NASA-CR-194878

Report Documentation Page

IN-61-CR

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No. *203557*
4. Title and Subtitle Image Pattern Recognition Supporting Interactive Analysis and Graphical Visualization		**5. Report Date** *31 P* January 1992
		6. Performing Organization Code
7. Author(s) James M. Coggins		**8. Performing Organization Report No.**
		10. Work Unit No. 506-59-11
9. Performing Organization Name and Address Computer Science Department Sitterson Hall University of North Carolina Chapel Hill, NC 27599-3175		**11. Contract or Grant No.** NAS5-30428
		13. Type of Report and Period Covered
12. Sponsoring Agency Name and Address CESDIS (Center of Excellence in Space Data and Information Sciences) Code 930.5 Goddard Space Flight Center/Greenbelt, MD 20771		Final
		14. Sponsoring Agency Code

15. Supplementary Notes

CESDIS (The Center of Excellence in Space Data and Information Sciences) is operated by Universities Space Research Association (USRA), American City Building, Suite 212, 10227 Wincopin Circle, Columbia, MD 21044.

16. Abstract

Image Pattern Recognition attempts to infer properties of the world from image data. Such capabilities are crucial for making measurements from satellite or telescope images related to Earth and space science problems. Such measurements can be the required product itself, or the measurements can be used as input to a computer graphics systems for visualization purposes. At present, the field of image pattern recognition lacks a unified scientific structure for developing and evaluating image pattern recognition applications. The overall goal of this project is to begin developing such a structure. This report summarizes results of a 3-year research effort in image pattern recognition addressing the following three principal aims: 1) to create a software foundation for the research and identify image pattern recognition problems in Earth and space science; 2) to develop image measurement operations based on Artificial Visual Systems; and, 3) to develop multiscale image descriptions for use in interactive image analysis.

17. Key Words (Suggested by Author(s)) image processsing; image pattern recognition; computer vision; multiscale; object-oriented programming; C++ class library; artificial visual systems; Earth and space science applications		18. Distribution Statement Unclassified - Unlimited	
19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of pages 30	22. Price

CENTER OF EXCELLENCE IN SPACE DATA AND INFORMATION SCIENCES

GSFCMail: *CESDIS*; Internet: *CESDIS@amarna.gsfc.nasa.gov*

PREFACE

The attached document represents the final report for Task 5 on
Universities Space Research Association (USRA) contract NAS5-30428. Although
a preface is not normally included in such a document, it is necessary in this
case to provide some background information to offset the potentially negative
impression which may result from some of the author's statements.

When the contract project CESDIS was put in place in July 1988, USRA
executed subcontracts with four universities to fund research projects which
had been selected from 86 proposals through a thorough peer review process.
These four projects were each funded for a three-year period with the
possibility for a two-year extension with mutual consent by both parties.
Additionally, a subcontract with the University of Maryland's Department of
Computer Science was executed to provide partial support for the CESDIS
Director and, in 1989, two assistant professors who were to serve as staff
scientists. CESDIS core operations, therefore, supported a small
administrative staff, the research activities through the University of
Maryland, and the four university peer-reviewed projects.

As time progressed, it became evident that NASA core funding of CESDIS
would be less than original estimates. Thus, in the first two years the
administrative staff and the University of Maryland operations were held below
planned levels to conserve funds for the subcontracted research activities.
In the second year of operation it was seen that this funding shortfall was of
such an extent that budget cuts in the four peer-reviewed projects would also
be required in the third year. The project investigators were notified of
this and conservative spending encouraged until complete details of impending
cuts were known.

Reductions were made on an equitable basis across all four projects in an
effort to control the negative impact on the level of effort that each project
would experience as much as possible. Even so, some projects could not absorb

the cuts as well as others as noted, for example, in the attached report by the project investigator at the University of North Carolina.

In the third year, each project investigator was asked to submit an informal plan of work for the fourth and fifth years to be reviewed for the possible extension of funding in the optional 2-year period. These plans, as well as the perceived accomplishments during the first three years, were then reviewed to determine which projects would be extended. The final decision was to continue the project at Duke University as well as the work at Stanford University. The other two project investigators (George Washington University and University of North Carolina) were notified that they would not be funded for the additional period.

Unfortunately, the total research activities for CESDIS were thus reduced. It is hoped that CESDIS can recover from this loss of research activity and again put out a call for proposals as additional research funding from NASA becomes available.

Raymond E. Miller
CESDIS Director

April 22, 1992

Image Pattern Recognition Supporting Interactive Analysis and Graphical Visualization

Final Report
February 26, 1992

James M. Coggins
Computer Science Department
University of North Carolina at Chapel Hill

Table of Contents

1. Introduction 2
2. Project Objectives 2
3. Background 4
    3.1 Requirements for Image Segmentation 5
    3.2 Multiscale Image Representations 6
    3.3 Multiscale Decompositions for Inference 6
    3.4 Software 10
    3.5 Literature Cited 10
4. Accomplishments of This Project 14
    4.1 Foundations 14
    4.2 Image Analysis Methods 18
    4.3 Interactive Tools 20
    4.4 Other Accomplishments 21
5. Limitations of This Project 24
    5.1 Implementation of Budget Cuts 24
    5.2 Effect on Year 3 Project Accomplishments 25
    5.3 Effect on Contract Extension 26
6. A Review of our CESDIS Experience 26
    6.1 Characteristics of NASA Image Analysis Problems 27
    6.2 CESDIS: The Vision and the Reality 28
    6.3 Conclusion 28

Image Pattern Recognition Supporting Interactive Analysis and Graphical Visualization

Final Report
February 26, 1992

James M. Coggins
Computer Science Department
University of North Carolina at Chapel Hill

1. Introduction

This is the final report of a three-year research contract funded through the Center of Excellence in Space Data and Information Science. This report summarizes the goals, accomplishments, and limitations of our research during this contract.

This research has been carried out at both the University of North Carolina at Chapel Hill and at the Goddard Space Flight Center. At UNC, participants in the contract have been from the Departments of Computer Science, Physics and Astronomy, and Geography. At Goddard, contacts were made with over thirty NASA scientists who have problems related to the topic area of this research. Several of our contacts spent time discussing their research with us and helped us to define specific image analysis problems in earth and space science that will drive our research for several years to come. Some have provided data to be explored and used to test algorithms for various tasks.

Section 2 will describe the objectives of this project as set forth in the original proposal. Section 3 will supply some background and motivation for this effort provided in the original proposal and explain the approach we have pursued. Section 4 details our accomplishments in each of the areas listed as principal objectives in Section 2. Section 5 discusses limitations of the research output of this project. Section 6 then provides evaluative comments from our project's experience.

2. Project Objectives

Computer vision algorithms construct inferences about the world from image data. The inferences produced by computer vision algorithms include measurements, classifications, descriptions, models, labels, and so on. The overall goal of this research is to develop a *unified, scientific framework for developing and evaluating computer vision algorithms.*

The goal of this research implies that the current state of the art in computer vision does not include a scientific framework for developing such applications, and that such a framework would be a desirable advance for the field. Instead of building a well-founded scientific application framework, research in computer vision has accumulated anecdotal evidence for an array of *ad hoc* methods without effective evaluation, comparison, or often understanding of those methods. Techniques are continually developed and tentatively evaluated on limited application data, but no methodology for comparing rival techniques exists other than comparison of empirical results. The ability to create inference mechanisms within a framework that permits one to understand and evaluate the inference process would dramatically increase the utility of such tools.

The approach we have investigated involves a kind of multiscale image decomposition implemented in an Artificial Visual System (AVS). The AVS structure was guided by models of the early human visual system, the structure of successful algorithms from statistical pattern recognition, and an analysis of the requirements of a broad set of computer vision problems. Inferences are derived from such systems according to a few basic principles related to the foundations of the science of spectroscopy.

Since the models and features that are of interest to a human analyst change too rapidly and with too much subtlety to allow encapsulation in an Artificial Intelligence system, we propose that the most productive approach for dealing with real problems is to develop powerful tools for supporting human analysts. Our aim is to develop human interfaces and pattern recognition tools that allow the user to perform impromptu visualization and analysis in which the criteria and objectives of the interaction are continuously redefined in the course of the interaction.

The specific aims of this project were as follows:

1. To create a foundation for development, evaluation, and testing of the methods to be developed in this research program.
    A. To identify specific image analysis applications in the earth and space sciences to use as driving problems for the research program.
    B. To design and implement a software base for the experimental research using modern software engineering practices.

2. To develop image measurement operations based on artificial visual systems.
    A. To develop new artificial visual systems for use in measurement tasks.
    B. To investigate the sensitivity of such measurement techniques to noise and distortion.

3.To develop multiscale image descriptions for use in interactive image analysis.
    A. To develop automatic methods for creating image descriptions in terms of regions.

B. To develop methods for user-guided interaction based on the automatically identified regions using interactive computer graphics techniques.

The review of the accomplishments of this project in Section 4 will be organized according to this list of specific aims.

This program was proposed as a five-year research effort, funded for three years with optional extension for two additional years. Funding cuts severely limited the project in the third year, and the two-year extension was not funded. Nevertheless, we have made substantial progress on all of the proposed specific objectives and as the project ended were, in some areas, ahead of the schedule set for the five-year plan. This substantial progress toward challenging goals was made possible by the hard work of several graduate students and the funding support provided for the first two years of this contract. Progress in the third contract year was limited, but we were able to consolidate results of the first two years and make some advances in specific techniques.

3. Background

The problems of graphic visualization of scientific data have attracted national attention [McCormick 1987]. This attention has focused on computer graphics tools for interactive, impromptu visualization, all but ignoring the mathematical and statistical tools for exploratory data analysis used in the field of pattern recognition [Jain and Dubes, 1988]. The pattern recognition methods can help build alternate representations of data that maintain fidelity while reducing the data's complexity (*viz.* dimensionality). Reconstructions based on pattern recognition methods are often difficult to envision geometrically or graphically, but sound mathematical and statistical principles guide them.

Algorithms are available for exploratory analysis of data represented as points in multidimensional spaces, one-dimensional functions, strings, and graphs. This project concerns the development of data reduction, measurement, representation, and visualization tools for images.

Because images are acquired using multiple modalities under varying conditions over long spans of time, investigators face a serious problem of information overload: accurate but irrelevant data, relevant but ambiguous data, or simply too much data. Researchers and analysts need software tools for warping, registering, measuring, recognizing, matching, detecting, enhancing, editing, and, most important, for simplifying image data. Abstractions of image sequences based on automatically computed image representations can help select and manipulate items of interest, providing a more natural basis for an analyst's interaction with the image data base.

Detecting, segmenting, and measuring objects in an image and then comparing those objects with known normal and abnormal patterns is the essence of image analysis. Software tools can contribute to this task: a human analyst can use

them to enhance the contrast of the image data, specify a region of interest, label regions or groups of regions as objects, and measure the objects based on spatial, radiometric, temporal, or spectral criteria. The following subsections describe the development of multiresolution methods for defining, describing, and measuring objects.

3.1 Requirements for Image Segmentation

An early attempt at segmenting an image into regions classifies each pixel according to the object class that is most likely to be associated with the pixel's intensity. We refer to this method as Bayesian Pixel Classification (BPC) [see Duda and Hart, 1973, Chapter 2]. BPC is computationally simple since each pixel is processed independently of its surroundings. When the object classes are associated with unique ranges of intensity (as are bone regions in computed tomography images), BPC works. In most cases, however, BPC fails because the *a priori* class-conditional probability densities overlap, yielding an unacceptably high probability of error. Intuitively, we may say that BPC fails when a given intensity is likely to appear in the images of many different objects. In these cases, it is necessary to use information from spatial relationships among pixels to define image segments.

Variations on BPC add heuristic constraints such as requiring all regions to have a minimum size or requiring each region with a particular label to have particular intensity statistics. Such heuristics also fail because they do not address the four essential issues underlying object definition: spatial structure, context, constraints, and invariants.

We identify objects by the **spatial structure** of the intensities of the pixels that compose them. This structure is not captured by statistics of intensity distributions.

The **context** of a pixel determines its interpretation. The intensity at a single pixel is of little importance in most applications due to noise processes, systematic (e.g. optical) distortions, and spatial variations in brightness due to differences in illumination or system gain [Fay, *et al.*, 1985]. Furthermore, segmentation by "edges" or region growing is inadequate because such methods capture only a limited scope of a pixel's spatial context. Both proximal and distal contexts jointly determine a pixel's meaning.

Constraints from the problem domain often provide powerful cues to guide the interpretation of an image [Zucker 1981]. Such constraints usually involve restrictions on plausible interpretations based on relationships among (problem-domain specific) objects, not context-free pixels. Constraints concerning containment and adjacency of objects are particularly common and effective. Paradoxically, interpretation of large, complex images can be easier than interpretation of small, isolated regions because more constraints can be applied. This drastically reduces the number of consistent interpretations.

Images of a particular object acquired at different times or under different conditions are usually different. Measuring these differences is easy but irrelevant. A key to making useful inferences is to characterize the **invariants** of the ensemble of images of the object. This requires a specification of allowable variation. For example, one might want to recognize an object in spite of translations from the ideal position of less than 1/4 of its width, or scale changes of less than one octave from ideal, or rotations of less than 20 degrees from the ideal. A good object representation should permit specification of permitted variation as well as capturing the invariants of the object's appearance. Differential invariants are functions whose values do not change when the image undergoes transformations such as translation, scaling, and rotation. Such functions can be derived from multiscale image decompositions.

3.2 Multiscale Image Representations

We have investigated multiscale, structure-sensitive image descriptions that take into account spatial structure, context, constraints, and invariants. This approach is based on the key insight that the information content of images exists at multiple scales, so image descriptions must be based on data structures that embody multi-scale information. Such methods can form the basis for automatic object recognition, user-directed interactive editing, and efficient image descriptions based on inferences derived from a multiscale decomposition.

Studies of human visual psychophysics suggest that the human visual system obtains information about a range of context scales in a direct way: by acquiring information at multiple scales. Several ``multiple channel" or ``multiscale" models of information processing in the early human visual system have appeared [Robson 1983; Wilson and Bergen 1979; Ginsburg 1978, 1980; Koenderink and van Doorn 1978; Koenderink 1986, 1987]. These models have been supported by neurophysiological studies that are consistent with multi-channel filtering hypotheses [Pollen 1983].

The multiscale models of human vision sparked efforts to apply the same principles to computer vision systems [Marr 1980; Rosenfeld 1984; Burt *et al*. 1981]. These approaches used various methods for extracting the levels of resolution or scale, including frequency-domain bandpass filters, the results of unweighted pixel averaging, and the results of Gaussian blurring and differences thereof. Since the implementation of the multiresolution concept can be made more efficient by decreasing the spatial sampling as the spatial resolution decreases [Crowley 1984], the resulting computer vision systems have been termed "pyramid" approaches.

3.3 Multiscale Decompositions for Inference

Multiscale decompositions can be used as the basis of inferences about the objects in the image. Three kinds of inferences will be described in this section, texture, measurement, and shape.

Texture

The problem of characterizing visual texture consumed almost two decades before the solution fell out of a simple application of the multiple-channel model of early human vision. Coggins [1982; Coggins and Jain 1985; Ginsburg and Coggins 1981] implemented the multiple-channel model of Ginsburg [1978, 1980] and found a simple feature space in which distance correlated closely with human texture perception. The feature space was interpretable as "average local energy in a sequence of channels" and was the first definition of texture independent of human visual performance. This result was extended later by using the same feature space to measure the orientation of tilted, textured planes [Coggins and Jain 1986].

These results are consistent with the psychophysical studies of Harvey and Gervais [1981] and Richards and Polit [1974] and an *ad hoc* spatial-domain representation for texture developed by Laws [1980].

Measurement

The studies of Coggins and Fay (a physiologist) [Coggins *et al.* 1985, 1986a, 1986b; Fay *et al.* 1984, 1985] measured the locations, lengths, and orientations of protein bodies in a single cell. Three-dimensional fluorescence images were acquired with a computer-controlled microscope. Three separate filter decompositions mapped the images of the protein bodies into an abstract feature space in which the axes of the space were filter outputs. Measurements were computed in the feature space by interpolation of filter output intensities and by matching energy outputs of the filters to known patterns. The measurements computed were the locations (x,y,z), the length (in voxel units) and the 3D orientation (phi and theta angles) of the bodies. Serious image distortions due to the high gain and high magnification in the optical system were corrected by scaling the feature space to compensate for the distortions. The measurements of position, size, and orientation were made directly in the 3-D data, and the performance of the artificial visual system far exceeded the capabilities of human observers who had been working on the same data for months.

The measurements produced by the artificial visual systems were used as input to a graphics system that allowed interactive manipulation of the graphical model so that the investigators could explore the 3-D protein structure in a single cell. The artificial visual system's measurements simplified the data so that interactive computer graphics tools could give the researchers new insights and new viewpoints that are inaccessible in the real cells but available in the virtual world of the graphics model.

Shape

Attempts to characterize shape from multiresolution image representations have been based on graph representations of image features, followed by graph matching between the observed and prototype images. Marr and Hildreth [1980] use an edge detector and heuristic combination rules to characterize the image,

but the resulting representation in terms of blobs and edge segments is bulky and error-sensitive. Crowley [1984] proposes an object definition scheme that uses intensity extrema and ridge lines along with heuristic combination rules. This approach is promising, but it has been criticized as being too elaborate and too dependent on *ad hoc* heuristics. Coggins and Poole [1986] created a simpler graph structure based on directional intensity extrema in scale-filtered images for use in character recognition. Their study suggested that powerful inferences could be obtained from a single channel if one could identify the channel containing the most relevant structural information. Koenderink [1984] suggests that a reasonable segmentation can be obtained from a hierarchical nesting of regions. The hierarchy is defined by following intensity extrema through continuously decreasing resolution to annihilation. Pizer *et al.* [1987] extends this idea to following any ``essential structure" such as symmetric axes through blurring to annihilation. The image description hierarchy provides an ordering of image components by ``structural significance" which involves both size and contrast of the image regions.

Multiscale Communication for Display

The bandwidth problems faced in image communication can be alleviated by implementing quick and deep analysis systems and by using a successive refinement strategy for transmitting image data in real time.

Successive Refinement

Successive refinement is a strategy developed for interactive computer graphics by which a display is created in stages [Bergman *et al.* 1986]. An approximation to the image is displayed first, then if time is available, the display is refined to portray more detail or to incorporate enhancements such as antialiasing. If a user requests a different view before the refinement is complete, the process restarts immediately with a rough approximation of the new view. This technique allows a relatively slow display device to respond quickly to interactive user commands, providing the user with feedback on the status of the graphical model under real-time control. In image communication, a successive refinement strategy transmits the multiresolution description in order of increasing resolution. A processor in an image workstation constructs an image from brief descriptions transmitted over the communication channel. Transmission continues until interrupted by a control signal. Successive refinement minimizes the bandwidth required to transmit image data since the transmission can be cut off as soon as sufficient data to support a decision is acquired.

Interactive Object Definition

In his doctoral dissertation research, Lawrence Lifshitz [Pizer *et al.* 1986a; Lifshitz 1987] investigated the behavior of intensity extrema (local intensity maxima and minima) under continuous gaussian blurring. As the degree of blurring increases, an extremum moves through the image, and its intensity converges toward the mean intensity of the surrounding image region until the extremum annihilates with a saddle point. Each non-extremum pixel is

associated with the extremum to which the pixel's iso-intensity curve merges. This iso-intensity, or level, curve is formed by all connected locations having the same intensity, with interpolations when necessary. The set of all pixels associated with a particular extremum constitute the "extremal region". This "image component" can be labeled by the degree of blurring required to achieve annihilation, the intensity at annihilation, the location at annihilation, and whether the extremum was a maximum or a minimum.

The iso-intensity curve of the extremum at annihilation can be followed through further blurring to determine its containing component: the extremal region of the extremum to which it merges. Lifshitz has shown that the components appearing relatively high in the image structure tree often correspond to semantically meaningful objects, or if not, a semantically meaningful object will consist of only a few such components. Based on these observations, Lifshitz has designed an interactive system that presents the image structure tree in a form that displays the position of annihilation, degree of blurring required for annihilation, and the containment relation. The user can quickly specify image objects by selecting image components [nodes] in the image structure tree. Lifshitz' demonstration system operates on a 2D or 3D images, but the methodology has not been fully tested in 2D nor seriously tested in 3D.

Lifshitz found that the image description based on multiresolution intensity extrema would segment objects into unnatural components in some cases. Apparently, the structure of objects in the image is not adequately characterized by the annihilation paths of intensity extrema alone. Students John Gauch and William Oliver have developed image descriptions based more on the shape of image components and in particular on the symmetric (or medial) axis [Pizer *et al.* 1986b; Pizer *et al.* 1987; Gauch *et al.* 1987]. In this approach, level curves at each intensity are represented by the symmetric axis of the curve. The collection of symmetric axes from all level curves, when stacked one atop the other, forms a Symmetric Axis Pile (SAP) in which the branching structures form sheets that represent the symmetric axis skeletons of regions making up objects in the image. The SAP can be computed after the image has been preprocessed by a contrast enhancement operator to increase the likelihood that an object is surrounded by a single level curve, but this step is not required for the method to work.

The end of every branch of a SAP can be associated with a curve of vertices (maximum curvature points) of intensity level curves. Unlike the SAP branches, each of these components is a simple curve on the (2D or 3D) image. These vertex curves can be followed through blurring, while using the SAP at only the original resolution to define the hierarchy. When a vertex curve annihilates with blurring, one can determined if the version of that curve at original resolution is associated with a SAP branch, and the connection of that SAP branch can be used to define the containing SAP branch and thus the vertex curve that is the ``parent" of the vertex curve under consideration. Image regions, now in grey scale, are defined by the SAP branches and their radius functions. Information about the shape of the image component can be obtained from the curvature and length of the

symmetric axis and from the magnitude and curvature of the radius function of the symmetric axes.

This approach provides image descriptions that are based on a hierarchy of regions. Such a description provides a natural language for human observers and a rich data structure for computer manipulation. Interactive methods for manipulating these descriptions have been developed in this project and related research at UNC.

3.4 Software

The software environment for supporting this research began before the CESDIS Call for Proposals was received [Coggins 1987a, 1987b]. Dr. Coggins' research software library has been developed using the ``Object-Oriented" design methodology which has appeared over the last several years [Cox, 1986; Peterson, 1987]. In particular, the environment is written in the C++ programming language [Stroustrup, 1986].

Much of the research library has now been used to support several offerings of graduate courses in computer graphics and image processing. The library has grown to include complex pattern recognition programs that we have converted from FORTRAN and integrated into the library. Portions of the library have been rewritten as the hardware environment in our laboratory has changed. The user interface classes have been rewritten several times as our laboratory changed over to the X Windows System and as further experience with C++ revealed better, easier ways to accomplish our design objectives in the user interface code.

Our relationship with our Department's SoftLab Software Systems Laboratory has developed into a real collaboration. SoftLab maintains and distributes a version of the library to outside inquirers and has benefitted from procedures developed for maintenance of this research library and by using this library as a test system for their own procedures.

3.5 Literature Cited

L. Bergman, H. Fuchs, E. Grant and S. Spach, "Image Rendering by Adaptive Refinement," *Proceedings of ACM SIGGRAPH 1986*, Dallas, August 18--22, 1986, 29--37.

P. Burt, T. Hong and A. Rosenfeld, "Segmentation and Estimation of Image Region Properties through Cooperative Hierarchical Computation," *IEEE Transactions on Systems, Man, and Cybernetics*, 11(12), 802--809, 1981.

J. M. Coggins, "Integrated Class Structures for Computer Graphics and Image Pattern Recognition," *Proceedings of the First USENIX C++ Workshop*, Santa Fe, NM, Nov. 8--10, 1987a, pp. 240--245.

J. M. Coggins, "Object-Oriented Programming in C++," *Proceedings of the 1987 IEEE Atlanta Software Technology Conference*, Atlanta, Nov. 11--13, 1987b.

J. M. Coggins and J. T. Poole, "Printed Character Recognition Using an Artificial Visual System," *Proceedings of the 1986 IEEE International Conference on Systems, Man, and Cybernetics*, Oct. 14--17, 1986, Atlanta, GA, pp. 1612--1616.

J. M. Coggins and A. K. Jain, "Surface Orientation from Texture," *Proceedings of the 1986 IEEE International Conference on Systems, Man, and Cybernetics*, Oct. 14--17, 1986, Atlanta, GA, pp. 1617--1620.

J. M. Coggins, K. E. Fogarty, and F. S. Fay, "Interfacing Image Processing and Computer Graphics Systems Using an Artificial Visual System," *Proceedings of Graphics Interface -- Vision Interface '86*, May 26--30, 1986a, Vancouver, B.C. Canada, pp. 229--234.

J. M. Coggins, K. E. Fogarty and F. S. Fay, "Development and Application of a Three-Dimensional Artificial Visual System," *Journal of Computer Methods and Programs in Biomedicine*, special issue of papers from the Ninth SCAMC, vol.22, 1986b, pp. 69--77.

J. M. Coggins, K. E. Fogarty, and F. S. Fay, "Development and Application of a Three-Dimensional Artificial Visual System," *Proceedings of the Ninth Symposium on Computer Applications in Medical Care (SCAMC)*, November 10--13, 1985, Baltimore, MD, pp. 686--690.

J. M. Coggins and A. K. Jain, "A Filtering Approach to Texture Analysis," *Pattern Recognition Letters*, v. 3, 1985, pp. 195--203.

J. M. Coggins, *A Framework for Texture Analysis Based on Spatial Filtering*, dissertation for Ph.D., Michigan State University, 168 pages, published by University Microfilms, Ann Arbor, 1982.

B. Cox, *Object-Oriented Programming: An Evolutionary Approach*, Addison Wesley, 1986.

J. Crowley and A. Sanderson, "Multiple Resolution Representation and Probabilistic Matching of 2D Gray-Scale Shape." Technical Report CMU-RI-TR-85-2, Carnegie-Mellon University, 1984. Also see J. Crowley, A. Parker, "A Representation for Shape Based on Peaks and Ridges in the Difference of Low-Pass Transform," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(2), 156--169, 1984.

R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*, New York: Wiley, 1973.

F. S. Fay, K. E. Fogarty, J. M. Coggins, "Analysis of Molecular Distributions in Single Cells Using a Digital Imaging Microscope," in *Optical Methods in Cell*

Physiology, P. de Weer and B. Salzburg, editors, John Wiley and Sons, New York, 1985, pp. 51--64.

F. S. Fay, K. E. Fogarty, and J. M. Coggins, "Computerized 3-dimensional Immunocytochemistry in Single Cells," *Biophysical Journal*, 45:105a, 1984.

J. Gauch, W. Oliver and S. Pizer, "Multiresolution Shape Descriptions And Their Applications In Medical Imaging," Technical Report 87-018, University of North Carolina at Chapel Hill, also in Information Processing in Medical Imaging, 10th Conference, Utrecht, Plenum, 1987.

A. P. Ginsburg and J. M. Coggins, "Texture Analysis Based on Filtering Properties of the Human Visual System," *Proceedings of the IEEE International Conference on Cybernetics and Society*, Oct. 1981, Atlanta, pp. 112--117.

A. Ginsburg, *Visual Information Processing Based on Spatial Filters Constrained by Biological Data* , Doctoral dissertation, University of Cambridge, England, 1977. Also published as Air Force Aerospace Medical Research Laboratory Technical Report AMRL-TR-78-129, December 1978.

A. Ginsburg, "Specifying Relevant Information for Image Evaluation and Display Design: An Explanation of How We See Certain Objects," *Proceedings of the Society for Industrial Design*, vol. 21, 1980, 219--227.

L. O. Harvey and M. J. Gervais, "Internal Representation of Visual Texture as the Basis for the Judgement of Similarity," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 7, 1981, 741--753.

A. K. Jain and R. C. Dubes, *Algorithms for Clustering Data*, Englewood Cliffs, NJ: Prentice Hall, 1988.

J. Koenderink, A. van Doorn, "Visual Detection of Spatial Contrast; Influence of Location in the Visual Field, Target Extent and Illuminance Level," *Biological Cybernetics*, Vol. 30, 157--167, 1978.

J. Koenderink, "The Structure of Images," *Biological Cybernetics*, Vol. 50, 363--370, 1984.

J. Koenderink, "Representation of Local Geometry in the Visual System," *Biological Cybernetics*, 1986.

J. Koenderink, "Image Structure," *Mathematics and Computer Science in Medical Imaging*, (Proceedings, NATO Advanced Study Institute, Il Ciocco, 1986), Springer-Verlag, 1987.

K. I. Laws, Textured Image Segmentation, Technical Report TR-940, University of Southern California Image Processing Institute, 1980.

L. Lifshitz, *An Image Description for Object Definition Based on Extremal Regions in the Stack*, Doctoral dissertation, University of North Carolina, Department of Computer Science, 1987.

D. Marr, "Visual Information Processing: The Structure and Creation of Visual Representations," *Philosophical Transactions of the Royal Society of London*, B290, 199--217, 1980.

D. Marr and E. Hildreth, "Theory of Edge Detection," *Proceedings of the Royal Society of London*, B207, 187--217, 1980.

G. Nagy, "Candide's Practical Principles of Experimental Pattern Recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence* , vol. 5, 1983, 199-200.

G. E. Peterson, editor, *Tutorial: Object-Oriented Computing, Volume 1: Concepts and Volume 2: Implementations*, IEEE Computer Society Press, 1987.

S. Pizer, W. Oliver, J. Gauch and S. Bloomberg, "Hierarchical Figure-Based Shape Description for Medical Imaging," *Proc NATO Advanced Science Institute on Mathematics & Computer Science in Medical Imaging*, Springer-Verlag 1987.

S. Pizer, J. Koenderink, L. Lifshitz, L. Helmink and A. D. J. Kaasjager, "An Image Description for Object Definition Based on Extremal Regions in the Stack," *Information Processing in Medical Imaging*, S. L. Bacharach, ed., 24--37, Martinus Nijhoff, Dordrecht, Netherlands 1986a.

S. Pizer, W. Oliver and S. Bloomberg, "Hierarchical Shape Description Via The Multiresolution Symmetric Axis Transform," *IEEE Trans PAMI*, 1986b.

D. A. Pollen and S. F. Ronner, "Visual Cortical Neurons as Localized Spatial Frequency Filters," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-13, no. 5, 1983, 907--915.

W. Richards and A. Polit, "Texture Matching," *Kybernetic*, vol. 16, 1974, 155--162.

J. Robson, "Frequency Domain Visual Processing," *Physical and Biological Processing of Images*, O. J. Braddick and A. C. Sleigh, ed., 73--87, Springer-Verlag, 1983.

A. Rosenfeld, "Image Pattern Recognition," *Proceedings of the IEEE*, vol. 69, 1981, 569--605.

A. Rosenfeld, *Multiresolution Image Processing and Analysis*, Springer-Verlag: Berlin, 1984.

B. Stroustrup, *The C++ Programming Language*, Addison-Wesley, 1983.

H. Wilson and J. Bergen, "A Four-Mechanism Model for Threshold Spatial Vision," *Vision Research*, Vol. 19, 19--32, 1979.

S. W. Zucker, "Computer Vision and Human Perception: An Essay on the Discovery of Constraints," *Proceedings of the IEEE Conference on Pattern Recognition and Image Processing*, Dallas, 1981, 1102--1116.

4. Accomplishments of This Project

Almost four years ago, Dr. Coggins, with Dr. Frederick Brooks as senior investigator, proposed a five-year program of research in image pattern recognition with three principal objectives:

• to develop a research software base to support interactive image pattern recognition related to earth and space science;

• to develop artificial visual systems to solve measurement and detection problems and to evaluate their effectiveness and sensitivity; and

• to apply these techniques to earth and space science data with the objective of solving real problems faced by earth and space science investigators.

We have made progress on all three objectives and have advanced the overall goal of constructing a scientific foundation for engineering image analysis applications. This section reviews the achievements of our research group under the CESDIS contract. These achievements are organized into three subsections corresponding to the Specific Aims listed in Section 2. Section 4.1 Foundations summarizes the software tools and literature review that provided a foundation for the rest of the project. Section 4.2 Image Analysis Methods describes the advances we made in applying Dr. Coggins' Artificial Visual System approach to various image analysis problems. Section 4.3 Interactive Tools summarizes our accomplishments toward developing interactive software tools for image analysis based on the Artificial Visual System approach. Section 4.4 Other Accomplishments discusses other activities undertaken during this project, including additional projects in image restoration and software engineering, publications, miscellaneous accomplishments, and a list of the students supported by the contract.

4.1 Foundations

1. To create a foundation for development, evaluation, and testing of the methods developed in this research program.

A. To identify specific image analysis applications in the earth and space sciences to use as driving problems for the research program.

• Hired research assistants from the departments of Geography and Astronomy to perform literature reviews of image analysis techniques in those fields.

Research Assistants Laura Kellar (Astronomy), Jay Stewart and Jim Stephens (Geography) were supported by our project. Their insights into current visual analysis problems in astronomy and geography made substantial contributions to the key insights developed during the project.

• Prepared bibliographies from the astronomy and geography literature, emphasizing recent conference proceedings, covering image analysis techniques and needs in earth and space science.

These bibliographies provide a snapshot of the state-of-the-art in image pattern recognition in earth and space science at the beginning of this project. We found that in astronomy there are very sophisticated techniques used for image enhancement but there are missed opportunities to use computers for image segmentation and measurement. In geography the sophistication of computer vision methods is very inconsistent. Some labs use quite sophisticated analyses, but more prevalent is the kind of unmathematical, poorly-justified, *ad hoc* analysis that the state of computer vision research might be expected to inspire and our research is intended to improve.

• Interviewed earth and space scientists at UNC concerning image analysis and visualization problems relevant to their research.

Of particular importance was the participation of Dr. James Rose, Dr. Gerald Cecil, and Dr. Wayne Christiansen (Astronomy) and Dr. Stephen Walsh (Geography). Through these relationships, we also have met and worked with Dr. Chris Powell (Geology) and Dr. John Bane and Dr. Christopher Martens (Marine Science).

• Visited scientists at the Goddard Space Flight Center and interviewed them concerning image analysis problems, data management and manipulation.

During Spring Break at UNC in 1989, Dr. Coggins and two students interviewed over 30 NASA scientists about the computer vision and visualization problems they face in their research. These were selected from over 60 NASA scientists who were contacted in a selective mass mailing the month before our visit. These people were contacted based on short articles from the Annual Report of the Goddard Space Flight Center that suggested projects and people that might provide useful information for our research and who our research might help. Addresses of those people were obtained from the Goddard telephone directory.

• Exchanged data with earth and space science investigators at UNC and at Goddard in order to begin investigations relating our computer vision techniques to real scientific problems.

We now have earth and space science data and problem definitions that we can use in teaching and research that were acquired during this project. We also

acquired a much better appreciation for the data analysis challenges facing NASA in the next decade. These data and insights are being used now in courses and research activities in our department.

• Used some NASA data to define course projects in computer vision at UNC.

Much of the data we acquired has already been used (with the permission of the data providers) in course and research projects at UNC. The impact is far-reaching since the earth and space science data is being used in graphics and image analysis projects and in courses in computer vision and graphics.

B. To design and implement a software base for the experimental research using modern software engineering practices.

• Added C++ classes to Dr. Coggins' IGLOO software library

The most important additions implement pattern recognition operations such as intrinsic dimensionality estimation, minimum spanning trees, and clustering; sorting and searching; advanced matrix and image operations; command parsing; manipulation of computer networks, 2-D Fourier Transforms KNN classifier, computation of distance matrices and near-neighbor matrices, pixel labelling facilities. Recently, the user interface classes were redeveloped to enable them to be used more flexibly and to better separate user interface concerns from the rest of the application program.

Requests for information about IGLOO and for copies of IGLOO itself continue to come in at UNC. The Computer Science Department's SoftLab Systems Laboratory handles such requests, relieving Dr. Coggins of a heavy administrative burden. IGLOO recently became the most-requested product supported by SoftLab.

• Developed software for performing spatial filtering on a remote Convex mini-supercomputer with data and results being passed over the network.

The Convex 2-D FFT routines run an order of magnitude faster than our FFT procedures on Sun workstations. Unfortunately, the time required for transferring the data and results over the network consumes most of the savings. As faster networks are installed or as our research improves the ability to reduce the data by performing further proecssing on the Convex machine before transmitting results, this facility may become a practical alternative.

• Revised and enhanced portions of the library concerned with user interface, computer graphics, probability densities, and descriptive statistics.

As one gains experience with object-oriented design, techniques for more effectively using advanced techniques in the language such as inheritance continue to arise. As these techniques have been developed, many portions of the library have been revised to incorporate the new insights. Often, such insights result in a significant decrease in the number of source code lines in the library.

• Revised and enhanced the management structure of the library to be simpler to understand and to permit easier maintenance of the library for multiple computer architectures simultaneously.

When first developed, the management techniques we used required a special version of the `make` utility called `gmake`. Modifications in the management strategy now make possible the use of the standard `make` utility provided by all of the UNIX vendors. This development greatly enhances the utility of the management techniques and increases the interest from throughout the C and C++ communities.

• Developed software for visualization of results of computer vision research.

These methods include software for plotting scale traces (plots of filter responses through scale), performing color-wash labelling of regions, constructing feature space images from pairs of filtered images, mapping labelled pixels into feature space images, plotting regions of selected pixels into a scale trace plot, mapping pixels selected in a feature space image back into the original image using a color wash, and creating filter kernel images corresponding to Koenderink's Polar Receptive Field Family.

• Developed an image region hierarchy editor.

This program is used to manipulate image labellings and is based on Dr. Coggins' design of a color-labelling display scheme. This program implements the notion of a *visually sensible region*, automatically defined by a multiscale, geometric analysis (not edge detection), as a primitive unit for human-computer interaction.

• Developed a prototype of a program, `imutil`, that may be further developed into an image processing compute server and interpreter.

The `imutil` prototype is based on Dr. Coggins' IGLOO class library. We intend to develop it into a network-based image server operated using a simple language resembing the assembly language of some computers. The prototype was first developed while Dr. Coggins was visiting Goddard in the summer of 1990.

• Developed a testing suite for the IGLOO library.

This series of test programs exercises all classes and all member functions within IGLOO. The incorporation of these testing tools into the IGLOO distribution remains to be performed, but this constitutes the first complete testing suite to be provided with any C++ library.

• Began a port of IGLOO to the Cray Y-MP

Dr. Coggins began porting his research software to the Cray Y-MP at Goddard. The software is all compiled there, but there remains a problem: Dr. Coggins'

software cannot link and run without the C++ run-time support library, and its source code is not public domain. It would be necessary to obtain and port the source code of that library or to obtain and port C++ to the Cray. No further progress was made on this issue either at Goddard or after several attempts to work around the problems with the help of the North Carolina Supercomputing Center.

4.2 Image Analysis Methods

2. To develop image measurement and description operations based on artificial visual systems.

A. To develop new artificial visual systems for use in measurement and description tasks.

• Investigated the power of an AVS composed of multiscale Gaussians for image segmentation.

K-C Low, Lisa Baxter, and Dan Fritsch worked on unsupervised and supervised methods for exploring the feature space defined by multiscale Gaussian filters. K-C found that multiscale Gaussians are powerful by themselves, without additional orientation filters. Lisa developed the "scale trace" technique that is used throughout our research now. Together, their work demonstrated that there are several geometrical classes into which pixels can fall: exterior, exterior edge, interior edge, interior, and center. This was the first hint that the medial axis might be a real, statistically verifiable entity. Dan Fritsch found a connection between the medial axis and filters defined in Koenderink's Receptive Field Families. David Rudolph is now continuing investigations of multiscale Gaussian derivative filters for geometric landmark detection.

• Investigated the use of orientation filters to form a Multiscale Orientation Field.

The MOF is a vector field description of an image in which the vector at each pixel gives the orientation and eccentricity of image energy at the measurement scale. It turned out to be little more than the gradient, and was demonstrated to be insufficient to define all of the geometric structure we need. However, the MOF led to a new rationale for the appropriateness of the Koenderink receptive field families and contributed to key insights we have developed since then.

• Showed that a feature space defined by multiscale Gaussian filters causes pixels to cluster by their geometrical role in the image.

This is the most important and far-reaching result of this research. It justified the use of a variant of the medial axis for capturing the structure of grayscale images and points the way toward further analysis of the feature space formed by the outputs at each pixel of multiscale Gaussian filters.

• Applied the clustering results to segmentation of biomedical images and LANDSAT images of central North Carolina.

Besides performing clustering of pixels in images based on multiscale Gaussians, we began investigation of such clusters as a foundation for robust image registration. David Rudolph, a dental student, has taken this objective seriously. His study involves the use of multiscale Gaussians and Gaussian derivatives for locating predefined landmarks in images that can be used for image registration and geometric measurement. Since one of the main categories of earth and space science computer vision problems is image registration, his work holds great potential for impacting earth and space science research.

• Tested a feature space for texture analysis on classification of ice and snow cover in synthetic aperture radar images of the Greenland ice cap.

Lisa Baxter obtained good results during a class project in Fall 1990 at segmenting SAR images of the Greenland ice cap by image texture (but see Section 5 of this report).

• Developed methods for analyzing objective prism plates from telescopes to locate the images of stars, measure their position and brightness on the plate, and distinguish stars from noise on the plate.

The method, based on blurring using multiple filters and classifying pixels in the resulting feature space, is in use now in the UNC Department of Physics and Astronomy.

• Applied the same image pattern recognition method to identification of stars in images of Comet Halley and showed that we could eliminate stars from the images.

The method works well, and two variants were developed (but see Section 5 of this report).

B. To investigate the sensitivity of such measurement techniques to noise and distortion.

• Performed a detailed analysis of the accuracy of the MOF for measuring orientations of line segments.

The MOF was demonstrated to be capable of highly accurate orientation and position measurement, even at oblique angles and on small objects.

• Derived an analytic form for the falloff orientation filters must posess in order to map a delta function stimulus moving through the measurement dimension into a circle in the feature space.

This important study showed how the principles of spectroscopy can be applied to design spatial filters in an Artificial Visual System.

• Studied results from the laboratory of Jan Koenderink in the Netherlands indicating that the noise sensitivity to high-order derivatives decreases as the scale of the derivative operators becomes larger.

• Demonstrated that corresponding points in multiple images of an object can be identified based on similar image geometry at large scale.

This crucial result for image registration has since been extended by David Rudolph to include landmarks requiring small-scale structure as well as large-scale structure to be brought to bear to match visible objects.

4.3 Interactive Tools

3. To develop multiscale image descriptions for use in interactive image analysis.

This section of our original proposal was scheduled for the last two years of the contract (the 2-year extension), but we have made some progress on these tasks during the initial three-year contract.

A. To develop automatic methods for creating image descriptions in terms of regions.

• Discovered a method by which the overall responses of multiscale Gaussian filters can be analyzed to determine which scales contain "interesting" structure in the image being analyzed.

Lisa Baxter demonstrated a method whereby the right range of scales for performing detailed analyses can be found. This important result helps us to define in a complex image which scales contain interpretable information. This focuses our search in relevant scale channels. This result is related to Dr. Coggins' earlier work on visual texture based on a multiscale channel model of the early human visual system. This technique was part of a demonstration suite prepared in December 1990 (but see Section 5 of this report).

• Developed region definition procedures based on supervised pixel classification.

Unlike Bayesian Pixel Classification methods based on intensity, our methods are based on the multiscale structure of the image as revealed by a series of multiscale Gaussian filters. Thus, the requirement that spatial context be incorporated for valid image segmentation is satisfied.

• Developed a procedure for performing region segmentation in an image based on a feature space description of pixels.

K-C Low developed this technique, which involves a region-growing control strategy. Unlabelled pixels are added to regions until a criterion *defined in the feature space* is met.

• Investigated Koenderink's Polar Receptive Field Family for image description and analysis. Developed tools to compute the filters, apply the filters to an image, and display the results in several different modes.

• Discovered that the vector fields used in the Multiscale Orientation Field is only one kind of vector field result obtainable from Koenderink's Polar Receptive Field Family. The same technique can be used to compute vectors from other pairs of filter outputs in Koenderink's filter set, with equally useful results. We call the resulting vector and scalar fields the Multiscale Geometry Field.

B. To develop methods for user-guided interaction based on automatically identified regions using interactive computer graphics techniques.

• Under other support, the Image Hierarchy Viewer designed by Dr. Coggins and students Tim Cullip and Eric Fredericsen has been refined for use on workstations with 2-D or 3-D image data. The refinements include harnessing the PixelPlanes graphics engine developed in our department by Henry Fuchs and John Poulton to render images of 3-D regions selected using IHE.

C. To develop methods for interactive navigation through an image data base using simplified image descriptions based on multiscale representations of the image data.

• Simplified image descriptions of objective prism plates were developed and interactive methods for examining those results are now in use in the Department of Physics and Astronomy at UNC.

4.4 Other Accomplishments

• Space Telescope Deconvolution Project

Dr. Milt Halem introduced Dr. Coggins to Dr. Ron Downes, an astronomer who was then working in one of the Guaranteed Time Observer groups for the Hubble Space Telescope. Dr. Downes provided Dr. Coggins with some ground-based telescope imagery and some point spread function data. Dr. Coggins used the point spread function to blur the images and applied two algorithms to deblur them. One was the familiar Van Cittert iterative method. The other was an original modification Dr. Coggins developed at CESDIS in August 1990 called Iterative-Recursive Deconvolution. The Van Cittert method removed 55% of the blur after 50 iterations. Dr. Coggins' method removed 99.5% of the blur in 10 iterations. This attracted the interest of the Hubble astronomers who provided a couple of real Hubble images and a point spread function for Dr. Coggins to try. Experiments are being run at UNC to attempt to characterize the nature of the deconvolution performed by the method and to understand its behavior with noisy data. It appears that the method is linear and would therefore preserve flux in the restored image. If this is proven, this method will be of great importance in astronomy and other fields.

Dr. Coggins has discussed his iterative/recursive deconvolution method with knowledgeable scientists in image reconstruction and enhancement who have all indicated that the modification developed by Dr. Coggins is unique in their experience. Dr. Coggins has analyzed the mathematics of similar methods and now understands more about what his new method is doing. Work on this technique is continuing under other funding. (Also see section 5 of this report).

• Contributions to C++ community

One of the most surprising accomplishments of this contract was the enthusiastic interest in Dr. Coggins' work on his C++ library from the C++ community. Dr. Coggins has been invited to speak to both industry and academic audiences, and is now a regular speaker at the major C++ conferences on the topics of library design and management. Dr. Coggins' CESDIS technical reports related to the design and management of C++ libraries are the most often requested of his writings according to CESDIS records. His C++ software library is the most-requested product supported by the SoftLab Software Systems Laboratory in the Computer Science Department at UNC.

• Hughes Aircraft Interaction

One of the contacts we made at the October 1990 CESDIS workshop was John Heinrichs of Hughes Aircraft. John spent nearly an hour at the busy workshop talking to us about our research in detail. In April 1991, this contact resulted in a proposal by Hughes and UNC with John Heinrichs as PI and Dr. Coggins as PI of a UNC subcontract proposing to NASA OSSA to use Dr. Coggins' computer vision methods to perform cloud classification. The proposal was not funded in its first submission, but positive comments from the reviews were encouraging. The criticisms of the proposal are being used to revise the proposal for resubmission.

• Demo Interpreter Program

Dr. Coggins and Lisa Baxter travelled to Goddard for the CESDIS Science Council Meeting held on December 10, 1990. Over the previous weekend, Lisa generated and arranged a series of images that illustrate our most recent results, and Dr. Coggins implemented a "demo interpreter" program to present Lisa's images along with explanatory text. (But see section 5.) The demo interpreter and Lisa's images have been used as visitors have passed through our laboratories at UNC.

• Students Supported Directly by the Contract

Kah-Chan Low
Lisa Baxter
Robert Ingram
Jay Stewart
Jim Stephens
Laura Kellar
David Chen
Yuchin Fu

• Publications

J. M. Coggins and G. Bollella, "Managing C++ Libraries," SIGPLAN Notices, 24(6) 37-48 June 1989.

J. M. Coggins, "Computer-Aided Object Definitions from Digital Images," Proceedings of the Joint US-Scandinavian Symposium on Future Directions of Computer-Aided Radiotherapy, Radiation Research Program, Division of Cancer Treatment, National Cancer Institute.

S. M. Pizer, J. M. Gauch, J. M. Coggins, R. E. Fredericksen, T. J. Cullip, V. L. Interrante, "Multiscale, Geometric Image Descriptions for Interactive Object Definition," Mustererkennung 1989 (Proc. 11th Symposium of DAGM, the German Association for Pattern Recognition), Informatik-Fachberichte 219: 229-239, Springer-Verlag, 1989.

D. J. Rudolph, J. M. Coggins, and D. A. Tyndall, "Computerized Image Processing of Dental Radiographs," published as abstract in Proceedings of 40th Annual Scientific Session of Americal Academy of Dental Radiology, Honolulu, Hawaii, Oct 31-Nov 3, 1989, p. 21.

K-C Low and J. M. Coggins, "Multiscale Vector Fields for Image Pattern Recognition" Proceedings of SPIE Conference on Intelligent Robots and Computer Vision VIII: Algorithms and Techniques, November 6-10, 1989, Philadelphia, SPIE Proceedings Series, vol. 1192, pp. 159-169.

J. M. Coggins, "Anticipated Methodologies in Computer Vision" SouthCon '90 Conference Record, Orlando Fl., March 20-22 1990, pp. 144-148.

J. M. Coggins, "Design Criteria for C++ Libraries, Proceedings of the Second USENIX C++ Technical Conference, San Fransisco, CA, April 9-11, 1990, pp. 25-35.

R. E. Fredericksen, J. M. Coggins, T. J. Cullip, and S. M. Pizer, "Interactive Object Definition in Medical Images Using Multiscale, Geometric Image Descriptions," Proceedings of the First International Conference on Visualization in Biomedical Computing, Atlanta, May 22-25, 1990, pp. 108-114.

J. M. Coggins, "A Multiscale Description of Image Structure for Segmentation of Biomedical Images," First International Conference on Visualization in Biomedical Computing, Atlanta, May 22-25, 1990, pp. 123-130.

K-C. Low and J. M. Coggins, "Biomedical Image Segmentation Using Multiscale Orientation Fields," First International Conference on Visualization in Biomedical Computing, Atlanta, May 22-25, 1990, pp. 378-384.

J. M. Coggins, "Measuring Image Structures Using a Multiscale Orientation Field," Proceedings of the Tenth International Conference on Pattern Recognition, Atlantic City, NJ, June 17-21, 1990, pp. 720-722.

J. M. Coggins, "Designing C++ Libraries", The C++ Journal, vol. 1, no. 1, June 1990.

J. M. Coggins, R. E. Fredericksen, and S. M. Pizer, "Image Structure Analysis Supporting Interactive Object Defintiion," Proceedings of the 12th International Conference of the IEEE Engineering in Biology and Medicine Society, Philadelphia, PA, Nov. 1-4, 1990.

L . C. Baxter and J. M. Coggins, "Supervised Pixel Classification Using a Feature Space Derived from an Artificial Visual System," OE/Boston '90, Applications of Optical Science and Engineering, Boston, November 4-9, 1990.

5. Limitations of This Project

While we are proud of the research accomplishments made possible by the first two years of the CESDIS contract support, important activities were terminated by funding cuts implemented in the contract's third year. This section describes how the funding cuts were implemented and how they affected the conduct of this project.

UNC and CESDIS administrators have different interpretations of the timing and extent of the budget cuts implemented in the third contract year. The purpose of this report is neither to advocate either interpretation nor to arbitrate between them. This section reports on how and why the research accomplishments of this project were so severely affected by the funding reductions imposed midway through the third contract year.

5.1 Implementation of Budget Cuts

During Year 2 of our contracts, CESDIS contractors were warned that budget cuts were coming. In response to this advance warning, UNC began curtailing discretionary expenditures. Several months into Year 3 of the contract, cuts were implemented. These cuts had devastating effects on our contract's activity in Year 3. These effects were amplified by several factors:

   • The cuts were implemented by allocating to the project a percentage of funds *actually spent* in years 1 and 2. Thus, our frugality in year 2 in response to the warning that budget cuts were coming not only limited our Year 2 activity but made the cuts implemented in Year 3 worse.

   • Unspent funds from our Year 2 approved budget were retained, so the cushion we were hoping to draw upon to lessen the impact of anticipated Year 3 cuts disappeared. This was surprising because no mention was made in the warnings of impending cuts of any effect on our Year 2 approved budget.

• While we were warned of cuts in our Year 3 budget, the cuts actually implemented were deeper than we anticipated or were led to expect, and were implemented after we had already committed to graduate student support for the Spring 1991 semester.

When the budget cuts were announced in January 1991 (Year 3 of the contract began September 1, 1990), UNC administrators found that the funding available for the rest of Year 3 barely covered our commitments to on-board students for the Spring 1991 semester. Since we were not willing to fail to meet our commitments to on-board students, UNC administrators had no choice but to impose a hard cutoff of all expenditures under this contract. Thus, we entered the Spring semester of 1991 with student support but with zero budgets for travel, telephone, supplies, and services.

Part of the expenses for travel to CESDIS for the December 1990 CESDIS Science Council meeting were funded by Dr. Coggins' personal funds. All expenses for the three short visits Dr. Coggins made to CESDIS in the spring and summer of 1991 were funded by Dr. Coggins.

5.2 Effect on Year 3 Project Accomplishments

The main effect of the budget cuts was to cut off contact with collaborating scientists at Goddard. Five projects we undertook in the summer and fall of 1990 using data obtained at Goddard all had results ready to be returned to our Goddard collaborators at the end of 1990. These projects were as follows:

• location of stars in images of Comet Halley;

• classification of ice, snow, rock and slush regions in SAR images of the Greenland ice cap;

• identification of critical points (landmarks) for registration in SAR images of sea ice;

• pixel classification on LANDSAT images using a multiscale, geometric feature space;

• analysis of a new iterative/recursive method for image deblurring and its application to Hubble Space Telescope images.

Our plan was to visit Goddard during the Spring semester 1991 (during spring break as we had done the previous year) to return these results to the scientists who provided us the data. As of January 1991, since travel or any effective remote collaboration was not possible, we spent the Spring 1991 semester enhancing the software library and documenting the software and projects we had completed in the hope that they might be communicated back to Goddard or carried forward later under other funding. As of the date of the last student paycheck in May 1991,

there were no further expenditures or activities on this contract. The contract itself expired on August 31, 1991.

Since then, only the image deblurring method has been investigated further, on Dr. Coggins' own initiative.

5.3 Effect on Contract Extension

In July 1991, we were informed by letter that the two-year extension of our initial 3-year contract would not be funded. The reasons for not funding the extension provided with the letter included several criticisms related to the failure of the project to follow-through with the Goddard scientists who were contacted during Year 2 of the contract. Also cited was a failure to apply the techniques developed in years 1 and 2 to NASA data.

While the failure to follow-through with our NASA contacts was real, it was a direct result of the hard budget cutoff we experienced in January 1991. We have been told that the peer review of our contract's extension proposal was made without knowledge of the impact of the Year 3 funding cuts on our project.

The criticism that we failed to apply our techniques to NASA data is false. Our techniques were applied to LANDSAT data in the summer of 1990 and were shown in Dr. Coggins' August 16, 1990 CESDIS Science Seminar. Results of the other projects listed above were ready to return to Goddard at the end of the Fall 1990 semester. At the December 1990 CESDIS Science Council meeting, Dr. Coggins and Lisa Baxter prepared a demonstration suite including images from the LANDSAT pixel classification project, the Comet Halley project and the Greenland ice cap project. The Science Council never came to see the demonstrations. Budget cuts cancelled our plan to return to Goddard in the Spring semester 1991 after some further work with the Goddard data.

6. A Review of our CESDIS Experience

Our CESDIS contract has supported research activity and infrastructure development that will continue to have an impact at UNC into the future. Through this contract, we were able to educate graduate students, purchase equipment, and collect data for use in algorithm testing. We have become aware of image analysis, restoration, and data storage problems faced by NASA in supporting earth and space science research and have been able to relate the image analysis needs in earth and space science to image analysis problems from other disciplines. Our experience in this project led to a new formulation of our overall objective in which we see our approach to computer vision as a kind of spectroscopy. This analogy, which has proven to be a crucial insight in our subsequent work, was a direct result of our CESDIS-supported research.

When this contract began, CESDIS did not yet physically exist. It was almost a year into the contract before a laboratory suitable for conducting collaborative

research existed at Goddard. That first year was spent mainly at UNC working on software foundations and working with geographers and astronomers locally. During the second contract year, the CESDIS Science Center was available so that students and investigators could be productive during visits to Goddard. We began serious outreach to NASA scientists, collected data, defined image analysis problems, and tried to define what our direct contribution back to NASA might be. Progress was made on several of these problems in the first half of the third contract year.

6.1 Characteristics of NASA Image Analysis Problems

During the data gathering stages of the project, we noticed two interesting properties of NASA-related image analysis problems.

First, NASA-related images are huge. One SAR image consists of about 8000x8000 pixels; each LANDSAT image is of comparable sample size but is repeated for each of seven bands. The sheer size of the NASA data meant that new procedures had to be developed to extract and condense components of these images to a size that could be analyzed using the workstations on which our experimental software runs. In order to view on a workstation screen the original image, some intermediate results, and the final results of an analysis, the original image size needs to be no larger than 512x512 pixels. Digital images from biomedical studies normally come in approximately this size or smaller. Thus, we found that additional processing in the form of subimage extraction and reduction by pixel averaging is required to prepare NASA data for computer vision experiments. By going through the exercise of acquiring and converting several such images for use in our experiments, we are now familiar with the sources of image data available in NASA and will be able to increasingly incorporate such data in our research and teaching efforts in the future.

Second, the image analysis tasks that need to be performed on NASA data are the same tasks required on biomedical data. Geometric correction, image enhancement, registration of images from different modalities, registration of grid lines on images, pixel classification based on intensity and image geometry, identification of landmark points in images, and color labelling of pixels according to various criteria are all common problems in both NASA-related earth and space science data and in biomedical data. Furthermore, both NASA and biomedical images are accompanied by geometric data describing image acquisition parameters that are important for calibrating quantitative analyses of the images. We conclude from these observations that research on biomedical image analysis is entirely relevant to NASA's image analysis interests. We further conclude that our vision of a unified scientific foundation for the engineering of particular artificial visual systems tuned to solve specific image analysis problems is supported by our new appreciation of image analysis problems in earth and space science.

6.2 CESDIS: The Vision and the Reality

The vision of CESDIS conveyed early in the contract period was one of an "open institute" or an "institute without walls". This is a vision that we supported by traveling to CESDIS whenever it was possible during the first two years of our contract. These visits were longer and more frequent in the second contract year after the CESDIS Science Center was established. We conclude from our experience that the notion of teaching faculty and graduate students traveling frequently to work at Goddard is not feasible for faculty and students outside the immediate Washington, D.C. area. We were able to travel to Goddard during term breaks, for CESDIS Workshops and Science Council meetings and for more extended periods during the summer. Additionally, short visits to CESDIS were coordinated with other trips, but these visits were too short for any serious collaborative work to be performed. Teaching faculty cannot leave for travel without disrupting courses, incurring administrative paperwork for travel authorizations and reimbursements, and interfering with course preparation. For graduate students, travel is an even greater burden since they have less control over their required activities and time. We spent as much time as we could at Goddard, but this was not nearly as much as the vision of the "open institute" implied. While we supported the concept by both words and actions, we did not find its full realization to be practical for either students or faculty.

On the other hand, the time we did spend at Goddard was useful and productive. Minimal distractions and a good working environment in the CESDIS Science Center made those periods particularly stimulating. Some of the best ideas we generated in our project arose during the CESDIS Workshops and while preparing demonstrations at CESDIS. NASA scientists at Goddard provided us with real data and background information on the objectives and importance of earth and space science research. We also saw how NASA runs earth and space science research. We were able to critically review computing and information science research at NASA. We encountered a spectrum of projects ranging from highly creative and sophisticated analyses to competent engineering efforts within the state-of-the-art to *ad hoc* collections of poorly-motivated, unproven methods that we hope our research will render obsolete, sooner rather than later. We are encouraged that our goals of a scientific foundation for engineering image analysis applications and smoother integration of image processing and computer graphics will have a significant impact on NASA-supported research.

6.3 Conclusion

We are proud of the accomplishments of our CESDIS contract, especially in view of what we were able to accomplish despite the decreased support in its final year. We would like to complete at some later date the research activities planned for Year 3 but cancelled due to funding cuts. We remain supportive of the vision of CESDIS as a center for collaborative research between the computer science community and the NASA earth and space science communities, and believe that with appropriately-scaled expectations and consistent funding it could become an important and productive contact between these research communities.