Reprinted from

Eighth International Symposium

Machine Processing of

Remotely Sensed Data

with special emphasis on

Crop Inventory and Monitoring

July 7-9, 1982

Proceedings

Purdue University
The Laboratory for Applications of Remote Sensing
West Lafayette, Indiana 47907 USA

ADVANCES IN INFORMATION EXTRACTION TECHNIQUES

G. NAGY

University of Nebraska/Department of
Computer Science
Lincoln, Nebraska

ABSTRACT

The analysis of satellite images may evolve from ad hoc methods of utilizing spatial and temporal context to the application of artificial intelligence oriented procedures of hierarchical scene analysis. Pattern representations more abstract than Euclidian vector spaces offer some hope of unifying structural and decision theoretical approaches. The estimation of expected classification error rates is becoming more sophisticated and rigorous, but useful finite-sample results for non-parametric distributions appear unobtainable. Focus on computational complexity allows comparison of algorithms, while software engineering techniques reduce the effort necessary to develop and maintain complex image processing systems. Advances in computer systems architecture, commercial database technology, and man-machine communications should be closely monitored by the remote sensing community. A NASA-sponsored task force offers comprehensive recommendations for research directions in mathematical pattern recognition.

I. INTRODUCTION

The objective of this paper is to present sundry recent developments which show some potential for affecting the automatic extraction of information from remotely sensed data. Although no attempt is made at exhaustive coverage of all potentially useful items, both theoretical concepts and technological developments are mentioned. The concluding section summarizes the author's impressions from a NASA-sponsored task force on suitable directions for pattern recognition research for remote sensing.

II. THEORETICAL DEVELOPMENTS

With a few exceptions, the ideas referenced below have not yet been applied to satellite remote sensing, and it is not clear that they ever will be. All of them, however, are related to classification.

A. EXPLOITATION OF SPATIO-TEMPORAL CONTEXT

A potentially far-reaching method of integrating the spatial, temporal and spectral components of multiple observations of the same scene has been proposed by Haralick and Shapiro. The consistent labeling procedure that they advocate is a hierarchic extension of decision theoretic classification. In addition to the customary conditional probabilities for each pixel given the class, a "world model" is postulated that specifies, in either probabilistic or deterministic terms, the permissible spatial or temporal configuration of pixels bearing each class label.

Configurations that satisfy the category constraints of the world model are determined by dynamic programming, Viterbi algorithm, relaxation, or other search techniques. World models may be specified by semantic networks, prediction rules, or Markov models. The process begins at the pixel level, combining pixels into homogeneous units, but these units are then assembled into higher-level units in the same manner.

The approach attempts to unify structural and statistical classification methods on the one hand, and pattern recognition and artificial intelligence techniques on the other. The transition from local to consistent global interpretation may correspond, in practice, to obtaining land use, eco-type, or habitat information from landcover in a systematic manner. Much remains to be worked out in detail, particularly with regard to the specification of higher-level constraints, before the method is likely to prove applicable to practical problems. [1,2]

Levine approaches scene analysis of color photographs of outdoor scenes from the point of view of computer vision and artificial intelligence. His modular, rule-controlled system is data directed and knowledge based. The knowledge base (rules) are in a permanent relational database, while the results of current operations, including hypotheses, are in a similarly organized short-term memory. Boundary and region analysis and relaxation labeling algorithms are used to partition the scene into interpreted components. These cooperative processes or algorithms operate at various levels in the system hierarchy. Although there is no unified theoretical formulation that encompasses the entire system, and the different components interact in such a complex manner that it may be difficult to find systematic means to improve performance, this approach represents a distinct departure from current methods of processing remotely sensed data. [3]

The notion of "spectral-temporal" trajectories is demonstrated by Wheeler and Misra. A trajectory for a specific surface patch is characterized by the angles of the vectors between pairs of two-dimensional points in the brightness-greenness plane obtained at different dates. The results are, however, difficult to assess, since no ground-truth measurements for the actual pixels was available. [4]

The practical problems of integrating spatial and spectral information are discussed in an article by Landgrebe that offers an unusual perspective on the development, over a nine-year period, of set of pattern-recognition procedures ("ECHO") for image data obtained through satellite remote sensing. The alternatives available at the various stages of development are discussed and perspicacious arguments are advanced in support of the final approach. The presentation of half a dozen different criteria for the comparison of various algorithms is itself instructive, particularly considering that every criterion was evaluated in hundreds of separate experiments, each requiring the assignment of class labels to tens of thousands of observation vectors. [5]

B. IMBEDDING OBSERVATIONS IN ABSTRACT SPACES

A bold attempt to unify structural and statistical pattern recognition is the abstract representation of Goldfarb. Goldfarb's principal argument against the customary Euclidian representation of patterns is that, contrary to widespread misconception, Euclidian spaces cannot accomodate arbitrary measures of similarity between patterns.

Consider, for example, four patterns A, B, C, and D such that

$$d(A,B) = d(B,C) = d(A,C) = 1.00$$

and $d(A,D) = d(B,D) = d(C,D) = 0.55$.

The dissimilarity measure d actually satisfies the metric inequality, yet there is no Euclidian space, of any dimensionality, wherein the patterns can be embedded in such a way as to preserve the pairwise similarities.

Since the notion of vector representation has proved so important in pattern analysis, the question is whether there exists any type of vector space that preserves the original structure of the data as obtained through the notion of pairwise similarities - considered as the fundamental property - yet lends itself to the machinery of discriminant analysis, statistical decision theory, clustering, and other standard analytical tools of pattern recognition.

Goldfarb's solution is to imbed the patterns in a "pseudo-Euclidian" vector space. A pseudo-Euclidian space is a pseudo-metric space with a non negative, real valued mapping which satisfies the symmetry and reflexivity axioms, and on which is defined a symmetric bilinear form or generalized inner product. The measure of dissimilarity or "distance" between two points in such a space can, therefore, take on negative values. A widely-known example is the Minkowsky space of special relativity.

Goldfarb develops the mathematical properties of this representation and demonstrates the correspondence necessary to implement the standard pattern analysis techniques on a given set of data. He argues that the approach accomodates mixed-mode observations, including structural or syntactic features, in a manner inherently impossible in Euclidian space. The appropriate pseudo-Euclidian space is defined by the relations of the patterns themselves, resulting in a parsimonious representation of minimum dimensionality. The addition of any new pattern may modify the space and affect the representation of all other patterns. If, however, the original sample size is sufficiently large, then a new pattern introduces no additional constraint and can be located in the existing space. [6]

To the author's knowledge, these ideas have not yet been subjected to any significant experimental test.

C. FINITE SAMPLE ERROR ESTIMATES

At the 1966 IEEE Workshop on Pattern Recognition, Louis Fein made a plea for "Impotence Principles for Machine Intelligence," citing inspiring examples from other disciplines such as Heisenberg's Uncertainty Principle, the Postulate of Relativity, the Second Law of Thermodynamics, and Godel's Theorem. [7]

A fine instance of an impotence condition in pattern recognition that may prevent some futile effort is a recent result by Devroye to the effect that it is impossible to guarantee the finite-sample classification performance (i.e., the error rate based on a fixed training set) for any nonparametric discrimination rule. The rate of convergence of the probability of error to the asymptotic probability of error may be slower than that of any prespecified sequence. Thus knowledge of the asymptotic Bayes risk - which can be calculated for nearest-neighbor and k-nearest-neighbor classification in terms of the asymptotic Bayes risk, regardless of the distribution - does not tell us anything about finite-sample performance. Furthermore, putting restrictions on the distribution of the unlabeled samples is insufficient; some information about the class-conditional pattern distributions is required to estimate the finite-sample error rate.[8]

There are, nevertheless, a number of methods for estimating the expected probability of error on new data using only the samples on the training data. Glick compares a number of estimators, including the popular leave-one-out, bootstrap, and posterior probability estimators, with regard to bias (optimistic prediction because the classifier is tuned to the training set), variance (uncertainty of the estimate), robustness (influence of "abnormal" samples), and computational cost. He argues that smoothing the data reduces both the bias and the variance of the estimator and advocates a smoothed modification of the sample success proportion. [9] A still valuable bibliography on the estimation of error probabilities is that of Toussaint. [10]

D. EFFICIENT NEAREST NEIGHBORS CLASSIFICATION

In "difficult" classification problems, the optimal discriminant function between classes is generally neither a hyperplane nor a quadratic surface. An adaptive algorithm for piecewise linear boundaries has been demonstrated long ago by Duda and Fossum, and there have been numerous attempts to cluster the patterns in each class in order to obtain more suitable discriminants. Nearest-neighbor and k-nearest neighbor techniques have agreeable asymptotic properties and have been shown to yield excellent practical results, but require storage and computation proportional to the number of patterns in the training set. [11]

It is clear, however, that the "optimal" discriminant or boundary is influenced most heavily by the patterns close to the boundary; portions of the boundary that are far from any pattern are not critical to classification performance. This is precisely the weakness of clustering techniques, where patterns within each category are grouped without regard to their relations to patterns in other categories.

Toussaint, Bhattacharya and Poulsen summarize previous methods for thinning or editing the training set and propose a method based on the Gabriel graph (a subgraph of the Voronoi graph or Thiessen diagram) for retaining only precisely those patterns that affect the nearest-neighbor decision boundary. A nearest-neighbor classification based on the retained points will thus yield exactly the same error rate on new data as nearest-neighbor classification based on the original training data. The preprocessing required is of order N log N, where N is the number of patterns in the training data. [12]

This result is an example of possible contributions to statistical decision theory from studies in computational geometry, a current and exciting field of research with much to offer to remote sensing and geographic data processing.

E. CLUSTERING

Our understanding of the manifold aspects of clustering or unsupervised classification - surely one of the most fundamental of cognitive operations - continues to grow. More than a dozen books specifically dedicated to clustering have appeared (see, for instance, Hartigan's [13]), and new contributions continue to appear from diverse disciplines such as graph theory, statistics, linear algebra, and computational geometry. A number of global objective criteria have been advanced for comparing algorithms, but cluster validity remains an elusive concept that seems difficult to define except in terms of a spcific application.

F. SIZE OF EXPERIMENTS

The size of experiments, in terms of pixels or acreage, has continued to grow throughout the decade, but the machinery to process entire LANDSAT frames is still available only in very few laboratories. The importance of large experiments lies in the fact that they render it increasingly difficult to hand-tailor the analysis procedure to the peculiarities of the data set. Multi-temporal experiments have also become more practicable with improvements in registration procedures. The verification of the results against independently obtained "ground truth" remains a major difficulty.

III. OTHER DEVELOPMENTS

In contrast to the above, the items mentionned in this section are broadly applicable to various aspects of remote sensing, not just classification.

A. COMPUTATIONAL COMPLEXITY

Computational complexity is not, of course, restricted to pattern recognition or remote sensing, but some notions may be worth summarizing here. A major result of the past decade is the recognition that there is a class of important problems which cannot be solved in a time proportional to the number of data points raised to a fixed power. This class of problems is called NP complete. If a solution based on a deterministic model of computation were found for one member of the class, then a solution would be known for all such problems. Thus the preferred tool for showing that a given problem is NP-complete is to demonstrate that it is computationally equivalent to a known member of the class. Examples of NP-complete problems are the traveling salesman problem and the knapsack problem. [14]

More recently, it has been shown that even if the worst-case situation requires an exponentially large number of operations for solution, for an arbitrary set of data points the correct solution may be obtained in polynomial time except for a very small number of specific cases. The probability of achieving the correct solution for a randomly selected example is thus arbitrarily close to 1.

The notion of a model of computation has been refined to allow accurate comparison between algorithms developed in different notations, but the choice of model does not affect the question of NP-completeness or other asymptotic results. Models of parallel computation are playing an increasingly important role as processes which were once considered to be of an essentially sequential nature, such as bounded searches of game trees, are implemented on networks of processors.

In computational geometry, three kinds of complexity are usually defined: storage complexity, run-time complexity, and preprocessing complexity. In supervised pattern classification, storage complexity would refer to the storage requirements of the classification algorithm rather than of the training algorithm; preprocessing complexity is the number of operations required to design the classifier.

B. SOFTWARE ENGINEERING

Software engineering is the methodical specificationdesign, development, testing and documentation of large programs. It extends over the entire life cycle of a program, from conception to last use. Useful developments include management techniques (programming teams, chief programmer), notational techniques (pseudo-code, HYPO-charts, display layouts), computer aids (the Programmer's Workbench, PSL-PSA), and stylistic programming rules (modular programs, indentations, naming conventions). Programs developed under such discipline are, it is widely agreed, easier to understand, debug, modify, and transport than programs written in a highly individualistic manner concerned mainly with machine efficiency. For the majority of current computer science graduates such techniques are second nature and good programming techniques will eventually permeate the remote sensing community as they already have permeated much of the business world. "New" languages, such as PASCAL and ADA, will help in this endeavour. [15]

C. SMALL IS BEAUTIFUL

In the late sixties some image processing groups shifted to minicomputers, and in the late seventies they started shifting to microcomputers. It is increasingly recognized that intelligent input/output devices, such as displays, plotters, and optical scanners, can substantially reduce the load on the central processor; in fact, the central processor has lost its centrality. Distributed processing systems, where the processing power can be more flexibly balanced against the load, are gradually appearing.

There are, however, more than forty major manufacturers of microprocessors in the United States. Because of the high cost of software development, most processors support only a very limited number of languages, and image input/output operations must still often be programmed in a low level language. Consequently, software portability remains as much of a problem as ever. Furthermore, in mix-and-match academic installations depending largely on student maintenance, reliability is almost impossible to ensure.

D. INTERACTIVE METHODS AND HUMAN FACTORS

At the Human Factors in Computing Conference in Gaithersburg in April 1982, attendance exceeded the expected two-hundred by more than fourfold. The development of experimental techniques and a body of observations on man-machine interaction is occurring simultaneously with increased interest in interactive techniques in the analysis and utilization of remotely sensed data. High-quality image display and image entry devices are still extremely expensive, but perhaps by the time they become generally available we will develop a consensus on how to use them effectively.

E. IMAGE DATABASES

The last five years have seen a number of attempts to extend database techniques to images. In some projects the images are stored as integral entities that can be inspected or retrieved but not modified. Indeed, tools for image modification are still essentially non-existent, except for overlays. Nor has there been much software developed to combine images and geographical information from other sources; it is, at best, a laborious effort. [16]

It is this author's conjecture that real progress will come only when commercial database systems can be expanded to two-dimensional applications. The overhead costs of developing communications protocols, back-up and archiving facilities, privacy and security measures, query languages, I/O interfaces, programming language interfaces, adequate maintenance and user documentation, and so on, are just too high in proportion to the restricted volume of geographic applications. Furthermore, even if the bulk of the storage available must be reserved for the image components, for most applications it will be essential to provide appropriate facilities for textual and formatted alphanumeric data.

At this time most installations can keep only a very limited amount of image data "on line." Video disk technology may eventually provide a more economical storage medium than magnetic disk.

F. NEW COMPUTER ARCHITECTURES

The race between general purpose sequential computers and special purpose parallel architectures for image processing has been on for more than twenty years. Because of the cost of manufacturing processors using discrete components, the early machines exhibited only a relatively low degree of parallelism. Current VLSI designs may process thousands of pixels simultaneously. The recent resurgence of experimental parallel systems for image analysis is cogently reviewed by Danielson and Levialdi.[17] In spite of the abundance of ideas and even of commercially available processors, most installations still depend largely on old-fashioned uniprocessors whose cost per operation continues to decrease dramatically. Although image-processing operations have been embedded in extensions of standard high-level programming languages, lagging software development and lack of uniform interfaces to standard processors appear to continue to retard the wide-scale acceptance of specialized systems.

G. IMPROVED SOURCES OF DATA

???

IV. RESEARCH GOALS IN MATHEMATICAL PATTERN RECOGNITION

Seeking to delineate the most promising areas for research in pattern recognition techniques applicable to the management of renewable natural resources, a group of ten scientists with extensive experience in remote sensing joined forces under the leadership of Professor L.F. Guseman, Jr. of Texas A&M University. Among the disciplines represented by the participants were mathematics, statistics, photogrammetry, electrical engineering, physics, and geography. The group functioned within the framework of a comprehensive study commissioned by R.B. MacDonald, Chief Scientist for Earth Resources Programs, NASA Johnson Space Center, which included other similarly constituted groups charged with making recommendations for "basic" research on Scene Radiation and Atmospheric Effects, Electromagnetic Radiation and Data Handling, and Information Utilization and Evaluation.

The Working Group on Mathematical Pattern Recognition and Image Analysis met eight times altogether, for two to three days at a time. In addition to drawing on the Working Group's own experience, crystallized through lengthy soul-searching discussions, three special workshops were organized on the topics of Preprocessing, Image Modeling, and Classification Techniques. A dozen specialists in the area covered by each workshop were invited from universities, other non-profit research institutions, and industry, to make overview presentations and guide discussion.

A. REGISTRATION AND RECTIFICATION

Because of the widely divergent background of the participants, it took several meetings to reach an acceptable level of communication. The first major topic of discussion, image registration and rectification, was fortunately the best defined, and a consensus developed that there indeed was a much neglected area. Topics singled out for further attention included the registration (overlay) of digital image arrays obtained at different times and possibly from different sensors (satellite MSS and RBV, radar, airphoto); the precise earth-location of images including consideration of topographic effects; the definition and automatic extraction of control points; the exploitation of USGS-prepared digital terrain models; the development of meaningful measures of accuracy for both registration and rectification; the relation between, scale, orientation, photometric quantization, spatial sampling and coordinate system; and the need for increased understanding of the resampling process.

B. IMAGE REPRESENTATION

The concept of digital image representation took considerably more effort to define. It was eventually agreed, however, that the core of the matter was the relation of observed features to the class of objects or items defined by an application dependent taxonomy. Topics selected for further investigation included texture information from multiple (multispectral, multitemporal, multisource) images; the relation between spatial (shape, texture, topology) and spectral features; image segmentation techniques; the role and appropriate representation of ancillary (non-image) information; syntactic ("structural") techniques, spatial context and temporal context; and the integration of non-image information (atmospheric, illumination, and sensor correction) into the generation and definition of primitives in application-independent scene models. These topics are often included under the heading of image restoration. Also sought is the extension of spectral dimensionality-reduction techniques to spatial and temporal dependencies.

C. CLASSIFICATION

The incorporation of previously defined digital image representations into systematic methods of determining the required attributes of object scenes is the subject of classification. The primary objectives of the classification process are considered to be mapping, inventory, and monitoring of natural resources. Mapping shows the location of classes, objects, items, or other types of interest; it includes both hardcopy and display. Inventory is concerned with counting, aggregation, census, or planimetry of items without specific retention of spatial coordinate information. Monitoring refers to change detection, discovery of unusual conditions, and other operations of limited spatial and temporal scope.

Classification includes concepts such as categorization, identification, recognition, clustering, partitioning, taxonomy, and segmentation. Of concern are supervised and unsuperivsed learning, teaching, or training; estimation of parameters, distributions, and error rates; the assignment of identities, labels, or symbols by either automatic or interactive means; and the general evaluation of the accuracy, dependability, and robustness of the entire process. Of particular interest is the role of the human and the contributions extracted from ancillary data. Techniques based on statistical as well as structural, syntactic, relational, and other deterministic approaches are germane. Algorithms need to be developed for multisource data, including multisensor observations, multitemporal observations, and combinations of multi-image data with non-image data.

When classification is not performed in a single step, the intermediate variables are often called features, components, signatures, dimensions, transformations, factors, primitives, characteristics, or measurements. Methods must be developed for obtaining these intermediate variables from the digital images and ancillary data and for incorporating them into the classification process.

The information sought by the end-user may be in a form different from the simple, non-overlapping mutually exclusive and totally exhaustive model provided by standard pattern-recognition texts. An example of the fuzzy taxonomies of possible interest is the class of grizzly-bear habitats, for which a complete specification may not even conceptually exist.

Considerable attention was devoted to mathematical techniques of proportion estimation. Among approaches worth pursuing are enumeration through classification, stratified area estimators, regression estimators, and direct estimators. Further progress is dependent on the development of algorithms which require only a small number of training samples, can deal with a large number of object classes, are responsive to non-stationary distributions, and can account for "mixed-pixel" measurements resulting from finite sensor resolution.

D. OTHER TOOLS

In contradistinction to map displays or statistical inventory information which forms the final product of the classification process and benefits the "end-user", data displays are intermediate products intended to improve the classification process itself. Specifically, they provide the opportunity for human interaction. The scope of the displays ranges from simple histograms, which allow judgment of the overlap between statistical distributions, to digital images providing photointerpreters a means of assigning labels to representative samples.

Data structures, data compression techniques, and special parallel computer architectures are of interest to the extent that they impact the classification process. Among data structures to be investigated are pixel-by-pixel storage, bit-plane structures, vector (polygon) methods, chain encoding, contour coding, hierarchical pyramid and quad-tree structures, and various two-dimensional polynomial approximations. Architectures to be investigated (though at a relatively low priority) are pipeline, multiple-instruction single-data-stream, and multiple-instruction-multiple-data-stream machines. Given the relative economies of special-purpose VLSI chip development and the rapidly decreasing cost of general purpose processors (particularly bit-slice architectures), it is expected that networks of processors, possibly with common memory access, will predominate. The development of adequate operating systems for these configurations is, however, a monumental task that cannot be borne by the remote-sensing community alone. In the expectation that parallel machines will be available in the next decade, however, increased attention to models of computation not based on the single-instruction, single-data-stream model is recommended.

E. RESEARCH ORGANIZATION

The Working Group discussed the possibility that forthcoming research endeavours be administered by universities or other research-oriented organizations rather than directly by the National Aeronautics and Space Administration. Specially designated Research Centers could prepare research plans, solicit and evaluate proposals, and hold semi-annual or annual symposia on each major topic.

This summary represents the perceptions of one of the participants of the Working Group on Mathematical Pattern Recognition and may differ in substance and emphasis from the final report of the Group. The official report of the Task Force (a document of about 200 pages) is available from the NASA Earth Observation Directorate, Johnson Space Center, Houston, Texas.

V. ACKNOWLEDGMENT

The support of the University of Nebraska Conservation and Survey Division through NASA Grant No. 28-004-020 is gratefully acknowledged.

VI. REFERENCES

1. Haralick, Robert and L. Shapiro "The Consistent Labeling Problem I" IEEE Trans. Pattern Analysis and Machine Intelligence 1, #2, March 1979, pp. 173 184.

2. Haralick, Robert and L. Shapiro "The Consistent Labeling Problem II" IEEE Trans. Pattern Analysis and Machine Intelligence 2, #3, May 1980, pp. 263-313.

3. Levine, Martin and S. Shaheen "A Modular Computer Vision System for Picture Segmentation and Interpretation" IEEE Trans. Pattern Analysis and Machine Intelligence 3, #5, September 1981, pp. 540-556.

4. Wheeler, S. and Misra, P.N. "Crop Classification with LANDSAT Multispectral Scanner Data II" Pattern Recognition 12, #4, 1980, pp. 219-228.

5. Landgrebe, David "The Development of a Spectral Spatial Classifier for Earth Observational Data" Pattern Recognition 12, #3, 1980, pp. 165-176.

6. Goldfarb, Lev A New Approach to Data Analysis, Ph.D. Dissertation, University of Waterloo, Department of Systems Design, 1980 (to be

published as a monograph); see also Wong, A.K.C., and L. Goldfarb, "Pattern Recognition of Relational Structures," in Pattern Recognition Theory and Applications (J. Kittler, K.S. Fu, and L.F. Pau, eds.) D. Reidel Publishing Company, Dordrecht, Holland, 1981, pp. 157-176.

7. Fein, Louis "Impotence Principles for Machine Intelligence" in Pattern Recognition (L. Kanal, ed.), Thompson Book Company, Washington, 1968, pp.443-447.

8. Devroye, Luc "Any Discrimination Rule Can Have an Arbitrarily Bad Probability of Error" to appear in IEEE Trans. Pattern Analysis and Machine Intelligence.

9. Glick, Ned "Additive Estimators for Probabilities of Correct Classification" Pattern Recognition 10, 1978, pp. 211-222.

10. Toussaint, Godfried "Bibliography on Estimation of Misclassification," IEEE Trans. Information Theory 20, 1974, pp. 472-479.

11. Cover, Thomas and P.E. Hart "Nearest Neighbor Pattern Classification," IEEE Trans. Inf'n Theory 13, January 1967, pp.21-27.

12. Toussaint, Godfried "Pattern Recognition and Geometrical Complexity," Proc. Fifth Int'l Conf. on Pattern Recognition, Miami Beach, December 1980, pp. 1324-1346.

13. Hartigan, John Clustering Algorithms, John Wiley and Sons, New York, 1975.

14. Horowitz, E. and S. Sahni Fundamentals of Computer Algorithms, Computer Science Press, 1978.

15. Wasserman, Anthony and S. Gutz "The Future of Programming," Comm. ACM 25 , #3, March 1982, pp. 196-206.

16. Zobrist, Al and G. Nagy "Pictorial Information Processing of Landsat Data", IEEE Computer 14, #11, November 1981, pp. 34-42.

17. Danielson, Per-Erik and S. Levialdi "Computer Architectures for Pictorial Information Systems," IEEE Computer 14, #1, November 1981, pp. 53-67.

18. Nagy, George and S. Wagle "Geographic Data Processing" ACM Computing Surveys 11, #2, June 1979, pp. 139-181.

George NAGY received the B.Eng. (Physics) and M.Eng. (EE) degrees from McGill University and the Ph.D. degree in Electrical Engineering from Cornell University (1962). He subsequently served ten years on the staff of the IBM T.J. Watson Research Center, developing pattern classification techniques for optical character recognition, speech processing, data compression, and remote sensing. He was appointed chairman of the Computer Science Department of the University of Nebraska in 1972 and held this position until 1981. He has spent leaves ranging from two months to a year at Cornell University, Universite de Montreal, IBM, Bell Laboratories, the Italian Research Council's laboratories at Genoa and Naples, and INRS Telecommunications - Bell-Northern Research in Montreal. He has served as a research consultant also for Tektronix, Compression Laboratories, Caere Corporation, and NASA. He has given lectures at over one hundred universities and technical conferences in the United States and abroad, and is the author of numerous research and survey articles. Currently a Professor at the University of Nebraska-Lincoln, he teaches courses on computer organization, pattern recognition, and remote sensing. His research areas are geographic data processing, digital image registration, and quantitative evaluation of the man-computer interface.