

1998-05

ARTMAP-FTR: A Neural Network For Fusion Target Recognition, With Application To Sonar Classification

<https://hdl.handle.net/2144/2350>

"Downloaded from OpenBU. Boston University's institutional repository."

**ARTMAP-FTR: A neural network for fusion target
recognition, with application to sonar classification**

Gail A. Carpenter and William W. Streilein

April 1998

Technical Report CAS/CNS-98-016

Permission to copy without fee all or part of this material is granted provided that: 1. The copies are not made or distributed for direct commercial advantage; 2. the report title, author, document number, and release date appear, and notice is given that copying is by permission of the BOSTON UNIVERSITY CENTER FOR ADAPTIVE SYSTEMS AND DEPARTMENT OF COGNITIVE AND NEURAL SYSTEMS. To copy otherwise, or to republish, requires a fee and / or special permission.

Copyright © 1998

Boston University Center for Adaptive Systems and
Department of Cognitive and Neural Systems
677 Beacon Street
Boston, MA 02215

ARTMAP-FTR: A neural network for fusion target recognition, with application to sonar classification

Gail A. Carpenter and William W. Streilein

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

ABSTRACT

ART (Adaptive Resonance Theory) neural networks for fast, stable learning and prediction have been applied in a variety of areas. Applications include automatic mapping from satellite remote sensing data, machine tool monitoring, medical prediction, digital circuit design, chemical analysis, and robot vision. Supervised ART architectures, called ARTMAP systems, feature internal control mechanisms that create stable recognition categories of optimal size by maximizing code compression while minimizing predictive error in an on-line setting. Special-purpose requirements of various application domains have led to a number of ARTMAP variants, including fuzzy ARTMAP, ART-EMAP, ARTMAP-IC, Gaussian ARTMAP, and distributed ARTMAP. A new ARTMAP variant, called ARTMAP-FTR (fusion target recognition), has been developed for the problem of multi-ping sonar target classification. The development data set, which lists sonar returns from underwater objects, was provided by the Naval Surface Warfare Center (NSWC) Coastal Systems Station (CSS), Dahlgren Division. The ARTMAP-FTR network has proven to be an effective tool for classifying objects from sonar returns. The system also provides a procedure for solving more general sensor fusion problems.

Keywords: sonar classification, sensor fusion, target recognition, ARTMAP-FTR, ART, ARTMAP, fuzzy ARTMAP, adaptive resonance, neural network

1. SONAR TARGET RECOGNITION

A Naval Surface Warfare Center (NSWC) data set records reflections of six types of targets: two mine-like objects, a water-filled drum, a smooth granite rock, a limestone rock, and a water-saturated log with a mine-like shape (Table 1). Each object was suspended in water and rotated during sonar testing, yielding 72 aspects per target for each of seven frequency bands of 40 kHz bandwidth. Researchers at the ORINCON Corporation, under a Phase II SBIR contract from the Office of Naval Research led by Dr. Larry Burton, have used the NSWC data set to develop a neural network system for automated sonar recognition. ORINCON simulation studies focused on the 20-60 kHz band. Sonar measurements were processed using both a standard matched filter and the Spectrogram Correlation And Transformation (SCAT) algorithm,¹ which emulates signal processing in the bat echolocation system. Preprocessing yielded input vectors with components representing 492 time-series points for each echo. Data records were also varied by the addition of simulated reverberations.

Object number	Abbreviation	Description
1	bullet	Bullet-shaped, metallic
2	cone	Truncated-cone-shaped, plastic
3	drum	Water-filled 50-gallon drum
4	limestone	Rough limestone rock
5	granite	Smooth granite rock
6	log	Water-saturated log

Table 1: Objects represented by the NSWC sonar data set. Objects 1 and 2 are considered mine-like and objects 1-3 are man-made.

The ORINCON project has now established a benchmark paradigm for development and comparison of classifier systems for target recognition. The NSWC sonar data set has been tested on several recognition tasks such as man-made vs. non-man-made object discrimination, mine-like vs. non-mine-like discrimination, identification of the man-made objects, and identification of all six objects. ORINCON researchers have examined the recognition capabilities of various neural network systems, particularly multi-layer perceptrons and ellipsoidal basis functions, the latter giving slightly better results. The classifier system protocol trains and tests on alternating aspects (Figure 1).

This paper introduces a neural network architecture that has been developed for data fusion tasks such as those encountered in sonar discrimination problems. The new classification system, called ARTMAP-FTR (ARTMAP - Fusion Target Recognition), is based on the ARTMAP family of neural networks (Section 2). Simulations focus on benchmark paradigms that have shown optimal performance in the ORINCON studies.² In particular, on a set of three-ping fusion tasks with matched filter preprocessing, ARTMAP-FTR performance compares favorably with performance measures reported by the ORINCON group (Section 3). Additional ARTMAP-FTR system studies show that performance improves with increasing numbers of pings; that inter-ping aspects do not need to be spaced at regular intervals for good performance; and that matched filter preprocessing gives substantially better classification results than the SCAT algorithm (Section 4). Section 5 shows how the ARTMAP-FTR fusion system is constructed as a hierarchy of network subsystems.

2. ART AND ARTMAP NEURAL NETWORKS

Researchers at the Boston University Department of Cognitive and Neural Systems / Center for Adaptive Systems (CNS/CAS) have introduced and analyzed the ART (adaptive resonance theory) family of neural network architectures for self-organizing category learning, recognition, and prediction.³⁻¹⁶ Capabilities of these systems include stable incremental learning, if-then rule extraction, and large-scale database interpretation.

This research program is now advancing state-of-the-art engineering, moving from neural network models to application prototypes and fielded systems. Examples of ART technology transfer that have been reported in publications include: a Boeing parts design retrieval system,¹⁷ an autonomous vision system, also being developed at Boeing,¹⁸ robot sensory-motor control,¹⁹⁻²² robot navigation,²³⁻²⁴ active vision,²⁵ 3-D object recognition,²⁶ face recognition,²⁷ medical imaging,²⁸ satellite remote sensing,²⁹⁻³² Macintosh operating system software,³³ automatic target recognition,³⁴⁻³⁷ electrocardiogram classification,³⁸⁻³⁹ air quality monitoring,⁴⁰ weather prediction,⁴¹ strength prediction for concrete mixes,⁴² signature verification,⁴³ decision making and intelligent agents,⁴⁴ document retrieval,⁴⁵ analysis of musical scores,⁴⁶ character classification,⁴⁷⁻⁴⁹ machine condition monitoring and failure forecasting,⁵⁰⁻⁵³ chemical analysis from UV and IR spectra,⁵⁴ multi-sensor chemical analysis,⁵⁵ combinatorial optimization,⁵⁶ detection of cancerous cells,⁵⁷ sorting of recycled materials,⁵⁸ frequency selective surface design for electromagnetic system devices,⁵⁹ and digital circuit design.⁶⁰

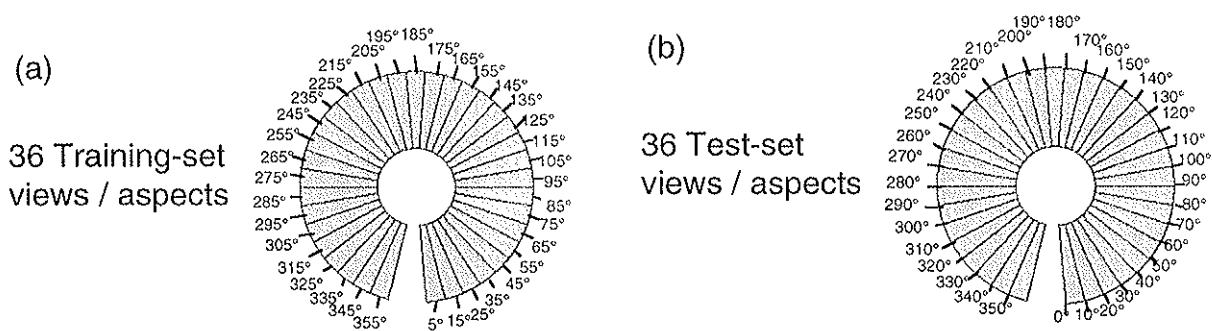


Figure 1: Sonar returns were recorded at 5° intervals for each of the six objects represented in the NSWC data set. (a) The classification training set includes returns at 10° intervals, beginning with the 5° aspect. (b) System performance was evaluated on returns from the remaining 36 aspects: 0°, 10°, ... 350°.

The ARTMAP-FTR sensor fusion system is outlined in Section 5 below.

3. THREE-PING SONAR RECOGNITION BY AN ARTMAP-FTR SYSTEM

The three-ping sonar identification problem presents a classifier system with three sonar returns at 30° aspect separation. Figure 2 illustrates a typical aspect set for this paradigm. Returns are preprocessed to produce matched filtered and threshold-centered data, which serve as classifier inputs. This benchmark problem facilitates comparisons between different classifiers applied to a set of common tasks. In particular, ORINCON Phase II results are reported largely for this paradigm, which gave the best overall performance in their studies.

Table 2 lists the rate of correct classification by an ARTMAP-FTR neural network on the three-ping sonar identification problem. These results compare favorably with optimal performance reported by ORINCON researchers. For example, in the ORINCON ONR SBIR Phase II final report, the best six-class recognition rate is 87.0%, using an ellipsoidal basis function network.^{2 (p. 29)} On the same task with the same preprocessing steps applied to the input data, ARTMAP-FTR performance was 91.6% (Table 2).




Figure 2: For three-ping fusion with 30° aspect separation, test-set input trials present the classifier system with returns from aspects {0°, 30°, 60°}, {10°, 40°, 70°} ... {350°, 20°, 50°}.

3-PING SONAR TASK 30° SEPARATION	ARTMAP-FTR Pcc
Man-made vs. non-man-made objects {1,2,3} vs. {4,5,6}	93.1%
Mine-like vs. non-mine-like objects {1,2} vs. {3,4,5,6}	93.1%
Discrimination among man-made objects {1},{2},{3}	100%
Discrimination among non-man-made objects {4},{5},{6}	97.2%
Six-class discrimination {1},{2},{3},{4},{5},{6}	91.6%

Table 2: ARTMAP-FTR test-set percent correct classification (Pcc) rate on three-ping sonar tasks.

All ARTMAP-FTR performance rates are reported for a test set that is not used in the training process. The network requires the selection of one free parameter, vigilance, which determines the degree of cluster granularity and code compression. Some studies also require a choice of decision threshold. All operating parameters for testing are chosen using a validation subset of the training set (Section 5.4). Sections 3.1 - 3.3 further describe the three-ping simulation studies.

3.1. Man-made vs. non-man-made object discrimination

A system making a two-class decision can vary the test-set output mix by varying an output decision threshold. With a threshold $\gamma \in [0,1]$, the system chooses class k if the fraction σ_k of the system output favoring that class is greater than γ . For the man-made / non-man-made task, setting $\gamma = 0.0$ causes system output for 100% of all man-made objects to meet the threshold criterion (Pcc = 1.0), but causes 100% of all non-man-made objects to meet the criterion as well (Pfc=1.0). In Figure 3, the case $\gamma = 0.0$ corresponds to the point in the upper right-hand corner of the graph, while the case $\gamma = 1.0$ corresponds to the point in the lower left-hand corner of the graph. As γ increases from 0 to 1, the graph plots the percent correct classification rate as a function of the false classification rate. An ideal point would lie in the upper left-hand corner. For the ARTMAP-FTR classifier, setting $\gamma = 0.5$ produces correct classifications of 93.1% of the man-made targets, with false classifications for 6.9% of non-man-made targets.




Figure 3: As the output decision threshold γ varies from 0 to 1, a graph plots the correct classification rate (Pcc) as a function of the rate at which objects are falsely classified (Pfc). For $\gamma = 0.5$, Pcc = 93.1% and Pfc = 6.9%. The optimal ARTMAP-FTR vigilance parameter (ρ^*) was determined using a validation subset of the training set.

		Actual class		
		Man-made	Non-man-made	Totals
Predicted class	Man-made	102	9	111
	Non-man-made	6	99	105
Totals		108	108	201
				93.1%

Table 3: ARTMAP-FTR confusion matrix for the man-made / non-man-made discrimination task.

Table 3 and Figure 4 provide more detailed information about the nature of ARTMAP-FTR errors on the man-made / non-man-made discrimination task. The test set includes records of 36 views for each of 6 objects (Figure 1b). Thus a total of 108 three-ping inputs (Figure 2) correspond to man-made objects and 108 to non-man-made objects. The confusion matrix (Table 3) indicates, for example, that 6 man-made inputs are mistakenly identified as coming from non-man-made objects. Figure 4 shows the exact locations of each error. For example, the aspect set $\{10^\circ, 40^\circ, 70^\circ\}$ from object 1 (the mine-like bullet-shaped metallic object) is mistakenly classified as from a non-man-made object. Similarly, the inputs with initial aspects 50° and 70° from object 1 and inputs with initial aspects $170^\circ, 180^\circ,$ and 210° from object 2 (the mine-like truncated-cone-shaped plastic object) are also labeled non-man-made. The other 102 three-ping sets from objects 1, 2, and 3 are correctly classified as man-made. Table 3 also indicates that 9 inputs from non-man-made objects are misclassified. Figure 4 shows that all inputs from objects 4 and 6 are correctly classified, with the 9 errors coming from object 5 (the smooth granite rock), from the inputs with initial aspects $50^\circ, 60^\circ, 70^\circ, 100^\circ, 220^\circ, 240^\circ, 250^\circ, 290^\circ,$ and 320° .

3.2. Mine-like vs. non-mine-like object discrimination

A confusion matrix (Table 4) shows error patterns for the mine-like / non-mine-like discrimination task. ARTMAP-FTR performance on this task is identical to performance on the man-made / non-man-made task. This similarity is due to the fact that the man-made but non-mine-like object (3) gives 100% correct discrimination (Figure 4). Thus the error patterns for the two tasks are the same.




Figure 4: Test-set error locations for the three-ping fusion, man-made / non-man-made discrimination task. Dark areas show the locations of man-made predictions for each of the six objects from each initial test-set aspect (key). Three inputs from object 1 and three from object 2 are mislabeled as non-man-made; and nine inputs from object 5 are mislabeled as man-made. The remaining 201 inputs (93.1%) are correctly classified.

		Actual class		
		Mine-like	Non-mine-like	Totals
Predicted class	Mine-like	66	9	75
	Non-mine-like	6	135	141
Totals		72	144	201
				93.1%

Table 4: ARTMAP-FTR confusion matrix for the mine-like / non-mine-like discrimination task.

3.3. Three-class and six-class object discrimination

The ARTMAP-FTR system is able to perform six-class object discrimination nearly as well as it performs two-class discrimination. This is because identification is nearly perfect within the subsets of man-made and non-man-made objects. In fact, a network trained to discriminate among the man-made object {1,2,3} achieves 100% correct test-set performance (Table 5a); and a network trained on the non-man-made objects {4, 5, 6} achieves 97.2% correct test-set performance (Table 5b).

The six-class object recognition task is performed in two processing stages. First, an ARTMAP-FTR network is trained to predict whether returns are from man-made or non-man-made objects (Table 3). Depending on the outcome, the input is then presented either to a network trained only on the man-made objects {1,2,3} (Table 5a) or to a network trained only on the non-man-made objects {4,5,6} (Table 5b). The six-class test-set confusion matrix (Table 6) includes the 15 man-made / non-man-made errors introduced at stage 1 (Table 3), plus three more errors from stage 2 (Table 5b), giving an overall six-class discrimination rate of 91.6%. Figure 5 shows the aspect locations of each confusion error listed in Table 6. Compare this gray-scale six-class map with the man-made / non-man-made aspect map (Figure 4).

(a) Man-made objects		Actual class			
		1 - bullet	2 - cone	3 - drum	Totals
Predicted class	1 - bullet	36			36
	2 - cone		36		36
	3 - drum			36	36
	Totals	36	36	36	108
					100%

(b) Non-man-made objects		Actual class			
		4 - limestone	5 - granite	6 - log	Totals
Predicted class	4 - limestone	33			33
	5 - granite		36		36
	6 - log	3		36	39
	Totals	36	36	36	105
					97.2%

Table 5: ARTMAP-FTR test-set confusion matrices after separate three-ping training of (a) the man-made objects {1,2,3} and (b) the non-man-made objects {4,5,6}.

4. ADDITIONAL SONAR TARGET RECOGNITION STUDIES

All ARTMAP-FTR results reported in Section 3 focus on the benchmark three-ping fusion problem with interping aspect separations always set equal to 30°. Additional studies that varied this set of tasks indicate how changes in the training paradigm affect recognition accuracy. These results indicate the robustness of ARTMAP-FTR network performance.

6-class discrimination		Actual class						
		1-bullet	2-cone	3-drum	4 - limestone	5 - granite	6 - log	Totals
Predicted class	1 - bullet	33						33
	2 - cone		33			9		42
	3 - drum			36				36
	4 - limestone	2			33			35
	5 - granite		3			27		30
	6 - log	1			3		36	40
	Totals	36	36	36	36	36	36	198
								91.6%

Table 6. ARTMAP-FTR six-class object recognition from the sonar test set.




Figure 5: Test-set error locations for the three-ping fusion, six-class recognition task. The darkest areas show the locations of object 1 predictions and the lightest areas show the locations of object 6 predictions (key). The circular map shows the initial aspects of the 18 confusion errors listed in Table 6. The other 198 aspect inputs identify the correct object class.

4.1 Matched filter vs. SCAT preprocessing

Simulation studies confirmed the ORINCON conclusion that matched filter preprocessing with threshold-centered data is the best of the paradigms considered for this task. For example, on sonar data preprocessed by the SCAT algorithm,¹ performance was significantly worse than with matched filter preprocessing. For example, on the three-ping, man-made / non-man-made task, the best ARTMAP-FTR performance across all values of the vigilance parameter was just 82% for SCAT data, compared to 93.1% for the matched filter data.

4.2 Variable aspect separations

Simulations were carried out to examine how varying the size of the aspect spacing interval affects performance on the three-ping, man-made / non-man-made discrimination task. Figure 6 shows that the 30° input aspect spacing employed in the ORINCON classification paradigm also gives the best ARTMAP-FTR performance, compared to shorter or longer intervals. On the other hand, accuracy does not deteriorate drastically as the size of the inter-ping aspect interval varies.

The observation that ARTMAP-FTR performance is stable with respect to inter-ping aspect separation intervals is reinforced by results of a study in which inter-ping aspect intervals were chosen randomly, between 10° and 50° (Figure 7). With random separation of inputs, ARTMAP-FTR performance on the three-ping task increases to 94.3%, which is better than the performance produced by any single fixed aspect interval (Figure 6). The fact that accuracy is not sensitive to aspect spacing indicates the reliability of ARTMAP-FTR performance on the sonar recognition task.

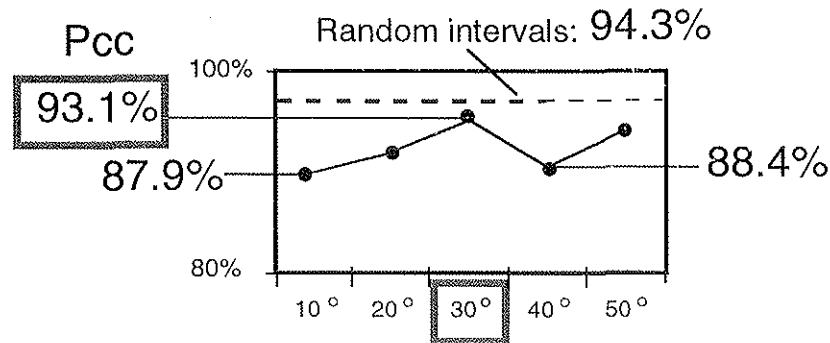


Figure 6: ARTMAP-FTR classification rate as a function of the inter-ping aspect interval, for the three-ping, man-made / non-man-made task.




Figure 7: Sample test-set aspect spacing for a typical three-ping fusion input with randomly chosen inter-ping intervals.

4.3 Variable numbers of pings

All ARTMAP-FTR simulation studies described so far have considered system performance on three-ping tasks. The effect of a varying the number of pings was also investigated by presenting N pings, with 30° aspect separations on each input trial. On the N -ping man-made / non-man-made discrimination task, ARTMAP-FTR performance was observed to improve with increasing numbers of pings. Table 7 shows that the percent correct classification rate increasing monotonically from 83.3% (1 ping) to 95.4% (6 pings).

5. THE ARTMAP-FTR NEURAL NETWORK

The goal of the ARTMAP-FTR network is to combine information from a sequence of sonar returns in order to achieve object identification that is more accurate than single-ping identification. The resulting system is designed as a neural network hierarchy. Coding, training, testing, and parameter selection for the ARTMAP-FTR hierarchy are now summarized.

5.1. Distributed coding by ARTMAP networks

Since the introduction of fuzzy ARTMAP,⁸ which features a winner-take-all (WTA) coding scheme, a number of ARTMAP variants that employ distributed coding have been developed. In many application domains, these variants have been shown to improve computational capabilities of the basic network. One such network is ART-EMAP (Stage 1),⁶¹⁻⁶² which, during testing, simply distributes activation across the coding nodes of a trained fuzzy ARTMAP network. The ARTMAP-IC network¹² adds instance counting to the coding scheme to track frequency of use of internal coding nodes. During testing, the system uses this information to bias distributed predictions. ARTMAP-IC also changes the original search algorithm slightly, which allows the system to encode inconsistent cases and also improves code compression in general.

Finally, distributed ARTMAP (dARTMAP)^{3,13} introduces computational elements that retain essential ARTMAP design principles, including stable coding with fast learning, while permitting distributed code representations during training as well as testing. One of these new computational elements is the increased gradient CAM (content addressable memory) rule for distributed activation at the coding field. A CAM rule models the steady-state activation pattern of a coding field in response to a given input vector. Compared to other coding algorithms, including power rules, the increased gradient CAM rule enhances differences among input components, which is useful for systems that tend to exhibit a compressed dynamic range of activation values. The increased gradient CAM rule has been found to be useful in networks, such as ARTMAP, that employ distributed coding only during testing, in addition to fully distributed coding networks such as dARTMAP.

N - PING SONAR TASK 30° SEPARATION	ARTMAP-FTR Pcc
1-ping	83.3%
2-ping	88.9%
3-ping	93.1%
4-ping	93.1%
5-ping	94.0%
6-ping	95.4%

Table 7: For man-made / non-man-made discrimination, ARTMAP-FTR performance improves with increasing numbers of pings.

5.2. ARTMAP-FTR training

The basic ARTMAP-FTR system incorporates two ARTMAP networks (Net 1 and Net 2), which are arranged hierarchically (Figure 8). For these networks, sonar simulations described in this paper use a variant of the ARTMAP-IC algorithm. This ARTMAP-IC variant employs the dARTMAP increased gradient CAM rule during testing.

During ARTMAP-FTR training for a multi-ping sonar task, Net 1 first uses WTA coding to learn to identify the six objects from single-ping inputs. The same training set is then re-presented to the trained network, except coding field activation is now distributed, yielding a distributed output pattern for each single-ping return. For multi-ping training, a concatenated set of these vectors is presented to a second ARTMAP system (Net 2). The training stage requires setting one free parameter, which corresponds to an ARTMAP baseline vigilance parameter. Given a vigilance value $\rho \in [0,1]$, Net 1 and Net 2 are trained to identify the six objects, based on the concatenated multi-ping input. The baseline vigilance parameter sets code granularity, with low ρ values permitting coarse categories and high ρ values creating fine categories. Following six-class target recognition, system outputs can be merged for further classification tasks, such as man-made / non-man-made discrimination.

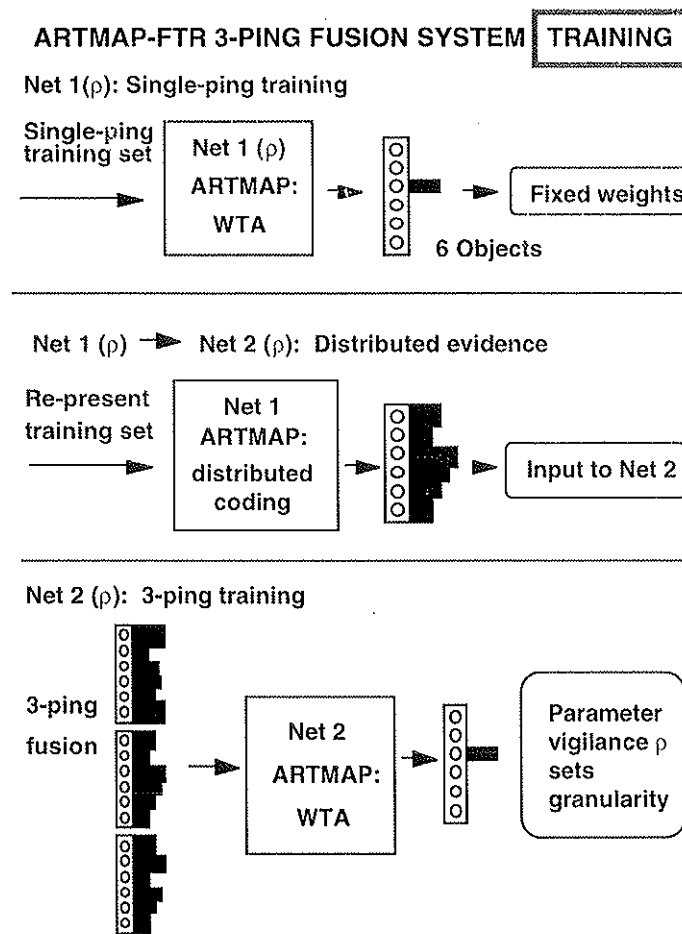


Figure 8: Multi-ping training by an ARTMAP-FTR hierarchy. Net 1 encodes single-ping inputs, then Net 2 encodes concatenated outputs from the previous stage. Sonar simulations use an ARTMAP-IC network with the increased gradient CAM rule for Net 1 and Net 2.

5.3. ARTMAP-FTR testing

Test set inputs are presented to an ARTMAP-FTR network that has previously been trained with some fixed value of ρ . During testing, the baseline vigilance parameter is set equal to 0, which forces the network to make a feedforward prediction for each input. As during training, distributed outputs from Net 1 are concatenated to form the input vectors to Net 2 (Figure 9).

As is typical of fast-learn systems, ARTMAP network coding varies somewhat with the order of input presentation. To factor out variations due to input ordering, each performance measurement reported here represents the result of voting across five orderings of the Net 1 input set. Voting provides a stable indicator of network accuracy. However, the performance of an ARTMAP-FTR system is still sensitive to the choice of the trained network's free parameter, ρ . For example, Figure 9 shows that the percent correct classification rate on a man-made / non-man-made discrimination task achieves a maximal accuracy rate of 93.1% for ρ equal to an optimal value, ρ^* , but that system accuracy is lower for other values of the baseline vigilance parameter. In order to be able to report the optimal value as the test-set performance measure (Figure 3), it is necessary to infer the optimal parameter value from training set data alone. The ARTMAP-FTR network does, in fact, achieve robust parameter estimation by training set validation, as follows.

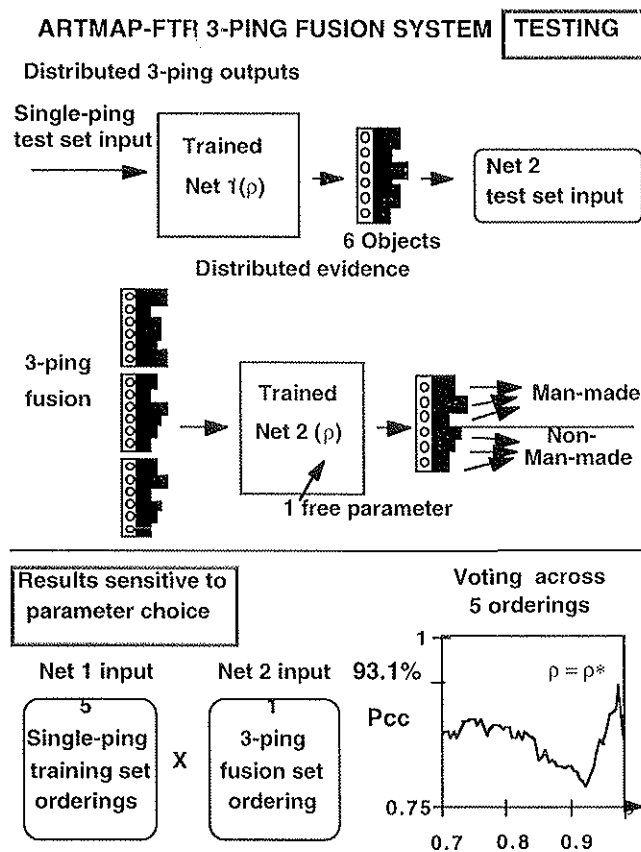


Figure 9: Multi-ping ARTMAP-FTR testing. Net 1 encodes single-ping inputs, then Net 2 encodes concatenated Net 1 outputs. A given network has been previously trained using a fixed value of the baseline vigilance, ρ , and system performance is sensitive to the value of this free parameter.

5.4. Parameter selection by validation

Figure 10 summarizes the validation procedure used for sonar recognition tasks. During training, a validation subset of the training set is reserved, and the network trained using inputs from the remaining aspects. For each study, five validation subsets are selected and, for each, the network is trained on five orderings of the remaining training set inputs. This procedure produces, for each value of the free parameter, a performance measure representing an average across the five validation subsets, each using five voters. The parameter ρ^* was taken to be the one with the best average validation set performance. An ARTMAP-FTR network is then retrained on the whole training set (with $\rho = \rho^*$) to produce the reported test-set accuracy measurements.

Some of the sonar recognition tasks also require the selection of a second free parameter, namely, an output decision threshold γ . This occurs, for example, in the three-class object discrimination tasks (Section 3.3). In this case, the validation set procedure would first be applied to choose an optimal threshold $\gamma(\rho)$ for each fixed ρ . The previously described steps would be followed to choose ρ , keeping each $\gamma = \gamma(\rho)$. In all cases, validation set choice of network parameters proved to be robust.

6. CONCLUSION: ARTMAP-FTR FOR TARGET RECOGNITION BY SENSOR FUSION

The ARTMAP-FTR network successfully performs target recognition by fusing sonar data. Although the system was designed specifically for object classification from multi-ping sonar returns, it could also be applied to a more general class of sensor fusion problems. For example, because of the modular nature of the network hierarchy, the output from several Net 1 systems, each classifying inputs from a different type of sensor, could be combined to form the input to Net 2 (Figure 8).

ARTMAP-FTR provides robust performance which remains reliable across many simulation trials. System parameters are stably selected by validation subsets of the training set, and performance accuracy remains high as the size of the inter-ping aspect intervals varies. Accuracy also increases with the number of pings.

Having been tested on the NSW benchmark sonar data set, the ARTMAP-FTR system can now be scaled up for larger problems, including those with more target object. Future studies could also include investigations of system performance with test set objects that are disjoint from objects in the training set and with different preprocessors.

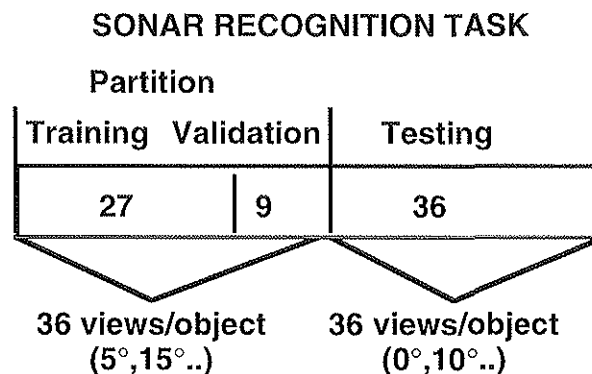


Figure 10: A validation subset of the training set is used to choose ARTMAP-FTR parameters.

7. ACKNOWLEDGMENTS

This research was supported in part by the Office of Naval Research (ONR N00014-95-1-0409 and ONR N00014-95-1-0657). The authors thank Dr. Gerald Dobeck from the NSWC Coastal Systems Station for providing the sonar data set; Drs. Larry Burton and Hung Lai from the ORINCON Corporation for helpful discussions; and Dr. Harold Hawkins from the Office of Naval Research for suggesting the project.

Technical Report CAS/CNS-TR-98-016, Boston, MA: Boston University.

8. REFERENCES

1. Simmons, J.A., Saillant, P.A., Wotton, J.M., Haresign, T., Ferragamo, M.J., & Moss, C.F. (1995). Composition of biosonar images for target recognition by echolocating bats. *Neural Networks*, **8**, 1239-1261.
2. Burton, L.L. (1997). ONR Phase II: Active target imaging and classification system. Topic N94-142, Final Report.
3. Carpenter, G.A. (1997). Distributed learning, recognition, and prediction by ART and ARTMAP neural networks. *Neural Networks*, **10**, 1473-1494.
4. 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.
5. Carpenter, G.A., & Grossberg, S. (1987b). ART 2: Self-organization of stable category recognition codes for analog input patterns. *Applied Optics*, **26**, 4919-4930.
6. Carpenter, G.A., & Grossberg, S. (1990). ART 3: Hierarchical search using chemical transmitters in self-organizing pattern recognition architectures. *Neural Networks*, **3**, 129-152.
7. Carpenter, G.A., & Grossberg, S. (1991). *Pattern Recognition by Self-Organizing Neural Networks*. Cambridge, MA: MIT Press.
8. 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.
9. 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.
10. Carpenter, G.A., Grossberg, S., & Rosen, D.B. (1991a). Fuzzy ART: Fast stable learning and categorization of analog patterns by an Adaptive Resonance system. *Neural Networks*, **4**, 759-771.
11. Carpenter, G.A., Grossberg, S., & Rosen, D.B. (1991b). ART 2-A: An adaptive resonance algorithm for rapid category learning and recognition. *Neural Networks*, **4**, 493-504.
12. Carpenter, G.A., & Markuzon, N. (1998). ARTMAP-IC and medical diagnosis: Instance counting and inconsistent cases. *Neural Networks*, **11**. Technical Report CAS/CNS TR-96-017, Boston, MA: Boston University.
13. Carpenter, G.A., Milenova, B.L., & Nocske, B.W. (1998). dARTMAP: A neural network for fast distributed supervised learning. *Neural Networks*, **11**. Technical Report CAS/CNS TR-97-026, Boston, MA: Boston University.
14. Grossberg, S. (1976). Adaptive pattern classification and universal recoding, II: Feedback, expectation, olfaction, and illusions. *Biological Cybernetics*, **23**, 187-202.
15. Grossberg, S. (1980). How does a brain build a cognitive code? *Psychological Review*, **87**, 1-51.
16. Williamson, J.R. (1996). Gaussian ARTMAP: A neural network for fast incremental learning of noisy multidimensional maps. *Neural Networks*, **9**, 881-897.
17. 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.
18. Caudell, T.P., & Healy, M.J. (1994). Adaptive Resonance Theory networks in the Encephalon autonomous vision system. In *Proceedings of the 1994 IEEE International Conference on Neural Networks*, pp. II-1235-1240. Piscataway, NJ: IEEE.
19. Bachelder, I.A., & Waxman, A.M. (1994). Mobile robot visual mapping and localization: A view-based neurocomputational architecture that emulates hippocampal place learning. *Neural Networks*, **7**, 1083-1099.

20. Baloch, A.A., & Waxman, A.M. (1991). Visual learning, adaptive expectations, and behavioral conditioning of the mobile robot MAVIN. *Neural Networks*, **4**, 271-302.
21. Bachelder, I.A., Waxman, A.M., & Seibert, M. (1993). A neural system for mobile robot visual place learning and recognition. *Proceedings of the World Congress on Neural Networks (WCNN'93)*, pp. I-512-517. Hillsdale, NJ: Lawrence Erlbaum Associates.
22. Dubrawski, A., & Crowley, J.L. (1994a). Learning locomotion reflexes: A self-supervised neural system for a mobile robot. *Robotics and Autonomous Systems*, **12**, 133-142.
23. Dubrawski, A., & Crowley, J.L. (1994b). Self-supervised neural system for reactive navigation. *Proceedings of the IEEE International Conference on Robotics and Automation*, San Diego, May, 1994, pp. 2076 - 2081. Los Alamitos, CA: IEEE Computer Society Press.
24. Racz, J., & Dubrawski, A. (1995). Artificial neural network for mobile robot topological localization. *Robotics and Autonomous Systems*, **16**, 73-80.
25. Srinivasa, N., & Sharma, R. (1996). A self-organizing invertible map for active vision applications. *Proceedings of the World Congress on Neural Networks (WCNN'96)*, pp. 121-124. Hillsdale, NJ: Lawrence Erlbaum Associates.
26. Seibert, M., & Waxman, A.M. (1992). Adaptive 3D object recognition from multiple views. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **14**, 107-124.
27. Seibert, M., & Waxman, A.M. (1993). An approach to face recognition using saliency maps and caricatures. In *Proceedings of the World Congress on Neural Networks (WCNN'93)*, pp. III-661-664. Hillsdale, NJ: Lawrence Erlbaum Associates.
28. Soliz, P., & Donohoe, G.W. (1996). Adaptive resonance theory neural network for fundus image segmentation. *Proceedings of the World Congress on Neural Networks (WCNN'96)*, pp. 1180-1183. Hillsdale, NJ: Lawrence Erlbaum Associates.
29. Baraldi, A., & Parmiggiani, F. (1995). A neural network for unsupervised categorization of multivalued input patterns: An application to satellite image clustering. *IEEE Transactions on Geoscience and Remote Sensing*, **33**, 305-316.
30. Carpenter, G.A., Gajja, M.N., Gopal, S., & Woodcock, C.E. (1997). ART neural networks for remote sensing: Vegetation classification from Landsat TM and terrain data. *IEEE Transactions on Geoscience and Remote Sensing*, **35**, 308-325.
31. Castelli, V., Li, C.-S., Turek, J.J., & Kontoyiannis, I. (1996). Progressive classification in the compressed domain for large EOS satellite databases. *Proceedings of the 1996 IEEE International Conference on Acoustics, Speech and Signal Processing*, **4**, 2201-2204.
32. Gopal, S., Sklarew, D. M., & Lambin, E. (1993). Fuzzy-neural network classification of landcover change in the Sahel. *Proceedings of the DOSES/EUROSTAT Workshop on New Tools for Spatial Analysis*, Lisbon, Portugal, November 18-20, 1993.
33. Johnson, C. (1993). Agent learns user's behavior. *Electrical Engineering Times*, June 28, pp. 43, 46.
34. Bernardon, A.M., & Carrick, J.E. (1995). A neural system for automatic target learning and recognition applied to bare and camouflaged SAR targets. *Neural Networks*, **8**, 1103-1108.
35. Koch, M.W., Moya, M.M., Hostetler, L.D., & Fogler, R.J. (1995). Cueing, feature discovery, and one-class learning for synthetic aperture radar automatic target recognition. *Neural Networks*, **8**, 1081-1102.
36. Rubin, M.A. (1995). Application of fuzzy ARTMAP and ART-EMAP to automatic target recognition using radar range profiles. *Neural Networks*, **8**, 1109-1116.
37. Waxman, A.M., Seibert, M.C., Gove, A., Fay, D.A., Bernardon, A.M., Lazott, C., Steele, W.R., & Cunningham, R.K. (1995). Neural processing of targets in visible, multispectral IR and SAR imagery. *Neural Networks*, **8**, 1029-1051.
38. Ham, F.M., & Han, S.W. (1993). Quantitative study of the QRS complex using fuzzy ARTMAP and the MIT/BIH arrhythmia database. In *Proceedings of the World Congress on Neural Networks (WCNN'93)*, pp. I-207-211. Hillsdale, NJ: Lawrence Erlbaum Associates.
39. Suzuki, Y. (1995). Self-organizing QRS-wave recognition in ECG using neural networks. *IEEE Transactions on Neural Networks*, **6**, 1469-1477.
40. Wienke, D., Xie, Y., & Hopke, P.K. (1994). An Adaptive Resonance Theory based artificial neural network (ART 2-A) for rapid identification of airborne particle shapes from their scanning electron microscopy images. *Chemometrics and Intelligent Laboratory Systems*.
41. Soliz, P., & Caudell, T.P. (1996). Inferring future states of the atmosphere with a laterally primed adaptive resonance theory (LAPART) neural network. *Proceedings of the World Congress on Neural Networks (WCNN'96)*, pp. 822-825. Hillsdale, NJ: Lawrence Erlbaum Associates.

42. Kasperkiewicz, J., Racz, J., & Dubrawski, A. (1995). HPC strength prediction using artificial neural network. *Journal of Computing in Civil Engineering*, **9**, 279-284.
43. Murshed, N.A., Bortolozzi, F., & Sabourin, R. (1995). Off-line signature verification, without *a priori* knowledge of class ω_2 . A new approach. *Proceedings of ICDAR 95: The Third International Conference on Document Analysis and Recognition*. Montreal, Canada, pp. 191-196.
44. Ruda, H., & Snorrason, M. (1996). Automated construction of a hierarchy of self-organized neural network classifiers. *Proceedings of the Southeastern Simulation Conference*.
45. MacLeod, K.J., & Surkan, A.J. (1992). Algorithm performance in ART I-like clustering of descriptor subsets for document retrieval. *Proceedings of the International Joint Conference on Neural Networks*, **I**, pp. 685-689. Piscataway, NJ: IEEE.
46. Gjerdingen, R.O. (1990). Categorization of musical patterns by self-organizing neuronlike networks. *Music Perception*, **7**, 339-370.
47. Gan, K.W., & Lua, K.T. (1992). Chinese character classification using an Adaptive Resonance network. *Pattern Recognition*, **25**, 877-88.
48. Kim, J.W., Jung, K.C., Kim, S.K., & Kim, H.J. (1995). Shape classification of on-line Chinese character strokes using ART I neural network. *Proceedings of the World Congress on Neural Networks (WCNN'95)*, pp. II-191-194. Hillsdale, NJ: Lawrence Erlbaum Associates.
49. Wang, T., Xu, Q., & Ziaoliang, X. (1992). Character recognition through improving adaptive resonance theory (ART). *Proceedings of the International Joint Conference on Neural Networks*, pp. I-761-764. Piscataway, NJ: IEEE.
50. Choi, J., Ly, S., Healy, M., & Smith, S. (1996). Prediction of cutter condition using LAPART. *Proceedings of the International Conference on Neural Networks: Plenary, Panel, and Special Sessions*, 226-230.
51. Ly, S., & Choi, J.J. (1994). Drill condition monitoring using ART-1. *Proceedings of the 1994 IEEE International Conference on Neural Networks*, pp. II-1226-1229. Piscataway, NJ: IEEE.
52. Tarng, Y.S., Li, T.C., & Chen, M.C. (1994) Tool failure monitoring for drilling processes. *Proceedings of the 3rd International Conference on Fuzzy Logic, Neural Nets and Soft Computing*, Iizuka, Japan, pp. 109-111.
53. Tse, P., & Wang, D.D. (1996). A hybrid neural networks based machine condition forecaster and classifier by using multiple vibration parameters. *Proceedings of the 1994 IEEE International Conference on Neural Networks*, **IV**, pp. 2096-2100. Piscataway, NJ: IEEE.
54. Wienke, D. (1994). Neural resonance and adaption - Towards nature's principles in artificial pattern recognition. In L. Buydens and W. Melssen (Eds.), *Chemometrics: Exploring and Exploiting Chemical Information*. Nijmegen, NL: University Press.
55. Whiteley, J.R., Davis, J.F., Mehrotra, A., & Ahalt, S.C. (1996). Observations and problems applying ART 2 for dynamic sensor pattern interpretation. *IEEE Transactions on Neural Networks*, **26**, 423-437.
56. Burke, L.I. (1994). Neural methods of the traveling salesman problem: Insights from operations research. *Neural Networks*, **7**, 681-690.
57. Murshed, N.A., Bortolozzi, F., & Sabourin, R. (1996). A fuzzy ARTMAP-based classification system for detecting cancerous cells, based on the one-class problem approach. *Proceedings of the 13th International Conference on Pattern Recognition (ICPR'96)*.
58. Wienke, D., & Kateman, G. (1994). Adaptive Resonance Theory based artificial neural networks for treatment of open-category problems in chemical pattern recognition - Application to UV-Vis and IR spectroscopy. *Chemometrics and Intelligent Laboratory Systems*.
59. Christodoulou, C.G., Huang, J., Georgiopoulos, M., & Liou, J.J. (1995). Design of gratings and frequency selective surfaces using fuzzy ARTMAP neural networks. *Journal of Electromagnetic Waves and Applications*, **9**, 17-36.
60. Kalkunte, S.S., Kumar, J.M., & Patnaik, L.M. (1992). A neural network approach for high resolution fault diagnosis in digital circuits. *Proceedings of the International Joint Conference on Neural Networks*, **I**, pp. 83-88. Piscataway, NJ: IEEE.
61. Carpenter, G.A., & Ross, W.D. (1993). ART-EMAP: A neural network architecture for learning and prediction by evidence accumulation. *Proceedings of the World Congress on Neural Networks (WCNN'94)*, pp. III - 649-656. Hillsdale, NJ: Lawrence Erlbaum Associates.
62. 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.