Fusion ARTMAP: A Neural Network Architecture for Multi-channel Data Fusion and Classification

Yousif R. Asfour[1], Gail A. Carpenter[2], Stephen Grossberg[3], and Gregory W. Lesher[4]

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

Abstract

Fusion ARTMAP is a self-organizing neural network architecture for multi-channel, or multi-sensor, data fusion. Single-channel Fusion ARTMAP is functionally equivalent to Fuzzy ART during unsupervised learning and to Fuzzy ARTMAP during supervised learning. The network has a symmetric organization such that each channel can be dynamically configured to serve as either a data input or a teaching input to the system. An ART module forms a compressed recognition code within each channel. These codes, in turn, become inputs to a single ART system that organizes the global recognition code. When a predictive error occurs, a process called parallel match tracking simultaneously raises vigilances in multiple ART modules until reset is triggered in one of them. Parallel match tracking hereby resets only that portion of the recognition code with the poorest match, or minimum predictive confidence. This internally controlled selective reset process is a type of credit assignment that creates a parsimoniously connected learned network. Fusion ARTMAP's multi-channel coding is illustrated by simulations of the Quadruped Mammal database.

Multi-Channel Data Fusion

A variety of pattern recognition applications require a system to fuse input data from multiple independent channels or sensors. One straightforward approach to this problem is vector concatenation. That is, inputs from each channel are joined to form one large vector that then becomes the input to a single-channel supervised learning system. This approach is used, for example, by Chu and Aggarwal (1992) to train a back propagation system on inputs from multiple sensors. One problem with concatenation is that network connectivity tends to grow multiplicatively with the size of the input vector.

Fusion ARTMAP uses the multi-channel structure of the input data to streamline the network design. One intra-channel code can contribute to several global codes, leading to reduced network connectivity. In addition, teacher and data input channels are dynamically defined via gain control, so each channel can play either role at different times (Figure 1a). Gain control also allows the system to function correctly even if input data to certain channels is missing at various times.

[1]Supported in part by DARPA (ONR N00014-92-J-401J), the National Science Foundation (NSF IRI 90-00530), and the Office of Naval Research (ONR N00014-91-J-4100).

[2]Supported in part by British Petroleum (BP 89-A-1204), DARPA (AFOSR 90-0083 and ONR N00014-92-J-4015), the National Science Foundation (NSF IRI-90-00530), and the Office of Naval Research (ONR N00014-91-J-4100).

[3]Supported in part by DARPA (ONR N00014-92-J-401J), the National Science Foundation (NSF IRI-90-24877), and the Office of Naval Research (ONR N00014-91-J-4100).

[4]Supported in part by the Air Force Office of Scientific Research (AFOSR F49620-92-J-0334), a National Science Foundation Graduate Fellowship, and the Office of Naval Research (ONR N0014-91-J-4100).

During unsupervised learning of single-channel inputs, Fusion ARTMAP is functionally equivalent to ART1 (Carpenter and Grossberg, 1987) for binary inputs and to Fuzzy ART (Carpenter, Grossberg, and Rosen, 1991) for analog inputs. During supervised learning of single-channel signal and teaching inputs, Fusion ARTMAP is functionally equivalent to ARTMAP (Carpenter, Grossberg, and Reynolds, 1991) for binary inputs and to Fuzzy ARTMAP (Carpenter, Grossberg, Markuzon, Reynolds, and Rosen, 1992) for analog inputs, as illustrated with a simulation (Circle-in-the-Square) below.

Parallel Match Tracking

Before Fusion ARTMAP activates a global recognition code, input to each channel activates a compressed recognition code in that channel's own Fuzzy ART module. Then, one global ART module, which receives compressed categorical input from each channel separately, organizes the multi-channel recognition code. The global ART system internally controls code formation via a nonspecific feedback signal sent in parallel to the ART systems of individual channels. This process is called *parallel match tracking* because it generalizes ARTMAP match tracking (Carpenter, Grossberg, and Reynolds, 1991), as follows.

In ARTMAP or Fuzzy ARTMAP (Figure 1b), match tracking implements internal dynamic control of search and reset when an input to ART_a makes an erroneous prediction at ART_b. In ART systems, search is triggered when an active input pattern fails to match the top-down expectation, or prototype, of an active category according to a matching criterion that is defined in terms of a dimensionless parameter called *vigilance* (Carpenter and Grossberg, 1987). In the Fuzzy ART module ART_a, activity $\mathbf{x}^a$ at the field F_1^a is determined by the match between a bottom-up input ($\mathbf{A}$) and a top-down prototype from F_2^a. The degree of match is defined by the ratio $|\mathbf{x}^a|/|\mathbf{A}|$. This ratio is small when the top-down and bottom-up inputs to F_1^a are poorly matched. Search for another ART_a code in F_2^a is initiated at an *orienting subsystem* when:

$$\frac{|\mathbf{x}^a|}{|\mathbf{A}|} < \rho_a, \tag{1}$$

where ρ_a is the ART_a vigilance parameter. In an isolated ART module, vigilance is an independent parameter. In ARTMAP, ρ_a is a variable that is internally controlled via match tracking. Initially, ρ_a equals a *baseline vigilance* $\overline{\rho}_a$ that is typically kept low to maximize code compression. When a predictive error occurs at ART_b, ρ_a is increased just enough to violate (1) and thereby cause reset of F_2^a. Hence the term match tracking, since ρ_a tracks the F_1^a match ratio. The resulting ART_a search leads either to activation of a different F_2^a code that makes the correct prediction at ART_b, or to the formation of a new F_2^a category that then learns the correct ART_b prediction.

In Fusion ARTMAP, parallel match tracking simultaneously raises the vigilances of multiple ART modules (Figure 2). A search is thereby triggered in just one of the modules. By (1), that module has the poorest match between bottom-up input and top-down prototype. It is hereby judged by the system to be the most likely source of the predictive error. Search activates a new code in that module alone, preserving other portions of the previously active pattern. This process of credit assignment efficiently shares code subsets across categories in the learned network, since predictively effective channels are not reset to correct errors caused by ineffective channels. Fusion ARTMAP thus creates more parsimonious codes, with fewer paths and weights, than would be needed by single-channel recognition systems. Connectivity of single-channel (ARTMAP) and multi-channel (Fusion ARTMAP) systems are illustrated below with simulations of the Quadruped Mammal database (Ginnari, 1992). The importance of sparse network connectivity increases multiplicatively with the dimension of the input vectors.

Circle-in-the-Square Simulations

A single-channel Fusion ARTMAP system was trained to recognize whether a point within a unit square was inside or outside a circle of one-half unit area. The results of the simulations were compared with benchmark Fuzzy ARTMAP simulations (Carpenter *et al.*, 1992). The performance of the two systems was identical, as expected. However, the two systems differed in terms of the total number of modifiable connections and in terms of the fan-in and fan-out at each node. Fusion ARTMAP produced more ART_{ab} category nodes than did Fuzzy ARTMAP at F^{ab} (Figure 1). However, the average fan-in and fan-out at each node in Fusion ARTMAP was significantly less.

Quadruped Mammal Database Simulations

Single-channel and multi-channel Fusion ARTMAP systems were simulated using the Quadruped Mammal database (Ginnari, 1992), which represents four mammals (dog, cat, giraffe, and horse) in terms of eight components (head, tail, four legs, torso, and neck). Each component is described by nine attributes (three location variables, three orientation variables, height, radius, and texture), for a total of 72 attributes. Each attribute is modeled as a Gaussian process with mean and variance depending on mammal and component. For example, the radius of a giraffe's neck is modeled by a different Gaussian from that of a cat's neck.

The first set of simulations configured Fusion ARTMAP to be functionally equivalent to an unsupervised Fuzzy ART system, with the entire attribute vector presented to a single channel, without a teacher. Fusion ARTMAP categorized the inputs into four stable categories corresponding to the four mammals.

The next set of simulations presented each of the eight component vectors to a different ART_a module (Figure 1a), and presented the target animal's identity to ART_b. Fusion ARTMAP achieved 100% prediction rates on both the training and testing sets within a single presentation when 1000 training exemplars were used. The resulting network was compared with that of a single-channel Fusion ARTMAP system trained on the same data sets, except with a merged attribute vector. Performance was identical, but the single-channel case required about 1.5 times as many path connections and weights as the multi-channel case.

References

Carpenter, G.A. and 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., Markuzon, N., Reynolds, J.H., and 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., 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.

Chu, C.C. and Aggarwal, J.K. (1992) Image interpolation using multiple sensing modalities. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **14**, 840-846.

Ginnari, J.H. (1992) Quadruped mammals. *UCI Repository of machine learning databases* (Machine-readable data repository). Department of Information and Computer Science, University of California, Irvine, CA.

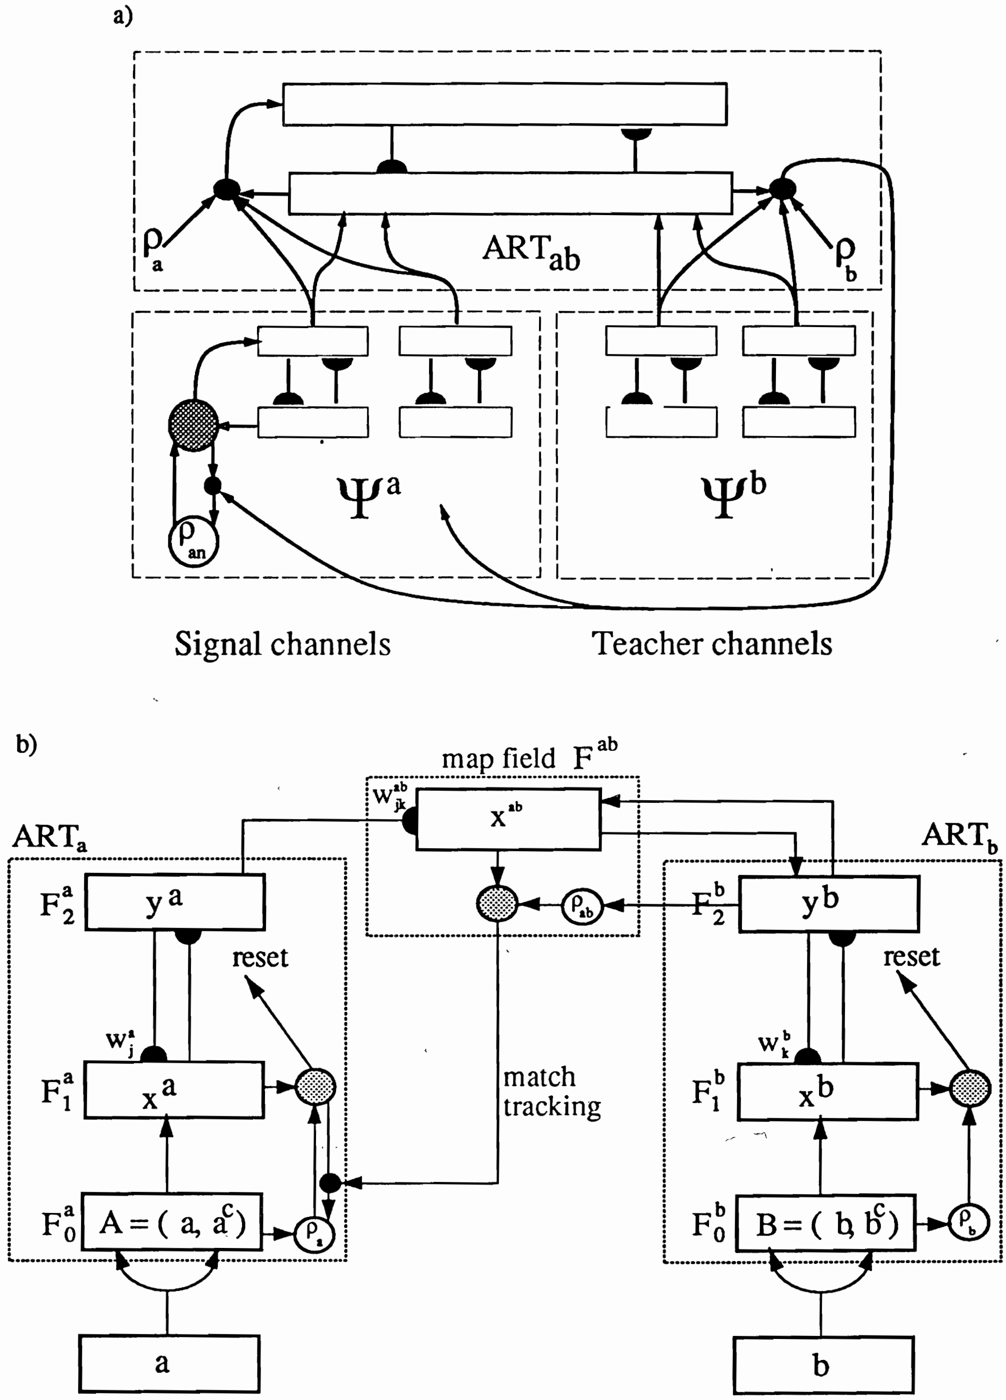

Figure 1: (a) Fusion ARTMAP generalizes Fuzzy ARTMAP, learning multi-channel maps from one dynamically configured subset of the input space to another. (b) During supervised learning, Fuzzy ARTMAP learns a predictive single-channel map from signal to ART_a to teaching inputs to ART_b (Carpenter *et al.*, 1992).

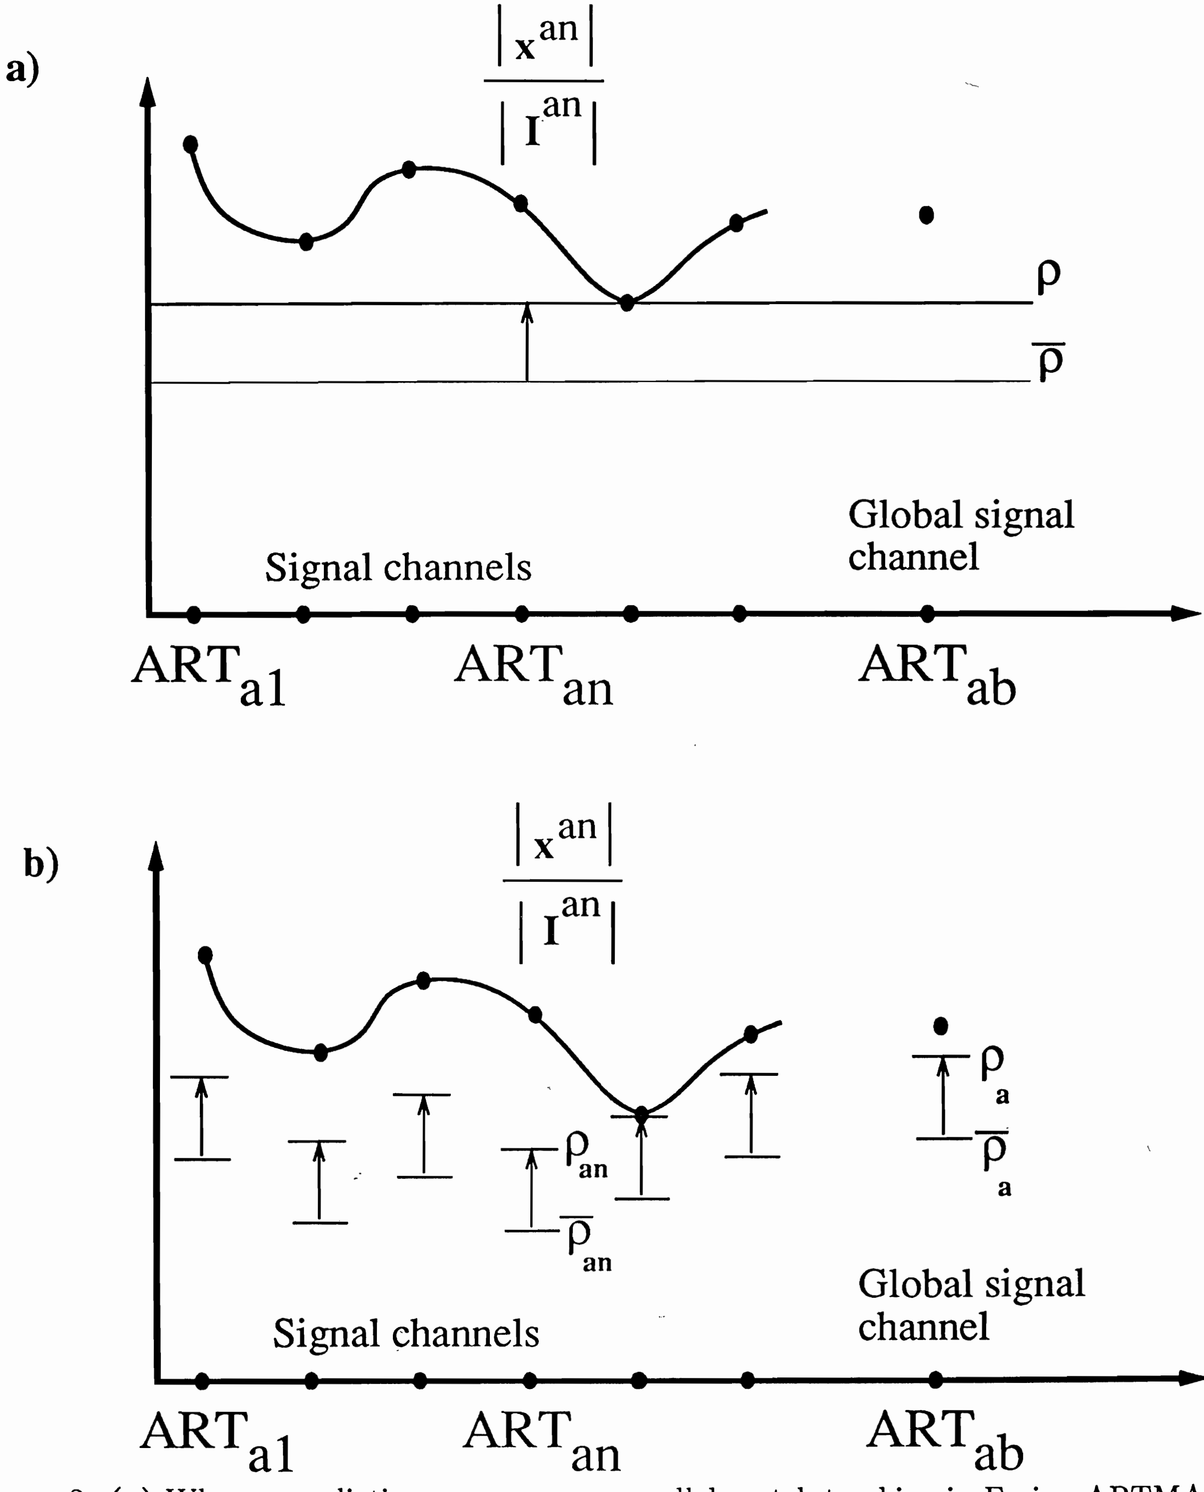

Figure 2: (a) When a predictive error occurs, parallel match tracking in Fusion ARTMAP raises multiple vigilance values simultaneously until reset occurs in the ART module most likely to have caused the error. (b) Parallel match tracking can simultaneously raise vigilances in independent Fusion ARTMAP modules each with its own baseline matching criterion.