JB 94-22

Entitled:

"Wavelet Based Progressive Classification with Learning: Applications to Radar Signals"

Authors:

(with S.I. Wolk)

Conference

SPIE 1995 International Symposium on Aerospace/Defense Sensing and Dual-Use Photonics,
Orlando, Florida
April 17-21, 1995

94-22

PROCEEDINGS

SPIE—The International Society for Optical Engineering

Wavelet Applications II

Harold H. Szu
Chair/Editor

17–21 April 1995
Orlando, Florida

SPIE PROCEEDINGS SERIES

Volume 2491
Part One of Two Parts

Wavelet based progressive classification with learning: applications to radar signals

John S. Baras*
Department of Electrical Engineering
and Institute for Systems Research
University of Maryland
College Park, MD 20742

Sheldon I. Wolk
Code 5750
Tactical Electronic Warfare Division
Naval Research Laboratory
Washington, D.C. 20375

ABSTRACT

In this paper we investigate the problem of fast and accurate classification of naval targets from radar returns. The algorithms can be applied to both 1-D and 2-D data (i.e. high range resolution and imaging radar returns). We describe the structure of these algorithms and report experimental results on their performance with synthetic returns from ships.

We have successfully addressed the problem of reducing the target model representations with respect to viewpoint variations and other sensor parametric variations. Our method can be viewed as a quantization of the space of sensing operations. The resulting multiresolution aspect graph is a (relational) graph representation of this quantization. Aspect graphs of target radar returns are generated algorithmically. Since our off-line model/parameter tuning methods are based on general vector quantization, our methods extend naturally and efficiently to multi-sensor data: LADAR, TV, mmWave, SAR, etc.. We describe new results on the "continuity" of the aspect graph, new properties and improvements of our algorithmic constructions.

Our basic classification algorithm utilizes a cascade of a wavelet preprocesor followed by a tree-structured clustering algorithm; a learning mode can be also added. We develop a high performance parallel progressive classification algorithm and report on its performance and complexity. We show experiments illustrating that the parallel algorithm outperforms a compund version (which is the more intuitive choice) and that it has performance close to Bayes optimal classification (via comparison to Learning Vector Quantization). We ouline an analytical framework for establishing these results theoretically. We also discuss similar experiments from face recognition and medical image classification problems.

Keywords: Wavelets, Tree-structured Vector Quantization. Progressive Classification, ATR, RADAR.

1 INTRODUCTION

High range resolution radar returns contain in their structure substantial information about the target which can be used to better identify complex targets consisting of many scatterers. This applies to many forms of radar signatures, including the amplitude of pulsed radar (PR) returns, the phase of pulsed radar returns, Doppler radars (DR), synthetic aperture radar (SAR) returns, inverse synthetic aperture radar (ISAR) returns, millimeter-wave (MM-wave) radar returns. With the increasing resolution of modern radars it is *at least theoretically* possible to store many of the possible returns (i.e. returns organized according to aspect, elevation, pulsewidth etc.) of a complex target and use them in the field for target identification. The advantage of the increasing radar resolution is the availability of more detailed information, and ultimately of *specific features*, characteristic of the radar return from a specific target. The disadvantage is that these very detailed characteristics require an ever increasing amount of computer memory to be stored. The latter not only results in unfeasible memory requirements but it also slows down the search time in real field operations.

It is therefore important to develop extremely efficient ways to compress the representations of high resolution data returns from real targets, and to design efficient search schemes which operate in a progressive manner on the compressed representations to recover the target identity. Wavelet theory [1]—[9] offers an attractive means for the development of the required multi-resolution representations. This can be roughly explained by the fundamental property of wavelet representations of signals to uncover the superposition of these signals in terms of different structures occurring on different time scales at different times (see [14, 17, 19, 20] for details). Successful methods to provide effective compression of radar returns must address also the substantial variability of the returns with respect to aspect, elevation and radar pulsewidth. It is therefore physically meaningful to cluster the radar returns from various viewpoints into equivalence classes using a measure of similarity. The resulting quantization of the signal space (i.e. of the radar returns) characterizes the limits of discriminating between returns from different targets using information about the viewpoint; in essence if we insist on extremely fine quantization cells we are modeling the radar sensor noise and not the underlying complex target. A systematic application of these two fundamental techniques leads to the development of economic radar target models.

The problem of automatic target recognition based on high range resolution (HRR) radar returns when a large number of targets is possible, presents formidable algorithmic and computational difficulties. A key step in the

*Also with AIMS, Inc.. Research supported by AIMS, Inc., 6110 Executive Blvd., Suite 850, Rockville, MD 20852.

design and implementation of high performance ATR algorithms is the organization and construction of efficient and economic target models which will result in significant search speed-up and memory requirements reduction. In this paper we use as target models scale space aspect graphs constructed using the two basic techniques mentioned above (i.e. wavelets and hierarchical clustering) as describe in our earlier work [14, 17, 19, 20].

In addition to economic target models we must address the development of massively parallel algorithms for ATR which efficiently utilize these target models. In this paper we develop further such algorithms, first announced in our earlier work [19, 20], operating in parallel, where the target data are processed concurrently, as they are received, by a parallel array of processors each tuned to a particular target model. We provide further experimental evidence that such massively parallel algorithms have superior performance than algorithms based on a compound model of the entire target database constructed with the same efficiency.

Our target representations lead naturally to classification schemes that are progressive; and this is an essential feature of our algorithms. That is, a small amount of information, in the form of a coarse approximation to the return, is used first to provide partial classification and progressively finer details are added until satisfactory performance is obtained. This results in a scheme where small amounts of computation are used initially and additional computations are performed as needed, resulting in extremely fast searches while preserving high fidelity in the search. Each target is reperesented by its multiresolution aspect graph [15], which is a quantization (produced by clustering) of the space of HRR returns and view points. Using an efficient Tree-Structured Vector Quantization (TSVQ) algorithm we cluster the returns from the various viewpoints into equivalence classes according to an appropriate discrimination measure. This approach automatically accounts for the discrimination capability of the sensor and in effect it performs a quantization of the sensory data which reduces the data input to the classification algorithm by orders of magnitude. In each equivalent class a "paradigm" is selected and the collection of these typical pulses arranged in a multi-scale tree constitute the target model that is guiding the on-line classification search. In its parallel version our scheme allows for the comparison of the sensor data to many hypothetical targets at the coarse resolution and to a progressively smaller candidate set as the resolution is refined and more detailed features are revealed. We therefore achieve progressively finer classification with a constant computational resource since the coarse comparisons are faster per target hypothesis as compared to the finer ones.

We demonstrate our results by experiments with highly accurate synthetic radar returns from the Naval Research Laboratory code 5750 ship radar return simulator. We show performance comparisons between typical classification results obtained with our algorithm vs results obtained with a Learning Vector Quantization (LVQ) algorithm which approaches the optimal Bayes decision rules [27]. The performance of our parallel algorithm is indeed very close to that of the LVQ algorithm and therefore very close to the optimal Bayes classification. This parallel multi-target model feature of our algorithm also solves the so called "new target" insertion problem. Indeed, the addition of targets is handled by adding the corresponding aspect graphs and by passing the sensor data through the processors representing these additional models; no new expensive off-line training of the algorithm is required.

2 MULTIRESOLUTION ASPECT GRAPHS OF SHIP RADAR RETURNS

A key construct in our approach to economical target model hierarchies is that of an aspect graph [15], which is a hierarchical data structure, indexed by sections of viewpoint, that stores compressed target formats. A general definition of the aspect graph [15] is that it is a graph structure in which there is a node for each *general view* of the object as seen from some maximal connected cell of *viewpoint space*, and there is an arc for each possible transition across the boundary between the cells of two neighboring general views. A general viewpoint is defined as one from which an infinitesimal movement in any possible direction in viewpoint space results in a view that is indistinguishable (by the sensor) from the original. Under this definition, the aspect graph is complete in that it provides an enumeration of the fundamentally different views of an object, yet is minimal in the sense that the cells of general viewpoint are disjoint. Thus the aspect graph is equivalent to a parcellation (tesselation) of viewpoint space into general views. Considerable research has been performed in recent years on algorithms that compute the aspect graph and its related representations [15]. However, todate these conventional methods have addressed only the ideal case of perfect resolution in object shape, in the viewpoint, and the projected image, leading to a set of important practical difficulties [15], [19, 20].

Almost exclusively, previous work on aspect graphs has focused on computer vision and object geometry [11, 15]. Our notion of the aspect graph as developed in [14, 17, 19, 20] is an extension of these concepts to sensors other than cameras such as radar. We present here further improvements and results from our work on the automatic construction of multiresolution aspect graphs of ship radar returns. Our earlier results reported in [14, 17, 19, 20] and in [18] incorporate scale (resolution) in the construction of the aspect graph in a manner consistent with the sensor considered. Indeed we have developed an algorithmic construction of *scale space aspect graphs (or multiresolution aspect graphs*. Scale space aspect graphs are equivalent to families of viewpoint space tesselations parameterized by scale. It is precisely these scale space aspect graphs that we use in this paper as efficient target models to develop high performance model based ATR algorithms. In our automatic construction of these multiresolution aspect graphs we have employed wavelets and Tree Structured Vector Quantization. We refer to [2], [3], [9] for wavelet fundamentals.

We generated radar return databases for various different ships, utilizing the NRL Code 5750 ship radar return simulator. In generating the synthetic data we kept the radar fixed and turned the ship. We varied the aspect angle from 0° to 360° in increments of 0.5°. This allowed for large variation in the number and appearance of dominant scatterers.

Let $\mathcal{S}$ denote the set of discretized radar pulses that we use to construct the multiresolution aspect graph. The fine resolution data will be denoted by $S^0 f(n)$, $n \in I^0$, where $I^0 = \{1, 2, \cdots, 2^7\}$, is the index set of the fine resolution data. We shall let $N = 2^J$ denote the number of samples in the fine resolution data, where J is the maximum possible number of scales that we can consider. In practice one considers scales up to J^* where $J^* < J$. Respectively for each resolution m we denote by I^m the subset of I^0 where sampled values of the m^{th} resolution pulse representation $S^m f$ are computed. I^m is obtained from I^{m-1} by decimation. We used $N = 128$, and $J^* = 3$. This gives us four scales (including the given fine scale) $m = 0, 1, 2, 3$, with vector lengths $128, 64, 32, 16$ and resolutions 10 ns, 20 ns, 40 ns, 80 ns, respectively. We identify $\mathbf{a^0}$ with the vector of sampled data $S^0 f$. Then we use the *pyramid* scheme [2] to recursively compute the successive approximations $S^m f$ to the pulse f at various scales m and the residual pulses $W^m f$. As we proceed with this analysis step from scale m to the coarser scale $m+1$, the space of signals becomes smaller, and the length of vectors is halved. Thus the algorithm recursively splits the initial vector $\mathbf{a^0}$ representing the sampled pulse $S^0 f$ to its components $\mathbf{c^m}$ at different scales indexed by m representing the wavelet residuals $W^m f$; the multiresolution scheme replaces the information in each pulse $f = S^0 f$ with the set $\{W^m f,\ m = 1, 2, \cdots, J^*,\ S^{J^*} f\}$.

We then construct the scale space aspect graph for each ship target by developing a hierarchical, tree-structured organization of radar returns, which utilizes the multiresolution representations provided by wavelets. Vector Quantization (VQ) is primarily used as a data compression method [10]. VQ in addition is a clustering algorithm. Indeed the codewords, represented by the centroids, can be thought of as representatives of the equivalence class represented by each cell of the VQ (each Voronoi cell). It is in this sense that we use VQ in our approach to the problem of hierarchical representations for HRR radar returns.

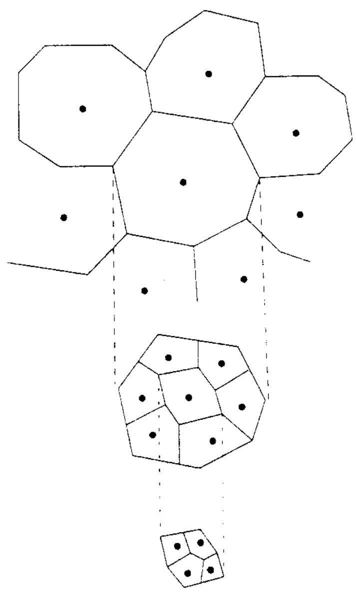

Figure 1: Illustrating a multiresolution TSVQ by splitting Voronoi cells based on different resolution data.

A useful method for designing the tree structure is based on the application the Linde–Buzo–Gray (LBG) algorithm [10] to successive stages using a training set. We have used a variant of this method which is of the "greedy" [13] variation. We first perform a multiresolution wavelet representation of the radar pulses, based on the selection of a mother wavelet. This allows us to consider each pulse reconstructed at different resolutions $S^0 f$, $S^1 f, \cdots, S^{J^*} f$. We then proceed by splitting the signal space at various resolutions in cells as indicated pictorially in Figure 1. The data vector space (signal space) is partitioned into cells, or collections of data vectors which are determined by the repeated application of the Linde-Buzo-Gray (LBG) algorithm. LBG is first applied to the coarsest resolution representation of the data vectors $\{S^{J^*} f,\ f \in \mathcal{S}\}$. The resultant distortion is determined based on a mean squared distance metric, and is computed using the finest resolution representation of the data vectors. The cell (equivalence class of coarse resolution representations) which is the greatest contributor to the total average distortion for the entire partition is the cell which is split in the next application of LBG. A new Voronoi vector is found near the Voronoi vector for the cell to be split and is added to the Voronoi vectors previously used for LBG. LBG is then applied to the entire population of data vectors, again using the coarsest representation of each vector. These steps are repeated until the percentage reduction in distortion for the entire population falls below a predetermined threshold. The partition in the coarsest resolution is then fixed, and further partitioning continues by splitting the cells already obtained based on finer resolution representations of the data vectors in the cell. The algorithm then iterates through these steps

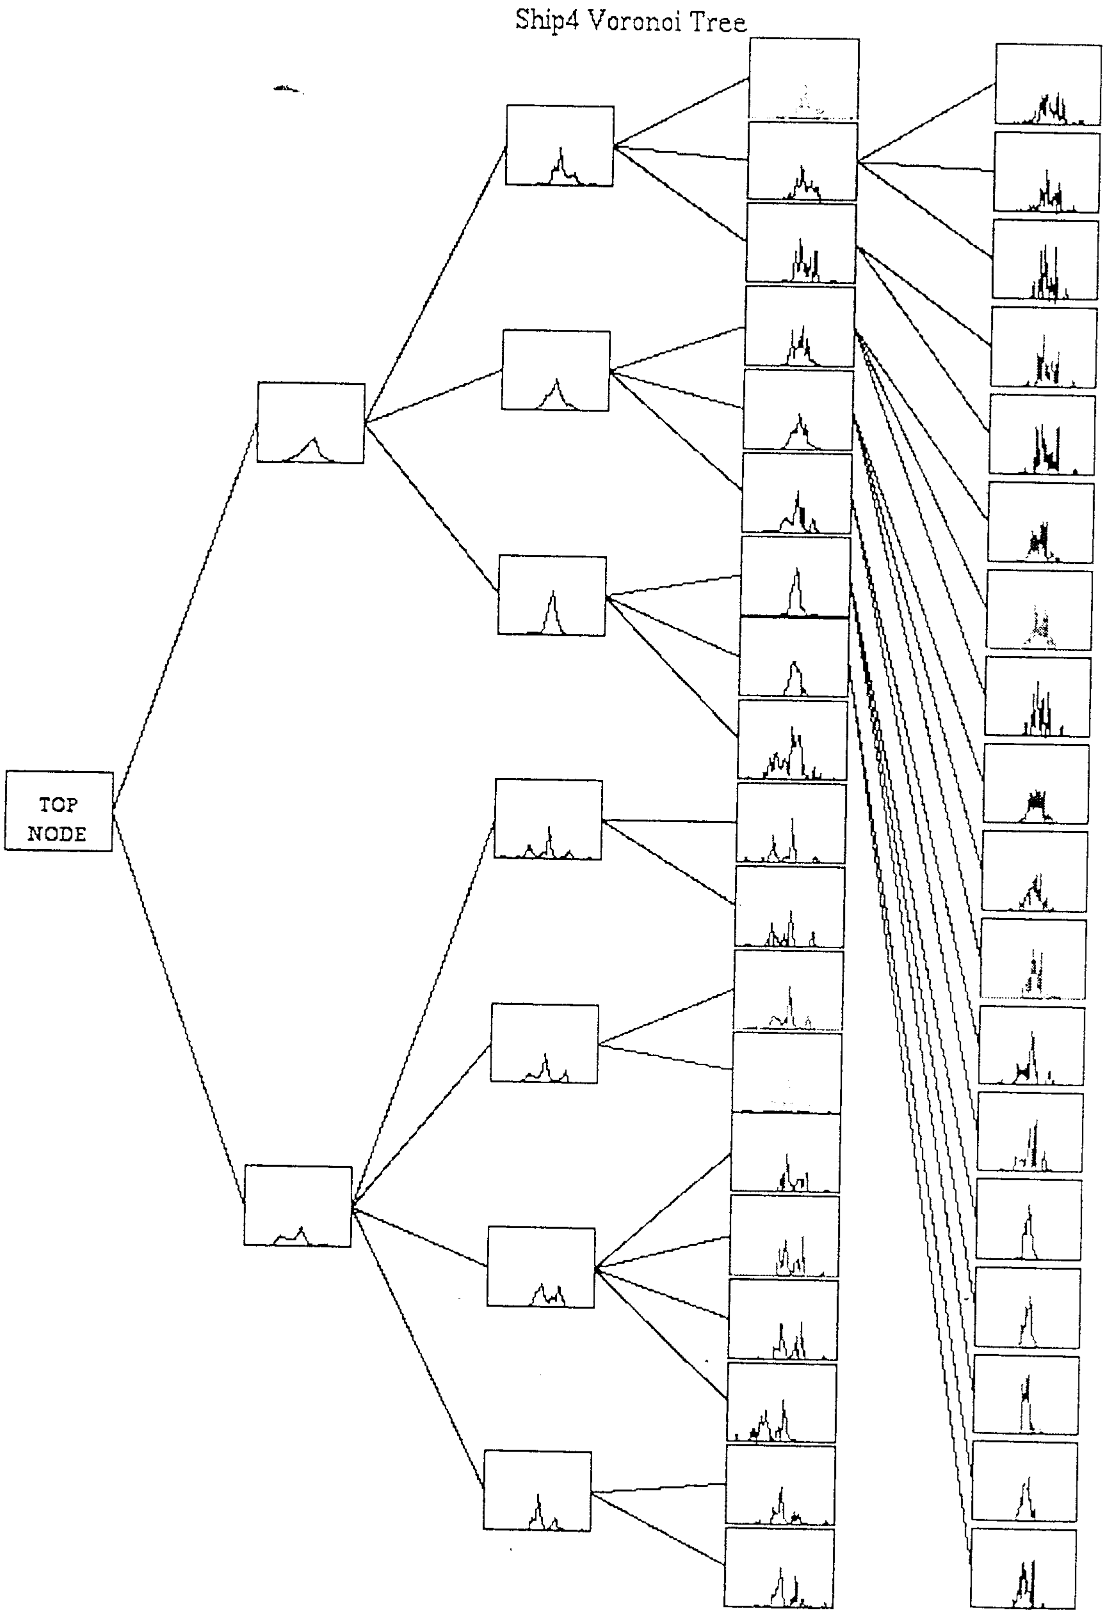

Figure 2: Multiresolution aspect graph (as a tree) for the Ship4 radar data set

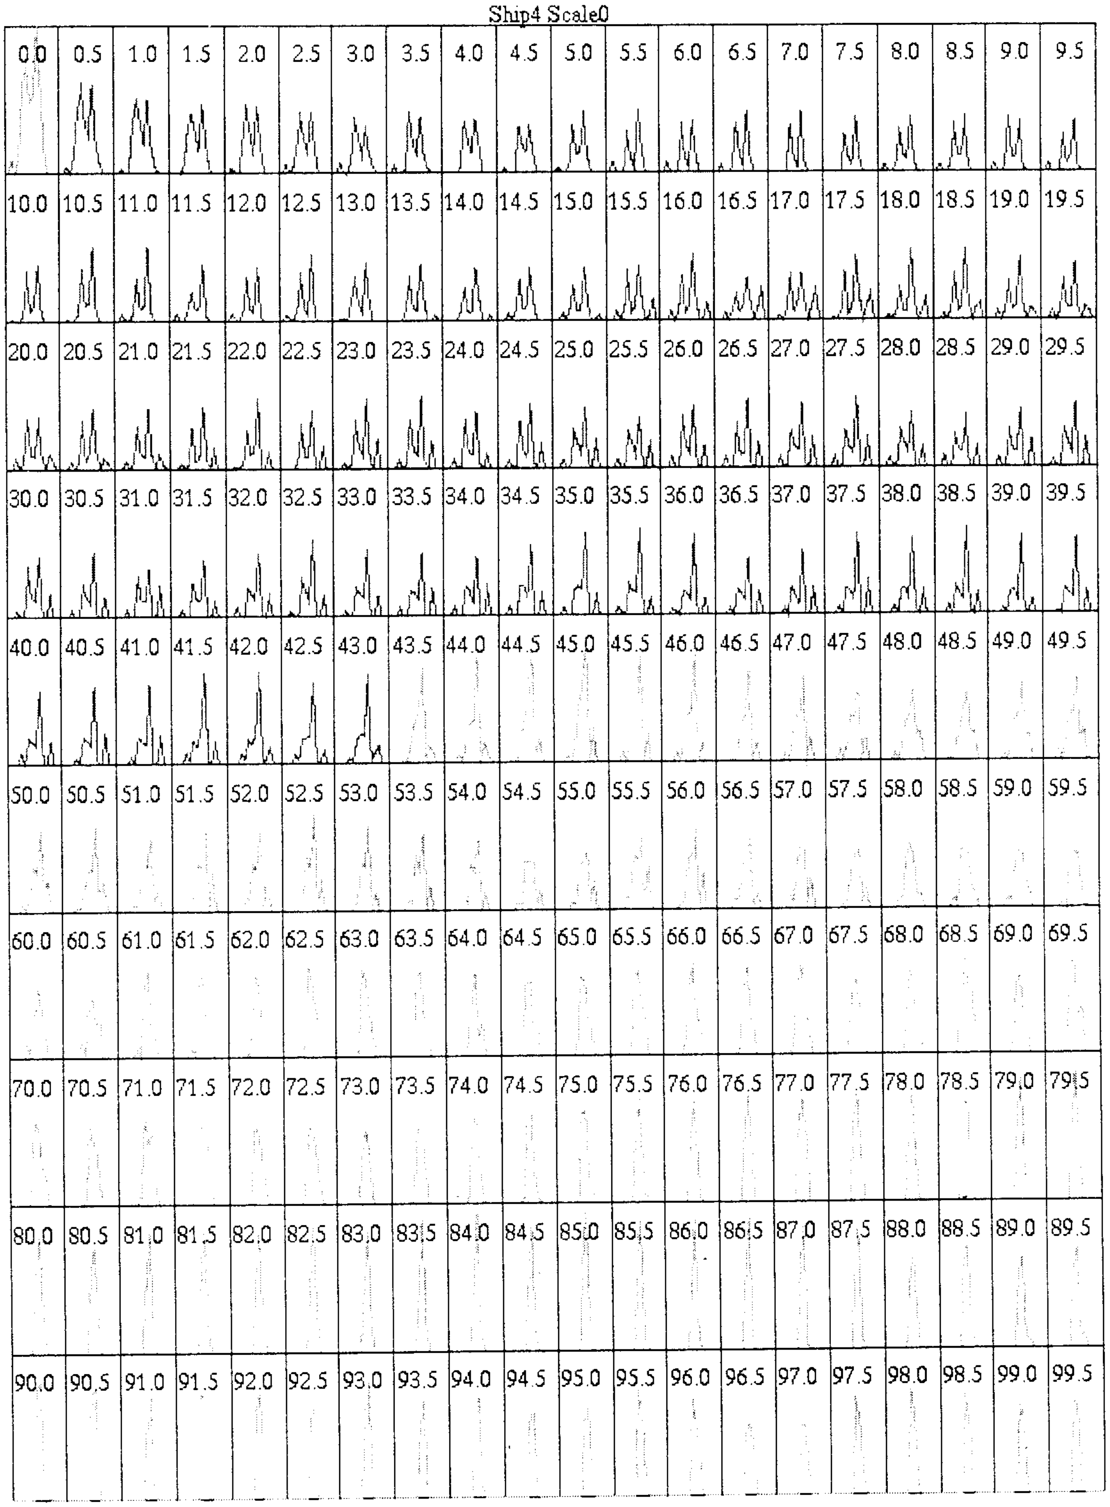

Figure 3: Average returned pulse variation at the coarse level vs aspect angle; Ship4

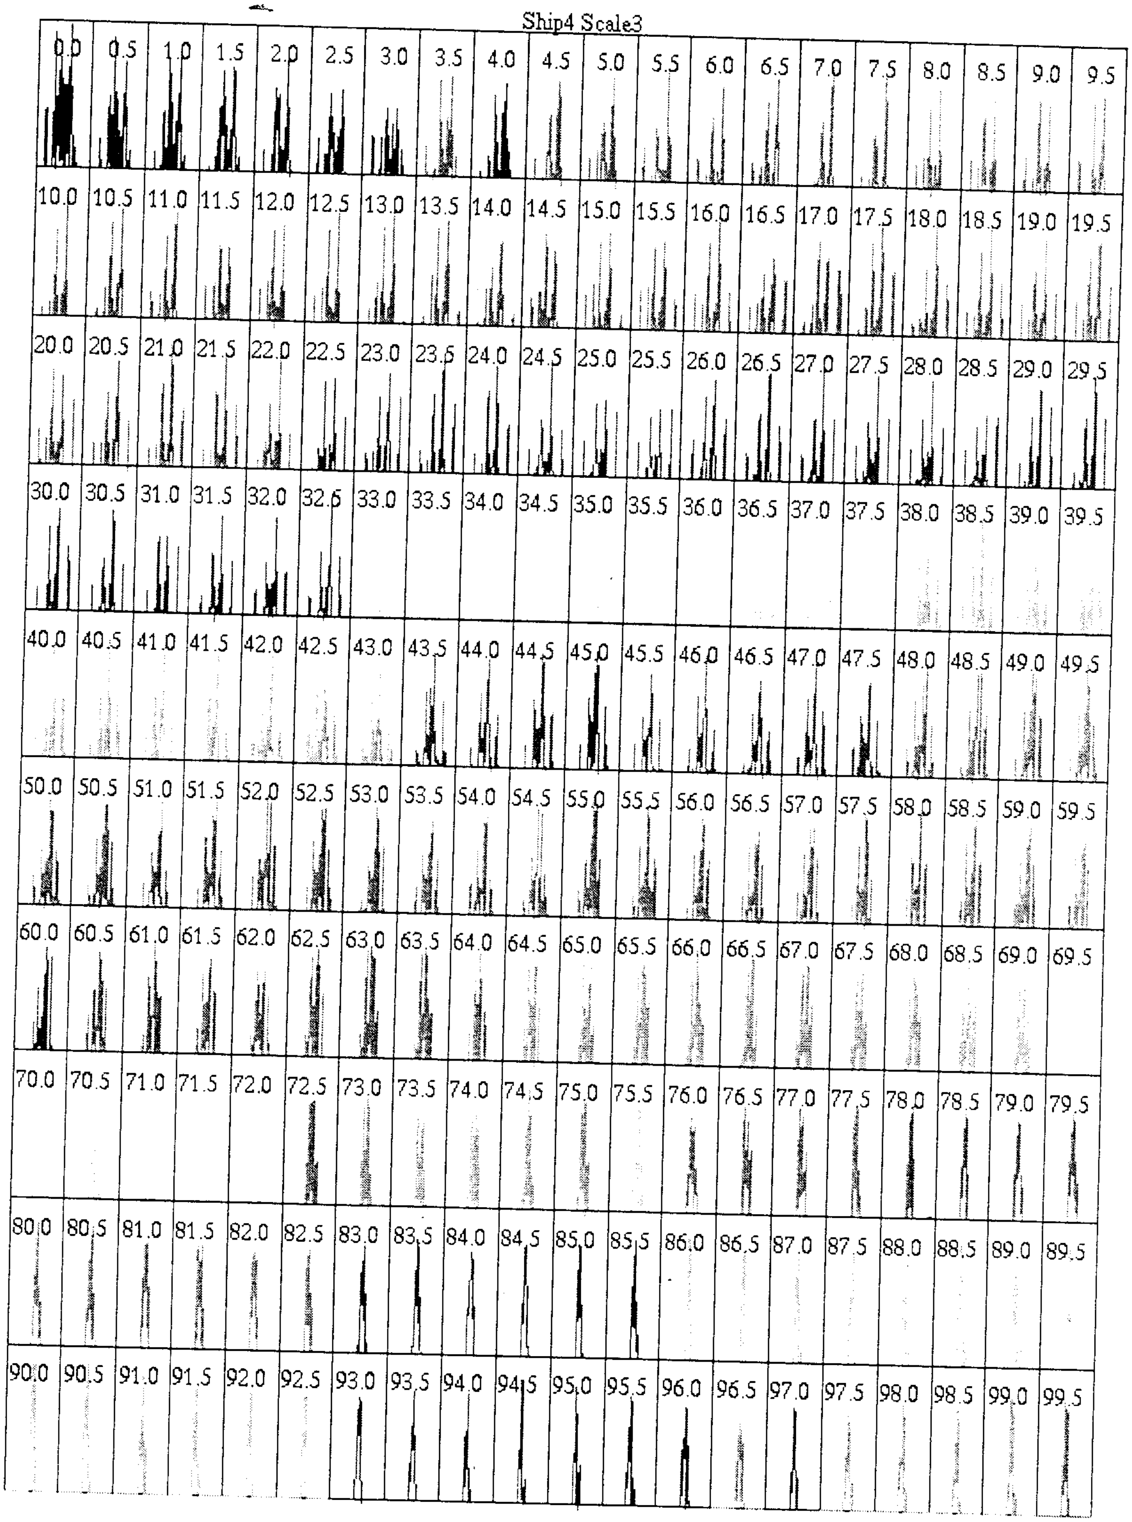

Figure 4: Average returned pulse variation at the fine level vs aspect angle; Ship4

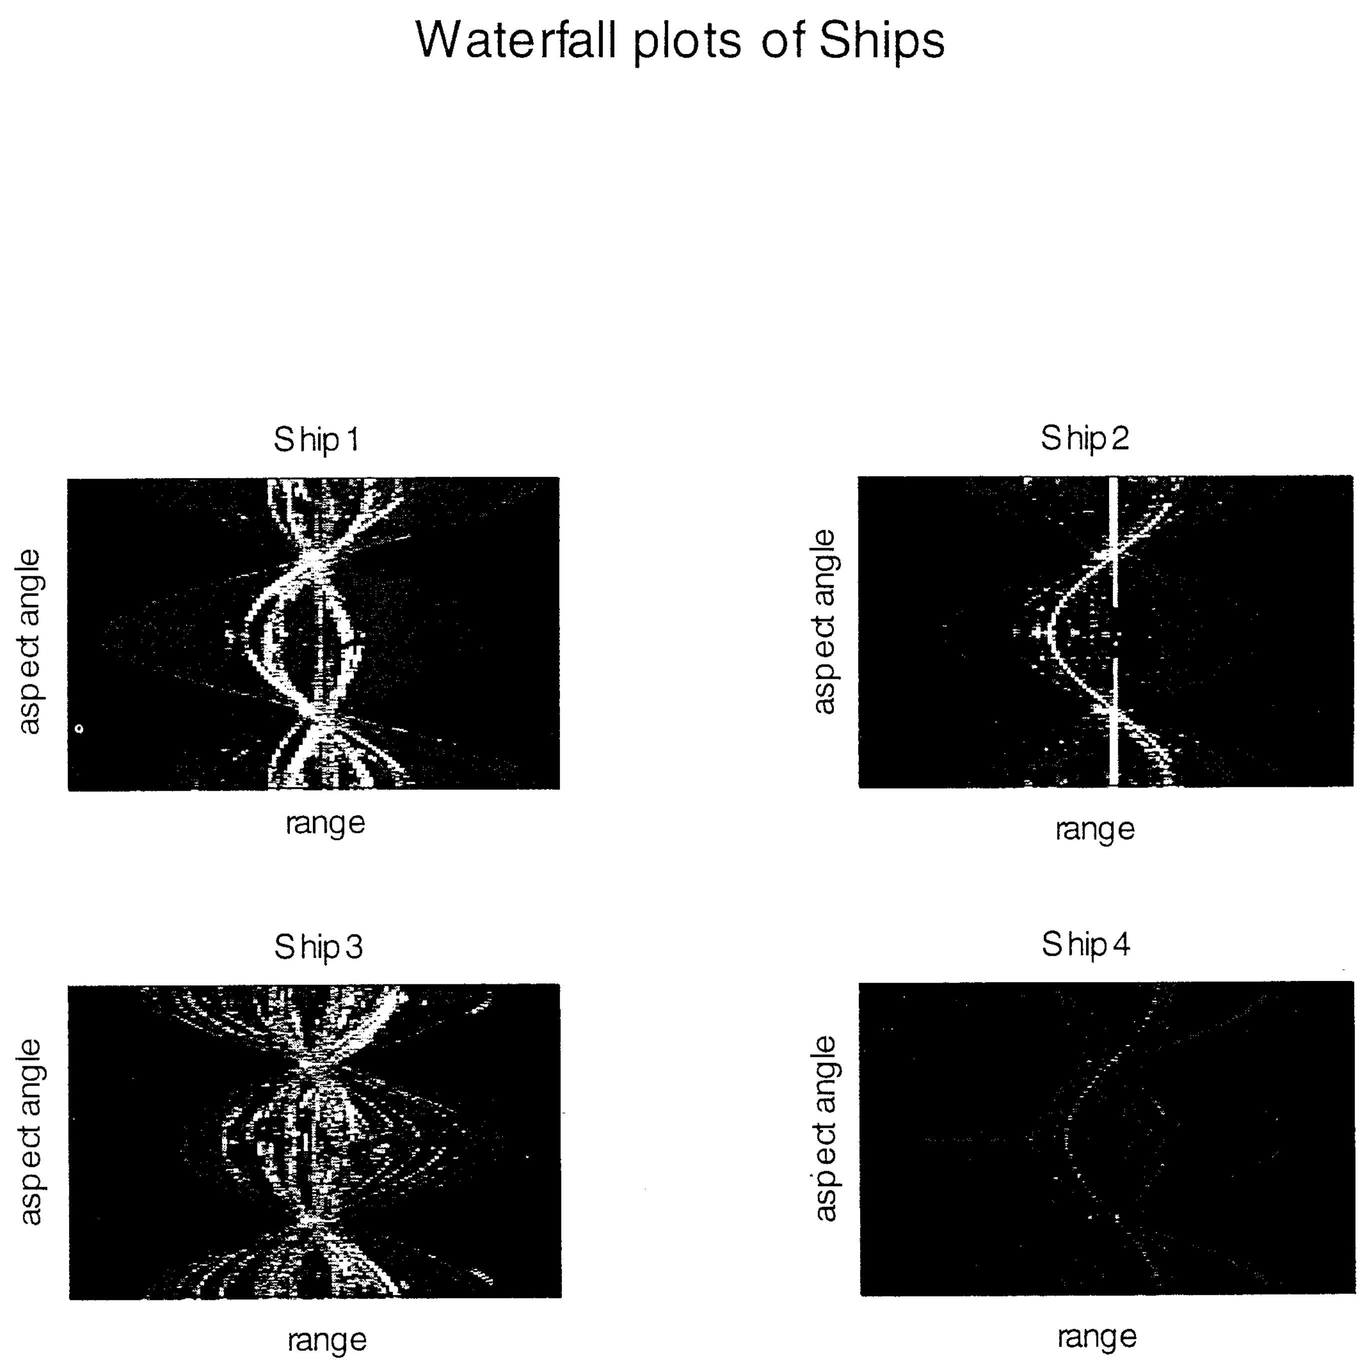

Figure 5: Waterfall diagrams for four ships used in our experiments

until the allotted number of cells have been allocated, or until total average distortion has been reduced to a requisite level. Each new layer in the tree corresponds exactly to partitions based on the next finer resolution representation of the data.

The algorithm constructs a hierarchical organization of the radar return data as a tree, which is conformant with the wavelet multiresolution data representations. This representation can be constructed for a single target or for a collection of targets. In the former case this construction produces a "model" for the target as viewed by the radar sensor. In the later this construction organizes an entire database of target radar returns. The resulting tree is our *aspect graph of the target(s)*. As illustrated in Figure 2 a multiresolution (or scale space) aspect graph for radar ship data results naturally from our construction. Here we show a typical aspect graph for one of the ships used in our experiments. Coarse resolution is represented in the nodes to the left, while as we move to the right we encounter nodes with progressively finer resolution. In each node we show the Voronoi vector (here the centroid pulse of the cell/node).

In our experiments we varied only the aspect angle. This method then produces aspect angle equivalence classes for the radar signals (returned pulses). These equivalence classes mean that the pulses in these clusters are difficult to discriminate due to their similarity. As we move from coarse to fine resolution the aspect angle cells split and can eventually get further characteristics of the target based on finer resolution information on the pulse. These characteristics are related to dominant scattering centers. One important concept in this construction is a notion of "continuity"; in other words these aspect angle equivalent classes should not be too "brittle" (should not be very small for coarser scales). We have analyzed extensively this problem from several points of view: frequency agile vs non frequency agile radar returns, ship motion, multipath conditions. Our experiments indicate that the results obtained are more or less similar, provided we are not operating in extreme ranges. As a result we typically construct the aspect graph for each ship using high resolution frequency agile pulses. We have obtained the best results (from the point of view of "less brittle" aspect classes) with such pulses. Indeed as we can see from the typical variation of average pulse returns vs aspect angle constructed for one of the ships, and shown for coarse scale in Figure 3 , and at the finest scale in Figure 4 , we obtain reasonably sized aspect angle cells, with non-impulsive changes among neighboring aspect angle cells. In these last figures different gray level designates different aspect cells. We also see from Figures 2, 3, and 4, that as we progress from coarse resolution to finer resolution the aspect cells subdivide locally in most cases; a sign that we have captured the esentials of quantizing the space of sensor viewpoints.

It is clear from this discussion that the aspect graph is a reduced but accurate model of the target and can be used to guide the ATR process in model-based ATR. In such an application the received pulse is compared with the "canonical" pulse at each node sequentially as the ATR process evolves.

As explained and demonstrated in detail in [14, 17, 19] these scale space aspect graph radar target models are not only efficient but very accurate.The resulting collection of centroid pulses at each resolution can be considered as a compressed representation of the data set. It is actually a quite good approximation. To test the accuracy of the approximation we computed "waterfall" diagrams of the ship radar returns as functions of aspect angle. In these diagrams, shown in Figure 5, the horizontal axis shows the range extent (range bins), while the vertical axis shows the aspect angle ($0°$ to $360°$). The value of each pixel is the radar return amplitude from this bin at the corresponding aspect angle. These are color coded but here are shown in gray scale. We have compared such waterfall plots at different resolutions for several synthetic ship data and we have verified that they form a sequence of approximations to the fine resolution waterfall plots. In Figure 5 we show such waterfall plots for four ships used in our experiments. The plots illustrate graphically the differences in the scatterers distribution in each ship; the ATR algorithm utilizes these differences to classify the targets.

3 HIGH PERFORMANCE PROGRESIVE CLASSIFICATION

The efficient radar target modeling techniques and model database organization described in the earlier sections will be used in this section to design high performance progressive classification algorithms. As explained in detail in [16, 18, 19, 20] the multiresolution aspect aspect graph constructed by our wavelet-TSVQ method provides an extremely efficient hypothesis generation and indexing mechanism to guide the on-line classification process. Our method constructs these efficient indexing schemes algorithmically, without resorting to *ad hoc* heuristics typical of previous work [23, 24, 25, 26]. Our methodology has distinct advantages over other indexing schemes such as the *Geometric Hashing* technique of Lamdan *et al* [22], or the methods of Knoll and Jain [21]. Our techniques hold great promise for ATR algorithms involving large libraries of target models, without commensurate increase in computational complexity. The key idea we are exploring with our aspect graphs is the use of clusters of features, represented in terms of geometric properties that are invariant under projection, as keys for indexing into a hashed library.

As described first in [19, 20]there are two distinct ways that we can use these efficient and progressive target models for ATR.

We assume we are given data which represent HRR returns from several ships, *Ship1* to *ShipM*.The first method uses half of the combined *Ship1* to *ShipM* radar data to construct a scale space aspect graph, representing all ships. In other words we construct a hierarchical tree structured representation of the radar target database. We next examine the leaf nodes of the tree constructed. If the majority of the data in the leaf node is from *Shipi* we label the leaf node "*Shipi*", if the majority is from *Shipj* we label the leaf node "*Shipj*". The resulting tree is ready to be used now as an on-line ATR algorithm. It is an efficient structure since we only need to store the nodes and their relationships (parent to children), the centroid pulse for each node, and the labeling of the leaf nodes (which was done by simple majority vote here). During on-line execution this ATR algorithm operates as follows. The incoming sensor data are first processed by a wavelet preprocessor and then are input to the tree. At each node of the tree, the distance between the sensor data and the Voronoi vectors (i.e. centroid pulses) of the "children nodes" is computed and the input pulse progresses to the node corresponding to the "child" with the smaller distance and the process repeats. When the incoming pulse reaches a leaf node it is declared as *Shipi* according to the labeling that has been assigned to that leaf node.

We named this method [19] the **Combined Aspect Graph (CAG) Method** in order to emphasize the fact that a single aspect graph is constructed for the entire database. It is a reasonable approximation to the more systematic method described and proposed in [19]. Its advantages are: (i) It includes all relations between the data; (ii) Due to its tree hierarchy it produces decisions in time logarithmic in the number of leaf nodes (roughly); (iii) It is based on an approximation of the multidimensional distribution of the data from all targets. Its disadvantages are: (i) The multidimensional distribution may "hide" the distinguishing characteristics of individual targets; (ii) The leaf labeling is *ad hoc*; (iii) When a new target needs to be encountered a costly redesign and retraining must be performed. The most serious drawback is (iii) which is often unacceptable despite the fact that it is performed off-line.

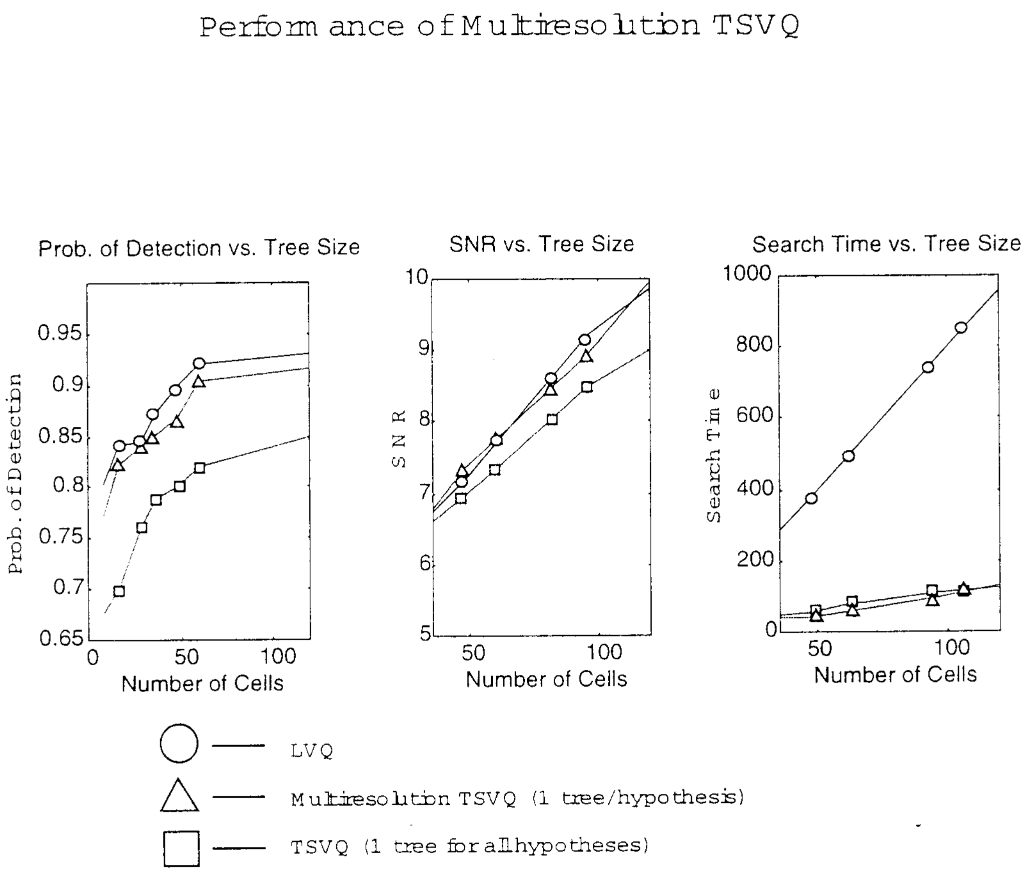

Figure 6: Comparison of performance vs complexity for LVQ, PAG and CAG algorithms

We have tested this CAG method on high accuracy synthetic data. The NRL Code 5750 digital simulation model is a flexible tool for experimentation, and it has been used as the basic data generation source for the studies reported here. This model has been validated against field returns and provides high accuracy simulations. The digitally simulated ship model consists of over 800 scatterers (for each viewpoint) of a variety of types, including flat plates, point scatterers and dihedrals. These scatterers are distributed in both range and space in accordance with their actual locations on a ship. To capture safely all ship pulses we used a range gate of 128 bins corresponding to a returned signal time duration of 1280 ns. At the finer resolution of 10 ns, and sampling at the corresponding rate produces 2^7 samples per pulse. We generated radar return databases for four different ships, *Ship1* to *Ship4*, utilizing the NRL Code 5750 ship radar return simulator. We varied the aspect angle from 0° to 360° in increments of 0.5°. This allowed for large variation in the number and appearance of dominant scatterers.The parameters in the synthetic data were as follows. Elevation angle: .023°, Sea state: 3, Pulsewidth: 10 ns, Pulse Repetition Interval:

1 s. Each database contained 720 pulses at fine resolution (720 for training and 720 for testing for each ship). We did not change the range gate and therefore each pulse in the database has the same time duration. We used different data for designing the aspect graphs and different data for testing the resulting ATR algorithms.

The second method [19] uses separately the data from each ship to construct an independent scale space aspect graph model for each target. The resulting trees are then used in parallel for the on-line ATR algorithm as follows. The incoming sensor pulse is processed by each tree independently in the usual fashion progressing through the parent to children nodes on the basis of minimum distance. At the end the incoming sensor pulse will reach a leaf node in each tree. If the distance from the Voronoi vector of the leaf node where the pulse ended is smaller in the tree representing *Shipi*, then a *Shipi* classification is declared.

In [19] we named this method the **Parallel Aspect Graph (PAG) Method**. Its advantages are: (i) It is extremely fast and maximally parallel; (ii) On-line training can be easily implemented; (iii) It provides an acceptable solution to the "new target" insertion problem. Its disadvantages are: (i) It does not account for correlations between target models; (ii) The decision scheme is *ad hoc*.

We have tested the ATR (classification) algorithms constructed with this PAG method using the same data as in the CAG method. We also used the same data with a Learning Vector Quantization (LVQ) algorithm which we have shown in [27] to converge to the optimal Bayes classsifications. The resulting Probability of correct classification, Signal to Noise Ratio and Search Time are shown for each method vs number of cells (leafs) in the apsect graph tress in Figure 6. The circles represent the performance of LVQ, the triangles the performance of PAG and the squares the perfoprmance of CAG.

Our measures of performance are average probability of correct classification, average signal-to-noise ratio (SNR), and average search time. Average probability of correct classification is defined in the usual way by constructing a confusion matrix and then averaging the diagonal entries:

$$\frac{1}{4}\sum_{i=1}^{4} p(i|i).$$

where $p(i|i)$ is the experimentaly derived probability of deciding that the i'th target is present given that it is indeed present. Average signal-to-noise ratio is computed as

$$\sum p_i 10\log_{10}\left\{\frac{\|v_i\|^2}{1/n_i\sum(v_i - x_j)^2}\right\}$$

where

p_i = the number of observations in cell i/the total number of observations
v_i = Voronoi vector for cell i
n_i = number of observations in cell i

and the first summation is over all leaf cells i, while the second summation is over all observations $x_j \in$ cell i. Our search time metric incorporated the binary nature of the tree structures. We applied the following formula to the resultant trees

$$\sum\sum_{j=0}^{l_i} 2^j p_i s_{ij}$$

where

l_i = the scale in which leaf node i is found
p_i = the population of cell i /the total population
s_{ij} = number of siblings of leaf node i in scale j

and the first summation is over all leaf nodes i.

In order to properly compare search times for the three algorithms CAG, PAG, and LVQ, we evaluated each algorithm as if it were to be executed by a single processor. This normalizes the results and it allows a direct extrapolation for parallel implementations. In each of the cases of CAG and LVQ this means that the expression for search time was evaluated for a single tree. In the case of PAG, the sum of the search times for each of the four trees was reported as the total search time.

A normalization of the results with respect to the number of cells was also done. In each case, the number of cells reported was the total number of cells in all trees involved.

It is clear, and all our experiments have verified this, that the parallel algorithm outperforms significantly the compound algorithm, and that the performance of the parallel algorithm is almost indistinguishable from that of the much more computationally costly LVQ algorithm. We are currently investigating theoretical proofs of these experimental findings which are very significant. The results will be reported elsewhere.

These findings indicate that we have discovered a fairly general and systematic method for constructing high performance, maximally parallel model based ATR algorithms, based on progressive classification and learning. Generalizations and extensions to other sensors and to multi sensor problems will follow easily.

Acknowledgement

We are indebted to Rajiv Patel of Kaman Sciences for his expert help with the graphical representations and computer support.

References

[1] I. Daubechies, "Orthonormal Bases of Compactly Supported Wavelets", *Comm. Pure Appl. Math. 41*, pp. 909—996, 1988.

[2] S.G. Mallat, "A Theory for Multiresolution Signal Decomposition: The Wavelet Representation", *IEEE Trans. on Pattern Analysis and Mach. Intel.*, Vol. 11, No. 7, pp. 674—693, 1989.

[3] Y. Meyer, *Ondelettes et Operateurs I, II*, Hermann, Paris, 1990.

[4] J.G. Daugman, "Complete Discrete 2-D Gabor Transforms by Neural Networks for Image Analysis and Compression", *IEEE Trans. on Ac. Sp. and Sign. Process.*, Vol 36, No. 7, pp. 1169—1179, July 1988.

[5] Teuvo Kohonen, *Self-organization and Associative Memory*, Springer-Verlag, New York, 1989.

[6] S.G. Mallat and S. Zhong, " Characterization of Signals From Multiscale Edges", *NYU, Computer Science Technical Report, Nov. 1991*; also in *IEEE Trans. Pattern Anal. Machine Intell*, Vol. 14, No. 7, pp. 710-732, July 1992.

[7] S. Zhong and S.G. Mallat, "Compact Image Representation From Multiscale Edges", *Proc. of 3rd International Conf. on Comp. Vision*, Dec. 1990.

[8] O. Rioul and M. Vetterli, "Wavelets and Signal Processing", *IEEE Signal Processing Magazine*, Vol. 8, No.4, pp.14—38, October 1991.

[9] I. Daubechies, *Ten Lectures on Wavelets*, SIAM, 1992.

[10] A. Gersho and R.M. Gray, *Vector Quantization and Signal Compression*, Kluwer Academic Press, 1991.

[11] K. Bowyer, D. Eggert, J. Stewman, and L. Stark, "Developing the Aspect Graph Representation for Use in Image Understanding", *Proc. DARPA Image Understanding Workshop 1989*, pp.831—849.

[12] L. Breiman, J.H. Friedman, R.A. Olshen and C.J. Stone, *Classification and Regression Trees*, Wadsworth and Brooks, 1984.

[13] E. Riskin, "A Greedy Tree Growing Algorithm for the Design of Variable rate Vector Quantizers", *IEEE Trans. on Signal Process.*, Vol. 39, No. 11, pp. 2500-2507, Nov. 1991.

[14] J.S. Baras and S.I. Wolk, "Hierarchical Wavelet Representations of Ship Radar Returns", submitted for publication to *IEEE Trans. on Signal Process.*, 1993; also *NRL Technical Report* NRL/FR/5755-93-9593.

[15] D.W. Eggert, K.V. Bowyer, C.R. Dyer, H.I. Christensen and D.B. Goldgof, "The Scale Space Aspect Graph", *IEEE Trans. on Pattern Anal. and Mach. Intel.*, Vol. 15, No. 11, pp. 1114-1130, Nov. 1993.

[16] J.S. Baras and J.L. Preston, "Wavelet and Learning Clustering. Algorithms for Automatic Target Recognition," AIMS Technical report AIMS-TR-93-04, Final progress report on DARPA SBIR Phase I Contract 92-61, September 1993.

[17] J.S. Baras and S.I. Wolk, "Efficient Organization of Large Ship Radar Databases Using Wavelets and Structured Vector Qantization", *Proc. of the 27th Annual Asilomar Conference on Signals, Systems, and Computers*, Vol 1, pp 491-498, November 1-3, 1993, Pacific Grove, California.

[18] J.S. Baras and D.C. MacEnany, "Model-Based ATR: Algorithms Based on Reduced Target Models, Learning and Probing," *Proceedings of the Second ATR Systems and Technology Conference*, Feb. 1992, Vol. 1, pp. 277-300.

[19] J.S. Baras and S.I. Wolk, "Model based automatic target recognition from high range resolution radar returns", *Proc. of SPIE International Symposium on Intelligent Information Systems*, Orlando, FL., April 4-8, 1994.

[20] J.S. Baras and S.I. Wolk, "Wavelet based progressive classification of high range resolution radar returns", *Proc. of SPIE International Symposium on Intelligent Information Systems*, Orlando, FL., April 4-8, 1994.

[21] T.F. Knoll and R.C. Jain, " Recognizing Partially Visible Objects Using Feature Indexed Hypotheses", *IEEE J. Robotics and Automation"*, 2(1), 1986, pp. 3-13.

[22] Y. Lamdan and H.J. Wolfson, "Geometric Hashing: A General and Efficient Model-Based Recognition Scheme," *Proc. Second Int. Conf. Comp. Vision,* 1988, pp. 238-249.

[23] K. Ikeuchi and T. Kanade, "Automatic Generation of Object Recognition Programs", *Proc. IEEE*, Vol. 76, No. 8, Aug. 1988, pp. 1016- 1035.

[24] D.E. Dudgeon, J.G. Verly, and R.L. Delanoy, "An Experimental Target Recognition System for Laser Radar Imagery", *Proc. DARPA Image Understanding Workshop 1989*, pp. 479-506.

[25] J.G. Verly and R.L. Delanoy, "Appearance-Model-Based Representation and Matching of 3-D Objects", Proc. of IEEE Computer Society Workshop on Computer Vision, 1990, pp. 248-256.

[26] M. Seibert and A. Waxman, "Adaptive 3-D Object Recognition from Multiple Views", *IEEE Trans. on Pattern Match. and Mach. Intell.*, Vol. 14, 2, pp. 107-124, Feb. 1992.

[27] J.S. Baras and A. LaVigna, "Convergence of a Neural Network Classifier", *Proceedings of the 29th CDC Conference*, pp. 1735-1740. Honolulu, Hawaii, December, 1990.