FUSION ARTMAP: AN ADAPTIVE FUZZY NETWORK FOR MULTI-CHANNEL CLASSIFICATION

Yousif R. Asfour, Gail A. Carpenter, Stephen Grossberg,
and Gregory W. Lesher

August, 1993

Technical Report CAS/CNS-93-052

Fusion ARTMAP:
An Adaptive Fuzzy Network for Multi-channel Classification

Yousif R. Asfour*, Gail A. Carpenter †, Stephen Grossberg ‡, Gregory W. Lesher §
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. Fusion ARTMAP generalizes the fuzzy ARTMAP architecture in order to adaptively classify multi-channel data. 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.

I. MULTI-CHANNEL DATA FUSION

Fusion ARTMAP is a neural network architec-

ture designed to adaptively classify objects using multiple sources of information, regardless of its source or type. An example of the fusion problem is the classification of trucks based on inputs from different types of sensors such as range, doppler, and camera. Alternatively, multiple input sources could represent different views of the truck, such as top, front, and side views. Trucks can also be classified using different spatial scales by combining information from cameras that zoom in on the tires and information from cameras that provide a view of the whole truck. In general, Fusion ARTMAP is designed to classify objects using information from multiple sources of any type.

One straightforward approach to the fusion 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 [7] to train a back-propagation neural network on inputs from infrared, range, and visual sensors.

Whenever the classifier makes a wrong prediction during training, it is desirable to modify some system parameters in order to improve the total system performance. Deciding which parameters to modify is known as the the *credit assignment problem*. Since the information from the different sensors is concatenated into a single feature vector, the predictive power of each individual sensor is unknown to the classifier. Therefore, the credit assignment problem is solved by assigning blame non-specifically to all input channels. Failure to account for the individual channels' predictive power leads to connectivity that tends to grow multiplicatively with the size of the input vector.

Fusion ARTMAP utilizes a modular approach to sensor fusion. Each sensor is assigned an indi-

vidual classifier, the outputs of which serve as the inputs to a global classifier which makes a global prediction. For example, information from a range sensor is first classified into depth codes while information from a doppler sensor is classified into speed codes. The compressed depth and speed codes become inputs to a global classifier, which predicts the type of truck. By assigning an individual classifier to each sensor channel, blame can be assigned selectively to the channels with lowest predictive confidence. Such an approach retains system predictive accuracy while reducing total network connectivity by maximizing compression within each channel.

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 the role of a teacher or the role of an input at different times. Gain control also allows the system to function correctly even if input data to certain channels is missing at various times. Thus, faulty sensors may be deleted or new sensors added as the need arises.

II. FUZZY ARTMAP:
A FUSION BUILDING BLOCK

Fuzzy ARTMAP is a supervised neural network classifier that learns to classify inputs by a fuzzy set of features, or a pattern of fuzzy membership values between 0 and 1 indicating the extent to which each feature is present. Fuzzy ARTMAP differs from many other fuzzy pattern recognition algorithms [2],[9] in that it learns each input as it is received on-line, rather than by performing an off-line optimization of a criterion function.

Each fuzzy ARTMAP system consists of a pair of fuzzy ART classifiers (ART_a, and ART_b) that create stable recognition categories in response to arbitrary sequences of input patterns (Fig. 1). During supervised learning, ART_a receives a stream $a^{(p)}$ of input patterns, and ART_b receives a stream $b^{(p)}$ of input patterns, where $b^{(p)}$ is the correct prediction given $a^{(p)}$. These modules are linked by an associative learning network and an internal controller that ensures autonomous system operation in real time. The controller is designed to create the minimal number of ART_a recognition categories, or "hidden units", needed to meet accuracy criteria. It does this by realizing a minimax learning rule that enables the fuzzy ARTMAP system to

learn quickly, efficiently, and accurately as it conjointly *minimizes* predictive error and *maximizes* predictive generalization. This scheme automatically links predictive success to category size on a trial-by-trial basis using only local operations. It works by increasing the vigilance parameter ρ_a of ART_a by the minimal amount needed to correct a predictive error at ART_b.

When the ART_a classifier is presented with an input vector a, the bottom-up activation from F_1^a causes the F_2^a layer to choose a category node based on the input's fuzzy membership in that category's fuzzy set. The chosen category then sends information back to the F_1^a layer which is compared to the input vector a. The fuzzy intersection of top-down activation with the input vector produces a *match value* that indicates the classifier's confidence in its category choice. Parameter ρ_a calibrates the minimum confidence that ART_a must have in a recognition category, or hypothesis, activated by an input a^p in order for ART_a to accept that category, rather than search for a better one through an automatically controlled process of hypothesis testing. Lower values of ρ_a enable larger categories to form leading to broader generalization and higher code compression. A predictive failure at ART_b increases ρ_a by the minimum amount needed to trigger hypothesis testing at ART_a, using a mechanism called *match tracking* [5]. Match tracking sacrifices the minimum amount of generalization necessary to correct a predictive error. Hypothesis testing leads to the selection of a new ART_a category, which focuses attention on a new cluster of $a^{(p)}$ input features that is better able to predict $b^{(p)}$.

Fuzzy ARTMAP can itself be used for multi-sensor fusion, by concatenating the information from all sensors into a single input vector. However, whenever a predictive error occurs during training, the match tracking signal resets the ART_a classifier without regard to the predictive confidence in the individual channel information.

III. FUSION ARTMAP
GENERALIZES FUZZY ARTMAP

Fusion ARTMAP extends the fuzzy ARTMAP classifier by incorporating an individual sensor classifier for each input channel, and by extending the match tracking technique in a manner that assigns blame only to the channels with least confidence in their predictions (Fig. 2).

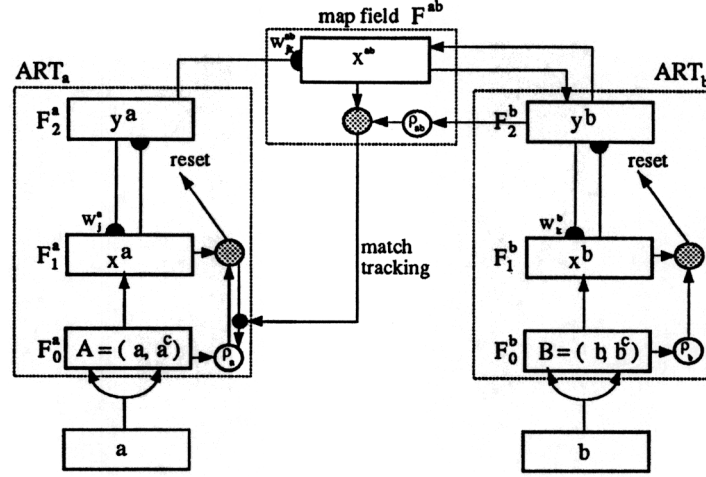Before a global recognition code is activated in

Figure 1: Fuzzy ARTMAP architecture. The ART_a complement coding preprocessor transforms the vector **a** into the vector $\mathbf{A} = (\mathbf{a}, \mathbf{a}^c)$ at the ART_a field F_0^a. **A** is the input to the ART_a field F_1^a. Similarly, in the supervised mode, the input to the ART_b field F_1^b is the vector $(\mathbf{b}, \mathbf{b}^c)$. When a prediction by ART_a is disconfirmed at ART_b, inhibition of the map field F^{ab} induces the match tracking process. Match tracking raises the ART_a vigilance (ρ_a) to just above the ART_a match value $|\mathbf{x}^a|/|\mathbf{A}|$. This triggers an ART_a search, which leads to activation of either an ART_a category that correctly predicts **b** or to a previously uncommitted ART_a category node.
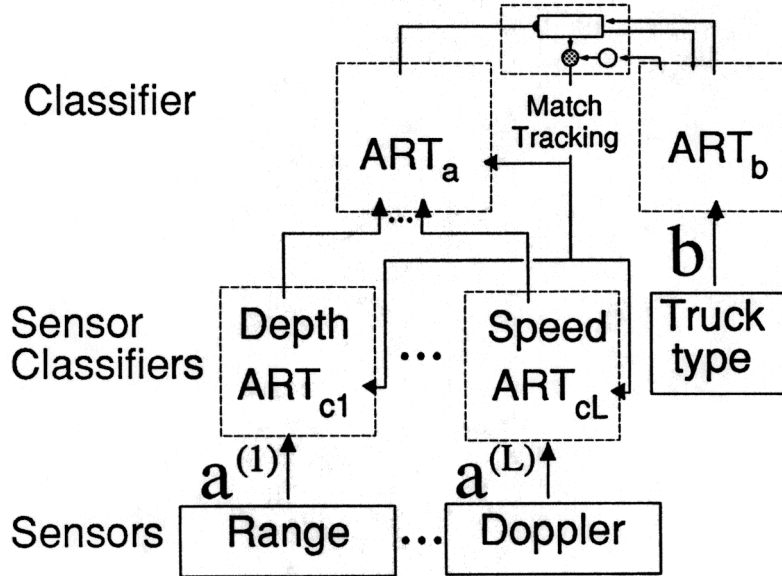

Figure 2: Fusion ARTMAP associates a single ART classifier to each input sensor. The outputs of these classifiers are used as inputs to a fuzzy ARTMAP system that makes a global prediction. Parallel match tracking raises the vigilance of all sensor classifiers simultaneously until the module with the least predictive confidence is reset.

3

Fusion ARTMAP, input to each channel activates a compressed recognition code in that channel's own fuzzy ART module. Then, one global fuzzy ARTMAP module, which receives compressed categorical input from each channel separately, organizes the multi-channel recognition code. The fuzzy ARTMAP system internally controls code formation via a nonspecific feedback signal sent in parallel to the ART systems of individual channels. This control process, called *parallel match tracking*, generalizes ARTMAP match tracking.

In Fusion ARTMAP, parallel match tracking simultaneously raises the vigilances of multiple sensor ART modules. A search is thereby triggered in just the one module that has the poorest match between bottom-up input and top-down prototype. It is hereby judged to be the most likely source of the predictive error and is defined to be the channel with the least predictive confidence. Search activates a new code in that module alone, preserving other input channel categories 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 unnecessarily reset to correct errors caused by ineffective channels. Fusion ARTMAP thus creates more parsimonious codes, than single-channel recognition systems, with fewer paths and weights.

IV. INTRODUCING SYMMETRY

Fusion ARTMAP can be generalized by replacing the inter-ART map field F^{ab} with a modified global ART module (Fig. 3). The outputs of all the sensor and teacher classifiers are used as inputs to a global ART module that self-organizes its inputs into stable categories.

Channels are designated to input or teacher status by a set of parametric biases. Changing the bias on a particular channel can change its function from that of an input sensor to that of a target teacher. This symmetrical approach allows the use of multiple teacher channels.

System analysis shows that the generalized symmetric Fusion ARTMAP architecture reduces functionally to the system described in section III in the case of a "fixed single teacher channel" with multiple input sensors.

V. QUADRUPED MAMMAL DATABASE SIMULATIONS

Single-channel fuzzy ARTMAP and multi-channel Fusion ARTMAP systems were simulated using the Quadruped Mammal database [8], 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 dependent on the 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 or a giraffe's tail.

In the first set of simulations, Fusion ARTMAP was configured 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 was allowed to self-organize the input vectors in categories. Fusion ARTMAP categorized the inputs into four stable categories corresponding to the four mammals.

In the next set of simulations, each of the eight component vectors was presented to a different ART_a module (Fig. 2), and the target animal's identity was presented 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 fuzzy 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

[1] Y.R. Asfour, G.A Carpenter, S. Grossberg, and G.W. Lesher, "Fusion ARTMAP: A neural network architecture for multi-channel data fusion and classification." *WCNN Conference Proceedings*, vol. 2, pp.210-215, July 1993.
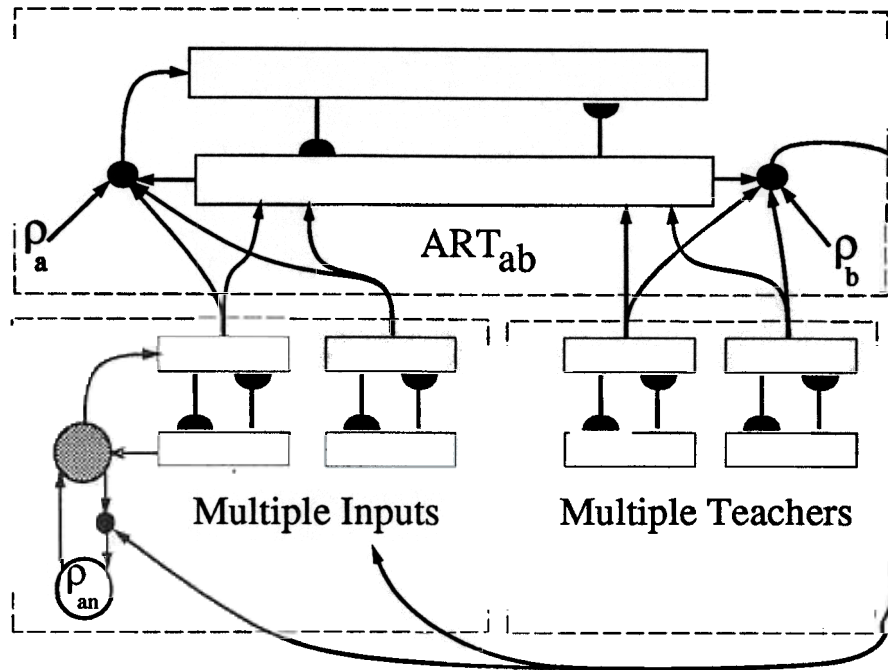
Figure 3: Replacing the inter-ART map field with a global ART module introduces symmetry between input and teacher classifiers, and allows multi-sensory fusion using multiple teachers.

[2] J.C. Bezdek, *Pattern Recognition with Fuzzy Objective Function Algorithms.* New York: Plenum Press, 1981.

[3] G.A. Carpenter, S. and Grossberg, S. "A massively parallel architecture for a self-organizing neural pattern recognition machine," *Computer Vision, Graphics, and Image Processing*, vol. 37, pp. 54-115, 1987.

[4] G.A Carpenter, S. Grossberg, N. Markuzon, J.H. Reynolds, and D.B. Rosen, "Fuzzy ARTMAP: A neural network architecture for incremental supervised learning of analog multidimensional maps." *IEEE Transactions on Neural Networks*, vol. 3, pp. 698-713, 1992

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

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

[7] C.C. Chu, and J.K. Aggarwal, "Image interpolation using multiple sensing modalities." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, pp. 840-846.

[8] J.H. Ginnari, "Quadruped mammals." *UCI Repository of machine learning databases* (Machine-readable data repository at ics.uci.edu:/usr2/spool/ftp/pub/machine-learning-databases). Department of Information and Computer Science, University of California, Irvine, CA, 1992.

[9] S. Pal, and D.K. Dutta Majumder, *Fuzzy Mathematical Approach to Pattern Recognition.* New York: Wiley, 1986.

a)

$$\frac{\left|\mathbf{x}^{an}\right|}{\left|\mathbf{I}^{an}\right|}$$

ρ

$\overline{\rho}$

Global signal
channel

Signal channels

ART_{a1} ART_{an} ART_{ab}

b)

$$\frac{\left|\mathbf{x}^{an}\right|}{\left|\mathbf{I}^{an}\right|}$$

ρ_{an}

$\overline{\rho}_{an}$

ρ_a

$\overline{\rho}_a$

Global signal
channel

Signal channels

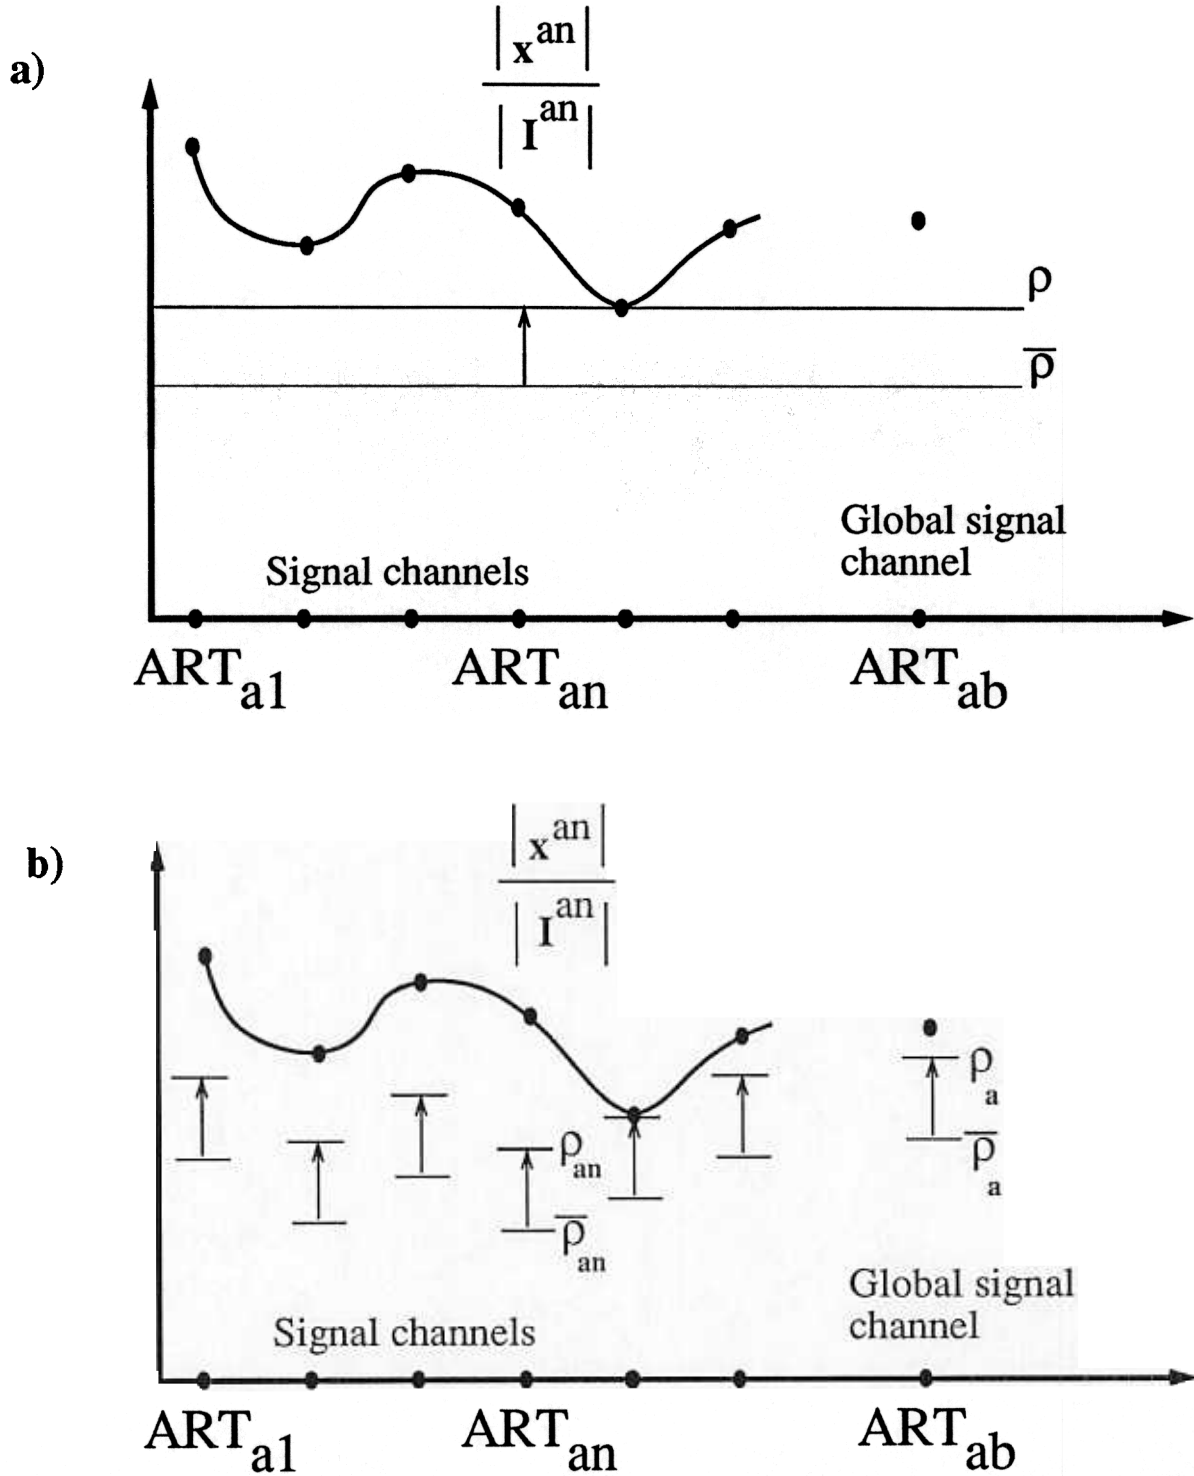ART_{a1} ART_{an} ART_{ab}

Figure 4: (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 initial match criterion value.