

Final Report

Unsupervised, Robust Estimation-Based Clustering for Multispectral Images

P. I.: Nathan S. Netanyahu
Computer Vision Laboratory, Center for Automation Research
University of Maryland College Park

ABSTRACT

To prepare for the challenge of handling the archiving and querying of terabyte-sized scientific spatial databases, the NASA Goddard Space Flight Center's Applied Information Sciences Branch (AISB, Code 935) developed a number of characterization algorithms that rely on supervised clustering techniques. The research reported upon here has been aimed at continuing the evolution of some of these supervised techniques, namely the neural network and decision tree-based classifiers, plus extending the approach to incorporating unsupervised clustering algorithms, such as those based on robust estimation (RE) techniques. The algorithms developed under this task should be suited for use by the Intelligent Information Fusion System (IIFS) metadata extraction modules, and as such these algorithms must be fast, robust, and anytime in nature. Finally, so that the planner/schedule module of the IIFS can oversee the use and execution of these algorithms, all information required by the planner/scheduler must be provided to the IIFS development team to ensure the timely integration of these algorithms into the overall system.

SUMMARY OF RESEARCH AND DEVELOPMENT

My work focused on a number of research and development efforts that were aimed, first and foremost, at the further advancement of the Intelligent Information Fusion System (IIFS) and the Regional Validation Center (RVC) projects under Code 935. Efforts concentrated on topics like unsupervised data clustering, supervised data classification, and image registration. The following text summarizes work accomplished in these domains.

In addition, I have continued to foster my professional contacts and related research interests at the University of Maryland and elsewhere. In particular, I have continued to pursue theoretical and practical aspects of computationally efficient algorithms for robust estimation and other research topics, e.g., robotic navigation/estimation. For details, see the attached list of publications.

Unsupervised Clustering of Remotely Sensed Imagery

Background

Numerous (unsupervised) clustering algorithms have been proposed and applied to various problem domains. Unfortunately, many of the strategies for partitional clustering may result in methods that depend heavily on assumptions made about the underlying probability density functions, the a priori number of clusters, etc. Also, these methods can be severely affected by outliers, i.e., data contaminated due to noise, coding errors, etc. A mode-seeking approach, however, can alleviate most of the above phenomena. In particular, processing (hyper)histograms in a discrete, multidimensional (feature) space seems especially suitable for mode-seeking-based clustering of multispectral images. Such an approach requires, on the other hand, careful consideration as far as the memory and run-time constraints of storing and searching multidimensional histograms are concerned.

Alternatively, we can adopt principles based on RE to construct a mode-seeking clustering scheme in a continuous domain. Consider a d -dimensional set of data which are hypothesized to belong to the same population. Estimating the parameters of the "point cloud" (e.g., its center and covariance matrix) depends heavily on the method employed. It is well known that in the presence of noisy data, a classical maximum-likelihood estimator (MLE), for example, can deviate arbitrarily from the underlying data pattern. On the other hand, an MVE estimator, which computes the center of the minimum volume (hyper)ellipsoid covering at least a fraction h of the data (where h can be taken equal to (approximately) $1/2$), is robust and can tolerate data contamination of up to 50%. It should be obvious, by definition, that the MVE estimator is a mode-seeking estimator. (See Figure 1.)

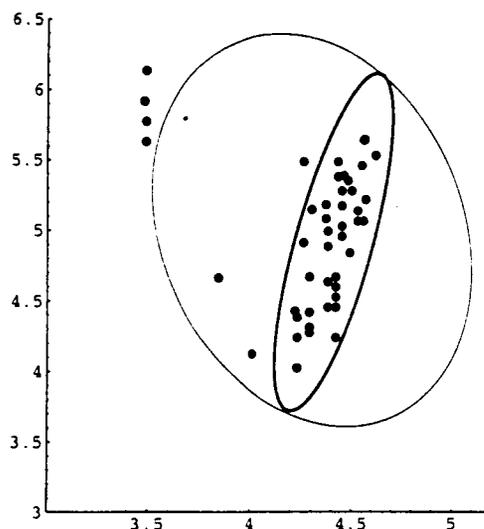


Figure 1: MVE (bold) vs. MLE cluster estimates for a contaminated set of points. Source: Rousseeuw and Leroy, 1987.

The Methodology Pursued

An algorithm for computing MVE is based on the (squared) Mahalanobis distance (MD) metric, often used to detect "leverage" points in multivariate data sets. Repeated computations using relatively small numbers of (random) data subsamples of size $d + 1$ yield an initial estimate of a minimum volume (hyper)ellipsoid containing exactly 50% of the data. This estimate is then refined by assigning a weight to every datum (depending on its proximity to the initial estimate) and computing a reweighted center of mass and a corresponding d -by- d covariance matrix which determines the shape of the hyperellipsoid fit. (See Rousseeuw and Leroy 1987.)

The MVE principle can be extended to detect several clusters in a multidimensional space by considering various (decreasing) $h \leq 0.5$ values per data set. The idea is to delete, at each step, a subset of points enclosed by the minimum volume hyperellipsoid for that value of h , such that the points in this hyperellipsoid "best" fit a cluster generated by an assumed underlying unimodal distribution. A "cluster hypothesis" is thus accepted for each deleted subset, and this process is repeated until the entire data set is exhausted. To examine cluster hypotheses, appropriate statistical tests can be employed (e.g., Chi-square, Kolmogorov-Smirnov (KS), etc.). Also, meeting a certain level of cluster validity may require, at times, merging clusters that initially were distinct. (See Jolion et al. 1991, Netanyahu 1994, for details.)

Another scheme based on the above approach is presented in Jolion et al. (1991) for low-dimensional ($d \leq 3$) data, with promising early results. Applying it, however, to remotely-sensed (multispectral) images having a larger number of bands, and studying its performance versus GT, requires further research.

Empirical Results

We have applied the above scheme to several remotely-sensed images. In particular, we have tested the algorithm on a 128 X 128 3-band Ridgely, Maryland Landsat/Thematic Mapper (TM) subimage. (Bands 2, 4, and 5 were used.) The image scene consists of various classes such as water, wetlands, shrub, trees, and various types of agricultural land. Figure 2(a) depicts band 5 of the input image. Also, Figure 2(b) shows a GT segmentation of the corresponding scene. The segmented image contains ten classes. It was arrived at by combining original ground truth with a hierarchical segmentation obtained by iterative parallel region growing (IPRG) (Tilton 1989). A typical result obtained from the clustering algorithm is illustrated in Figure 2(c).

We have also combined the current clustering approach in feature space with Tilton's spatial region growing approach. Specifically, a partially segmented image can be provided by the IPRG scheme as input to the RE-based clustering algorithm. This enables the clustering scheme to constrain the sampling, so that the number of subsamples that ensures a reasonable accuracy in the computation of the MVE can be greatly reduced. The clustering algorithm proceeds to complete the segmentation of the image. Figure 2(d) illustrates a segmented image arrived at using this combined approach.

Deriving contingency tables for the images displayed in Figures 2(c)-2(d) versus available GT reveals overall pixel "hits" of 65.7% and 69.6%, respectively. Furthermore, when aggregating into three classes (e.g., water/ wetlands, shrubs/ trees, and only one type of agricultural land), the hit ratios for the above images become 92.7% and 93.8%, respectively. This enhanced performance suggests that the unsupervised clustering scheme could be exploited in the context of hierarchical segmentation. Its output could serve as an approximation to be later refined by other classification methods.

Table 1 summarizes the "hit ratio" ranges obtained for the Landsat images that we have experimented with. The results for the Washington, D. C. image reflect only 13,274 pixels with available ground reference. At most 4 out of 7 bands of the TM images were used. The hit ratio depends, to some extent, on certain free parameters that are associated with the clustering scheme. (See Netanyahu 1994 for details.)

Supervised Classification of Remotely Sensed Images

Background

Under the IIFS project, a number of novel techniques have been developed for handling the automatic archiving and querying of terabyte-sized scientific spatial databases in a sophisticated manner. To allow for efficient (repetitive) querying of the resulting voluminous databases, archiving should involve characterization of the data and extraction of its content from remotely sensed imagery. The intent is to provide a variety of (image) classification tools that would meet the needs of multiple users with diverse (research) interests. In other words, the "pool" of image classification schemes that run under the IIFS should permit selection of a method according to the required level of understanding of the data's content.

For example, at the lowest level, a method might categorize every individual pixel, whereas at the highest level, a more global representative characterization could be provided, e.g., the list of classes found in an entire image (Chettri and Crompton 1993). In any event, providing good solutions, on different

hierarchical levels, to the problem of classification of multispectral images is of primary importance, and should prove extremely useful in many remote sensing/EOS-related applications (e.g., locating mineral deposits, vegetation/deforestation mapping, fire detection, etc.).

Work Accomplished

To pursue metadata extraction from remotely sensed imagery (i.e., obtain a "characterization vector" or a high level description of the contents of the data), in the global context of the IIFS, we have focused on the following subtasks, which are mainly supervised classification oriented.

- **Data Accumulation for Development/Experimentation**

To establish reliable, accurate classification schemes that operate on a large-scale basis, it is imperative that prototypes be tested as extensively as possible. Thus, as part of the infrastructure developed for the IIFS, we have populated its mass storage device mainly with scores of Landsat/TM images (courtesy of the Distributed Active Archive Center (DAAC), Code 902.2, NASA Goddard). Our complete pool of Landsat/TM data contains about 250 images (acquired over roughly 60-70 scenes around the globe) and occupies a total memory of ~25 Gbytes.

The concentrated gathering of such a diverse data set should be considered unique in a research and development environment, especially in view of the underlying accessibility constraints on Landsat/TM data. Such a data set should prove invaluable in further developing our classification schemes as it enables experimentation with large numbers of images, generation of training sets (see the next subsection), testing the notion of "signature extension", etc.

- **(Semi)automated Procedures for Training Set Generation**

Given the technical difficulties associated with ground truth (GT), particularly its lack of availability for the most part, we have taken an alternative approach to generating training sets that are essential for supervised classification. Our approach is based on photointerpretation (PI) of a given scene, whereby a user assigns labels to a relatively small sample of pixels. It is desirable, of course, that such pixel labeling be carried out by users who are familiar with the scene or by remote sensing specialists. Once a labeled set of pixels is available (associated with each labeled pixel are the classification [i.e., label], spectral readings, ancillary data, etc.), training subsets can be derived and experimented with for image classification.

In an attempt to automate the above procedure(s), we have specified an interactive photointerpretation tool (PIT) which allows a user to select, label, and record a desired sample of pixels. (The set of "committed" pixels can easily be modified at a subsequent PIT session, depending on the user's preferences, interests, etc.) Also, we have generated various interactive conversion modules that transform a "photointerpreted image" (i.e., PIT's output) to a training set in a format that meets the input specifications of the classification schemes.

- **Efficient Algorithms for Supervised Classification**

Previous studies carried out in the Applied Information Sciences Branch suggest that extracting metadata from remotely sensed images in a fast and relatively accurate manner can be adequately achieved through a combination of probabilistic neural networks (PNN) and networks that are backpropagation-trained. (Both techniques require a user-supplied training set.) Given its fast training time, the PNN initially serves to establish an "optimal" training set that is representative of all

of the classes in a given (set of) scene(s), after which a trained backpropagation network is invoked in its feed-forward mode to classify these scenes.

The PNN's learning process is based on nonparametric density estimation techniques. Specifically, the network estimates the probability density function of a newly introduced test pattern by computing for each class (that is present in the training set) a sum of Gaussians centered at each individual training pattern that belongs to the class, and evaluated at the test pattern. The pixel (or test pattern) is assigned to that class for which the above computation is the highest.

From the description above, it should be apparent that the PNN becomes computationally intensive if it is used to classify a very large number of patterns (e.g., millions of pixels). On the other hand, the PNN lends itself naturally to single-instruction, multiple-data (SIMD) parallelization. To reduce its run time significantly, we have implemented a parallel version of the PNN on a massively parallel machine, the 16K processing element MasPar, MP-2. (As an example, it takes the parallel version approximately 30sec (net data processing unit (DPU) time) to classify a 7-band 512 X 512 Landsat image against a training set of ~1500 patterns, as opposed to the hours it takes for a sequential version to run on a SUN/SPARC workstation.)

We have also parallelized a nearest neighbor (NN) classification scheme and run it on the MasPar. (The run-time and classification results appear to be comparable to those of the PNN.) The module can be readily incorporated into the IIFS, thereby further enhancing the system's overall versatility.

Having generated training sets from a number of Landsat/TM images (in accordance with the previous subsection), we ran the PNN (and the NN searching module) to classify these and other Landsat/TM scenes. The characterization vectors obtained serve to populate the database of the IIFS and help meet relevant prospective queries with respect to the data's contents. Figure 3 depicts the spectral signatures of the training set patterns selected (via PIT) from a Landsat/TM scene over Ashdod, Israel. (The land cover types, i.e., the scheme, is in accordance with Anderson's USGS Level I Standard.)

Figure 4 provides examples of the classification maps obtained using the above-described supervised schemes. It shows (a) band 3 of a 512 X 512 subimage of the Landsat/TM image over Israel, (b) its classification using the PNN, and (c) its classification using NN searching. Note the good segmentation into "urban", "agriculture", "water", and "barren" (e.g., beach sand) classes. Likewise, Figure 5 demonstrates that the above procedures can be applied to AVHRR images, classifying them into cloud height categories. It shows (a) band 2 of a 512 X 512 AVHRR subimage, (b) its PNN-based classification, and (c) its classification using NN searching.

It should be noted that in cases where only an approximation of the contents of a data set is required, classifying only a relatively small sample of the image should suffice. Specifically, we have shown (theoretically and empirically) that for a typical 2984 X 4320 Landsat/TM scene quadrant, it suffices to classify only a sample of ~17000 pixels to obtain a characterization vector whose individual components are within 1% of the corresponding class frequencies in the entire image. This feature offers extended flexibility, as far as planning/scheduling of the concurrent tasks that the IIFS is expected to handle. (See Netanyahu and Cromp (1995) for a detailed discussion.)

- Integration with Wavelet Parameters

When using a straightforward classification scheme, most of the classification errors are expected to occur at the boundaries between classes. These errors seem to result from the fact that pixel-based (e.g., neural network-based) classification methods do not incorporate local, spatial information in the classification process. Based on this premise, we have conducted a preliminary study of the impact

that such information could have on the overall classification. Specifically, to improve the performance i.e., accuracy of the classification module(s) discussed previously, it is estimated that the PNN provides ~70-80% accuracy. We have attempted to exploit texture information (in addition to spectral intensities) in the overall framework of a neural network-based classification scheme. (See Szu et al. 1997 for a detailed discussion.)

In general, texture can be captured through statistical, spectral, structural, or model-based methods. In particular, statistical methods rely on the spatial distribution and spatial dependence among local gray tones. (See Haralick (1979), Davis et al. (1979) on the notion of a co-occurrence matrix as a means of capturing texture information.) Spectral methods, on the other hand, extract relevant texture information by computing the energy associated with different frequency bands through the use of the Fourier transform. Recently, there has been a growing interest in wavelet transforms for texture extraction. (In a nutshell, wavelet transforms provide better localization than traditional Fourier transforms, as well as better division of the time/space-frequency domain(s) than so-called windowed Fourier transforms.) Wavelet-based texture analysis can be regarded as both a statistical and spectral technique since an isotropic wavelet exploits the localization of a wavelet transform in computing (local) energy properties, and since it also provides a spatial density function of the co-occurrence texture. Indeed, various recent studies have demonstrated the usefulness of applying wavelets in texture analysis. (See, e.g., Rogers et al. 1992, Unser 1993, Manjunath and Ma 1996).

In Szu et al. (1997), we pursued the following combined approach. First the multiband image(s) were pre-processed, retaining the most significant bands obtained by a principal component analysis (PCA). (For example, Figures 6(a)-(b) depict the two most significant PCA bands obtained for a Landsat/TM image over Washington, D. C.) Then, an isotropic, composite wavelet filter (which preserves edge information while emphasizing texture components) was applied to these PCA bands. (The specific wavelet transform was realized as a combination of a Mexican Hat and a Morlet-type wavelet. Also, a number of scale values, a , were experimented with for each PCA band. See Figures 6(d)-(f).)

The augmented set of images, i.e., the PCA bands plus the additional images produced by the above wavelet transform, were then fed as new input to the PNN. (Training sets, too, were augmented accordingly.) Preliminary classification results suggest that additional wavelet information may enable refinement of the classification in regions where texture is well structured, e.g., urban areas. See, for example, Figures 6(g)-(h). On the other hand, this information may not be as useful in regions where spectral information is sufficient to determine the classification (e.g., water).

- **Mixed Pixel Classification**

Usually classification procedures assume a "winner take all" strategy with a single pixel being placed within one land use/land cover category. While this is a valid assumption over large homogeneous regions, there are many cases where a pixel actually consists of several ground cover elements. For example, the National Oceanic and Atmospheric Administration (NOAA) AVHRR platform series consists of sensors having a spatial resolution of 1.1 [km/pixel]. Given this (low) resolution, it is quite likely that several surface cover classes will be found within each pixel in the image. We refer to such a pixel as a mixed pixel. Thus, in classifying mixed pixels, an alternative approach would be to assume that a pixel is composed of all possible end-members (i.e., pure classes) of a given scheme, e.g., that its observed (spectral) reflectance comprises a linear combination of the (spectral) reflectances of the various end-members. Determining the abundance (i.e., the relative fraction) of each component is known as the (linear) mixture modeling problem. Providing a solution for this problem is considered more appropriate, in many cases, for the extraction of metadata/content from remotely sensed images.

Indeed, a number of approaches to the spectral unmixing problem have been proposed in the remote sensing literature. For example, Shimabukuro and Smith (1991) and Settle and Drake (1993) present solutions to the problem that are based on the conventional principle of constrained least squares. However, these solutions are not general (the number of end-members is assumed smaller, e.g. 3, than the number of spectral bands) and they are ad hoc in the way the desired coefficients are guaranteed to be non-negative. A more recent paper presents a similar approach, but it, too, does not provide a general solution to the problem. Extending the technique beyond the 4-class case considered in Bosdogianni et al. (1997) is not obvious, and maintaining non-negative coefficients in the general case could become highly inefficient.

In Chettri and Netanyahu (1996) we have proposed an alternative approach which is based on the principle of maximum entropy. We provide extensive justification, from an information-theoretic and a combinatorial perspective, for the appropriateness of this approach in the specific context of mixed-pixel classification. Our methodology yields a general solution to the linear mixture modeling problem by optimizing the entropy associated with the abundance vector. Also, it guarantees (by definition) non-negative fractional distributions.

The paper first introduces the mathematical formalism of linear mixture modeling and the principle of maximum entropy. It then presents algorithms for the solution of the linear mixture model both in the conventional case (Gebbinck and Schouten 1995) and in the case where a more sophisticated technique, the penalty function method (Zangwill 1967), is employed to optimize the entropy function. The issue of noise is dealt with in both cases. In addition, the paper makes an interesting connection to Bayesian methods, and compares the maximum entropy-based approach with standard regularization methods. Finally, empirical results are presented for real and simulated satellite data. Real data were obtained from a large repository of AVHRR data and simulated data were obtained via realistic models of reflectance properties of different land surface types and the inclusion of atmospheric effects. Our preliminary results suggest that the maximum entropy-based method yields higher accuracy. (Although an earlier paper by Pendock (1992) on the application of maximum entropy to the unmixing of mineral spectra appeared in the remote sensing literature, it did not analyze noisy data, did not provide a formulation of the method in a Bayesian context, and lacked a performance comparison to conventional methods.)

In an attempt to reduce the computational complexity of the maximum entropy method, we have proposed to extend the method by using the concept of multiresolution. The basic idea is to decompose (recursively) a given image into a set of lower-resolution images by applying, for example, well-known wavelet transforms, e.g., the Haar transform. At each level of decomposition, four new images are computed; each image represents low/high frequency information along the horizontal and vertical directions in the original/previous image. We then carry out the maximum entropy unmixing technique with respect to one of these images (e.g., the so-called low/low (LL) image). Applying our technique to this smaller-size image yields satisfactory fraction approximations as compared to those obtained for the original image while gaining considerable speedup. The resulting procedure is roughly 4 times faster. See Chettri et al. (1997) for a detailed discussion.

Image Registration

Background

Many of the analysis techniques that a Regional Validation Center (RVC) is expected to utilize will involve the integration of multiple data sources. For example, the analysis of global coverage by low-resolution data could be validated by using local, very-high resolution data. As image registration (e.g., image-to-image registration, absolute georegistration, etc.) comprises an important first step of such an operation—

to enable users to analyze large amounts of pertinent data more accurately and efficiently—a preliminary requirement is that a sound, automatic image registration scheme be incorporated into the RVC.

Although automatic image registration has been extensively studied in the domain of image processing, it continues to pose a formidable challenge in the realm of remote sensing. This can be attributed, in part, to the vast diversity of instruments/data sources that have become available, the large areas covered by these instruments, etc. Recently, a team of researchers in the AISB has proposed to develop an operational toolbox that will consist of various basic image registration modules and to carry out a comparative performance evaluation of these modules with respect to criteria such as accuracy, level of automation, computational efficiency, etc. (See Xia et al. 1997.) The idea is not only to provide a user with a variety of frequently used registration techniques (thereby eliminating, to some extent the redundancy that exists in the remote sensing community), but to enable him to select an optimal, end-to-end registration scheme, based on the above evaluation study.

We distinguish between two fundamental approaches to image registration. Essentially, one approach makes direct use of the original data (or edge gradient data), whereas the other is based on feature matching. (Features can be control points, corners, line segments, etc., and they are assumed available through the use of standard feature extraction algorithms.) The former approach, which is usually based on a correlation measure between the two images, could prove computationally expensive and sensitive to noise unless substantial preprocessing is employed. The latter approach, on the other hand, tends to yield more accurate results as features are usually more reliable than intensity or radiometric values. We note, though, that the matching process may run into difficulties if feature extraction results in missing/spurious points.

One disadvantage of feature-based methods is that they, too, can become computationally expensive, though perhaps not as much as correlation-based methods. (See Ton and Jain (1989) for a partial review of the computational complexity of point matching methods.) Thus, in an attempt to arrive at a sound registration scheme that would be both accurate and fast, we have begun to pursue an efficient algorithmic framework for a robust, feature-based matching module. We believe that the availability of such a generic module will extend the flexibility of our registration toolbox and enhance its overall functional capability.

Robust Image Matching

Given two images of roughly the same scene, image registration is the process of determining the transformation that most nearly maps one image to another. This problem is of particular interest in remote sensing applications, where it is known that two images correspond to roughly the same geographic region, but the exact alignment between the images is not known. There have been many techniques reported with respect to image registration. In Mount, Netanyahu, and Le Moigne (1997), we have considered an approach based on extracting a set of point features from each of the two images, thereby reducing the problem to a point pattern matching problem.

Because of measurement errors and the presence of outlying data points in either of the images, it is important that the distance measure between two point sets be robust to these effects. Specifically, we have used the generalized Hausdorff distance (Huttenlocher et al. 1992, 1993, 1997) to measure these distances. An important element of image registration applications is that the search begins with a priori information on the bounds of the desired transformation. Thus, a good algorithm should take advantage of this information.

As indicated, point matching can be a computationally intensive task. Indeed, there have been a number of algorithmic approaches proposed for solving this problem, both from theoretical and applied standpoints. One common approach is based on a geometric branch-and-bound search of

transformation space and another is based on using point alignments to derive the matching transformation. The former has the advantage that it can provide guarantees on the accuracy of the final match, and that it naturally uses any a priori information to bound the search. The latter has the advantage of simplicity and speed.

In Mount et al. (1997) we have introduced a new practical approximation algorithm, which we called bounded alignment, for robust point pattern matching. The algorithm derived comprises a novel combination of the above two approaches, in the sense that it operates within the framework of branch-and-bound, but employs point-to-point alignments to accelerate the search. Furthermore, we have shown that this combination retains many of the strengths of branch-and-bound search, but provides significantly faster search times by exploiting alignments. Finally, we have implemented the algorithm and demonstrated its performance on synthetically generated, as well as some remotely sensed data. (See Mount et al. (1997) for a detailed discussion.)

SEMINARS, LECTURES, PRESENTATIONS

Unsupervised, Robust Estimation-based Clustering of Remotely Sensed Images. IEEE International Geoscience and Remote Sensing Symposium", Florence, Italy, July 12, 1995.

Detection and Reclassification of Mixed Pixels for Enhanced Metadata Extraction. Internal talk, NASA Goddard Space Flight Center, Earth and Space Data Computing Division, Information Sciences and Technology Branch (Code 935), November 9, 1995.

Remote Sensing Products for the Intelligent Information Fusion System. NASA Goddard Space Flight Center, Regional Data Center Kick-Off Workshop, March 21, 1996.

Tools and (Parallel) Algorithms/Product Generation for the Regional Data Center System. NASA Goddard Space Flight Center, Regional Data Center Training Workshop, July 19, 1996.

Robust Estimation of Parameters for Fitting Circular Arcs (with Arnold J. Stromberg). Symposium on Statistics and the Sciences, Halifax, Nova Scotia, Canada, August 15, 1996.

Robust Detection of Road Segments in Noisy Aerial Images. Thirteenth IAPR International Conference on Pattern Recognition, Vienna, Austria, August 27, 1996.

Efficient Algorithms for Robust Point Matching. CESDIS Image Registration Workshop, NASA/GSFC, November 21, 1997.

PUBLICATIONS AND PAPERS

Netanyahu, N., & Weiss, I. (1994). Analytic line fitting with non-Gaussian noise. *Proceedings of the NSF/ARPA Workshop on Performance versus Methodology in Computer Vision.* (pp. 146-151). Seattle, Washington, June 24-25.

Netanyahu, N., Short, N. M., Jr., Crompton, R. F., Campbell, W. J., Tilton, J. C., Le Moigne, J., Fekete, G., and Sylvain, G. (1994). AI challenges within NASA's Mission to Planet Earth. *Proceedings of the AAAI Workshop on AI Technologies for Environmental Applications.* (pp. 1-15). Seattle, Washington, July 31-August 4.

Netanyahu, N., & Weiss, I. (1994) Analytic outlier removal in line fitting. *Proceedings of the Twelfth IAPR International Conference on Pattern Recognition: Vol. II.* (pp. 406-408). Jerusalem, Israel, October 9-13.

- Netanyahu, N., * Crompton, R. F. (1995). *Random Sampling of Remotely Sensed Images for Efficient Metadata Extraction*. Internal publication, NASA Goddard Space Flight Center, Information Sciences and Technology Branch (Code 935).
- Netanyahu, N., Tilton, J. C., & Gualtieri, J. A. (1995). Unsupervised, robust estimation-based clustering of remotely sensed images. *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium*. (pp. 1150-1152). Florence, Italy, July 10-14.
- Netanyahu, N., Mount, D. M., Silverman, R., & Wu, A. (1995). Chromatic approximate nearest neighbor searching: A query sensitive approach. *Proceedings of the Seventh Canadian Conference on Computational Geometry*. (pp. 261-266). Quebec City, Quebec, Canada, August 10-14. To appear in *Computational Geometry Theory and Applications*.
- Netanyahu, N., Short, N. M., Jr., Crompton, R. F., Campbell, W. J., Tilton, J. C., Le Moigne, J., Fekete, G., Ligon, W. B., III, & Wichmann, K. (1995). Mission to Planet Earth: AI views the world. *IEEE Expert*, 10(6), 24-34.
- Netanyahu, N., Stromberg, A. J., Philomin, V., & Rosenfeld, A. (1996). *Robust Estimation of Parameters for Fitting Circular Arcs*. Presented at the Minerva Symposium on Statistics and the Sciences, Halifax, Nova Scotia, Canada, August 12-16.
- Netanyahu, N., Philomin, V., Rosenfeld, A., & Stromberg, A. J. (1996). Robust detection of road segments in noisy aerial images. *Proceedings of the Thirteenth IAPR International Conference on Pattern Recognition, Vol. B*. (pp. 151-155). Vienna, Austria, August 25-30. Also in (1997) *Pattern Recognition*, 30(10), 1673-1686.
- Netanyahu, N., Cucka, P. & Rosenfeld, A. (1996). Learning in navigation: Goal finding in graphs. *International Journal of Pattern Recognition and Artificial Intelligence*, 10(5), 429-446. Special issue in memory of King-Sun Fu. Also in Freeman, H., (Ed.). (1996). *World Scientific Series on Machine Perception & Artificial Intelligence: Vol. 25. Studies in Pattern Recognition* (pp. 39-56).
- Netanyahu, N., Szu, H. H., Le Moigne, J., & Francis, M. (1996). Wavelet index of texture for artificial neural network classification of Landsat images. *Proceedings of the Twenty Fifth Annual SPIE Applied Imagery Pattern Recognition Workshop* (pp. 36-44). Washington, D. C., October 16-18.
- Netanyahu, N., & Chettri, S. (1996). Spectral unmixing of remotely sensed images using maximum entropy. *Proceedings of the Twenty-fifth Annual SPIE AIPR Workshop on Emerging Applications of Computer Vision* (pp. 55-62). Washington, D. C., October 16-18.
- Netanyahu, N., Cucka, P., & Rosenfeld, A. (1996). *Robotic Estimation*. College Park: University of Maryland, Center for Automation Research. (CS-TR-3699).
- Netanyahu, N., Mount, D. M., Romanik, K., Silverman, R., & Wu, A. (1997). A practical approximation algorithm for the LMS line estimator. *Proceedings of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms* (pp. 473-482). New Orleans, Louisiana, January 5-7.
- Netanyahu, N., Le Moigne, J., & Lyon, R. (1997). Remote sensing image processing at CESDIS. USRA Newsletter *CESDIS in Review*, 3(Winter).
- Netanyahu, N., Szu, H. H., Le Moigne, J., & Hsu, C. (1997). Integration of local texture information in the automatic classification of Landsat images. *Proceedings of the SPIE Conference on Wavelet Applications* (pp. 116-127). Orlando, Florida, April 22-24.

Netanyahu, N., Chettri, S., Garegnani, J., Robinson, J., Coronado, P., Cromp, R. F., & Campbell, W. J. (1997). Multiresolution maximum entropy spectral unmixing. *Proceedings of the International Symposium on Artificial Intelligence, Robotics, and Automation in Space* (pp. 347-352). Tokyo, Japan, July 14-16.

Netanyahu, N., Xia, W., Le Moigne, J., Tilton, J. C., Lerner, B-T., Kaymaz, E., Pierce, J., Raghavan, S., Chettri, S., El-Ghazawi, T, Mareboyana, M., Campbell, W. J., & Cromp, R. F. (1997). A registration toolbox for multi-source remote sensing applications. *Proceedings of the International Conference on Earth Observation and Environmental Information*. Alexandria, Egypt, October 13-16.

Netanyahu, N., Mount, D., & Le Moigne, J. (1997). An efficient algorithm for robust feature matching. *Proceedings of the CESDIS Image Registration Workshop* (pp. 247-256). NASA Goddard Space Flight Center, Greenbelt, Maryland, November 20-21.

Netanyahu, N., Le Moigne, J., Xia, W., Chettri, S., El-Ghazawi, T., Kaymaz, E., Lerner, B-T., Mareboyana, M., Pierce, J., Raghavan, S., Tilton, J. C., Campbell, W. J., & Cromp, R. F. (1997). Towards an intercomparison of automated registration algorithms for multiple source remote sensing data. *Proceedings of the CESDIS Image Registration Workshop* (pp. 307-316). NASA Goddard Space Flight Center, Greenbelt, Maryland, November 20-21.

Netanyahu, N., & Mount, D. M. (1997). *Efficient Randomized Algorithms for Robust Estimation of Circles and Aligned Ellipses*. College Park: University of Maryland, Center for Automation Research (CS-TR-3860).

Scene	Landsat	d	Size	Hit ratio
Ridgely, MD	TM	3/7	128×128	65.7–69.6%
Blackhills, SD	MSS	4/4	512×512	66.5%
Wash., DC	TM	4/7	256×256	73.1–78.1%

Table 1: Summary of results for various Landsat images

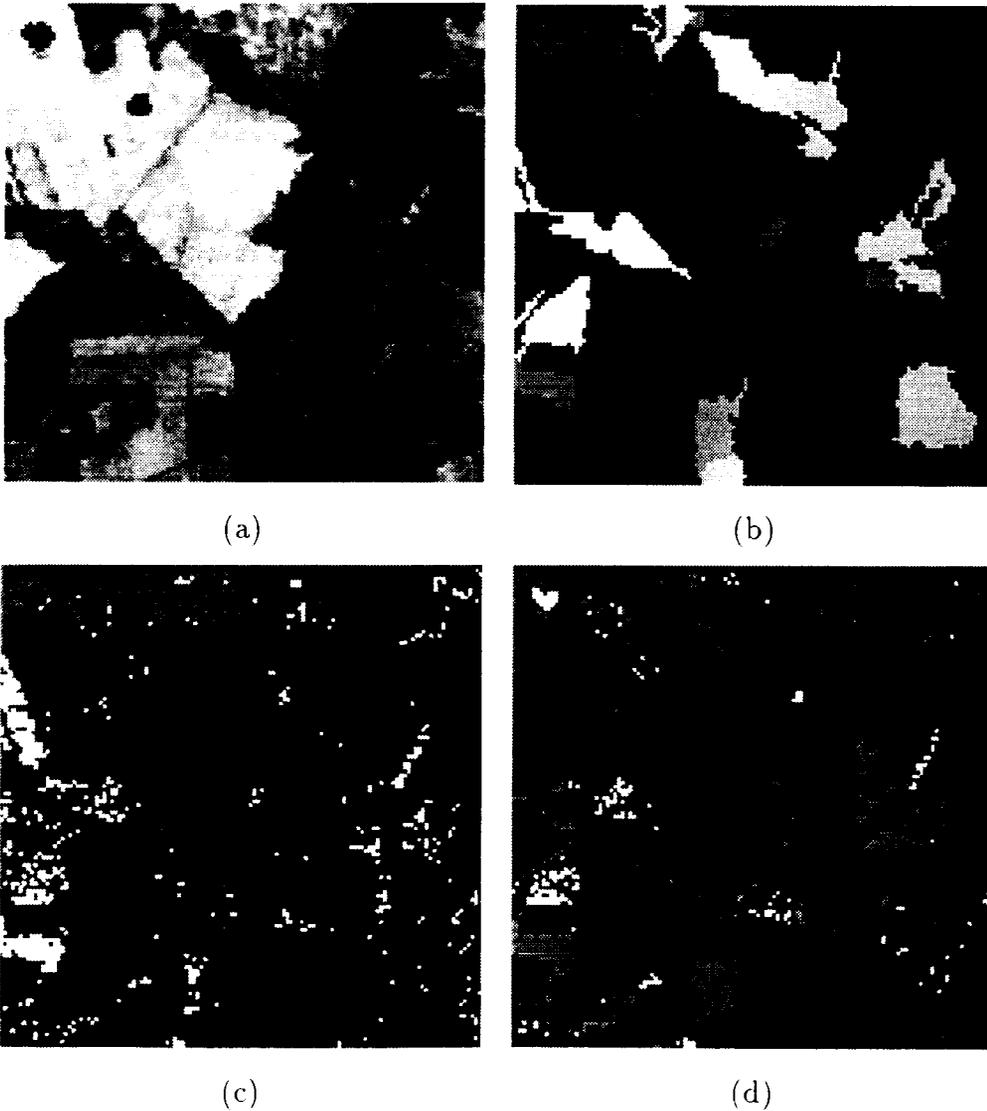


Figure 2: (a) Band 5 of a 128×128 Landsat/TM subimage of Ridgely, MD; (b) Ground truth; (c) RE-based clustering; (d) RE and IPRG combined

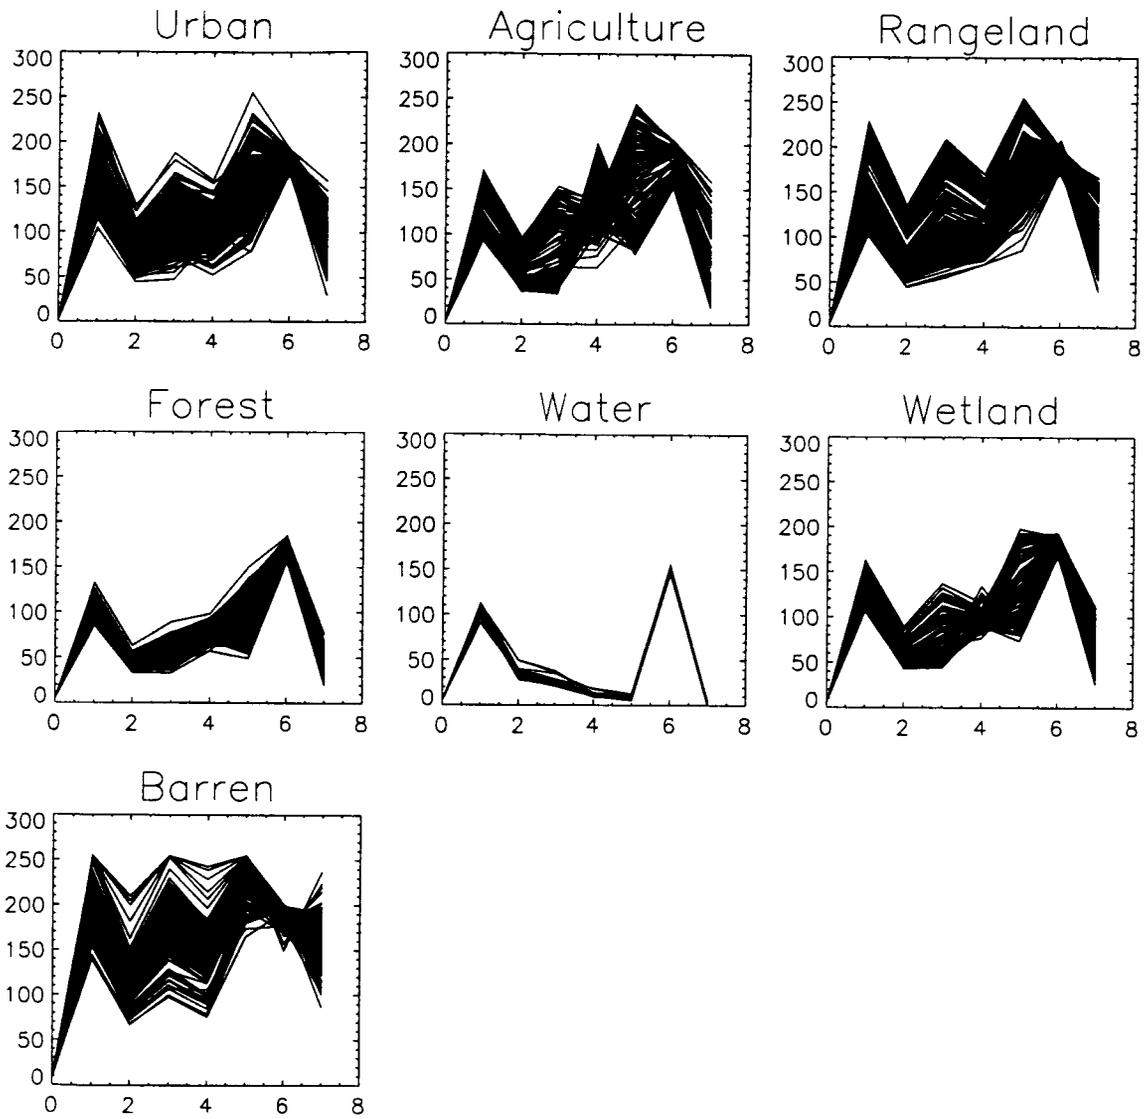
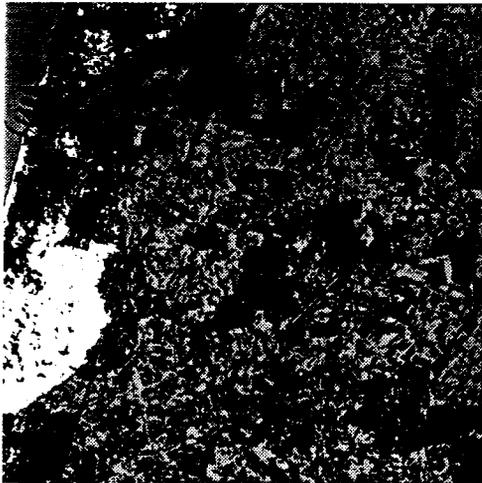


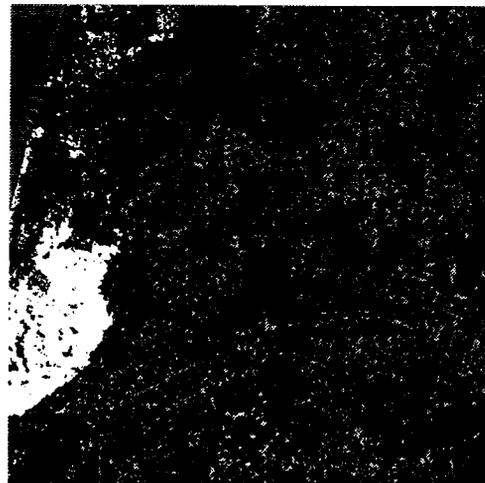
Figure 3: Training sets obtained by photo-interpretation for land use/land cover applications.



(a)

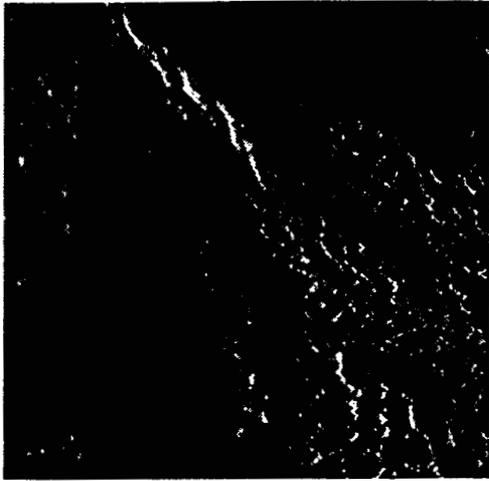


(b)

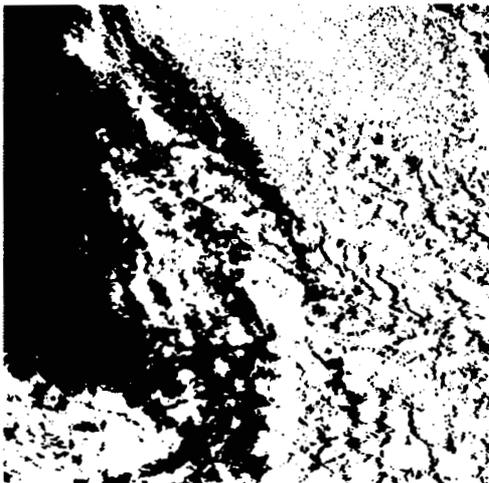


(c)

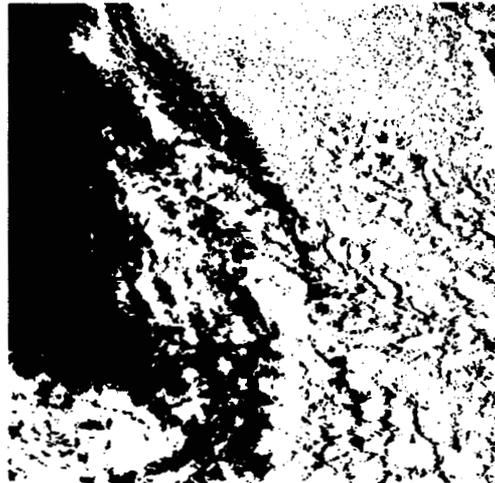
Figure 4: (a) Band 3 of a 512×512 Landsat/TM subimage of Ashdod, Israel; (b) land-cover classification via a probabilistic neural network; (c) ditto, using nearest-neighbor searching.



(a)



(b)



(c)

Figure 5: (a) Band 2 of a 512×512 AVHRR scene; (b) its, classification (into cloud height categories) via a probabilistic neural network; (c) ditto, using nearest-neighbor searching.

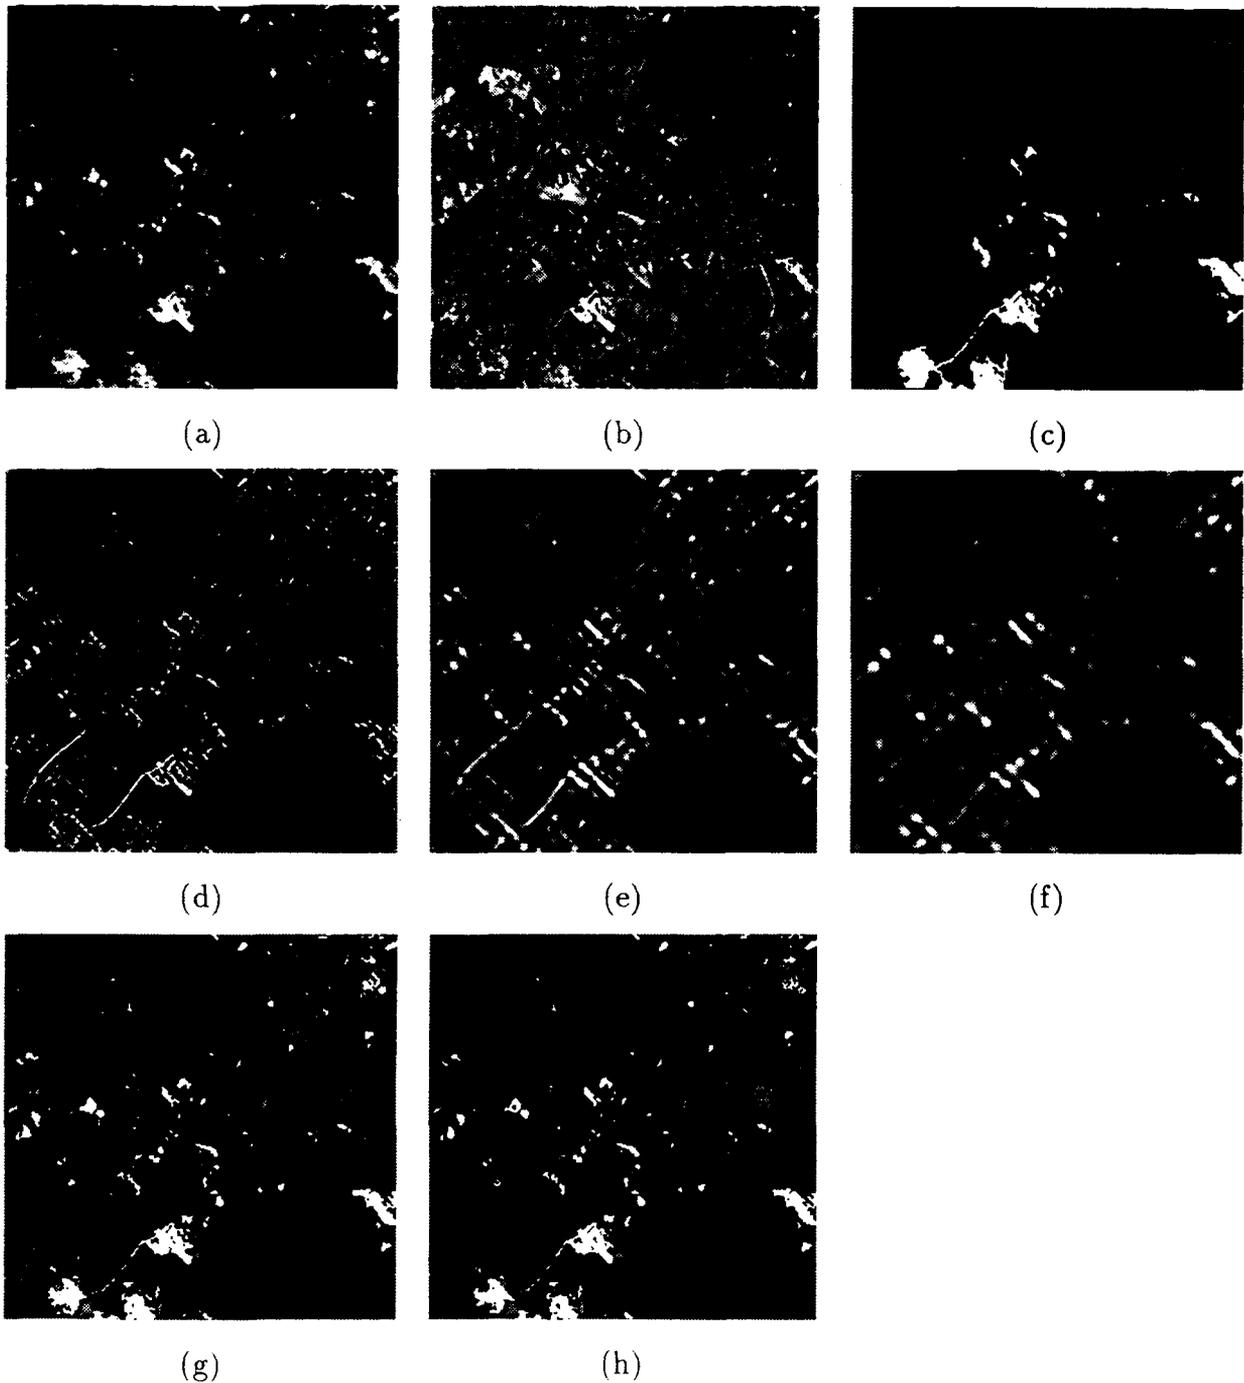


Figure 6: (a) Most significant band of a 256×256 Landsat/TM scene over Washington, DC; (b) second PCA band of the above scene; (c) ground truth image; (d) wavelet transform (applied to most significant PCA band) for $a = 0.25$; (e) ditto, for $a = 0.5$; (f) ditto, for $a = 0.75$; (g) PNN-based classification; (h) ditto, with spatial information (using a wavelet transform).

NASA			
Report Documentation Page			
1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.	
4. Title and Subtitle Unsupervised, Robust Estimation-based Clustering for Multispectral Images		5. Report Date	6. Performing Organization Code
7. Author(s) Nathan S. Netanyahu		8. Performing Organization Report No.	
9. Performing Organization Name and Address Computer Vision Laboratory Center for Automation Research University of Maryland College Park, MD 20742		10. Work Unit No.	
		11. Contract or Grant No. NAS5-32337 USRA subcontract No. 5555-37	
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Washington, DC 20546-0001 NASA Goddard Space Flight Center Greenbelt, MD 20771		13. Type of Report and Period Covered Final July 6, 1993 - December 31, 1997	
		14. Sponsoring Agency Code	
15. Supplementary Notes This work was performed under a subcontract issued by Universities Space Research Association 10227 Wincopin Circle, Suite 212 Columbia, MD 21044 Task 44			
16. Abstract To prepare for the challenge of handling the archiving and querying of terabyte-sized scientific spatial databases, NASA GSFC's Applied Information Sciences Branch developed a number of characterization algorithms that rely on supervised clustering techniques. The research reported upon here has been aimed at continuing the evolution of some of those supervised techniques, namely the neural network and decision tree-based classifiers, plus extending the approach to incorporating unsupervised clustering algorithms, such as those based on robust estimation techniques. The algorithms developed under this task should be suited for use by the Intelligent Information Fusion System metadata extraction modules, and as such these algorithms must be fast, robust, and anytime in nature. Finally, so that the planner/scheduler module of the IIFS can oversee the use and execution of these algorithms, all information required by the planner/scheduler must be provided to the IIFS development team to ensure the timely integration of these algorithms into the overall system.			
17. Key Words (Suggested by Author(s)) spatial databases remote sensing		18. Distribution Statement Unclassified--Unlimited	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 16	22. Price