

Progress on the CRCD Experiences
at the University of Central Florida: An NSF Project

Michael Georgiopoulos*, Erol Gelenbe**, Ronald Demara*, Avelino Gonzalez*, Marcella Kysilka*, Mansooreh Mollaghasemi*, Annie Wu*, Georgios Anagnostopoulos***, Ingrid Russell****, Jimmy Secretan*

(*) University of Central Florida

(**) Imperial College

(***) Florida Institute of Technology

(****) University of Hartford

Abstract

Machine Learning has traditionally been a topic of research and instruction in computer science and computer engineering programs. Yet, due to its wide applicability in a variety of fields, its research use has expanded in other disciplines, such as electrical engineering, industrial engineering, civil engineering, and mechanical engineering. Currently, many undergraduate and first-year graduate students in the aforementioned fields do not have exposure to recent research trends in Machine Learning. This paper reports on a project in progress, funded by the National Science Foundation under the program *Combined Research and Curriculum Development (CRCD)*, whose goal is to remedy this shortcoming. The project involves the development of a model for the integration of Machine Learning into the undergraduate curriculum of those engineering and science disciplines mentioned above. The goal is increased exposure to Machine Learning technology for a wider range of students in science and engineering than is currently available. Our approach of integrating Machine Learning research into the curriculum involves two components. The first component is the incorporation of Machine Learning modules into the first two years of the curriculum with the goal of sparking student interest in the field. The second is the development of new upper level Machine Learning courses for advanced undergraduate students. In the past, we have reported on our experiences of introducing Machine Learning modules in sophomore and junior undergraduate classes, in an effort to recruit students for our senior level classes (Current Topics in Machine Learning I (CTML-I) and Current Topics in Machine Learning II (CTML-II)). This paper focuses on discussing our experiences in teaching these senior level classes of CTML-I and CTML-II.

1. Introduction

In the last decade there has been an explosion of research in machine learning. A contributing factor is that traditionally independent research communities in symbolic machine learning,

computational learning theory, neural networks, genetic algorithms, statistics, and pattern recognition have achieved new levels of collaboration. The outcome has been a plethora of results in machine learning emerging from all of these research communities working synergistically. The second reason for the explosive growth is that machine learning has been applied successfully to a growing range of problems in science and engineering, such as speech recognition, handwritten recognition, medical data analysis, game playing, knowledge data discovery in databases, language processing, robot control, and others.

Machine Learning has traditionally been a topic of research and instruction in computer science and computer engineering programs. Yet, due to its wide applicability in a variety of fields, its research use has expanded in other disciplines, such as electrical engineering, industrial engineering, civil engineering and mechanical engineering. Currently, quite a few undergraduate and first-year graduate students in the aforementioned fields do not have access to coursework and exposure to recent research trends in Machine Learning. The effort in this CRCD project is attempting to remedy these shortcomings. By involving in this CRCD effort a strong team of faculty from a variety of disciplines, such as computer science, computer engineering, electrical engineering and industrial engineering with interest in Machine Learning, we hope to significantly increase the exposure of engineering and science students to machine learning technology. The faculty affiliated with the CRCD have significant research experiences in the theory and applications of machine learning, and are therefore qualified to transfer these research experiences into the undergraduate and first-year graduate curricula.

The approach that we have chosen to integrate machine-learning research into the curriculum involves two components. The *first component* inspires student interest in machine learning, while the second component introduces the student to current research results in machine learning. The first component is realized through the incorporation of machine learning modules in sophomore and junior level coursework. One reason that we can incorporate a wide variety of interesting and inspiring machine learning modules is because of our diverse team of researchers. Hence, our goal of reaching out to a wide variety of students from several engineering disciplines will be readily accomplished. The undergraduate students enrolled in classes with machine learning modules will create our primary pool of candidate students to take our proposed *new* machine learning classes, the Current Topics in Machine Learning I (CTML-I) class, and the Current Topics in Machine Learning II (CTML-II) class. These new machine-learning classes constitute the *second component* of our curriculum development efforts, and the focus of this paper. .

The CTML-I class will serve as the entrance point to our proposed curriculum. In this course a number of our PIs will instruct the students on machine-learning techniques developed through their current research and how it relates to other machine learning approaches. The approach taken here for CTML-I is to create a course that is enthusiastically taught by the principal investigators (since it will always be relevant to their machine learning research), will reflect the ever changing research interests of the machine learning community (the topics in this class will

frequently change), and will be a welcome alternative to undergraduate students (students are normally exposed to knowledge that is at least a few years old). The approach taken in CTML-II is to build on the fundamentals that the students have been exposed to in CTML-I in order to embrace the research aspects of a graspable Machine Learning project. In CTML-II the students will have the opportunity to interact with the PIs on a one-to-one basis. This way the students will be exposed to research that will positively affect their critical thinking, stimulate their interest, and improve their communication skills.

In this paper we are focusing on our experiences with the teaching of the Current Topics in Machine Learning I and II classes. In particular, we will explain how the research topics that are being taught in CTML-I lead us in a natural way to the research projects assigned in the CTML-II class. To achieve this goal we are going to emphasize one research topic, discussed in CTML-I and an associated project proposed to one of the students in CTML-II. The research topic is that of *ART Neural Networks*.

3. Current Topics in Machine Learning I and II Classes

The CTML-I and CTML-II classes were taught in the Fall of 2003 and Spring of 2004. The research topics discussed in the CTML-I class included: Introduction to Machine Learning, Inductive Learning, ART Neural Networks and Applications, Random Neural Network and Applications, Simulation, Experimental Design and Simulation Metamodeling, Genetic Algorithms, and Evolutionary Hardware. In particular a more detailed CTML-I course outline is provided below.

CTML-I Course Outline:

1. Introduction to CRCD and Machine Learning.
2. Inductive Learning (Learning from examples using inductive learning, Algorithms to build induction trees based on examples, Extraction of rules from induction trees, Design of minimal inductive trees).
3. ART Neural Networks (Rationale for Adaptive Resonance theory, Description of Fuzzy ART and Fuzzy ARTMAP, Geometrical Interpretation of ART Architectures, Learning in Fuzzy ART and Fuzzy ARTMAP, Applications [Mushroom Database, Letters Database]).
4. Random Neural Networks: Theory and Applications (The RNN – Inspiration from natural neurons, Basic theory of the RNN and computational model, Applications of the RNN to image texture learning and discrimination, Application of the RNN to decision making and packet network routing, Application of the RNN to modeling somatosensory circuits in the brain]).
5. Experimental Design and Simulation Metamodeling (input data modeling, random number and random variable generation, output analysis, variance reduction techniques, and experimental design).
6. Genetic Algorithms: Learning by Evolution (Introduction to the GA, Problem representation, Genetic operators, Selection method, Evaluation, Evolvable Hardware)

7. Description of Potential Research Projects (for CTML-II).

More details about each and every one of these topics can be found on our web-site (www.ece.ucf.edu/~crcd). The lectures for the class were based on individual notes that each one of the instructors prepared for this class (the notes can be found on our web-site under the link Machine Learning Courses, and Current Topics in Machine Learning I). The prerequisite for this class was a class a programming language class (such as C, or C++), or a basic statistics class, or the consent of the instructor.

In the next three sub-sections we discuss, in detail, the material taught under the topic of ART Neural Networks (CTML-I class), the ART homework assigned in the CTML-I class, and finally one of the project assignments that was proposed in the CTML-II class and carried out by one of the students in the CTML-II class. By elaborating on only one research topic we believe that we will communicate effectively, what has transpired with the rest of the research topics that other CRCDC PIs taught in the CTML-I class, and supervised projects for in the CTML-II class.

3.1 ART Neural Networks in the CTML-I class

The major sub-topics discussed under the ART Neural Networks topic were: Introduction to the Adaptive Resonance Theory (ART), ART literature review, the training and performance phases of the Fuzzy ARTMAP neural network, the geometrical interpretation of the ART neural network equations related to category choices and learning, and ART neural network applications.

In the Introduction to Adaptive Resonance Theory (ART) we emphasized the fact that adaptive resonance theory was introduced by Grossberg (see Grossberg, 1976) in an effort to solve the stability versus plasticity dilemma. This dilemma is of primary importance to any system that learns from experience. That is “how can we design a system that learns, which is stable enough not to forget important information that it has learnt, and at the same time, is plastic enough to learn new information”. Grossberg’s answer to this question was the theory referred to as adaptive resonance theory, on whose principles a number of ART learning systems have been built. Then, we focused on the relevant ART literature from the perspective of new ART neural network architectures, ART properties of learning, and ART applications. We also explained some of the good properties of ART neural networks, whose primary utility is in solving classification problems. These properties are: (a) they can solve complex classification problems, (b) they converge fast to a solution, after a few iterations through the collection of the data used for their training, (c) they possess on-line learning capabilities, thus not requiring re-training with old information as they are being exposed (learn) new information, (d) they offer easy to understand explanations for the answers that they provide, (e) they have been proven to be very competitive (compared to other techniques) classifiers on a variety of classification problems.

We then continued the lectures by emphasizing on one of the most celebrated ART neural network architectures, the Fuzzy ARTMAP (FAM) neural network architecture (see Carpenter, et al., 1992). Initially we talked about the basic layers of nodes in FAM (see Figure 1 for a block

diagram of the FAM neural network architecture). That is we emphasized that FAM consists of three layers of nodes: the *input layer* (F_1^a), the *category representation layer* (F_2^a), and the *output layer* (F_2^b). The input layer of Fuzzy ARTMAP is the layer where an input vector \mathbf{I} of dimensionality $2M_a$ of the following form is applied

$$\mathbf{I} = (\mathbf{a}, \mathbf{a}^c) = (a_1, \dots, a_{M_a}, a_1^c, \dots, a_1^c, \dots, a_{M_a}^c) \quad (1)$$

where

$$a_i^c = 1 - a_i; \quad 1 \leq i \leq M_a \quad (2)$$

The assumption here is that the input vector \mathbf{a} is such that each one of its components lies in the interval $[0, 1]$. The above form of an input pattern is referred to as a *complementary coded* form. The layer F_2^a of FAM is referred to as the *category representation layer*, because this is where categories (or groups) of input patterns are formed. Finally, the output layer is the layer that produces the outputs of the network. An output of the network represents the output to which the input applied at the input layer of Fuzzy ARTMAP is supposed to be mapped to.

Then, we proceeded by emphasizing the interconnection weights in FAM. Two types of interconnection weights were discussed. The weights from F_2^a to F_1^a , designated as w_{ji}^a ; ($1 \leq j \leq N_a$, $1 \leq i \leq 2M_a$), are referred to as top-down weights. Furthermore, the vector of weights $\mathbf{w}_j^a = (w_{j1}^a, w_{j2}^a, \dots, w_{j,2M_a}^a)$ is called a *template*. Its functionality is that it represents the group of input patterns that chose node j in the category representation layer of Fuzzy ARTMAP as their representative node and this node encoded them. Finally, there are weights that emanate from every node j in the category representation layer to every node k in the output layer. These weights are designated as W_{jk}^{ab} (called inter-ART weights). The vector of inter-ART weights emanating from every node j in Fuzzy ARTMAP (i.e., $\mathbf{W}_j^{ab} = (W_{j1}^a, W_{j2}^a, \dots, W_{j,N_b}^a)$) corresponds to the output pattern that this node j is mapped to.

Then, we emphasized the difference between the training and the performance phase in Fuzzy ARTMAP. The training phase of Fuzzy ARTMAP can be described as follows: Given a list of input/output pairs, $\{(\mathbf{I}^1, \mathbf{O}^1), \dots, (\mathbf{I}^r, \mathbf{O}^r), \dots, (\mathbf{I}^{PT}, \mathbf{O}^{PT})\}$, we want to train Fuzzy ARTMAP to map every input pattern of the training list to its corresponding output pattern. To achieve the aforementioned goal we present the training list to Fuzzy ARTMAP architecture repeatedly. That is, we present \mathbf{I}^1 to F_1^a , \mathbf{O}^1 to F_2^b , \mathbf{I}^2 to F_1^a , \mathbf{O}^2 to F_2^b , and finally \mathbf{I}^{PT} to F_1^a , and \mathbf{O}^{PT} to F_2^b . We present the training list to Fuzzy ARTMAP as many times as it is necessary for Fuzzy ARTMAP to correctly classify all these input patterns. The task is considered accomplished (i.e., the learning is complete) when the weights do not change during a list presentation. The

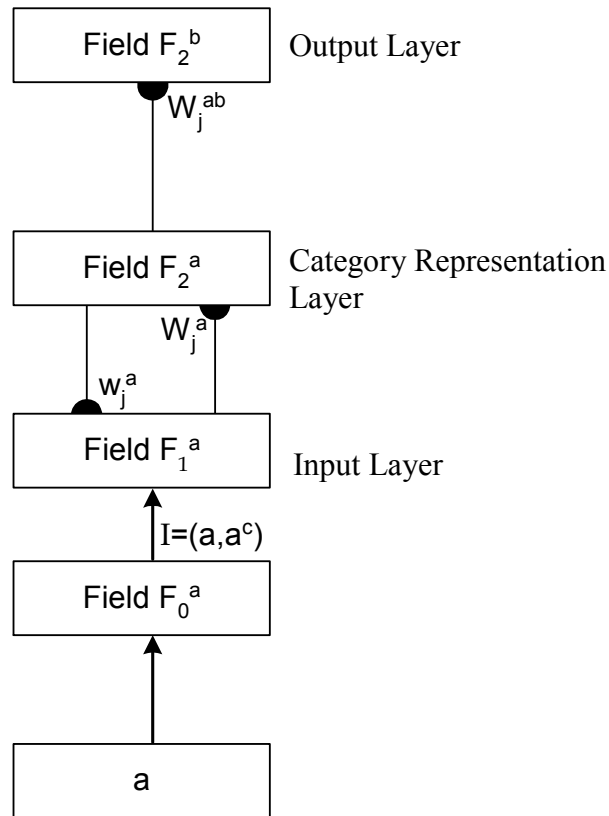


Figure 1: Block Diagram of the FAM Architecture

aforementioned training scenario is called *off-line learning*. The performance phase of Fuzzy ARTMAP works as follows: Given a list of input patterns, such as $\tilde{\mathbf{I}}^1, \tilde{\mathbf{I}}^2, \dots, \tilde{\mathbf{I}}^{PS}$, we want to find the Fuzzy ARTMAP output produced when each one of the aforementioned test patterns is presented at its F_1^a layer. In order to achieve the aforementioned goal we present the test list to the trained Fuzzy ARTMAP architecture and we observe the network's output.

Our next focal point was a short description of the FAM network parameters and their associated functionality. The operation of Fuzzy ARTMAP is affected by two network parameters, the choice parameter β_a , and the baseline vigilance parameter $\bar{\rho}_a$. The choice parameter β_a takes values in the interval $(0, \infty)$, while the baseline vigilance parameter $\bar{\rho}_a$ assumes values in the interval $[0, 1]$. Both of these parameters affect the number of nodes created in the category representation layer of Fuzzy ARTMAP. Higher values of β_a and $\bar{\rho}_a$ create more nodes in the category representation layer of Fuzzy ARTMAP, and consequently produce less compression of the input patterns.

After these initial and necessary comments about FAM it was time to elaborate on the training phase and performance phases of FAM. We believe that one of the most difficult points of the FAM lecture notes was the understanding of the training phase of FAM. So a considerable amount of time was spent explaining the FAM training phase. A complete description of the training phase of Fuzzy ARTMAP is included in the ART lecture notes provided to the CTML-I students and available on the web (www@ece.ucf.edu/~crcd, lecture notes link of the CTML-I class).

Once the training and performance phases were discussed, an example of input/output pairs was presented and FAM's training steps were followed faithfully until the architecture learned all the correct mappings. In this example, six patterns, designated by $\mathbf{I}^1, \mathbf{I}^2, \mathbf{I}^3, \mathbf{I}^4, \mathbf{I}^5, \mathbf{I}^6$ are presented to Fuzzy ARTMAP. Patterns $\mathbf{I}^1, \mathbf{I}^2, \mathbf{I}^3, \mathbf{I}^5$ belong to class 1, represented by label \mathbf{O}^1 , and patterns $\mathbf{I}^4, \mathbf{I}^6$ belong to class 2, designated by \mathbf{O}^2 . The input patterns and the output patterns are chosen as follows: $\mathbf{I}^1 = (0.3 \ 0.2 \ 0.7 \ 0.8)$, $\mathbf{I}^2 = (0.6 \ 0.7 \ 0.4 \ 0.3)$, $\mathbf{I}^3 = (0.2 \ 0.1 \ 0.8 \ 0.9)$, $\mathbf{I}^4 = (0.25 \ 0.25 \ 0.75 \ 0.75)$, $\mathbf{I}^5 = (0.8 \ 0.85 \ 0.2 \ 0.15)$, $\mathbf{I}^6 = (0.28 \ 0.5 \ 0.72 \ 0.5)$, $\mathbf{O}^1 = (1 \ 0)$, $\mathbf{O}^2 = (0 \ 1)$. The FAM parameters $\beta_a, \bar{\rho}_a$ were chosen equal to 0.01 and 0.7, respectively. Furthermore, $M_a = 2$, $N_b = 2$. This particular example was chosen for a purpose. After the first four input/output pairs presentations we would have had the chance to go through all the steps of the algorithm. So by going through this numerical example the students tested their understanding of all the steps that the FAM algorithm undergoes.

Fuzzy ARTMAP has an interesting geometrical interpretation and we felt that it was necessary to emphasize this interpretation to the students. The templates (\mathbf{w}_j^a 's) of nodes in the category representation layer of Fuzzy ARTMAP can be represented as hyperrectangles (rectangles in 2-D). The meaning of this hyperrectangle is that it encloses within each boundaries all the input patterns that chose this node (template) as their representative node (template) and were encoded by it (see Figure 2). This hyperrectangle starts from its trivial size of 0, corresponding to the case where it has encoded a single pattern and it grows as more patterns are encoded by it. The size of the hyperrectangle is limited by the value of the vigilance parameter (ρ_a). The maximum size hyperrectangle is equal to $M_a(1 - \rho_a)$.

Learning in FAM can also be explained geometrically. Figures 2 and 3 depict the two possible cases of learning where a category node j , with a corresponding rectangle representation, learns a new input pattern presented to FAM. In Figure 2, the input pattern represented by $\hat{\mathbf{a}}$ is enclosed within the rectangle's boundaries; so no new learning is required. In Figure 3, the input pattern represented by $\hat{\mathbf{a}}$ is outside the rectangle's boundaries. Then the FAM learning equations expand the rectangle in a way that includes within its boundaries the new input pattern.




Figure 2: Rectangular representation of a FAM category




Figure 3: FAM Learning (Pattern inside the rectangle)




Figure 4: FAM Learning (Pattern outside Rectangle)

In review, ART covers the input space, where the patterns reside, by rectangular regions, each one of which represents a group of similar input patterns. These rectangular regions are then mapped to appropriate output labels. The significance of the label of a rectangular region is that if this region attracts a new input pattern the label of this input pattern coincides with the label of the rectangular region.

We closed our lectures on ART by emphasizing two applications, both dealing with pattern recognition of mushrooms into the “edible” versus “poisonous” class, based on some observed mushroom features (such as cap problems. The first application dealt with the classification shape, cap color, etc.), called *Mushroom database*. The second application dealt with the problem of recognizing capital letters belonging to different fonts, called *Letters database*. In both applications the measures of performance for ART were the size of the architecture created (smaller is better) and the generalization performance of the trained network (bigger is better). Generalization performance is defined to be the networks’ performance on a set of data (called test set) different than the data used in the training of the network.

3.2 ART Neural Network Homework

Homework was assigned at the end of the ART lectures to reinforce the concepts covered in class and to help students understand topics that time did not allow to cover in class. An example

of the latter type of a homework problem is discussed below. The complete homework assignment can be found at our web-site (www.ece.ucf.edu/~crcd).

Homework Problem 1: Present to Fuzzy ARTMAP the following input/output pairs

$$\mathbf{I}^1 = 111010 \quad \mathbf{O}^1 = 01$$

$$\mathbf{I}^2 = 111100 \quad \mathbf{O}^2 = 01$$

$$\mathbf{I}^3 = 111000 \quad \mathbf{O}^3 = 10$$

as many times as it is necessary for the architecture to learn the required mapping. The input/output pairs are presented to Fuzzy ARTMAP in the order depicted above. The network parameters are chosen as follows:

$\beta_a = 0.01$, $\bar{\rho}_a = 0$. Furthermore, choose the initial top-down weight components $w_{ji}^a = 2$ (note that the typical choice for these initial top-down weights is 1). Do you detect any problems in the training process? Is Fuzzy ARTMAP going to learn the required mapping? (Note that the input patterns in the training collections are not complementary encoded).

The purpose of this problem is to demonstrate to the students the importance of complementary encoding the input patterns before presenting them to FAM. Without complementary encoding the input patterns (as it is the case of the problem above) FAM learning can run into problems.

3.3. Current Topics in Machine Learning II Class

The ultimate objective of the CTML-I class is to motivate students in this class to register for CTML-II, which is taught in the Spring semester, following the Fall semester when CTML-I is taught. CTML-II is project oriented and its purpose is to involve undergraduate students in machine learning research. The CTML-II class was structured as follows. At the beginning of the Spring 2004 semester, the CRCDC PIs gave 1 lecture each, covering the projects that were related to their research interests in Machine Learning (obviously these Machine Learning research interests were covered in more detail in the CTML-I class). Then the students were given a week to discuss these potential projects with the corresponding PIs. Hence, a month after the beginning of the semester each student or student group had a research project that they were working on. For the rest of the semester the students interacted on a weekly basis with the professors that were supervising their project. Furthermore, the students also collaborated frequently with graduate students helping them to understand better certain aspects of their project. All the CRCDC students had to present the results of their project at the end of month 2, 3, and 4 of the semester. The last presentation was their final presentation for the project. The students also had to write a final report, where they delineated in detail the results of their completed project. The students were also requested to keep a portfolio of their activities that explained in brief the difficulties they faced, how they resolved them, and the progress that they were experiencing towards the completion of their project. A number of potential projects related to the ART neural networks were explained to the students. One of these projects dealt with the parallel implementation of a variation of ART (see Anagnostopoulos, et al., 2003), using a ring of processors (see Manolakos 1998). One of the students in this class was interested in this project

and he carried it through, under the supervision of one of the CRCO PI faculties and his graduate student.

The project assigned to the student in the CTML-II class dealt with the implementation of a Fuzzy ARTMAP variant (called no-match tracking FAM) on a Beowulf cluster. One of the reasons that such an implementation is important is because Fuzzy ARTMAP, despite its impressive convergence to a solution (i.e., it takes only a few iterations through the training data to converge to a solution) slows down considerably when the size of the training set is large, as it is usually the case in a number of data-mining applications. To remedy this shortcoming, this project focused on an efficient, parallel implementation of the no-match tracking FAM algorithm on a cluster of workstations (Beowulf cluster). In this work, the CRCO student, in close collaboration with a Ph.D. student, proposed an implementation, and demonstrated its good properties theoretically (theorems were postulated and proved) and through experimentation (the parallel algorithm was tested on a number of large datasets). In particular, the no-match tracking FAM was tested on a real-database (Forrest Covertype database from the UCI repository) and on simulated databases (Gaussianly distributed data belonging to two different classes). The results on the Forrest Covertype database are shown in Figure 5. In this figure the speed-up attained by the parallel no-match tracking FAM compared to the sequential no-match tracking FAM implementation is depicted. The speed-up (shown in Figure 5) is demonstrated for training set sizes ranging from 32,000 patterns to 512,000 patterns (in steps of powers of 2), and for 1, 2, 4, 8, 16, and 32 processors of the Beowulf cluster. An obvious observation from this figure is that speed-up achieved using the parallel no-match tracking FAM grows linearly with the number of processors in the Beowulf cluster.


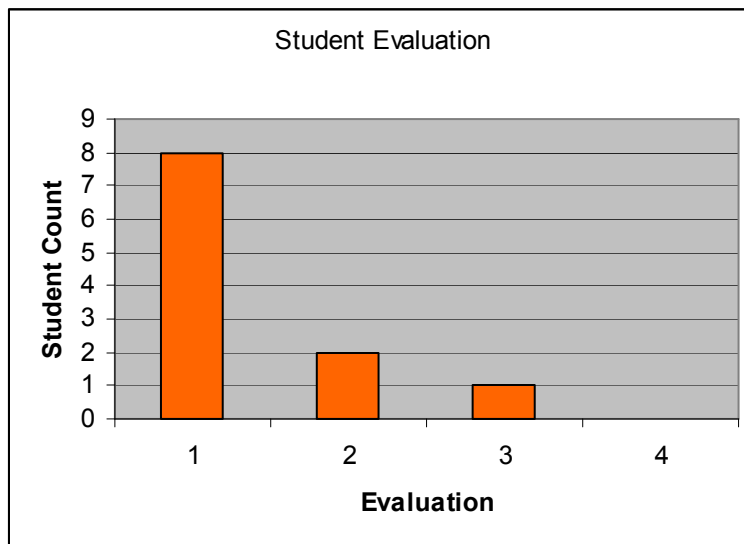


Figure 5: Speed-up performance of the no-match tracking FAM for the CoverType database for different pipeline sizes (1 through 32 in steps of 2) and different database sizes (32,000 to 512,000 in steps of power of 2).

5. Assessment and Evaluation of the Module


A questionnaire was handed out to the students at the end of the CTML-I class. Actually a questionnaire was handed out at the end of every research topic that was discussed with the students (ART Neural Networks, Genetic Algorithms, Evolutionary Hardware, Inductive Reasoning, Simulation Meta-modeling). The ART Neural network questionnaire is included in Appendix I. The responses of three of those students to this questionnaire are included in Appendix II. Each one of these questionnaires resulted in student feedback regarding their perception of how much they learned from the material covered in the class. Some of the students in the class had a problem with the tedious computations that were needed by one of the ART problems in the homework set. In a subsequent offering of the course we reduced the number of computations required by the students for this homework problem. The responses of the students to the rest of the homework assignments in the CTML-I class were of similar nature as the ones depicted for the ART homework set. One telling story of how well the students responded to the material covered in the CTML-I class is that 9 out of the 11 students registered for the class chose to take the CTML-II class as well.

A questionnaire was also handed to the students at the end of the CTML-II class. This questionnaire assessed the process that they have been involved with in the CTML-II class, and how successful this process was in completing the requirements of the project. In the evaluation charts that follow “1” means “STRONGLY AGREE”, “2” means “AGREE”, “3” means “DISAGREE”, and “4” means “STRONGLY DISAGREE”.




Question a: I found the individual time I spent with my professor helpful while I was working on my project


Question f: My professors were available to help me when I needed it as I worked on my project




Question b: I found I could get help from my fellow students while working on my project




Question c: The initial lectures in the CTML-II class helped me to decide on the project chosen



Question d: The CTML-I's lectures helped me as I worked on my project



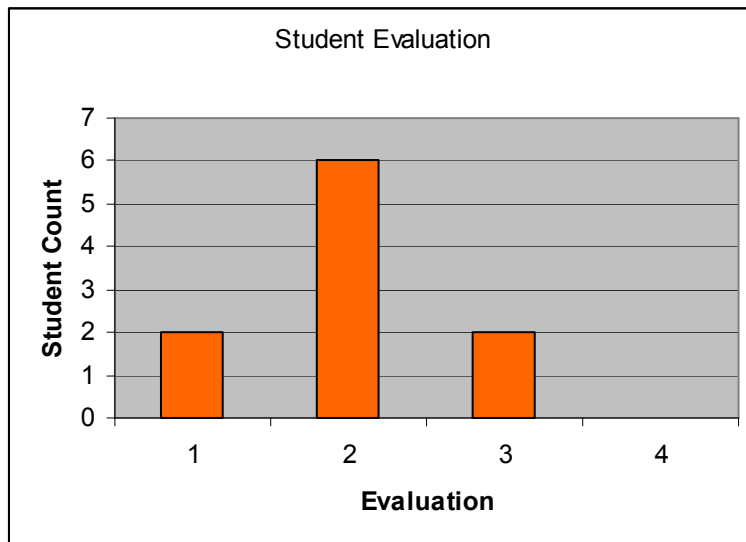
Question e: I found the idea of the portfolio of work helpful as I worked on my project




Question g: I like having to complete a project as part of the class assignments

Question h: I liked the way the class was structured (initially lectures, then work on project, monthly presentations etc.)


Question i: I found the idea of providing input about the progress of my work (on a monthly basis) helpful for me and my other classmates




Question j: I would have preferred if the project was assigned to me earlier



Question k: I found the papers that were assigned for me to read helpful in completing the project



Question l: I would be willing to take another class like this



Question m: I think all the professors were effective in explaining their ideas

As it can be seen from the above evaluations the students responded very positively to the research projects, and the process associated with completing them, assigned in the CTML-II class. It is worth noting that the students who participated in the CTML-I and CTML-II classes had some initial interest within the field of AI, Machine Learning, robotics and related fields. We believe that the CTML-I and II courses reinforced their interest in these fields. Of the nine undergraduate students that took the CTML-II class three of them were juniors and they are continuing their studies at UCF, and also continuing in performing machine learning research under the supervision of the CRCRD faculty that advised their work in the CTML-II class. Three of them decided to pursue graduate work at UCF (two at the Ph.D. level within the field of Machine Learning, one of them at the Masters level in the field of communications). One of them chose to work in the industry and he has plans of enrolling at the Masters program at UCF when the opportunity arises. The last undergraduate student was an Army Officer and he went back to his home base after graduating with a B.S. degree in the Spring of 2004. There were also two graduate students in the CTML-II class that they are continuing to pursue their graduate studies at UCF.

6. Summary and Conclusions

We have reported in this effort some of the work conducted towards the satisfaction of some of the deliverables of a funded NSF project entitled “Advances in Machine Learning for Engineering Education”. Our focus in this paper has been the teaching of a newly introduced class entitled Current Topics in Machine Learning I. In this class the PIs in this effort produce and teach material related to their current research in Machine Learning. The ultimate goal is to motivate these students to take the continuation of this class (Current Topics in Machine Learning II) where they will work, in consultation with the PIs, on a research project of their

own. In this paper we have extensively discussed the teaching material produced, and its delivery on one of the research topics of interest, the ART neural networks topic. Our assessment and evaluation results indicated that the students liked the sequence of these Machine Learning classes, they learnt from them, and they were (most of them) motivated by them to pursue graduate studies. More detailed feedback from the students regarding this CRCD experience is provided in the Appendices, and the charts provided.

References

1. Anagnostopoulos, G. C., and Georgiopoulos, M. "Putting the utility of match-tracking in Fuzzy ARTMAP to the test," In *Proceedings of the Seventh International Conference on Knowledge-Based Intelligent Information (KES)*, London, England, 2003
2. Anagnostopoulos, G. C., and Georgiopoulos, M., "Ellipsoid ART and ARTMAP for incremental unsupervised and supervised learning," In *Proceedings of the IEEE-INNS-ENNS conference*, Vol. 3, pp. 1221-1226, Washington-DC, 2001
3. Carpenter, G. A., Grossberg, S., Markuzon, N., Reynolds, J., and Rosen, D. B., "Fuzzy ARTMAP: A neural network architecture for incremental supervised learning of analog multi-dimensional maps," *IEEE Transactions on Neural Networks*, Vol. 3, pp. 698-713, 1992.
4. Grossberg, S., "Adaptive pattern classification and universal recoding I: Parallel development and coding of neural feature detectors," *Biological Cybernetics*, Vol. 23, pp. 121-134, 1976.
5. Manolakos, E. S. Parallel Implementation of ART1 Neural Networks on Processor Ring Architectures . Eds N. Sundararajan, P. Saratchandran. IEEE Computer Society Press, 1998.

Acknowledgement

The authors from the University of Central Florida acknowledge the partial support from NSF through a CRCD grant number 0203446. M. Georgiopoulos, G. Anagnostopoulos and Jimmy Secretan acknowledge also the partial support of the NSF under grant number #CCLI:0341601.

Biographical Information

MICHAEL GEORGIOPOULOS is a Professor of the School of Electrical Engineering and Computer Science at the University of Central Florida. His research interests lie in the areas of neural networks and applications of neural networks in pattern recognition, image processing, smart antennas and data-mining. He is an Associate Editor of the IEEE Transactions on Neural Networks since 2001.

EROL GELENBE is a Professor at the Imperial College (Dennis Gabor Chair and Head of Intelligent Systems and Networks) in London. He is also a Research Professor at the University of Central Florida. He is a Fellow of IEEE

and a Fellow of ACM. His research interests cover packet network design, computer performance analysis, artificial neural networks and simulation with enhanced reality.

RONALD DEMARA is an Associate Professor at the School of Electrical Engineering and Computer Science at the University of Central Florida. He has been a reviewer for National Science Foundation, Journal of Parallel and Distributed Computing, IEEE Transactions on Parallel and Distributed Computing. His interests lie in the areas of Parallel and distributed processing, self-timed architectures.

AVELINO GONZALEZ is a Professor of the School of Electrical Engineering and Computer Science at the University of Central Florida. He has co-authored a book entitled, "The Engineering of Knowledge-Based Systems: Theory and Practice". His research interests lie in the areas of artificial intelligence, context based behavior and representation, temporal reasoning, intelligent diagnostics and expert systems.

MARCELLA KYSILKA is a Professor and Assistant Chair of the Education Foundations Department at the University of Central Florida. She is active in her professional organizations and currently serves as Associate Editor of the "Journal of Curriculum and Supervision" (the scholarly journal of the Association for Supervision and Curriculum Development). Her research interests are in curriculum studies.

MANSOOREH MOLLAGHASEMI is an Associate Professor at the Industrial Engineering and Management Sciences (IEMS) Department at the University of Central Florida. She has co-authored three books in the area of Multiple Objective Decision Making. Her research interests lie in Simulation Modeling and Analysis, Optimization, Multiple Criteria Decision Making, Neural Networks and Scheduling.

ANNIE WU is an Assistant Professor at the School of Electrical Engineering and Computer Science at the University of Central Florida. Her research interests are in the areas of genetic algorithms, machine learning, biological modeling, and visualization

INGRID RUSSELL is a Professor of Computer Science at the University of Hartford. Her research interests are in the areas of artificial neural networks, pattern recognition, semantic web technologies, and computer science education. She has been involved in several computer science curriculum projects. Most recently she chaired the Intelligent Systems focus group of the IEEE-CS/ACM Task Force on Computing Curricula 2001.

GEORGIOS ANAGNOSTOPOULOS is an Assistant Professor in the Department of Electrical and Computer Engineering at the Florida Institute of Technology. His research interests lie in the areas of Neural Networks, Machine Learning and Pattern Recognition.

JIMMY SECRETAN is a Ph.D. student in the Department of Electrical and Computer Engineering at the University of Central Florida. His research interests lie in the areas of Machine Learning and cluster computing.

Appendix I
 Sample Evaluation
 ART Neural Networks
 EEL 4932
 Current Topics in Machine Learning I

Answer each question honestly by circling the number that best represents your ability to meet the objective. 1= strongly agree; 2 = agree; 3 = unsure; 4 = disagree and 5 = strongly disagree. Responses will not affect your grade. We are trying to evaluate the effectiveness of the course.

Content Knowledge

Question	Response
1. I can differentiate ART neural networks from other neural networks and can identify appropriate applications of ART NNs.	1 2 3 4 5
2. I can explain the functionality of FAM parameters and the training steps in FAM.	1 2 3 4 5
3. I can apply the major steps of FAM's performance phase to given examples.	1 2 3 4 5
4. I can apply the FAM training steps for input/output pairs of the FAM training example given in the notes, as well as other examples.	1 2 3 4 5
5. I can apply the equations that verify the geometry of FAM on any given example.	1 2 3 4 5
6. I feel comfortable in re-designing the training phase of the FAM example given in the homework, using the geometrical perspective of FAM.	1 2 3 4 5
7. I can explain the connection between the size of a new and old rectangle in FAM.	1 2 3 4 5
8. I can discriminate between the on-line and off-line training phase of FAM within the application domain of the mushroom database.	1 2 3 4 5
9. I can discriminate between the one-classifier and the multiple	1 2 3 4 5

classifier results within the application domain of the letters database

10. The homework was practice for what we learned in class. 1 2 3 4 5
11. Class discussions clarified concepts presented in class. 1 2 3 4 5

Homework Activities

1. Which homework activities were the easiest to comprehend and complete?
(Provide the number or name)

2. Which homework activities were the most difficult to complete? Explain why
(Provide the number or name)

3. Which of the homework activities did you like doing?
(Provide the number or name)

4. Which of the homework activities did you think were not necessary? Explain why.
(Provide the number or name)

5. Did the homework activities spark an interest in learning more about ART NNs?

6. What suggestions would you make to improve the homework activities?

Demographics:

Please provide the following information:

Gender Male Female

Level Junior Senior

What is your current GPA? In major _____ Overall _____

List the previous courses, skills that helped you in completing the assignments associated with this chapter.

What are your immediate goals after you graduate?

Has this class helped you rethink your goals after graduation? If so, how?

Appendix II

Evaluation 1

ART Neural Networks

EEL 4932

Current Topics in Machine Learning I

Answer each question honestly by circling the number that best represents your ability to meet the objective. 1= strongly agree; 2 = agree; 3 = unsure; 4 = disagree and 5 = strongly disagree. Responses will not affect your grade. We are trying to evaluate the effectiveness of the course.

Content Knowledge

Question	Response
1. I can differentiate ART neural networks from other neural networks and can identify appropriate applications of ART NNs.	1 (2) 3 4 5
2. I can explain the functionality of FAM parameters and the training steps in FAM.	1 (2) 3 4 5
3. I can apply the major steps of FAM's performance phase to given examples.	(1) 2 3 4 5
4. I can apply the FAM training steps for input/output pairs of the FAM training example given in the notes, as well as other examples.	(1) 2 3 4 5
5. I can apply the equations that verify the geometry of FAM on any given example.	1 (2) 3 4 5
6. I feel comfortable in re-designing the training phase of the FAM example given in the homework, using the geometrical perspective of FAM.	1 (2) 3 4 5
7. I can explain the connection between the size of a new and old rectangle in FAM.	1 (2) 3 4 5
8. I can discriminate between the on-line and off-line training phase of FAM within the application domain of the mushroom database.	1 (2) 3 4 5
9. I can discriminate between the one-classifier and the multiple	1 2 (3) 4 5

classifier results within the application domain of the letters database

10. The homework was practice for what we learned in class. 1 (2) 3 4 5
11. Class discussions clarified concepts presented in class. (1) 2 3 4 5

Homework Activities

7. Which homework activities were the easiest to comprehend and complete?
(Provide the number or name)

The first five questions were the easiest to understand but the first four were the easiest to complete.

8. Which homework activities were the most difficult to complete? Explain why
(Provide the number or name)

I thought problem 8 was the hardest to complete. It took me a while to find a set of input/output pairs when presented to the Fuzzy ARTMAP created rectangles that overlap.

9. Which of the homework activities did you like doing?
(Provide the number or name)

The questions that required me to work through the algorithm (5,6,7,8, and 10) were the questions I liked to do.

10. Which of the homework activities did you think were not necessary? Explain why.
(Provide the number or name)

If I had to pick I would say 1 and 2 but even then if you didn't pick up on those in class you wouldn't have learned it.

11. Did the homework activities spark an interest in learning more about ART NNs?

Yes, I found the whole topic very interesting.

12. What suggestions would you make to improve the homework activities?

I think all together it was a really good homework assignment.

Demographics:

Please provide the following information:

Gender Male Female

Level Junior Senior

What is your current GPA? In major 3.44 Overall 3.1

List the previous courses, skills that helped you in completing the assignments associated with this chapter.

- Intro to Computer Engineering
- Understanding algorithms
- Algebra

What are your immediate goals after you graduate?

Find a job in Computer Engineering and hopefully have the company pay for me to get my Masters degree.

Has this class helped you rethink your goals after graduation? If so, how?

I wasn't sure exactly what more I would like to study about in Graduate school but now I would like to know more about Machine Learning.

Evaluation 2
 ART Neural Networks
 EEL 4932
 Current Topics in Machine Learning I

Answer each question honestly by circling the number that best represents your ability to meet the objective. 1= strongly agree; 2 = agree; 3 = unsure; 4 = disagree and 5 = strongly disagree. Responses will not affect your grade. We are trying to evaluate the effectiveness of the course.

Content Knowledge

Question	Response
1. I can differentiate ART neural networks from other neural networks and can identify appropriate applications of ART NNs.	2
2. I can explain the functionality of FAM parameters and the training steps in FAM.	2
3. I can apply the major steps of FAM's performance phase to given examples.	1
4. I can apply the FAM training steps for input/output pairs of the FAM training example given in the notes, as well as other examples.	1
5. I can apply the equations that verify the geometry of FAM on any given example.	2
6. I feel comfortable in re-designing the training phase of the FAM example given in the homework, using the geometrical perspective of FAM.	2
7. I can explain the connection between the size of a new and old rectangle in FAM.	1
8. I can discriminate between the on-line and off-line training phase of FAM within the application domain of the mushroom database.	1
9. I can discriminate between the one-classifier and the multiple classifier results within the application domain of the letters database	3

10. The homework was practice for what we learned in class. 1
11. Class discussions clarified concepts presented in class. 1

Homework Activities

13. Which homework activities were the easiest to comprehend and complete?
(Provide the number or name)

The easiest to comprehend were definitely problems 1 – 4, as they were restatements or explanations of learned material about FAM.

14. Which homework activities were the most difficult to complete? Explain why
(Provide the number or name)

Problems 5 and 6 were definitely the most time consuming, as they required manually stepping through the algorithm and showing all the math. However, the most difficult was problem 9. First, even though my intuition told me it was true, I had to prove it to myself, which I wasn't entirely sure how to do. However, once I did this, I could use one of the examples that I demonstrated to myself for the solution.

15. Which of the homework activities did you like doing?
(Provide the number or name)

Problems 8 and 10 were both very interesting. I was suspicious that problem 8 could be done but had not had the chance to come up with an example on my own time. I had also wondered about number 10 in regards to if and how much the order could affect the trained weights or even the number of nodes.

16. Which of the homework activities did you think were not necessary? Explain why.
(Provide the number or name)

The only one I felt was unnecessary was problem 5. Since this was just “plugging-and-chugging,” there was not any critical thinking behind it. Problem 6 could clearly demonstrate that the student understood the training process, and it required some analysis thereof, where as 5 was more or less mindless computation

17. Did the homework activities spark an interest in learning more about ART NNs?

To some extent, problem 6 did. More than anything, it made me wonder about the different parameters and initial values. Specifically, what if they were given values different than the ones usually used or given in class?

18. What suggestions would you make to improve the homework activities?

The only suggestion I have would be to remove problem 5 and maybe one other problem. I found myself not spending as much time on each question because 5 and 6 took so long and I just wanted to be done with it. However, as boring as they are, questions like 5 and 6 are often necessary.

Demographics:

Please provide the following information:

Gender Male Female

Level Junior Senior

What is your current GPA? In major 3.9 Overall 3.9

List the previous courses, skills that helped you in completing the assignments associated with this chapter.

In my Topics in Computer Science class (COT 4810), I did a presentation on neural networks, which familiarized me with some of the very general ideas behind neural networks.

What are your immediate goals after you graduate?

Graduate school.

Has this class helped you rethink your goals after graduation? If so, how?

It has helped me realize that my choice to pursue research and academia over a professional career in the private sector is the right one for me. If I wasn't dealing with interesting, engrossing subject matter like this, I would not enjoy myself or my job.

Evaluation 3
 ART Neural Networks
 EEL 4932
 Current Topics in Machine Learning I

Answer each question honestly by circling the number that best represents your ability to meet the objective. 1= strongly agree; 2 = agree; 3 = unsure; 4 = disagree and 5 = strongly disagree. Responses will not affect your grade. We are trying to evaluate the effectiveness of the course.

Content Knowledge

Question	Response
1. I can differentiate ART neural networks from other neural networks and can identify appropriate applications of ART NNs.	1 2 3 4 5
2. I can explain the functionality of FAM parameters and the training steps in FAM.	1 2 3 4 5
3. I can apply the major steps of FAM's performance phase to given examples.	1 2 3 4 5
4. I can apply the FAM training steps for input/output pairs of the FAM training example given in the notes, as well as other examples.	1 2 3 4 5
5. I can apply the equations that verify the geometry of FAM on any given example.	1 2 3 4 5
6. I feel comfortable in re-designing the training phase of the FAM example given in the homework, using the geometrical perspective of FAM.	1 2 3 4 5
7. I can explain the connection between the size of a new and old rectangle in FAM.	1 2 3 4 5
8. I can discriminate between the on-line and off-line training phase of FAM within the application domain of the mushroom database.	1 2 3 4 5
9. I can discriminate between the one-classifier and the multiple classifier results within the application domain of the letters database	1 2 3 4 5

10. The homework was practice for what we learned in class. 1 2 3 4 5
11. Class discussions clarified concepts presented in class. 1 2 3 4 5

Homework Activities

19. Which homework activities were the easiest to comprehend and complete?
(Provide the number or name)
#3, Explaining the steps of the training phase was the easiest. We were well prepared to answer that question given the amount we covered it.
20. Which homework activities were the most difficult to complete? Explain why
(Provide the number or name)
#5 was not so much hard but time consuming. At least for me, completing all of the steps necessary to show that it functionally completed. That involved 2 epochs and took a bit to set up and show all nodes. I would have rather programmed it rather than show it by hand.
21. Which of the homework activities did you like doing?
(Provide the number or name)
#5 was also one of the more enjoyable. Completing it took awhile, but I was able to include elements of automation that made it more fun.
22. Which of the homework activities did you think were not necessary? Explain why.
(Provide the number or name)
#7 was mostly understood, so it was not necessary to cover in homework. If it falls within the bounding box of the element, then it is in that box, and no change to the box. Yet, if it is outside the box and close, then it will expand the box.
23. Did the homework activities spark an interest in learning more about ART NNs?
It sparked my interest in that I wanted to program a method for implementing the ART NN instead of manually doing the homework. I will most likely implement one just for fun or for the vehicle I am working on.
24. What suggestions would you make to improve the homework activities?
Like Dr. Gonzalez, have a section to implement the ART in software. I find that sometimes implementing the algorithm is more useful than doing it by hand, since then you must learn the ins and outs to fully complete it.

Demographics:

Please provide the following information:

Gender Male Female

Level Junior Senior

What is your current GPA? In major 3.957 Overall 3.972

List the previous courses, skills that helped you in completing the assignments associated with this chapter.

A lot of them....I'd have to say the AI class in CS and some of the data structures classes.

What are your immediate goals after you graduate?

Hopefully to get a Ph.D. in Computer Science and work with robotic vehicles for future unmanned military and civilian projects to organize human movement.

Has this class helped you rethink your goals after graduation? If so, how?

I will try to use this course to organize and implement more intelligence methods for my future study and profession instead of the more direct and tried approaches to robotics.