

# A Data Partitioning Approach to speed up the Fuzzy ARTMAP algorithm using the Hilbert space-filling Curve

José Castro  
Department of Computer Engineering  
University of Central Florida  
Orlando, FL 32816-2786  
e-mail: jcastro@pegasus.cc.ucf.edu

Michael Georgiopoulos  
Department of Electrical and  
Computer Engineering  
University of Central Florida  
Orlando, FL 32816-2786  
e-mail: michaelg@mail.ucf.edu

Ronald Demara  
Department of Computer Engineering  
University of Central Florida  
Orlando, FL 32816-2786  
e-mail: demara@pegasus.cc.ucf.edu

Ronald Demara  
Department of Computer Engineering  
University of Central Florida  
Orlando, FL 32816-2786  
e-mail: demara@pegasus.cc.ucf.edu

**Abstract**—One of the properties of FAM, which is a mixed blessing, is its capacity to produce new neurons (*templates*) on demand to represent classification categories. This property allows FAM to automatically adapt to the database without having to arbitrarily specify network structure, but it also has the undesirable side effect that on large databases it can produce a large network size that can dramatically slow down the algorithms training time. To address this problem we propose the use of the Hilbert space-filling curve. Our results indicate that the Hilbert space-filling curve can reduce the training time of FAM by partitioning the learning set without a significant effect on the classification performance or network size. Given that there is full data partitioning with the HSFC we implement and test a parallel implementation on a Beowulf cluster of workstations that further speeds up the training and classification time on large databases.

## I. INTRODUCTION

Neural Networks have been used extensively and successfully to attack a wide variety of problems. As computing capacity and electronic databases grow, there is an increasing need to process considerably larger databases. In this context, the algorithms of choice tend to be ad-hoc algorithms or tree based algorithms such as CART and C4.5 [9]. Variations of these tree learning algorithms, such as SPRINT (Shafer, et al., [10]) and SLIQ (Mehta, et al., [7]) have been successfully adapted to handle very large data sets.

Neural network algorithms have a higher computational complexity and for some applications have prohibitive convergence times. Even one of the fastest training time neural network algorithms, the Fuzzy ARTMAP (FAM) algorithm, tends to lag in convergence time as the size of the network grows.

The FAM algorithm corresponds to a family of neural

network architectures introduced by Carpenter, et al., 1991-1992 [3], [2] and has proven to be one of the premier neural network architectures for classification problems. Some of the advantages that FAM has, compared to other neural network classifiers, are that it learns the required task fast, it has the capability to do on-line learning, and its learning structure allows the explanation of the answers that the neural network produces.

One of its properties which is a mixed blessing, is its capacity to produce new neurons (*templates*) on demand to represent classification categories. This property allows FAM to automatically adapt to the database without having to arbitrarily specify network structure, but it also has the undesirable side effect that on large databases it can produce a large network size that can dramatically slow down the algorithms training time. It would be desirable to have a method capable of keeping the training set on a manageable size without seriously affecting FAMs convergence, classification and generalization properties.

In this paper we propose the use of the Hilbert space-filling curves for processing the training set to be used with FAM classification (hFAM). Our research on Hilbert space-filling curves has shown that they can dramatically reduce the training time of FAM by partitioning the training set without a significant effect on the classification performance or network size. Skopal et al. [4] analyze different space-filling curves, amongst them the Peano curve, Z curve and the Hilbert curve, and also provide measures for their appropriateness. Moon et al. [8] argue and prove that the Hilbert space-filling curve is the mapping that provides the *least number of splits* of compact sets from  $[0, 1]^n$  to  $[0, 1]$ . This can be interpreted as stating that points that are close on the mapping will also be

close on the  $n$ -dimensional hypercube. Lawder [6] has taken advantage of this and used the Hilbert space-filling curve to develop a multi-dimensional indexing technique.

This paper is organized as follows: First we review the Fuzzy ARTMAP architecture and parameters, then we examine the computational complexity of FAM and analyze how and why a partitioning approach can considerably reduce the algorithms training time. After that we discuss space-filling curves in general and the Hilbert space-filling curve in particular and why this curve can be instrumental in improving the FAM algorithm's convergence time. Furthermore, experimental results are presented on a sequential machine and on a Beowulf cluster of workstations that illustrate the merit of our approach. We close the paper with a summary of the findings and suggestions for further research.

## II. THE FUZZY ARTMAP ARCHITECTURE

The Fuzzy ARTMAP architecture consists of four layers or fields of nodes (see Figure 1). The layers that are worth describing are 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 of dimensionality  $2M_a$  of the following form is applied

$$\mathbf{I} = (\mathbf{a}, \mathbf{a}^c) = (a_1, a_2, \dots, a_{M_a}, a_1^c, a_2^c, \dots, a_{M_a}^c) \quad (1)$$

where:

$$a_i^c = 1 - a_i; \forall i \in \{1, 2, \dots, M_a\} \quad (2)$$

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 layer  $F_2^a$  of Fuzzy ARTMAP 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 FAM is supposed to be mapped to.

There are two sets of weights worth mentioning in FAM. The first set of weights are weights from  $F_2^a$  to  $F_1^a$ , designated as  $w_{ji}^a$ , ( $1 \leq j \leq N_a, 1 \leq i \leq 2M_a$ ), and referred to as top-down weights. 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 to represent the group of input patterns that chose node  $j$  in the category representation layer of Fuzzy ARTMAP as their representative node. The second set of weights, worth mentioning, 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.,  $\mathbf{W}_j^{ab} = [W_{j1}^{ab}, W_{j2}^{ab}, \dots, W_{j,N_b}^{ab}]$ ) corresponds to the output pattern that this node  $j$  is mapped to.

Fuzzy ARTMAP can operate in two distinct phases: the *training phase* and the *performance phase*. 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

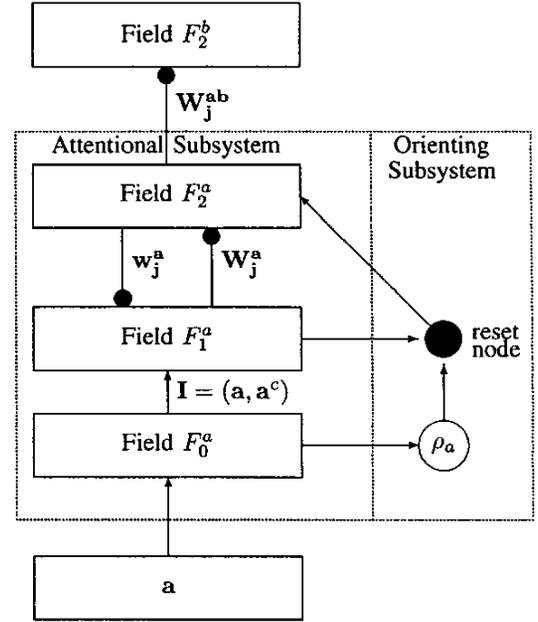


Fig. 1. Fuzzy ARTMAP Diagram

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 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.

The operation of Fuzzy ARTMAP is affected by two network parameters, the choice parameter  $\beta_a$ , and the baseline vigilance parameter  $\bar{\rho}_a$ . The choice parameter takes values in the interval  $(0, \infty)$ , while the baseline vigilance parameter 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  $\beta_a$  and  $\bar{\rho}_a$  create more nodes in the category representation layer of Fuzzy ARTMAP, and consequently produce less compression of the input patterns. There are two other network parameter values in Fuzzy ARTMAP that are worth mentioning. The vigilance parameter  $\rho_a$ , and the number of nodes  $N_a$  in the category representation layer of Fuzzy ARTMAP. The

vigilance parameter  $\rho_a$  takes value in the interval  $[\bar{\rho}_a, 1]$  and its initial value is set to be equal to  $\bar{\rho}_a$ . The number of nodes  $N_a$  in the category representation layer of Fuzzy ARTMAP increases while training the network and corresponds to the number of committed nodes in Fuzzy ARTMAP plus one uncommitted node.

Before training the top-down weights (the  $w_{ji}^a$ 's) of Fuzzy ARTMAP are set equal to 1, and the inter-ART weights (the  $W_{jk}^{ab}$ 's) are chosen equal to 0. One of the specific operands involved in all of these operations is the *fuzzy min operand*, designated by the symbol  $\wedge$ . Actually, the fuzzy min operation of two vectors  $x$ , and  $y$ , designated as  $x \wedge y$ , is a vector whose components are equal to the minimum of components of  $x$  and  $y$ . Another specific operand involved in these equations is designated by the symbol  $|\cdot|$ . In particular,  $|x|$  is the size of a vector  $x$  and is defined to be the sum of its components.

#### A. The Fuzzy ARTMAP Learning Algorithm

**Operation 1:** Calculation of bottom up inputs to every node  $j$  in  $F_2^a$ , as follows:

$$T_j^a = \frac{|\mathbf{I}^r \wedge \mathbf{w}_j^a|}{\beta_a + |\mathbf{w}_j^a|} \quad (3)$$

after calculation of the bottom up inputs the node  $j_{max}$  with the maximum bottom up input is chosen.

**Operation 2:** The node  $j_{max}$  with the maximum bottom up input is examined to determine whether it passes the vigilance criterion. A node passes the vigilance criterion if the following condition is met:

$$\frac{|\mathbf{I}^r \wedge \mathbf{w}_j^a|}{|\mathbf{I}^r|} \geq \rho_a \quad (4)$$

if the vigilance criterion is satisfied we proceed with operation 3 otherwise node  $j_{max}$  is disqualified and we find the next node in sequence in  $F_2^a$  that maximizes the bottom up input. Eventually we will end up with a node  $j_{max}$  that maximizes the bottom up input and passes the vigilance criterion.

**Operation 3:** This operation is implemented only after we have found a node  $j_{max}$  that maximizes the bottom-up input of the remaining nodes in competition and that passes the vigilance criterion. Operation 3 determines whether this node  $j_{max}$  passes the prediction test. The prediction test checks if the inter-ART weight vector emanating from node  $j_{max}$  (i.e.,  $\mathbf{W}_{j_{max}}^{ab} = [W_{j_{max}1}^{ab}, W_{j_{max}2}^{ab}, \dots, W_{j_{max},N_b}^{ab}]$ ) matches exactly the desired output vector  $\mathbf{O}^r$  (if it does this is referred to as *passing the prediction test*). If the node does not pass the prediction test, the vigilance parameter  $\rho_a$  is increased to the level of  $\frac{|\mathbf{I}^r \wedge \mathbf{w}_j^a|}{|\mathbf{I}^r|} + \epsilon$  where  $\epsilon$  is a very small number, node  $j_{max}$  is disqualified, and the next in sequence node that maximizes the bottom-up input and passes the vigilance is chosen. If, on the other hand, node  $j_{max}$  passes the predictability test, the weights in Fuzzy ARTMAP are modified as follows:

$$\mathbf{W}_{j_{max}}^a \leftarrow \mathbf{W}_{j_{max}}^a \wedge \mathbf{I}^r, \quad \mathbf{W}_{j_{max}}^{ab} \leftarrow \mathbf{O}^r \quad (5)$$

Fuzzy ARTMAP training is considered complete if and only if after repeated presentations of all training input/output pairs

to the network no weight changes are produced. In some databases noise in the data may create over-fitting so a single pass over the training set may be preferable.

In the performance phase of Fuzzy ARTMAP only Operations 1 and 2 are implemented for every input pattern presented to Fuzzy ARTMAP.

#### B. Fuzzy ARTMAP pseudocode

We will be interested in classification problems where we associate input patterns to category labels. The following pseudo code algorithm makes use of these assumptions and simplifies accordingly the FAM algorithm.

FAM-TRAINING-PHASE( $Patterns, \bar{\rho}, \beta_a, \epsilon$ )

```

1  templates ← {}
2  for each I in Patterns
3    do ρ ←  $\bar{\rho}$ 
4    repeat
5       $T_{max} \leftarrow \frac{M}{2M + \beta_a}$ 
6      status ← FoundNone
7      for each  $w_j$  in templates
8        do if  $\rho(I, w_j) \geq \rho$ 
9           and  $T(I, w_j, \beta_a) > T_{max}$ 
10          then
11             $T_{max} \leftarrow T(I, w_j, \beta_a)$ 
12             $j_{max} \leftarrow j$ 
13            status ← FoundOne
14          if status = FoundOne
15            then if  $class(I) = class(w_{j_{max}})$ 
16                  then status ← ThisIsIT
17                  else status ← TryAgain
18                   $\rho \leftarrow \rho(I, w_{j_{max}}) + \epsilon$ 
19          until status ≠ TryAgain
20          if status = ThisIsIT
21            then
22               $w_{j_{max}} \leftarrow w_{j_{max}} \wedge I$ 
23            else
24              templates ← templates  $\cup \{I\}$ 
25          return templates

```

Where  $T(I, w, \beta)$  is defined by equation 3 and  $\rho(I, w)$  is defined by the left hand side of inequality 4.

#### C. FAM time complexity analysis

If we call  $\Gamma$  the average number of times that the **repeat** loop is executed for each input pattern. Then the number of times that a given input pattern  $I$  passes through the code will be:

$$Time(I) = O(\Gamma \times |Templates|) \quad (6)$$

Also, under the artificial condition that the number of templates does not change during training it is easy to see that the time complexity of the algorithm is:

$$Time(FAM) = O(\Gamma \times PT \times |Templates|) \quad (7)$$



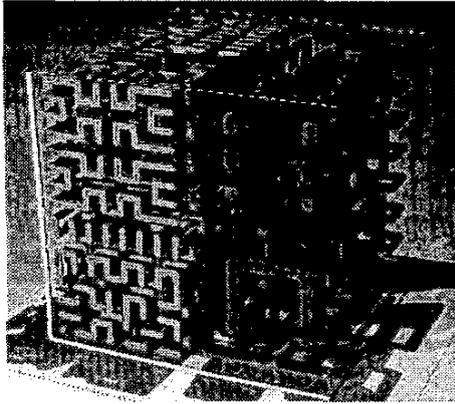


Fig. 5. Hilbert 3D

Machine Learning Repository. The number of attributes of each pattern is 54. There are no missing values on this data.

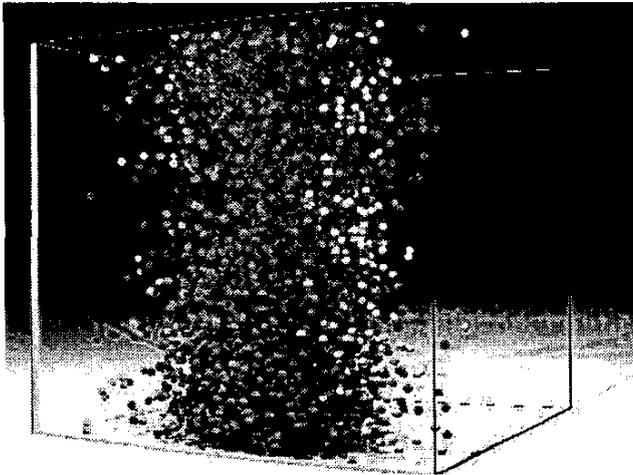


Fig. 6. Visualization of Forest CoverType data

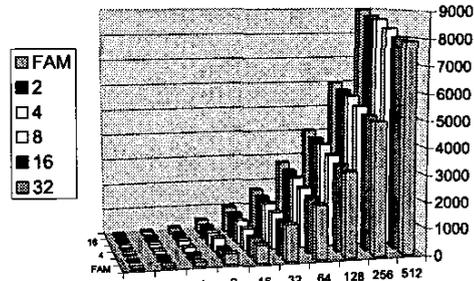
Patterns 1 through 512,000 were used for training, The testing set for all trials where patterns 561,001 to 581,000. A visualization of the first 3 dimensions of the Forest Covertype can be seen in figure 6, different tones correspond to different classes. Classification performance for different machine learning algorithms for this database usually hovers around 75%.

### B. Gaussian Databases

The Gaussian data was artificially generated using the polar form of the Box-Muller transform with the R250 random number generator by Kirkpatrick and Scholl [5]. We generated 2 class 16 dimensional data. 532,000 patterns were generated for each Gaussian database. 512,000 patterns were used for training, the last 20,000 were used for testing. One Gaussian database had a 15% overlap while the other had a 5% overlap.

## V. EXPERIMENTAL RESULTS

In figure 7 we can see a bar graph of the number of templates on the Y axis, the training set size on the X axis (in thousands of patterns), and on the Z axis the number of partitions. Differences on the Z indicate that the number of partitions do not affect considerably the compression ratio. This is also confirmed in figures 8.



	1	2	4	8	16	32	64	128	256	512
FAM	50	80.40625	158.78125	410.34375	863.90625	1347.875	2031.84375	3200.5625	5004.69625	7823.5
2	60	103.25	170.5	341.09375	715.90625	1257.3125	2120.4375	3303.84375	5134.8875	8037.46875
4	83	117.34375	188.40625	449.300625	790.375	1304.4375	2298.375	3480.6875	5384.34375	8322.05125
8	70	129.65625	211.5	437.86875	877.90625	1538.25	2505.84375	3754.96875	6860.75	10317.68375
16	74	138.1875	220.65625	483.5	920.21875	1592.3125	2809.28125	3953.96875	5789.8125	8750.9825
32	80	144.825	237.28125	487.4375	989.5	1696.71875	2781.5	4055.375	6013.0625	8888

Fig. 7. Number of Templates on Covertype Data

The classification of the partitioned data is very similar from the classification of the non-partitioned (1 partition) data figure 9. The Tree Covertype database classification continues to improve up to 512,000 patterns, clearly indicating that the database is sufficiently complex as to need a large training set size. This behavior is not observed on the Gaussian artificial data (5% or 15% data), which peaks performance at 32,000 patterns and makes the classification performance graph flat and uninteresting.

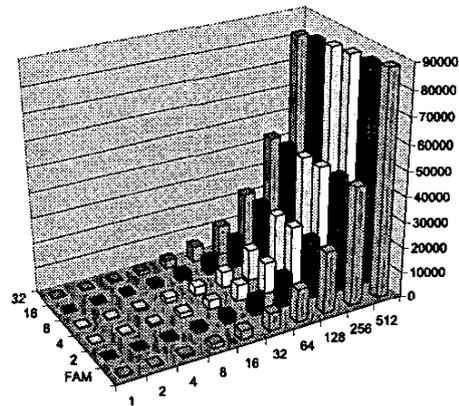


Fig. 8. Number of Templates on 5% Overlap Gaussian Data

Figure 10 shows the speedup of the partitioned FAM with  $p$  partitions running in parallel using equation ?? . The best speedups obtained were in the order of 100.

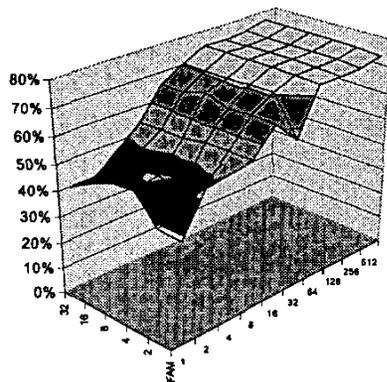


Fig. 9. Classification Performance of Forest Cover database

The speedup for the same data using a sequential processing machine can be seen in figure 11, we can see that the speedup for an single processor is in the order of  $p$  (17 for 32 processors) as indicated by equation 10 and the speedup of the parallel implementation is in the order of  $p^2$  (100 for 32 processors) as indicated by equation 11.

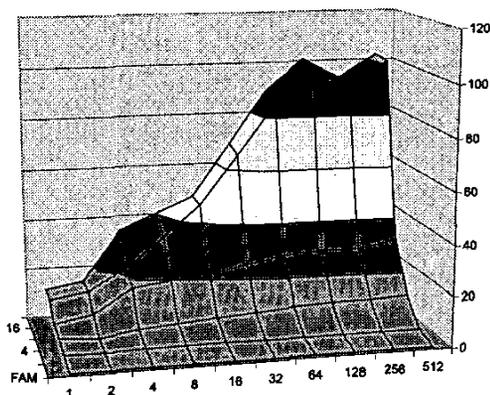


Fig. 10. Hilbert Parallel Partitioning Speedup on Covertyp Data

## VI. CONCLUSIONS

We presented the FAM algorithm applied to classification tasks on large databases and saw that its training time tends to slow considerably when the size of the database grows. By analyzing the algorithm we proposed the use of Hilbert space-filling curves to attack this problems, experimental results on 3 databases confirm our projections: classification performance is not affected, in fact, it is improved for the real database results, although this was not an objective in our study. Compression ratio is only slightly affected, and convergence time is improved linearly on the sequential machine and quadratically on the parallel machine. Nevertheless there is still room for improvement, in this study we applied a network partitioning approach to the training set. We believe that combining this

with a network partitioning approach is the next step to achieve optimal workload balance in the parallel implementation, this is one of the directions of our current research.

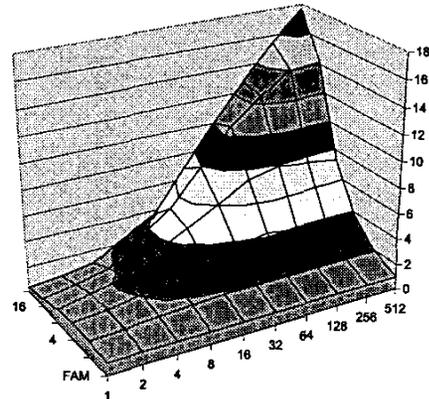


Fig. 11. Hilbert Sequential Partitioning Speedup on Covertyp Data

## ACKNOWLEDGMENT

The authors would like to thank the Institute of Simulation and Training and the Link Foundation Fellowship program for partially funding this project. This work was also supported in part by the National Science Foundation under grant no. 0203446.

## REFERENCES

- [1] J. A. Blackard, "Comparison of neural networks and discriminant analysis in predicting forest cover types," Ph.D. dissertation, Department of Forest Sciences, Colorado State University, 1999.
- [2] G. A. Carpenter, S. Grossberg, N. Markuzon, J. H. Reynolds, and D. B. Rosen, "Fuzzy ARTMAP: A neural network architecture for incremental learning of analog multidimensional maps," *IEEE Transactions on Neural Networks*, vol. 3, no. 5, pp. 698–713, September 1992.
- [3] G. A. Carpenter, S. Grossberg, and J. H. Reynolds, "Fuzzy ART: An adaptive resonance algorithm for rapid, stable classification of analog patterns," in *International Joint Conference on Neural Networks, IJCNN'91*, vol. II. Seattle, Washington: IEEE/INNS Inc, 1991, pp. 411–416.
- [4] *Properties of Space Filling Curves and Usage With UB-Trees*. High Fatra, Slovakia: ITAT 2002, 2002.
- [5] S. Kirkpatrick and E. Stoll, "A very fast shift-register sequence random number generator," *Journal of Computational Physics*, vol. 40, pp. 517–526, 1981.
- [6] J. K. Lawder and P. J. H. King, "Using space-filling curves for multi-dimensional indexing," *Lecture Notes in Computer Science*, vol. 1832, 2000.
- [7] M. Mehta, R. Agrawal, and J. Rissanen, "SLIQ: A fast scalable classifier for data mining," in *Extending Database Technology*, 1996, pp. 18–32. [Online]. Available: [citeseer.nj.nec.com/mehta96slmq.html](http://citeseer.nj.nec.com/mehta96slmq.html)
- [8] B. Moon, H. Jagadish, C. Faloutsos, and J. H. Saltz, "Analysis of the clustering properties of the hilbert space-filling curve," *IEEE Transactions on Knowledge and Data Engineering*, vol. 13, no. 1, JANUARY 2001.
- [9] J. R. Quinlan, *C4.5: Programs for Machine Learning*. San Mateo, California: Morgan Kaufmann, 1993.
- [10] J. C. Shafer, R. Agrawal, and M. Mehta, "SPRINT: A scalable parallel classifier for data mining," in *Proc. 22nd Int. Conf. Very Large Databases, VLDB*, T. M. Vijayaraman, A. P. Buchmann, C. Mohan, and N. L. Sarda, Eds. Morgan Kaufmann, 3–6 1996, pp. 544–555. [Online]. Available: [citeseer.nj.nec.com/shafer96sprint.html](http://citeseer.nj.nec.com/shafer96sprint.html)