4311

Conference Record

Twenty-Seventh Asilomar Conference on Signals, Systems & Computers

NOVEMBER 1–3, 1993
PACIFIC GROVE, CALIFORNIA

Volume 1 of 2

IEEE Computer Society Press
The Institute of Electrical and Electronics Engineers, Inc.

Efficient Organization of Large Ship Radar Databases Using Wavelets and Structured Vector Quantization

John S. Baras*
Department of Electrical Engineering
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

We investigate the problem of efficient representations of large databases of pulsed radar returns from naval vessels in order to economize memory and minimize search time. We use synthetic radar returns from ships as the experimental data. The results extend to real ISAR returns. We develop a novel algorithm for organizing the database, which utilizes a multiresolution wavelet representation working in synergy with a Tree Structured Vector Quantizer (TSVQ), utilized in its clustering mode. The tree structure is induced by the multiresolution decomposition of the pulses. The TSVQ design algorithm is of the "greedy" type. Our experiments todate indicate that the combined algorithm results in orders of magnitude faster data search time, with negligible performance degradation from the full search vector quantization. The combined algorithm provides an efficient indexing scheme (with respect to variations in aspect, elevation and pulsewidth) for radar data which can facilitate the development of ATR, surveillance and multi-sensor fusion systems.

1 Introduction

High resolution ship 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. This is true for naval targets in particular. 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 ship. 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 ships, and to design efficient coding schemes which operate in a hierarchical manner on the compressed representations to recover the ship identity. It is our contention that *multi-resolution representations* of the radar data, followed by properly designed hierarchical clustering are key means to achieve both objectives.

Wavelet theory [1]—[9] offers an attractive means for the development of such 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.

In this paper we describe recent results on hierarchical representations of high resolution radar returns from ships. This is accomplished by developing hierarchical clustering schemes using wavelet representations which are motivated from the physics of radar scattering [16], in the same way that wavelet analysis of sound is motivated by the physics of sound and speech signal generation. In addition we use sophisticated versions of Vector Quantization (VQ) [11] to further compress and cluster the wavelet representations of the radar signals, in a way that permits hierarchical search across resolutions and a progressive scheme for identification of the ship (target). For a detailed exposition we refer the reader to [16].

High range-resolution radar returns can be described as complex valued signals of finite duration. For a complete characterization of a complex target one can store the whole set of these two dimensional functions (pulses) for all possible values of radar pulsewidth δ (i.e. different resolution), aspect α and elevation ϵ. Even if one quantizes the three-dimensional space of δ, α, ϵ the required storage is enormous and impractical for real applications. Our efforts to date have concentrated on amplitude representations only. Given the amplitude of a high-range resolution radar return, several characteristics of the scatterer distribution of the target can be revealed. Varying the radar pulsewidth

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

δ changes the resolution of the returned pulse, in the sense that more (narrow pulse) or less (wide pulse) details can be distinguished. Varying the view-point (i.e. the aspect, elevation (α, ϵ) pair) changes the shape of the returned pulse, because dominant scatterers have typically highly directive returns (in space), and because small variations in aspect produce large variations in the phase of the signal returned from each scatterer. Successful methods to provide effective compression of radar returns must address the substantial variability of the returns. As a consequence, some sort of averaging (or clustering) is necessary in representing the more meaningful, slower variation of the radar return (or the RCS) as aspect and elevation are changing. 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.

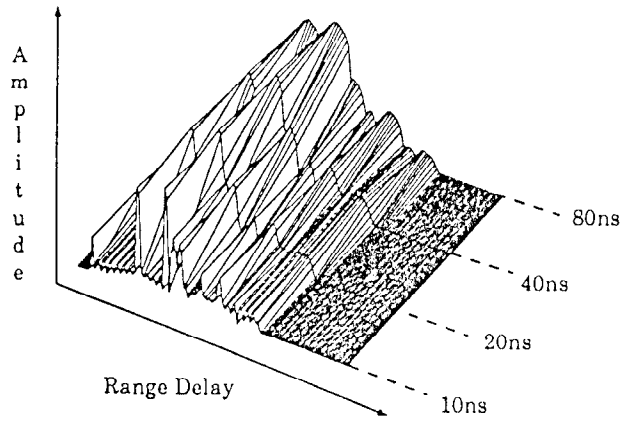

Figure 1: Variation of ship radar return vs pulsewidth.

Higher compression rates can only be achieved by developing coding techniques that are adapted to the information content of the signals, and to their physical nature. This requires the organization of the radar return data in such a way as to represent and separate the important features of the return. For the amplitude of ship radar pulse returns, this means the location of the local maxima and their geometrical characteristics. In particular we need ways to describe how these local maxima coalesce as the radar pulsewidth δ increases. This corresponds to changes in the return pulse as the resolution varies and corresponds to the physical property of coherently combining the returns from dominant scatterers as a function of range extent. Such a coding scheme results in a hierarchical clustering induced by the multiresolution representation, resulting in efficient storage and speedy recovery of the information [16].

Experiments with variable pulsewidths and real targets, in order to obtain a multiresolution representation of the ship is not a very practical solution. 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 400 scatterers 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. We are interested in the variability of pulsewidth from a minimum of 10 ns. At a bow-on aspect the ship is approximately 300 ft long, resulting in a range extended pulse return of approximately 600 ns plus a pulsewidth. 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. We also maintain pulse-to-pulse independence, by selecting the radar pulse repetition interval (PRI) to be long enough with respect to the correlation time of an individual scatterer.

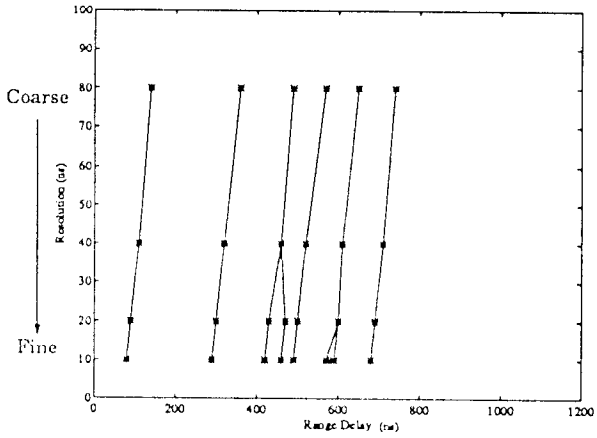

Figure 2: Scale space diagram for ship radar returns

In Figure 1 we show a typical result of a ship return with transmit pulse widths of 10 ns, 20 ns, 40 ns, 80 ns. In the three-dimensional representation shown, we see clearly the coalescence of the ship scatterers as we move from fine to coarse pulsewidths. In Figure 2 we have extracted the location of the maxima of the ship radar return as a function of pulsewidth, or equivalently resolution. Uniform localization corresponds to vertical lines. We call diagrams such as the one depicted in Figure 2 *scale space diagrams*, extending a notion originally introduced by Witkin in computer vision [13]. Our experiments to date regarding uniform localization of radar features

are not complete and this issue will be revisited elsewhere.

One of the objectives of the work reported in [16] is to show, by experiments with synthetic data, that our methodology (combining wavelet representations with clustering algorithms) provides a substitute which involves only processing the high resolution radar returns; recapturing in essence the effects of variable pulsewidth data. This result in itself is an important finding.

2 Hierarchical Clustering of Multiresolution Radar Returns

We refer to [2], [3], [9] for wavelet fundamentals. In such a multiresolution analysis [2] one has *two functions: the mother wavelet ψ and a scaling function ϕ*. We denote by f the generic radar pulse, by $\psi_{m,n} = 2^{-m/2}\psi(2^{-m}t - n)$, $\phi_{m,n}(x) = 2^{-m/2}\phi(2^{-m}x - n)$ the functions obtained by dilation and translation from ψ and ϕ. The coefficients of expanding f in terms of the $\psi_{m,n}$ are $c_{m,n}(f)$, while $a_{m,n}(f)$ are the coefficients of expanding f in terms of $\phi_{m,n}$. Usually one denotes by V_m the space spanned by the $\phi_{m,n}$. The spaces V_m describe successive approximation spaces, $\ldots V_2 \subset V_1 \subset V_0 \subset V_{-1} \subset V_{-2} \ldots$, each with resolution 2^m. This sequence of successive approximation spaces V_m constitutes a *multiresolution analysis* [2, 9]. W_m denotes the space which is exactly the orthogonal complement in V_{m-1} of V_m. These concepts result in a fast algorithm for the computation of the $c_{m,n}(f)$ (or $\mathbf{c^m}$) and $a_{m,n}(f)$ (or $\mathbf{a^m}$) [2]. The whole process can also be viewed as the computation of successively coarser approximations of f, together with the "difference in information" between every two successive levels.

We generated radar return databases for two different ships, *Ship 1* and *Ship 2*, utilizing the NRL Code 5750 ship radar return simulator. In generating the synthetic data we kept the radar fixed and turned the ship, including the motion induced by sea waves. We varied the aspect angle from 0° to 360° in increments of 0.05°. This allowed for large variation in the number and appearance of dominant scatterers.The parameters in the synthetic data were as follows. Radar frequency: 16.25 GHz, Elevation angle: .023°, Sea state: 3, Pulsewidth: 10 ns, Pulse Repetition Interval: 80 ms to 1 s. Each database contained 7,200 pulses at fine resolution. We did not change the range gate and therefore each pulse in the database has the same time duration.

Let $\mathcal{S}$ denote the set of discretized radar pulses. The fine resolution data will be denoted by $S^0 f(n)$, $n \in I^0$, where $I^0 \subseteq \{1, 2, \cdots, 2^7\}$, is the index set of the fine resolution data. We could subsample the given data to economize computations but in our experiments we took I^0 to be the full set $\{1, 2, \cdots, 2^7\}$. 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 $J^* < J$. We denote by $\Delta\tau^0$ the sampling interval for the fine resolution; it can be also thought as the resolution of the fine resolution data. 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; resulting in our case in $|I^m| = 2^{7-m}$. 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 $W^m f$. All signals are digitized and we identify the vector $\mathbf{a^m}$ with $S^m f$ and the vector $\mathbf{c^m}$ with $W^m f$. 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 at scale 1 we have $N/2 = 2^{J-1}$ samples and resolution $\Delta\tau^1 = 2\Delta\tau^0$. At scale m we have $N/(2^m) = 2^{J-m}$ samples and resolution $\Delta\tau^m = 2^m\Delta\tau^0$. 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 $W^m f$. Thus the multiresolution scheme replaces the information in each pulse $f = S^0 f$ with the set $\{W^m f,\ m = 1, 2, \cdots, J^*,\ S^{J^*} f\}$.

As was shown in [16] the peaks of the radar pulse coalesce as we vary the resolution from fine to coarse, in a manner similar to the one observed when we used variable pulsewidths of comparable resolution; c.f. Figure 1. This justifies the suitability of wavelet analysis for radar signals. It was shown in [16], that within the approximations involved, the "smoothed" return $S^J r$ at resolution J is equivalent to the return due to a "smoothed" transmit pulse $S^J p$ at the same resolution.

The principal result of this paper is the development of 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. By properly defining a rate-distortion measure between the respective sample distributions one can reinterpret the process of vector quantization in the context of optimal decision theory. In fact, this flexibility of the definition and interpretation of rate and distortion in Shannon's theory has recently lead to very beneficial cross-fertilization between these two areas, in particular between tree-structured vector quantizers [11] and classification (decision) trees [14]. 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 ship radar returns.

Implementations of the basic VQ algorithm in which the search is exhaustive are called *full-search*. In many VQ systems the time overhead associated with this search is too costly. Of particular interest is tree structured vector quantization (TSVQ) [11], which provides

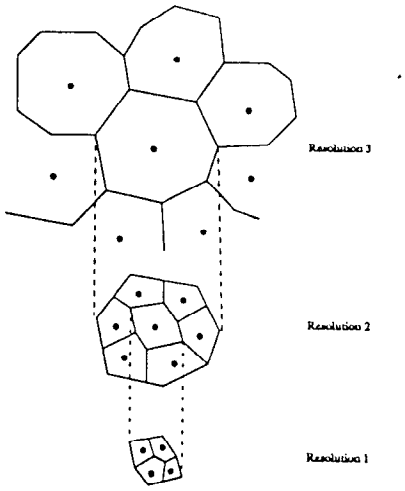

Figure 3: Illustrating a multiresolution TSVQ by splitting Voronoi cells based on different resolution data.

logarithmic (in the number of data) search time *vs* linear (in the number of data) search time provided by full search VQ. TSVQ is a special case of hierarchical VQ [11]. TSVQ is one of the most effective and widely used techniques for reducing the search complexity in VQ. A useful method for designing the tree structure is based on the application the Linde–Buzo–Gray (LBG) algorithm [11] to successive stages using a training set. We have used a variant of this method which is of the "greedy" [15] variation. More precisely our algorithm splits the cell which contributes the largest portion of the current overall distortion.

The algorithm which implements our overall approach is described in the sequel of this section. 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 3, and detailed below.

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 $\{S^{J^*} f,\ f \in \mathcal{S}\}$. Since it is the coarsest representation, the corresponding length of the data vectors is the shortest; in our experiments that length was 16. As a result this clustering is faster than a clustering performed on the much longer fine resolution representations of the data vectors. 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 the following steps 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.

An important step in the algorithm addresses the so-called "centering problem" for the various radar pulses. In real field databases the radar pulses will not be perfectly aligned with each other; there will be random time shifts. If we compute the L^2 distance of even very similar pulses, which have been shifted with respect to each other, large errors will result. These errors can distort the VQ computations, produce erroneous cell centroids and reduce the performance of the overall algorithm. The following adjustments were introduced and effectively corrected these problems. First, when we compute the distance between a pulse f and a centroid θ, we compute the $\inf_\tau \{\|f_\tau - \theta\|\}$, where f_τ is the shifted version of f by τ, and τ is within allowed bounds. Each pulse remaining in the cell is then shifted by the τ determined by this infimization. The new centroid is computed by averaging these shifted pulses. Second, the new centroid is centered so that its median point coincides with the middle point of the range gate interval. This process is iterative.

We have constructed a hierarchical organization of the radar return data as a tree. The tree is conformant with the wavelet multiresolution data representations. This representation can be constructed for a single ship or for a collection of ships. In the former case this construction produces a "model" for the ship as viewed by the radar sensor. In the later this construction organizes an entire database of ship radar returns.

3 Experiments, Interpretations and Conclusions

It is clear that the search of the hierarchically organized database (as a tree) will be much faster than the case by case search of the overall database. The question is: how much performance did we sacrifice? To answer this question we compared the results of the wavelet-TSVQ algorithm with those of full-search VQ applied to the finest resolution data, by means of the total (operational) distortion *vs* number of cells (i.e. the number of terminal nodes) performance curve. In all our experiments the two curves were very close showing that the performance of the wavelet-TSVQ algorithm is indeed excellent. Additional results can be found in [16]. This performance of our algorithm shows that we are getting high performance in an efficient

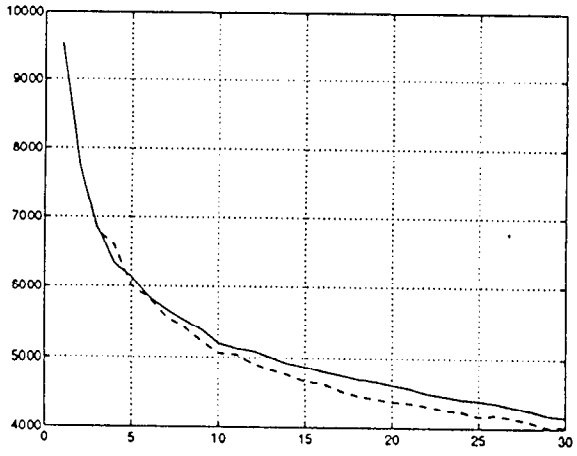

Figure 4: Performance comparison between full-search VQ on fine scale pulses vs wavelet-TSVQ on multiresolution pulses; *Ship 2* data set

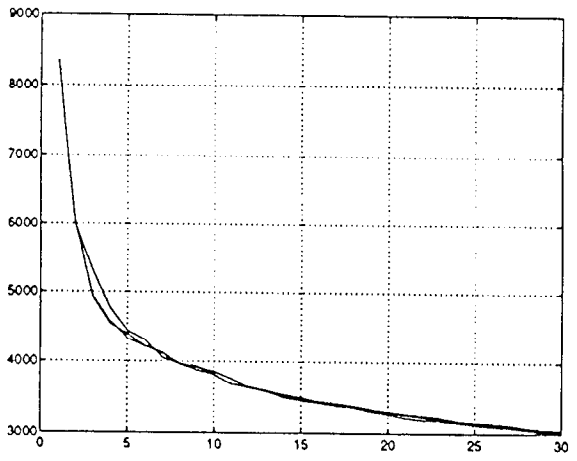

Figure 5: Performance of the wavelet-TSVQ algorithm on multiresolution pulses for three wavelets; *Ship 1* data set

way. A typical result is depicted in Figure 4. The results shown used the *Ship2* data set. The solid curve corresponds to the wavelet TSVQ algorithm, while the dashed curve corresponds to full-search VQ. There is a small discrepancy when the number of cells is very small, 3 to 5. These should be disregarded as it is not reasonable to assume such a small number of cells. The mother wavelet used in these results was the Haar wavelet.

We also examined the performance sensitivity of our algorithm with respect to the mother wavelet used in the multiresolution representation. Our experiments todate demonstrate that the tree itself is sensitive to the mother wavelet used, but the performance of the algorithm is not. In Figure 5 we show the total (operational) distortion *vs* number of cells (i.e. the number of terminal nodes) performance curves for the wavelet-TSVQ algorithm, operating on *Ship1* data set, for three different mother wavelets: the Haar wavelet, an orthonormal wavelet with highest number of vanishing moments from [9, chapter 6] constructed from ${}_N H$ with $N = 6$, and a biorthogonal symmetric wavelet, based on splines with less dissimilar lengths $l = 4 = k$, $\tilde{k} = 4$ from [10]. The three curves are almost indistinguishable, designating a robustness of our algorithm with respect to wavelt selection. Further research is needed in this direction however.

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 "water fall" diagrams of the ship radar returns as functions of aspect angle. In these diagrams, see for instance the bottom image in Figure 6, 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. In Figure 6, the bottom image shows the raw data, or the pulses at the finest scale. The middle (resp. upper) image shows the "water fall" diagram corresponding to the coarser resolution 1 (resp. resolution 2). These coarser diagrams are constructed by depicting only the centroid pulses, from each cell. These centroids correspond to the raw data and are shifted so as to be aligned with the corresponding finest resolution pulse, using the method described at the end of section 2. Figure 6 shows a typical multiresolutiuon "water fall" digram for the *Ship 1* data set, while Figure 7 for the *Ship2* data set. In both Figures we used the Haar wavelet. We conclude that our algorithm provides increasingly (with resolution) accurate representation of the ship radar returns.

We constructed "water fall" diagrams corresponding to our wavelet-TSVQ reperesentation in the same way. In particular we investigated the similarity of these diagrams with those of the original data as the number of cells increased. In Figure 8 we show such a typical comparison for the *Ship1* data set. The bottom diagram corresponds to the finest resolution data, the middle diagram corresponds to TSVQ with 30 cells, and the top to 5 cells.

As explained in detail in [16] the trees constructed by our wavelet-TSVQ method indeed achieve an indexing scheme for ship targets reminiscent of that obtained by varying the aspect angle and the pulsewidth. This is an indexing scheme because it provides a hierarchical organization of the multi-viewpoint (aspect and elevation) multi-pulsewidth radar data from a ship using significant clusterings in the δ, α, ϵ parameter space. We also examined carefully the resulting cells to discover on what characteristics of the pulse the clustering was based in an essential way. We concluded in [16] that at the coarse level the wavelet-TSVQ algorithm

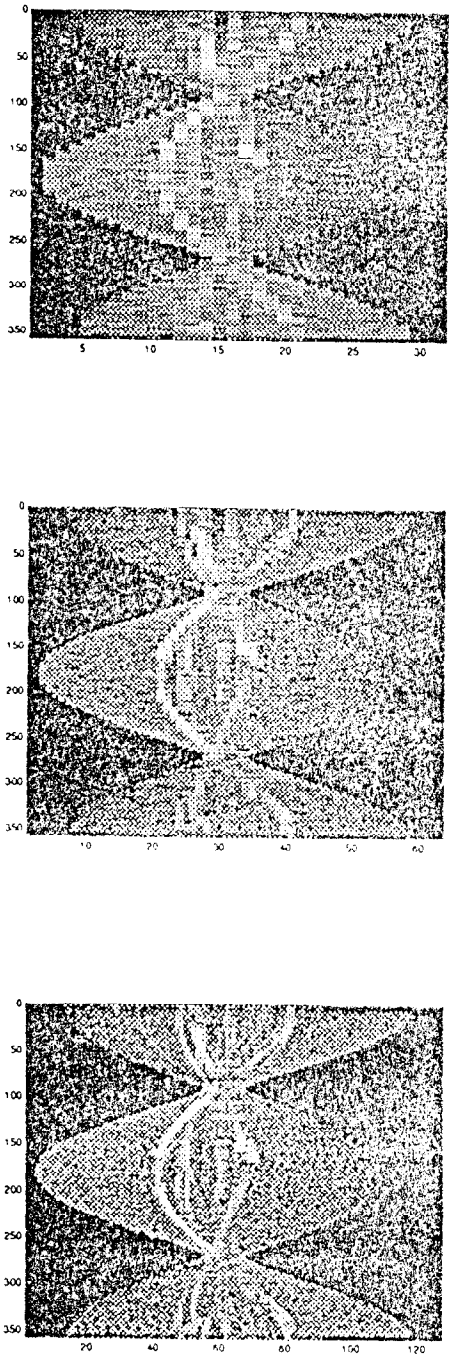

Figure 6: Multiresolution water fall diagrams; *Ship 1* data set

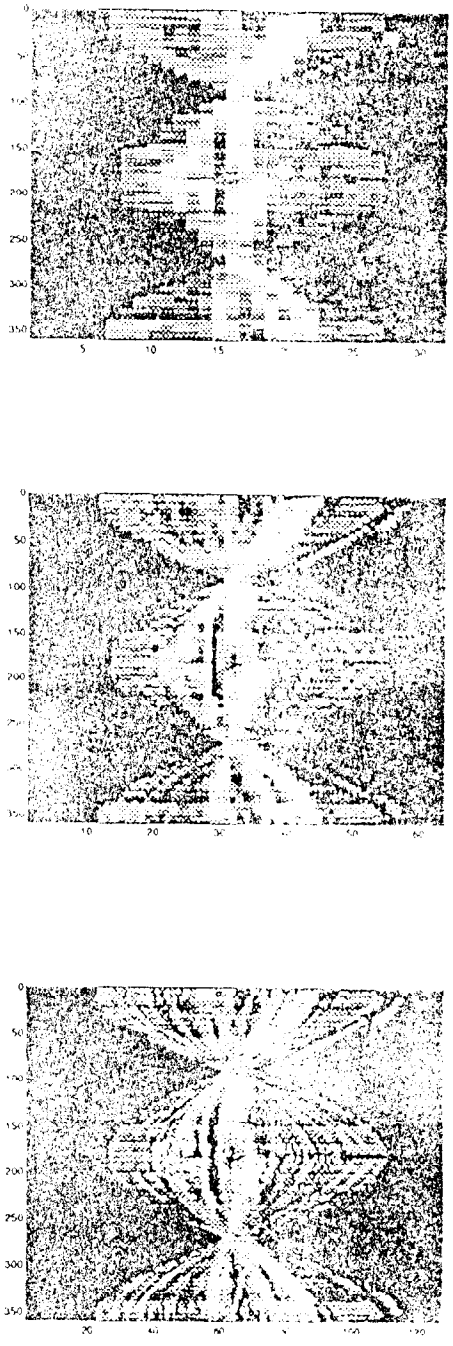

Figure 7: Multiresolution water fall diagrams; *Ship 2* data set

clusters the radar pulses according to aspect. As we move to finer resolutions the pulses cluster according to the location of their maxima. These maxima correspond to significant scatterers. This is a very nice and natural indexing of the radar returns from a ship. What is an important contribution is that we have developed a systematic, automatic method for constructing this indexing. To better realize the efficiency of our method it suffices to compare it with conventional methods of indexing radar pulses based on small aspect–elevation cells. The method also reflects the accuracy limitations of the sensor, in the sense that it does not attempt to separate the pulses more than

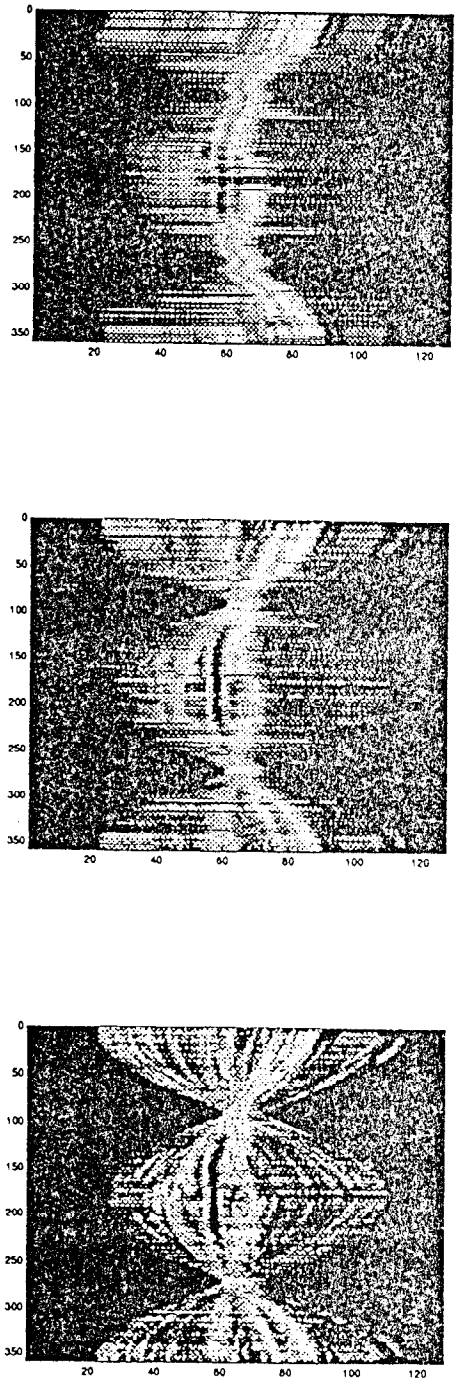

Figure 8: Multiresolution TSVQ water fall diagrams; *Ship 1* data set

the sensor noise will permit. Putting all this together we have discovered [16] in addition an extremely efficient indexing scheme for high range-resolution radar data, which is akin to the *aspect graph* widely used in computer vision based object recognition [12].

In Figures 9 and 10 we show the aspect graphs as trees constructed by our method for the data sets *Ship1* and *Ship2* resp.. We used the Haar wavelet. We show the cells that the algorithm created at each resolution (tree layer) as well as the percentages of pulses that were clustered in each cell. The notation we have used to designate the cells is as follows: *cell m,k* denotes the cell number k at resolution m. We note that the two trees are different, reflecting the different ship signatures.

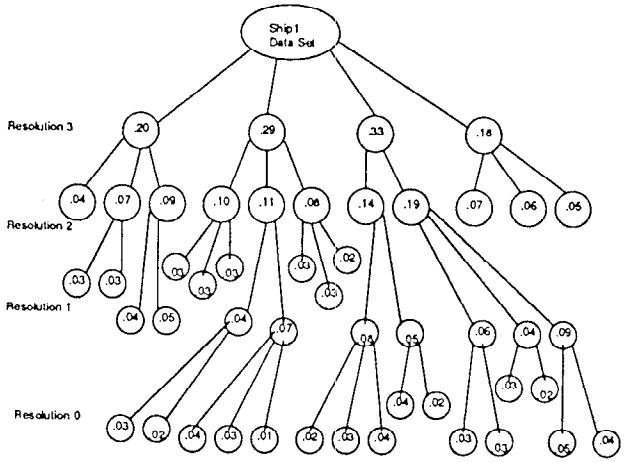

Figure 9: Apect graph (as a tree) for the *Ship1* data set

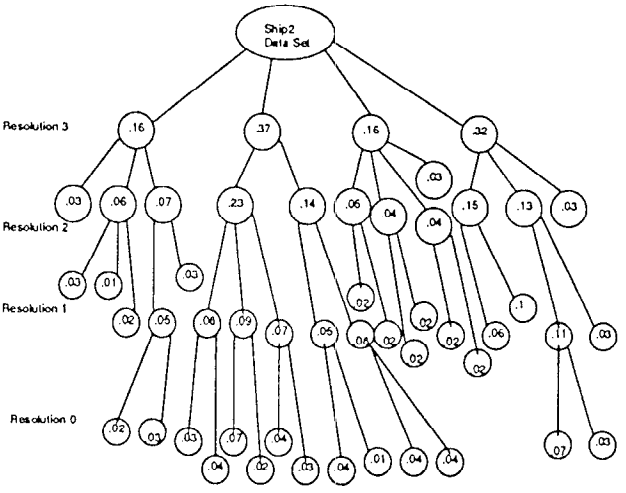

Figure 10: Apect graph (as a tree) for the *Ship1* data set

The nodes (cells in these figures) correspond to aspect-elevation neighborhoods for which the corresponding returned pulses are too similar to be separated. To each node there is associated a "canonical" pulse which corresponds to the centroid. The nodes

are given for various resolutions. Transitions from one node to the other indicate either a change in aspect-elevation, or in resolution, of adequate magnitude to cause changes in the pulse that can be discriminated by the sensor. In the present case these changes are due to grouping (or ungrouping) of scatterers, or scatterer visibility (or non visibility) from the particular aspect–elevation cell. 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. The aspect graph directs the search in an efficient and speedy manner; it is well known that tree based search is logarithmic in the number of terminal nodes, which is a substantial reduction from conventional methods. This search for identification can be done in two ways: (i) By passing the received pulse train through each tree in parallel and (ii) By constructing a large tree containing all models. The former method has substantial implementation advantages. These constructs are extremely useful in ATR, ship classification and in the retrieval of data from large radar pulse databases. The techniques are generic however and can be applied to a great variety of signal classification and hierarchical organization problems. We shall pursue several of these developments elsewhere.

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, IEEE Trans. Pattern Anal. Machine Intell.*

[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] M. Antonini, M. Barlaud, P. Mathieu and I. Daubechies, "Image Coding Using Vector Quantization in the Wavelet Transform Domain", *Proc. 1990 IEEE ICASSP*, Albuquerque USA, pp. 2297—2300, April 1990.

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

[12] 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.

[13] A. Witkin, "Scale Space Filtering", *Proc. Int. Joint Conf. Artificial Intell.*", 1983.

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

[15] 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.

[16] 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.