolds J. H., & Rosen D. B. (1992). Fuzzy ARTMAP: A neural network architecture for incremental supervised learning of analog multidimensional maps. *IEEE Transactions on Neural Networks*, *3*, 698–713.

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

Carpenter, G.A., & N. Markuzon, N. (1998). ARTMAP-IC and medical diagnosis: Instance counting and inconsistent cases. *Neural Networks*, *11*, 323-336.

Carpenter, G.A., Martens, S., & Ogas, O.J. (2005). Self-organizing information fusion and hierarchical knowledge discovery: a new framework using ARTMAP neural networks. *Neural Networks*, *18*(3), 287-295.

Carpenter, G.A., Milenova, B.L., & Noeske, B.W. (1998). Distributed ARTMAP: a neural network for fast distributed supervised learning. *Neural Networks, 11*, 793-813.

Carpenter, G.A., & Ravindran, A. (2008). Unifying multiple knowledge domains using the ARTMAP information fusion system. *Proceedings of the 11th International Conference on Information Fusion*, Cologne, Germany, June 30 – July 3, 2008.

Carpenter. G.A., & Ross, W.D. (1995). ART-EMAP: A neural network architecture for object recognition by evidence accumulation. *IEEE Transactions on Neural Networks*, *6*, 805-818, 1995.

Caudell, T.P., Smith, S.D.G., Escobedo, R., & Anderson, M. (1994). NIRS: Large scale ART 1 neural architectures for engineering design retrieval. *Neural Networks*, *7*, 1339-1350.

Grossberg, S. (1999). The link between brain, learning, attention, and consciousness. *Consciousness and Cognition*, *8*, 1-44.

Grossberg, S. (2003). How does the cerebral cortex work? Development, learning, attention, and 3D vision by laminar circuits of visual cortex. *Behavioral and Cognitive Neuroscience Reviews*, *2*, 47-76.

Hurvich, L. M., & Jameson, D. (1957). An opponent-process theory of color vision. *Psychological Review*, *64*, 384-390.

Lisboa, P. (2001). Industrial use of saftey-related artificial neural netoworks. *Contract Research Report 327/2001,* Liverpool John Moores University.

Parsons, O., & Carpenter, G. A. (2003). ARTMAP neural networks for information fusion and data mining: map production and target recognition methodologies. *Neural Networks*, *16*(7), 1075-1089.

Schiller, P.H. (1982). Central connections of the retinal ON and OFF pathways. *Nature*, *297*, 580-583.

Williamson, J.R. (1998). Gaussian ARTMAP: A neural network for fast incremental learning of noisy multidimensional maps. *Neural Networks*, *9*, 881-897.