1990

A New Approach to Analyzing High-Resolution Aerial Photographs of Urban Areas.

Timothy Dwayne Phillips
Louisiana State University and Agricultural & Mechanical College

Order Number 9112263

A new approach to analyzing high-resolution aerial photographs of urban areas

Phillips, Timothy Dwayne, Ph.D.

The Louisiana State University and Agricultural and Mechanical Col., 1990

U·M·I
300 N. Zeeb Rd.
Ann Arbor, MI 48106

A New Approach to Analyzing High Resolution Aerial Photographs of Urban Areas

A Dissertation

Submitted to the Graduate Faculty of the
Louisiana State University and
Agricultural and Mechanical College
in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

in

The Department of Electrical and Computer Engineering

by
Dwayne Phillips
B.S., Louisiana State University, 1980
M.S., Louisiana State University, 1984
August 1990

Dedication

In memory of Neal Phillips, November 6, 1930 - January 18, 1990

This is dedicated to all those who had the ability but were not blessed with the opportunity to have such an experience and education.

Most of all, this is dedicated to my three sons Seth, Nathan, and Adam because the future belongs to them.

Acknowledgements

I would like to thank the members of my committee: Dr. Charles Harlow, Dr. Subhash Kak, Dr. Jerry Trahan, Dr. Ahmed El-Amawy, and Dr. Donald Kraft. I would also like to thank Dr. Richard Conners and Dr. Mohan Trivedi who served during the first three years of my work.

Thanks go to the staff of the Remote Sensing and Image Processing Lab past and present. These include Don Middleton, Kevin Marshal, Don Dirosa, Dave Evans, and George Orhberg.

Thanks go to my brother-in-law Bruce Bundy for keeping me supplied with materials from the library.

Thanks go to my mother and father and my wife's mother and father for taking care of the children on many occasions and supporting and loving my wife and me.

Special thanks to my wife Karen. She lived with this work night and day. She never complained, always encouraged, and always loved.

Finally, thanks to God from whom all blessings flow.

Table of Contents

Abstract

This dissertation proposes a new approach for analyzing high resolution aerial photographs of urban areas. Analyzing aerial photographs is the process of constructing an overall description of a scene. It involves knowledge of visual sensors, computing systems, artificial intelligence, software engineering, and perceptual psychology. Researchers have had only limited success in this area. This dissertation considers a high level analysis approach. Most aerial photograph interpretation systems concentrate on analyzing an airport, roadway, or urban scene. Those systems, however, do not explain how they knew they were examining such a scene. This dissertation concentrates on how to reach that point. It begins with "this is an aerial photograph" and works its way down through a hierarchy of labels until it reaches the point of "this is an urban area - find and label the objects." The new analysis approach introduces a unique use of three basic ideas. These ideas are (1) the use of context, expectations, selective attention, and the perceptual cycle, (2) analyzing the image through a hierarchy of increasingly specific labels, and (3) the interplay between the segmentation and interpretation processes. These are developed in a computer vision system for analyzing aerial photographs. The system comprises (1) a control mechanism, (2) a knowledge base, (3) a belief maintenance system, and (4) an image processing interface. In general, the system uses the knowledge stored in frames to investigate areas in the image. The control mechanism calls low level routines in the image processing interface. They report the results back to the control mechanism which invokes the belief maintenance system. The belief maintenance system reports which frame is the most probable label for the area under investigation. To demonstrate the system, this dissertation presents the results of analyzing a high resolution, multi-spectral, aerial image of an urban area. It also presents the results of analyzing three different housing areas taken from a single channel, gray scale image of a color aerial photograph. These show the validity of the new approach and the power and portability of the system.

1 - The Problem - How to Analyze High Resolution Aerial Photographs

Analyzing aerial photographs is a problem that has received the attention of researchers for years. Automating the process of analyzing aerial photographs is an increasingly important social and scientific issue. This field has applications to diverse and sometimes opposing special interest groups such as map makers, tax assessors, environmentalists, treaty verifiers, law enforcement officials, farmers, and oil companies. Human photointerpreters perform with high levels of skill and expertise. They are, however, few in number, very expensive, and they require large amounts of time to carry out their tasks. As time passes, our ability to collect aerial photographs grows much faster than the number of human photointerpreters.

Automated systems are no longer a curiosity but a necessity. The special interest groups mentioned above need systems that can analyze hundreds of aerial photographs in a matter of hours without expert human intervention. The techniques and tools in the systems must be simple, to the point, and flexible. In many cases, aerial image analysis does not require labeling the image to the last detail. Special interest groups want systems that can sift through thousands of photographs and hand them the four or five that contain toxic waste dumps, military bases, or rain forests. At that point, the special interest groups can have an expert human photointerpreter detail the stacks of asbestos, barrels of chemical weapons, or logging camps. Researchers have made progress in creating new and better low level operators that locate and label objects in aerial photographs. These operators, however, only work when given specific instructions about the objects for which they are searching and they do not address the higher level analysis questions. (Note, while many of the aerial photographs analyzed today are not really photographs but are digitally scanned images, the following general discussions will refer to both photographs and images as photographs.)

A major problem that still exists is how to approach the analysis of an aerial photograph. Most work done to date uses the problem statement "given an aerial photograph of an airport, label the runways, buildings, and aircraft," or "given an aerial photograph of a housing area, label the streets, sidewalks, houses, yards, and vehicles."

The problem statement of this dissertation is "given an aerial photograph, discover what type of area it contains and label the objects that pertain to that type of area." This is a very broad and ambitious problem statement. The scope of this problem cannot be completely and thoroughly satisfied in a single dissertation. Nevertheless, this dissertation will outline some of the obstacles and propose solutions. A computer vision system was created to address these problems and demonstrate solutions (see chapter four). The system experimented with the analysis of two vastly different types of aerial photographs. This showed the validity of the approach.

One problem is the number of ambiguities in aerial photographs. Given a rectangle in the photograph, is it a building, swimming pool, car, or bale of hay? If the rectangle is a building, then is it a warehouse, hospital, or prison? If it is a warehouse, then a linear feature next to it is probably a road - or maybe a river or a railroad line. Now the interplay and interdependence of the objects becomes important and confusing.

The above relates to the problem of the explosion of possible scenes an aerial photograph can contain. When presented with an aerial photograph of a specific type, the number of possible objects is immense. When you multiply this by the number of different scene types, the number of objects becomes unmanageable. The analysis approach must reduce the size of the problem. The number of alternatives at any one time must be small. There should be under seven or eight - preferably only two or three.

Another problem is guiding low level operators. A system must select and direct properly even the best low level operators if they are to succeed. The general, all purpose operator has not yet been created. Each operator or tool works well in only select, specific situations. If applied in the wrong situation, an operator will return results that are wrong. Histogram analysis tools fail when applied to a texture image. Texture analysis tools fail when applied to a simple line drawing.

Once a system selects the proper low level operator, there is the problem of how to operate on only specific features of the image. Images are too large and operators are too complex to apply the latter to the former in their entirety. Images contain too much information for a system to process them fully. As an illustration, consider recognizing a person in a photograph. Humans can recognize a person in a color photograph. Humans can recognize the same person in a black and white photograph, so the color photograph contains more information than is necessary. Humans can recognize the same person in a simple line drawing, so the black and white photograph contains more information than is necessary. Finally, humans can recognize the same person in a partially covered, simple line drawing, so the full line drawing contains more information than is necessary. The point is there are certain salient features that contain the minimum information required to recognize something. The operators should only work on those features.

Another problem is allowing operators to function on their own without help. Operators often go astray because of noise, ambiguities, or occluded objects. The aerial photograph analysis approach should tie the operators together so they can feed information to and direct themselves and each other.

Another problem is attempting to explain and understand the results of less than perfect operators. Image processing operators often return faulty, incorrect, and contradictory results. Basing an interpretation of an aerial photograph on such results leads to unreliable results. In analyzing an aerial photograph, the understanding system might have to explain why a swimming pool is next to a prison or why a single bale of hay is larger than a warehouse. The understanding portion of the analysis approach needs to feed information to the operators to guide and redirect them. The approach could relax or restrict the parameters with the new results altering or reinforcing the previous explanation.

Related to guiding low level operators and focusing in on salient features is the overall problem of what to do and when to do it. Often the most difficult step is the first one. The situation should determine the action. Therefore, the system should always know the current situation, the alternatives, how to limit the number of alternatives, and how to select the best alternative.

These are the major problems facing the analysis of aerial photographs. This dissertation proposes attacking them with the following ideas (see chapter two for further explanation of these ideas). The use of a hierarchy of scene labels will significantly reduce the explosion of possible scenes that can confront an analysis system (see section 3.1 for a discussion of this hierarchy). The hierarchy will limit the number of possibilities facing the system to a manageable amount. Expectations can guide the low level operators. At any given time, the system can expect certain situations. The expectations guide and direct the operators properly. The operators can operate on only the salient features when they use selective attention. Humans use selective attention without thinking about it to function and solve simple, everyday problems. The system should use the perceptual cycle [Neisser 1976] to tie together operators so they can function with help from themselves and other operators. Tying together the segmentation and interpretation processes can improve dramatically the performance of low level operators. These processes cannot be independent. They must work together. The unifying principle is context. The context of the current situation drives the action, the interpretation, and the next action. Context limits the complexity, drives the operators, and interprets the results.

These concepts - a hierarchy of scene labels, expectations, selective attention, context, the perceptual cycle, and tying segmentation and interpretation together - are not new. Researchers have expressed them in many ways on many occasions in the perceptual psychology, artificial intelligence, and computer vision literature. Nevertheless, no one has ever tied them together as the basis for an approach to analyzing aerial photographs. The end result of this dissertation is a computer vision system that analyzes aerial photographs of urban areas. Figure 1.1 shows the block diagram of this system. The system's four basic parts are (1) the control mechanism, (2) the knowledge base, (3) the belief maintenance system, and (4) the image processing interface. The control mechanism implements the cycle and the interpretation part of the interpretation and segmentation interplay. The knowledge base is the hierarchy and also imbeds the ideas of expectations, selective attention, and context. The belief maintenance system works with the control mechanism to draw reasonable conclusions from the analysis. The image processing interface is the segmentation portion of the segmentation and interpretation interplay.

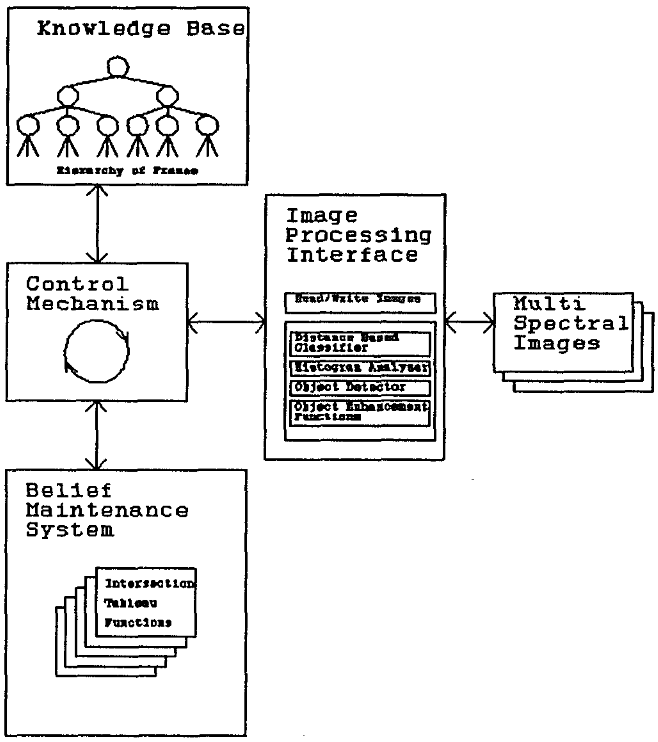

Figure 1.1 - Block Diagram of System

In general, the system uses the knowledge stored in frames to investigate areas in the image. The control mechanism calls low level routines in the image processing interface. They report the results back to the control mechanism which invokes the belief maintenance system. The belief maintenance system reports which frame is the most probable label for the area under investigation. The control mechanism either uses the most probable frame for further investigation or for labeling the area.

The knowledge base for the system is a hierarchy (as mentioned above) of frames. Figure 1.2 shows the top portion of the hierarchy of frames. Section 3.1 discusses the hierarchy in detail. Section 3.2 discusses frames and how they are well suited for a knowledge base in a computer vision system.

A key question before launching into the development of a major computer vision system deals with the advantages of such an effort. The question is "given an aerial photograph of a known urban area, what will be the difference between this approach and system and other systems (such as will be described in chapter 2)." The first major difference deals with the implicit knowledge imbedded into each point in the hierarchy. Each point in the hierarchy contains the context of the situation. This frees the low level operators from worrying with unnecessary details and complications. If the system is at a housing node, then the low level linear feature detector knows it is looking for sidewalks and roads. It does not need to consider linear features such as streams, rivers, or runways. This removes much of the complexity of the operator and allows it to concentrate on the specific problem at hand. It also limits the search area of the low level operators. The objects of interest are located in a small portion of the image. The context knows that small portion and limits the operator's search by focusing their attention.

Several other advantages of the approach concern reducing the computational complexity of the problem. At different levels in the hierarchy the system faces different problems and can use different operators. This translates into simple operators at the higher levels. The system delays using the fine detail, complex operators until later in the image analysis. Another reduction in complexity comes in the belief maintenance system. This system uses a form of the Dempster-Shafer theory of evidence (see section 3.3 for a discussion of this). There are several major simplifications to this theory when it is used with a hierarchy. Another advantage is the system will be able to label all types of urban areas - not just one or two. The hierarchy will encompass all types such as housing, transportation, manufacturing, government, etc. The initial implementation of the system may not have all the knowledge necessary for this, but it will contain the needed framework for such a task.

Several advantages of the approach concern working with large numbers of images requiring less than expert analysis. Since the approach is based on a hierarchy of labels, there is no reason why it cannot stop analysis at any mid-point in the hierarchy. This is a great departure from other analysis systems whose goals are to always work to the finest

```
world
  ├─ culture
  │    ├─ urban (built up areas)
  │    │    ├─ industry/utility
  │    │    ├─ commercial/residential
  │    │    └─ institutional/governmental
  │    ├─ transportation/navigation
  │    │    ├─ railroads
  │    │    ├─ roads
  │    │    ├─ aeronautical/aerospace
  │    │    ├─ naval/marine
  │    │    └─ associated transportation features
  │    └─ landmark/rural features
  │         ├─ communication/transportation
  │         ├─ storage
  │         ├─ agricultural
  │         ├─ recreational
  │         └─ miscellaneous
  └─ landscape
       ├─ hydrography
       │    ├─ water
       │    └─ snow/ice
       ├─ physiography
       │    ├─ exposed soils (surface composition)
       │    └─ landforms
       └─ phytography
            ├─ cropland
            ├─ rangeland
            ├─ woodland
            └─ wetland
```

Figure 1.2 - Top Portion of Hierarchy

detail possible. For example, if a special interest group wanted to analyze a thousand images and separate them according to man made or natural, then the system could do this using quick and simple operators. Another case is classifying an image as a target or non-target image. Suppose a special interest group wanted to find the one or two images out of a thousand that contained airports. The system would work its way down the hierarchy as usual. As soon as the system branched down a section of the hierarchy that was out of the airport path, it would stop and label the image as non-target. If it worked its way down to the airport label, then it would label that image as target and hand it over to an expert photointerpreter for final analysis.

The remainder of this dissertation will expand upon these ideas and demonstrate their use. It concentrates on the high level approach. Provan [Provan 1987], [Provan 1988] took a track similar to this. He used a simplified puppet world image. This allowed him to concentrate on reasoning and a truth maintenance system instead of signal processing. This dissertation concentrates on the high level approach instead of image processing. The author created a computer vision system to address these issues and demonstrate a solution. Chapter four of this dissertation describes the implemented system. Chapter five presents the results of several experiments and analyzes the advantages of the system.

Chapter two of this dissertation takes a step back and looks at the general computer vision problem. It discusses why computer vision encompasses knowledge from many different fields ranging from psychology to software engineering. Some successes have been achieved in computer vision. Labeling aerial photographs, however, has not yielded much success. The chapter also reviews some early computer vision systems. The basic bottom-up, data-driven, rule-based approach was predominant in these early systems. This approach completely separates the segmentation and interpretation processes. The results were not good because the segmentation process produced faulty results. This was because of the nature of the data. The chapter ends with a review of the current use of perceptual psychology in low level operators. Perceptual properties form the basis of much of the work of this dissertation. This chapter examines how other researches have used such properties. Most current work uses perceptual or geometric grouping in low level operators. The goal is to find perceptually significant low level image features. Many early vision operators found mathematical features in images. These features, however, did not correspond to objects in the real world. Perceptually significant image features correspond to objects in the real world. The researchers discussed in chapter two have created several excellent, generic, low level operators. These could be incorporated into general computer vision systems. The final part of chapter two presents the perceptual properties that are the basis of this dissertation. It discusses selective attention, expectations, context, and the perceptual cycle. Other researchers use these principles to some extent in other work. Nevertheless, no one has ever used them together for a high level analysis approach.

Chapter three details the new analysis approach that this dissertation proposes. This chapter introduces the Defense Mapping Agency's hierarchy for labeling aerial photographs. It discusses an example that shows how using a hierarchy leads one through the analysis of an image. This process employs the principles given above and avoids the obstacles in aerial image analysis presented earlier. The chapter also discusses knowledge representation in computer vision. It presents frames as the logical choice for knowledge representation. Frames have several advantages over other knowledge representation schemes for computer vision. They allow expectation driven processing, are one of the corners of the perceptual cycle (see chapter five), are well suited for expressing hierarchies in images, express knowledge in an explicit and modular manner, and allow both procedural and declarative knowledge. Finally, the chapter discusses belief maintenance and the Dempster-Shafer (D-S) theory of evidence. Computer vision needs a belief maintenance system. A characteristic of computer vision is that different pieces of evidence are often uncertain, inadequate, and contradictory. Computer vision needs a system to pool different pieces of evidence and draw logical conclusions. The D-S theory has gained acceptance from several sources during the completion of this dissertation. Several references in the literature agree with this dissertation that the D-S theory works, is sound theoretically, and is not too computationally complex.

Chapter four examines the implementation of the new approach. The computer vision system created using this approach is simple and modular, but it required a significant software effort. The system was written in C with the highest priority given to portability and modifiability. During its development, it was moved among four computers using three different operating systems and four different compilers. It has proven that it is portable and that it can easily accept new and complex operators developed by other researchers.

The knowledge base clearly separates the spatial and spectral information. This allows movement from one image type to another by modifying only the spectral properties of basic materials (e.g. concrete, roofing, etc.). Chapter four discusses details of each part of the system.

Chapter five presents the results of the analysis of several images. The system analyzed two vastly different types of images and discusses the advantages of the approach. The first is a high resolution, multi-spectral, aerial image of an urban area. The second is three housing areas of a single channel, gray scale scan of a color aerial photograph. This chapter details how the system works its way down through the hierarchy of labels using the principles given earlier. The chapter concludes with a discussion of the advantages of the approach and the system built around it.

Chapter six draws conclusions about this work. It reviews what has been done and how this dissertation contributes to the field of computer vision.

2 - Background Material

This chapter reviews the computer vision problem. Computer vision requires skills and techniques from many disciplines. The goal of computer vision research is to teach a computer to see. Researchers have succeeded in some limited problem domains, but labeling aerial photographs has had only limited triumphs. The basic bottom-up, data-driven, rule-based approach has been tried extensively with poor results. Newer research has used perceptual psychology properties in low level operations. These efforts are promising, but they are only in the low level parts of the problem. The chapter closes with a discussion of the perceptual properties that form the basis of this dissertation's approach.

2.1 - The General Computer Vision Problem

The computer vision problem is the task of teaching a computer to "see." The process involves connecting a visual sensor to a computing system and having a computer program "recognize" what is given in the input data. The problem is extremely difficult. It involves knowledge from many different fields. The computer vision researcher needs knowledge of visual sensors, computer hardware and architecture, artificial intelligence and expert systems, image processing tools, psychology, and software engineering. This dissertation work proved that and is in agreement with [Li, Kender] and [Nicolin, Gabler].

Knowledge of visual sensors is a must. Sensors have the task of measuring and recording the scene. Next, a system must digitize the measured image for the computer. A sensor can be a simple black and white camera or a satellite. This dissertation works with airborne sensors. These range from small airplanes to the space shuttle.

Computers host vision systems. The researcher must be able to assemble a powerful and flexible computing system. The system requires special peripherals for image input and display. Many researchers are trying to build parallel processing systems to better implement vision algorithms.

Image processing forms the base for computer vision work. Image processing spawned from digital signal processing and extended that field to two dimensions. Many of the early image processing algorithms are still useful, especially in preprocessing or filtering noisy images.

Many useful techniques have come to computer vision from artificial intelligence research. Implementing knowledge bases and belief maintenance systems are two major examples. Other tools include languages such as LISP and PROLOG and some of their associated workstations.

In recent years, researchers have added perceptual psychology to the already long list of tools. Early operators often failed on images because they detected features that did not correspond to real world objects. Most operators select shadows, glare spots, and defects in surface materials as objects. Newer operators that use results from perceptual psychology can remove these false objects because the false objects are not perceptually significant.

Finally, since computer vision algorithms and systems are written in software, knowledge of and strict adherence to software engineering is essential. This is one area where unfortunately most research efforts fail. Most research efforts produce software that proves an idea or concept. The software, however, is usually not modular, portable, maintainable, or understandable by anyone other than the author. This prevents others from using the results of the research effort. Others must attempt to implement the algorithm on their own. Seldom if ever is the resulting software the same as the original author's.

Researchers are years, perhaps decades, away from solving the general computer vision problem. Researchers have met major problems in vision work. This has led researchers to simplify the problem by working on different, limited applications areas. In some applications researchers restrict the inputs to the vision system enough to achieve success. Examples of this are small robot manufacturing stations, circuit board inspection, and simple x-ray analysis. The area of labeling aerial images, however, has not had much success.

Nagao and Matsuyama [Nagao, Matsuyama] describe the problem of analyzing aerial photographs as constructing an overall description of a scene. This differs from statistical methods of image analysis which seek to label each point in a scene. Nagao and Matsuyama claim that the core of image understanding is knowledge based symbolic processing. A system uses any knowledge which may help in the analysis. Specific sources of knowledge are photographic conditions, intrinsic properties of objects, and contextual and semantic information.

Nagao and Matsuyama list four difficulties which arise in the analysis process: (1) the number of possible combinations of objects in a scene are immense, (2) it is a formidable task to organize the knowledge needed to analyze the scene, (3) low-level image processing operators are often inept, and (4) how to resolve the issue of top-down or bottom-up processing. Nagao and Matsuyama created an early aerial photograph interpretation system [Nagao, Matsuyama]. That system used a blackboard architecture for a knowledge base that pointed to rules in a production system. Each rule in the production system pointed to an object detection subsystem. The object detection subsystems were the low level operators that would actually locate the individual objects. Nagao and Matsuyama had sixteen different object detection subsystems. A problem with this was that each object detection subsystem lacked generality and power. They could only detect very specific objects. Nagao and Matsuyama used their system on very limited images. This was a good early system and led to work in the same area by many other researchers.

In 1982 Binford [Binford] performed a survey of vision systems. He listed the goals of a vision system to be high performance, generality, completeness, intelligence, ease of use, and system support. High performance means that the system should be able to analyze complex, real world scenes - not just laboratory images. Generality relates to analyzing any images that can occur whether it be indoors, outdoors, or aerial photographs. Completeness means the system should span all tasks. The system should be self contained and complete in itself. Intelligence means that the system should be able to reason on its own without human intervention. Easy to use is a common requirement for any computing system. System support gets back to software engineering. One person cannot create vision systems in one day. They require many persons and many years. The system must be structured and constructed to allow new techniques to be incorporated.

Binford found general shortcomings of the systems he reviewed. These were that they had severely limited context, were image dependent, and the low-level operators, especially the texture operators, were not sufficient. Binford discussed several vision systems of that time. He repeatedly pointed out that weaknesses were in the low level operators. The advent of low level operators using perceptual and geometric properties (see section two of this dissertation) has corrected some of these weaknesses. This dissertation proposes overcoming the problem of weak operators by using a better analysis approach to guide, direct, and interpret. Another repeated weakness was that the systems worked on only specific types of scenes. This dissertation proposes overcoming that problem by using an all encompassing hierarchy of labels that would include any aerial scene.

Levine [Levine, Shaheen], [Levine, Nazif 1984], and [Levine, Nazif 1985] has worked on outlining the basic structure of a vision system. We note his work here because many systems were created along the basic lines he established. He defines the objective of a computer vision system as to outline the objects in a picture and label them with an appropriate interpretation. Levine lists three requirements for computer vision systems. They are (1) extensibility, (2) modularity, and (3) separability. The system must accept new model and control information easily. The development of the system will last years and involve many people working independently. The knowledge and control systems must be modular enough to allow additions, deletions, and corrections without harming other parts of the system. The knowledge base must be separate from the analysis program. This enforces generality and forces one to create functions that are scene independent.

Levine divides the system into three components: (1) the Long Term Memory (LTM), (2) the Short Term Memory (STM), and (3) the Analysis Processors. Figure 2.1 shows a block diagram of the system. The STM contains the image data and analysis results. The LTM contains the semantic knowledge or the model of the scenes which the system can analyze. The Analysis Processors are a group of processors that operate on the image. Each processor specializes on a particular task. The processors continually update and change the contents of the STM. They use information in both the LTM and STM to activate themselves.

The low level processor analyzes raw picture data. The first low level processor was a region segmentation algorithm. The feature analyzer computes a set of attributes for the segmented picture. It then sends the results to the STM. The hypothesis initializer takes region descriptions from STM, uses knowledge from LTM, and generates possible interpretations for each region. The hypothesis verifier verifies and interprets each region label based on confidences. The focus of attention processor controls the order of analysis of regions. It recognizes the situation and begins action. The scheduler is responsible for

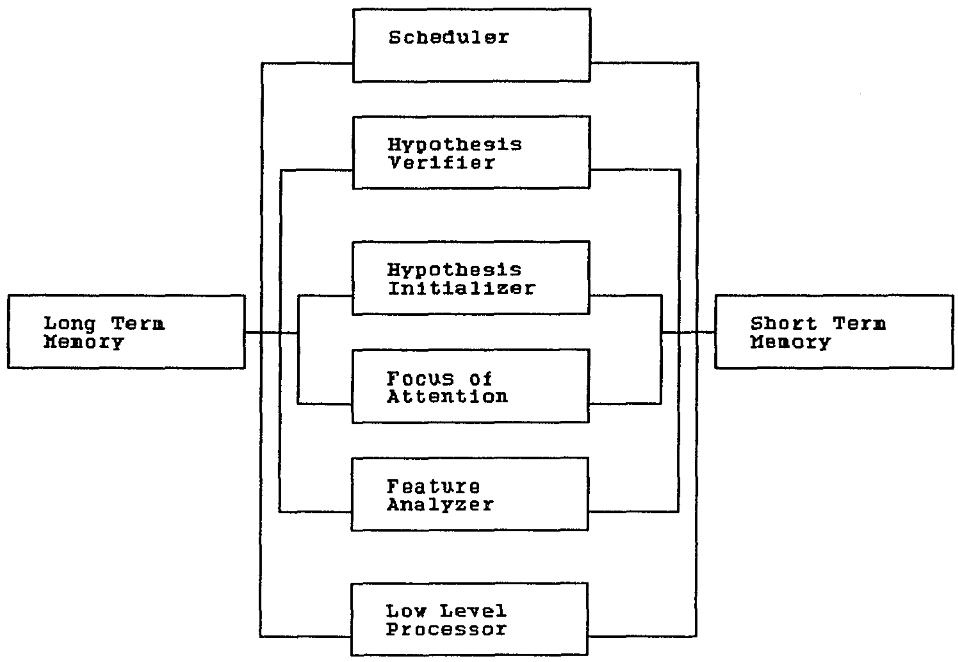

Figure 2.1 - Block Diagram of Levine's System

deciding which process should be started and when. In theory, the analysis processors could each be a separate physical computer processor. They would each have independent access to the memories and would examine the STM constantly and activate themselves. In reality, they are software and need a scheduler to determine which software process activates when.

2.2 - Early Computer Vision Systems

This section will describe several existing computer vision systems. The current literature describes many systems. This section will discuss only a small sample.

The Multi Spectral Image Analysis System (MSIAS) described in [Ferrante,Carlotto,Pomarede,Baim] is a rule-based system for labeling low-resolution satellite photographs. The goal is to label one pixel at a time as to its land use classification. One feature of MSIAS is that it organizes knowledge in a hierarchy as shown in figure 2.2. One advantage of a hierarchical structure is that the knowledge base is simpler. There only needs to be enough knowledge at each node to distinguish the children of that node. At the *vegetation* node, for example, there only needs to be enough knowledge to distinguish between *crops* and *other.* The node does not require knowledge about *silt* and *planted.* Another advantage is that each node has contextual information. When the decision reaches the *vegetation* node, this implies the pixel is not *soil* and is not *water.* The system does not express this information explicitly, but it does know and use it. MSIAS uses multi-spectral images. Each land class has unique properties in one or more spectral bands. A system can use multiple spectral bands to great advantage when analyzing aerial photographs.

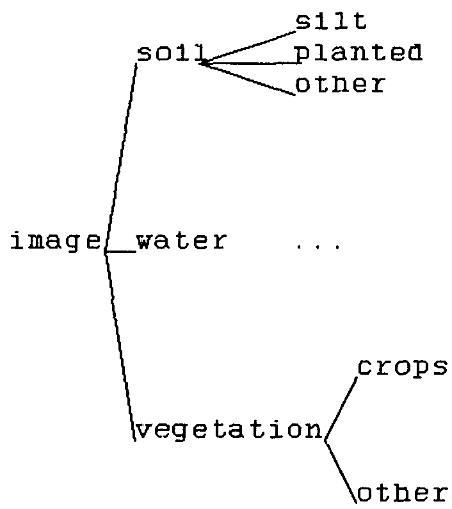

Figure 2.2 - MSIAS Hierarchy

The next three systems reviewed (LES, SPAM, and ANGY) all use the same basic data driven, rule-based, bottom-up approach. This technique is quite common among computer vision systems. There are two steps in this approach. The first is to apply scene independent image processing operators to segment the image and form a data base. The second is to apply a set of rules to the data base to interpret and label the image. The segmentation and interpretation processes are separate.

This approach is data driven because the data itself dictates the outcome of the segmentation in the first step. This outcome in turn dictates which rules to activate when interpreting and labeling the image. The data, therefore, drives or controls the entire process. This approach is rule-based because the knowledge base is a set of IF-THEN rules. One reason researchers use the rule-based approach is there are many rule-based expert system tools available. This approach is bottom-up because processing begins at the lowest level, the pixels. The system groups pixels into objects and then groups objects into regions during segmentation. This is the opposite of top-down processing which begins at the highest level, the image, and works its way down to individual objects and pixels.

Researchers at Lockheed have tried to analyze aerial photographs using their in-house, rule-based, general purpose expert system LES (Lockheed Expert System) [Perkins, Laffey, Nguyen]. The first stage of the system is a set of low level image processing operators which segment the image into atomic regions. The system then calculates a set of properties for each region. The rule-based stage of the system uses the properties of the atomic regions and tries to label each region. The system achieves only limited success.

SPAM [McKeown, Harvey, McDermott], [McKeown, Harvey], and [McKeown] is a system designed to label aerial photographs of airports. It is a rule-based system and uses mapping and airport design information as part of the knowledge base. One good point of this system is that it uses cartographic coordinates. Cartographic coordinates are of interest to map makers and map users. SPAM has a separate first stage that segments the image into atomic regions. The rule-based stage tries to label the regions. This system performs well if the image segmentation is done by human hand. The performance, however, degrades when it uses a computer segmentation.

ANGY [Stansfield] is a system which looks at angiograms. It is a rule-based system and operates like the two previously discussed systems. It has a low-level image segmentation section and a rule-based section which attempts to label the regions. The author of ANGY concludes with some candid and honest comments on this general approach to computer vision. Separating the segmentation operators from the rule-based interpretation system is a simple and fundamental idea. It comes from the principle of dividing a problem into smaller problems and solving them one at a time. Current segmentation operators, however, are not capable enough. The interpretation rule-base cannot work with the faulty results of the segmenters. In the future, the segmentation operators may have the expertise to allow a separate interpretation rule-base. Nevertheless, at this time this is not possible.

Niblack and others [Niblack, Petkovic, Damian] report another study of the rule-based approach. This study applied a rule-based system to the problems of circuit board inspection and analyzing satellite photos of ice flows. The approach worked well for the circuit boards, but performed poorly on the satellite photos. This is because of the regular, clean nature of circuit boards. The satellite photos proved too difficult. The authors concluded "Rule based methods can be useful but provide no fundamental breakthrough."

The data driven, rule-based, bottom-up approach does not work well. The low level image processing operators cannot process the image accurately enough to enable simple rules to label the image.

The strict bottom-up approach does not perform well when representing the knowledge by means other than simple rules. TESS, [Gilmore, Fox, Stevenson, Rabin], uses a hierarchy of frames as a knowledge base. This is an excellent method to represent knowledge, but TESS uses bottom-up, data driven processing before invoking the knowledge. The results are not promising.

Smyrniotis and Dutta [Smyrniotis, Dutta] described a system using mostly the top-down approach. The system directs a large set of image processing operators based upon overall knowledge of the image and the desired output. This system has several good qualities, but it is purely top-down or goal-driven.

The conclusion is that neither bottom-up processing nor top-down processing performs acceptably. Systems must use a combination of top-down and bottom-up processing. Uhr in [Uhr] and many other researchers agree with this. This relates to the relationship between segmentation and interpretation. Most research systems segment and then try to interpret. The results have not been good. Segmentation and interpretation cannot be separated. Several researchers that have addressed this question include [Kohl, Hanson, Riseman] and [Nicolin, Gabler]. The noisy nature of images [Haralick, Lee] and the limited ability of current operators [Huertas, Nevatia] preclude separating the segmentation and interpretation processes.

The segmentation processes lead astray the interpretation processes.

The question is - what type of combination of top-down and bottom-up processing will work? What is the proper relationship between segmentation and interpretation?

Chapter three will discuss this question and propose a solution.

2.3 - Current Low Level Use of Perceptual Psychology

This section discusses perceptual psychology and how some researchers use it in computer vision. Perceptual psychology deals with how a person's mind perceives or understands what he sees.

Recent publications have explored the importance and usage of perceptual psychology in computer vision. Researchers are working on discovering perceptually significant low level image features. The goal is to find practical, easy to use features and include them into general purpose computer vision systems. Another goal is to build systems that improve on the human visual and perceptual system [Hochberg] and [Hink, Woods]. Optical illusions easily fool humans. If a computer vision system is modeled exactly after a human, then optical illusions could fool it. The goal is to understand how illusions trick humans and use this knowledge to avoid having the computer vision system fooled. A common thread is to find properties that occur in the actual three dimensional object and also occur in the two dimensional image representation. A simple example is that a straight edge in a three dimensional object is represented by a straight line in a two dimensional image.

Biederman in [Biederman] was one of the first to approach the computer vision problem using perceptually significant features. He developed recognition-by-components (RBC). RBC is a proposal for a particular vocabulary of components. The key is how an arrangement of these components can access a representation of an object in memory. The goal in Biederman's system is to identify simple components of an object. Individual components are easier to identify than an entire object in a degraded image.

The properties on which Biederman concentrates are collinearity, curvature, symmetry, and cotermination.

> These ideas lead to Biederman's Principle of Componential Recovery. "If the components in their specified arrangement can be readily identified, object identification will be fast and accurate."

Walters [Walters 1986], [Walters 1987], [Walters, Krishnan] described another computer vision system built around perceptually significant features. A goal of Walters's research is to develop data driven, general purpose, generic algorithms that select perceptually significant features from an image. This is similar to the simple bottom-up approach described in the previous section. Walters, however, wants the results of the low level operators to agree with human perception. Walters's system works only with line drawings.

Walters has two uses for the generic algorithms. The first use is to enhance noisy or degraded line drawings. The second use is for segmentation of line drawings. Walters's system can divide line drawings into perceptually meaningful segments. The system is also able to group lines into objects. Walters performed several psychophysical experiments to determine significant low level features. The properties on which Walters's system concentrates are line length and types of connections between ends of lines.

Lowe created a computer vision system called SCERPO (Spatial Correspondence, Evidential Reasoning, and Perceptual Organization) [Lowe 1985] and [Lowe 1987]. SCERPO concentrates on analyzing images of three dimensional objects without the use of depth information. SCERPO is fairly successful at locating three dimensional objects in a line drawing.

The perceptual properties used by Lowe are collinearity, proximity, and parallelism. These properties meet two conditions set down by Lowe: the viewpoint invariance position; and the detection condition. The viewpoint invariance condition is that the perceptual features must remain stable over a wide range of viewpoints of some corresponding three-dimensional structure. The detection condition is to constrain the perceptual features so accidental instances are unlikely to arise.

Pentland in [Pentland 1986a], [Pentland 1986b], and [Pentland 1987] has developed a theory of part models. His goal is to find generic part models and use them to recognize the contents of an image as a combination of these primitives. Pentland's idea uses two concepts: representation and analysis. The representation concentrates on processes not models. The focus is on lumps of clay or superquadrics. There is a small number of primitives. The emphasis is on processes such as stretching, bending, twisting, and tapering. Given these and combinations of them, the system can represent the world by formative processes. The analysis is a simple global search of the models. The system can employ an exhaustive search because the number of primitives is small.

Chien and Aggarwal in [Chien, Aggarwal] attempted to recognize 3-D objects from single silhouettes. They concentrated on occluding contours and corners as the primary perceptual features. Perceptual psychology experiments showed that information concentrates in the occluding contour of a viewed object and at places where the contour changes most rapidly. The system used these features to recognize multiple objects with occlusion.

Weiss and Boldt in [Weiss, Boldt] apply perceptual psychology features in a low level generic operation. They use some of the perceptual organization techniques of Lowe in an edge detector. They begin with the traditional zero crossing method. They apply a hierarchy to join line segments using both geometric and intrinsic properties. The process is bottom-up. This general purpose low level operator uses the perceptual psychology properties of collinearity, symmetry, parallelism, proximity, repetition, and closure. This is an original. It is a generic, smart edge detector.

Mohan and Nevatia [Mohan, Nevatia] developed a system to detect buildings in aerial photographs. They concentrate on lines and perceptual grouping. The first step is a simple edge detector. The output of an edge detector usually contains too much false information. The system takes this output and groups the edge segments perceptually. The first grouping uses the edges. Next, the system groups parallel lines. Next comes 'U' shapes, and finally complete rectangles. This perception based, low level operator works quite well. It could be incorporated into large aerial image analysis systems.

Fua and Hanson [Fua, Hanson 1987] developed a system to detect several types of objects in aerial photographs. They use loosely defined shape models. Their operator starts with detecting edges. Next, the operator examines the area inside the edges. The loosely defined generic shape models possess predictive power. After an initial examination, the operator tries to verify the predictions. This allows a refining of parameters as the analysis progresses. This is quite similar to the perceptual cycle described later in this paper. Fua and Hanson do not specifically mention perceptual psychology properties. They do, nevertheless, employ them. The properties used are parallelism, perpendicular, closure, and collinearity. The result is a fine, generic, adaptive, smart object detector.

Similar work has been done by Harwood, Chang, and Davis [Harwood, Chang, Davis]. They first enhance an aerial image, then segment it. Next comes the similar process of adjusting parameters and optimizing on the objects they can find in the image. Then, they search for "missing" objects, i.e. objects that the initial parameters missed. Now that the parameters are optimal, they can locate objects missed during the initial analysis. The strategy employed in this low-level object detector improves the reliability.

Reynolds and Beveridge [Reynolds, Beveridge] worked on a class of algorithms for grouping collections of tokens into geometrically significant components. The first step is to segment the image and find the edges. The next step uses geometric segmentation and grouping algorithms. The output is a set of tokens that satisfy some geometric relations. They use the geometric or perceptual properties of collinearity, parallelism, relative angle, and spatial proximity. Spatial proximity divides into three parts. They are spatially proximate orthogonal, spatially proximate collinear, and spatially proximate parallel.

Huertas, Cole, and Nevatia [Huertas, Cole, Nevatia] applied their efforts to aerial photographs of airports. Like others, this group uses a hypothesize and verify low-level strategy. The system uses a four step process. The steps are (1) low-level segmentation, (2) hypotheses formation, (3) hypotheses verification, and (4) symbolic description. The results are good. This is just the lowest level of an overall system still under development.

Huertas and Nevatia [Huertas, Nevatia] worked on detecting buildings in aerial photographs. They used geometric (perceptual) models of buildings. The models included the properties of straight lines (collinearity), corners (perpendicular), sides (parallel), and box shape (closure). This low-level building detector is driven or triggered by a global goal, i.e. "find the buildings." It is, however, a data-driven operator. It performs well and would be a fine addition to any large system which analyzes aerial photographs.

Meisels and Bergman [Meisels, Bergman] describe their Rule-based Object Finder (ROF). As the title indicates, knowledge in the form of rules drives this low level segmenter. This is similar to [Levine, Nazif] with the addition that ROF uses context while [Levine, Nazif] was context independent.

This low-level operator (ROF) actually has several levels. These levels work back and forth on the image data.

Haralick and Lee [Haralick, Lee] employed a perceptual psychology property in a mathematical form. They used context in an edge detector. Their operator is different in that it is application or domain independent, but image dependent. Their operator performs a calculation over the entire test image and uses the result to find edges at different locations in the image. Each image, therefore, provides the context for local edge detection.

The research work described above is good. This is a young field and the results are promising. It is a little disturbing that so many different research efforts have discovered so many different "essential perceptual properties." This is no doubt a reflection of the subject - human perception. Every person sees with a unique set of eyes and we all have our own features that we notice.

Several papers have directly mentioned efforts at using combination top-down and bottom-up processing. These works are at the end of this section to better compare them to the ideas expressed in the next section.

Matsuyama [Matsuyama] reports on the SIGMA system. The SIGMA system is unique in that it uses reasoning at three levels. The system reasons by (1) the Low Level Vision Expert, (2) the Model Selection Expert, and (3) the Geometric Reasoning Expert. The Low Level Vision Expert reasons about the image segmentation processes. Given a goal such as "find a white rectangle 20x20," the Low Level Vision Expert reasons through the processes available to reach the goal. If one string of processes does not succeed, another string of processes tries. The Model Selection Expert reasons about the transformation between object models and their appearances. The Geometric Reasoning Expert reasons about structures of and spatial relations among objects. The Geometric Reasoning Expert applies a form of top-down and bottom-up processing. It uses top-down to direct the lower levels to search for objects. The low levels use bottom-up initially to find the objects. The Geometric Reasoning Expert then applies top-down processing again to search for "missing" objects. Missing objects are those objects not found by the low level operators, but all evidence indicates they should be present. For example, if the low level operator finds a row of houses with a large space between two houses, then there is probably a house in that space. The system calls the low level operator again using relaxed parameters. The SIGMA system, therefore, uses the combination processing to locate these missing objects.

Nicolin and Gabler [Nicolin, Gabler] produced a system that uses a little more top-down and bottom-up processing. They too use the combination processing to locate missing objects. They, however, use some combination processing in the early segmentation phase. The system performs some low level processing initially to bring out some cues. These cues guide or focus attention for initial region segmentation. This segmentation is somewhat goal-directed. After segmentation, structural analysis uses the Gestalt psychology properties of similarity, proximity, smooth continuation, symmetry, and familiarity. Structural analysis points out missing objects. Segmentation repeats itself with relaxed parameters to find the missing objects.

COBIUS (COnstraint Based Image Understanding System) is a system for aerial photograph interpretation [Kuan, Shariat, Dutta, Ransil]. The authors outline three major problems in the field. They are (1) generic domain object representation, (2) unreliable image segmentation, and (3) knowledge control. They attack the first problem by representing objects with a hierarchy of frames going from coarse to fine detail. They attack the second problem using a "multiple feature fusion approach with model-based feature verification capability." They attack the last problem using a form of the Dempster-Shafer theory to pursue the most probable hypothesis first. The multiple feature fusion approach is the combination top-down and bottom-up idea. First, a coarse segmenter provides an initial interpretation. Next, they resegment looking for expected objects and attempting to verify their object models. This is part of the perceptual cycle discussed in the next section. COBIUS uses some of the ideas described later in chapter three of this paper. It uses top-down and bottom-up processing, selective attention, and to some degree expectations.

Much of the work reported in this section concerns tactics. They are working on low level perceptually significant features. The next section deals with the main point of this research effort - *strategy.* What are the perceptually significant properties that guide the overall analysis of an image?

2.4 - Selective Attention, Expectations, Context, and the Perceptual Cycle

Certain perceptual psychology properties control the analysis of images. The projects described in the previous section covered the use of perceptual psychology in low level tactics. The subject of this research effort is the use of perceptual psychology in high level strategy. This section discusses selective attention, expectations, context, and the perceptual cycle.

A review of psychology books and papers has uncovered the critical points listed in the following paragraphs. These points influenced the creation of the computer vision system described later. Humans use selective attention [Goldstein], [Neisser 1967]. We direct our attention to only those items which interest us and ignore all else. Part of the reason for this is that humans can only handle seven bits of information [Miller] at a time. We simply cannot simultaneously perceive all of the objects and information in an image. We cannot handle all the information present. We selectively focus our attention on the things which interest us. Another reason is the physical graded resolution of the human eye [Browse, Rodrigues].

Expectations play an important role in perception because they guide our attention [Lindsay, Norman]. All of our experiences lead us to expect types of objects. An example is a new restaurant. You expect to see some type of chairs, tables, and menus and they draw your attention. The use of expectations is growing in the current literature. You will find instances in [Draper, Brolio, Collins, Hanson, Riseman], [Burt], and [Haralick, Lee]. Context works closely with expectations in driving human perception [Bruner, Minturn]. You do not expect to see tires, spark plugs, and motor oil in a restaurant. You, therefore, do not look for them and do not burden your mind with them.

Human perception operates in a cycle [Neisser 1976] and [Rao, Jain]. Figure 2.3 shows the perceptual cycle. The flow is counterclockwise. The object is what we are perceiving. Schema is the knowledge stored. Exploration is looking at the object or experimenting.

Start the cycle at the schema (the knowledge base). Knowledge directs exploration. You cannot initially explore without some known basis. The exploration takes samples or

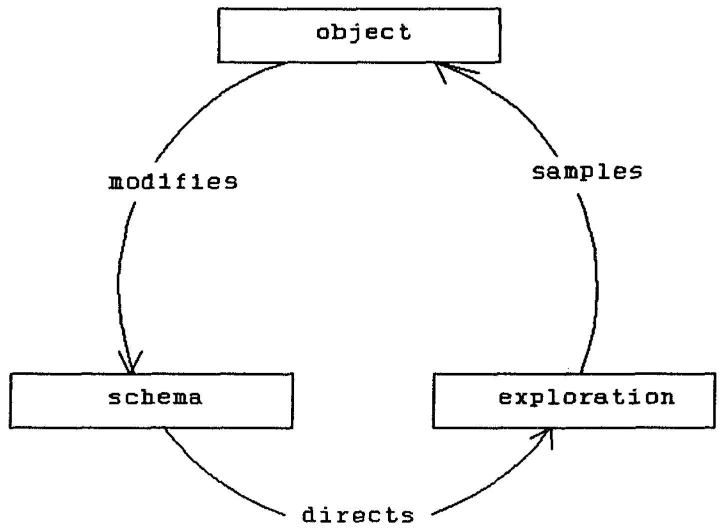

Figure 2.3 - Perceptual Cycle

observations of the object under study, i.e. it performs calculations on the image. Exploration of the object produces some result which modifies the knowledge. The modified knowledge directs the exploration in a different direction. The new exploration samples the object. The result of exploration modifies the knowledge... This is a simple closed loop feedback system. Though not specifically mentioned, several research efforts such as [Fua, Hanson] and [Kohl, Hanson, Riseman] use the perceptual cycle.

Because of these concepts, it is obvious that humans do not spend equal time, energy, and effort on all objects in a scene. Why should a computer vision system do that? That was the case in the data driven, bottom-up systems described in section 2.2. Those systems devoted equal processing time and effort to each pixel in the image. This is not how human visual perception works.

A computer vision system should be guided and controlled so it spends most of its processing time and effort on only the critical portions of the image. The knowledge base and control mechanism should focus attention and time [Koons, McCormick], [Arkin, Riseman, Hanson], [Lehrer, Reynolds, Griffith February 1987], and [Ballard, Ozcandarli]. The system should commit resources to only those objects and properties that are necessary for object and region identification. In agreement with this idea are [Clark, Ferrier], [Sha'ashua, Ullman], and [Kuan, Shariat, Dutta, Ransil].

Selective attention, context, expectations, and the perceptual cycle are not original ideas. Many papers mention these concepts and use them in various ways. It is, however, an original concept to use the four together as a global strategy and use them to tie together the segmentation and interpretation processes.

The next section discusses the issue of control strategy in more detail. That section discusses top-down and bottom-up processing. Those discussions are closely related to the above conclusions concerning perceptual psychology in computer vision.

3 - Introduction to the New Analysis Approach

This chapter introduces the new approach to analyzing high resolution aerial photographs. It talks about the approach in terms of perception, a hierarchy of labels, and the interplay between segmentation and interpretation processes. Next, the chapter discusses the knowledge base to be used by the approach. The chapter closes with a look at the Dempster-Shafer theory of evidence and how it is modified for use in the approach.

3.1 - A New Approach to Analyzing High Resolution Aerial Photographs

Researchers agree that some combination of bottom-up and top-down processing is appropriate for computer vision. The question is - what combination? This section proposes an answer to that question.

The approach presented in this section is the heart of the originality of this work. The approach uses three basic ideas in a unique manner. The three ideas are (1) the use of selective attention, expectations, context, and the perceptual cycle, (2) analyzing the image through a hierarchy of increasingly specific labels, and (3) the interplay between the segmentation and interpretation processes. These ideas translate into the computer vision system mentioned briefly in chapter one and described in detail in chapter four.

Computer vision systems should make use of selective attention, expectations, and context. The previous chapter discussed these perceptual properties. Context limits the number of possibilities which confront the computer vision system at a given moment. The system must use context at each moment in the analysis to describe the alternatives, limit the number of realistic alternatives, and select the proper option. Expectations guide the low level processing. When the system expects a group of buildings in an area, then it should use a specific building detector to search for the group. Selective attention focuses the processing to only the essential portions of the image. This ties closely to the use of expectations. There is no reason to process linear features while looking for buildings. Selective attention points the operators at only what is necessary.

Several of the systems described earlier in chapter two use these three concepts to some degree. Those systems use the concepts for low level operators such as edge detectors or building finding operators. This dissertation proposes to use them in the high level, overall, analysis approach. They are not used in the low level operators here. Using them in both the high and low level approaches is an excellent idea, but again developing high quality low level operators is not the point of the dissertation.

A key to the overall approach is the hierarchy of labels. The hierarchy of labels allows the system to work its way gradually to the point where it knows the type of scene in the image. At this point, it is ready to identify the pertinent objects. Each node in the hierarchy has only a few possible choices. This greatly reduces the complexity of the problem.

There are many types of hierarchies used in vision research. Several of the references mentioned earlier used hierarchies to reduce the complexity of different parts of the vision problem. An example hierarchy would be of the parts of a residential area when viewed from an aerial photograph. The residential area has sub parts roads and blocks. The blocks have sub parts lots and walkways. Each lot has house, yard, tree, sidewalk, driveway, and swimming pool.

The hierarchy used in this project gives labels for areas in an aerial photograph. Figure 3.1 shows the top portion of the hierarchy. This hierarchy was given by the Defense Mapping Agency (DMA) [Rusco]. The DMA uses this hierarchy to map the world and label aerial photographs. Chapter four repeats this figure and appendix 2 lists the entire hierarchy. The DMA hierarchy works well for most of the image analysis task. Nevertheless, it has problems at the bottom of the hierarchy. It does not include a typical vision hierarchy such as the ones mentioned above. At the node *commercial/residential*, the DMA hierarchy does not include a vision hierarchy of roads and blocks, then lots, then yard, house, sidewalk, driveway, and tree. This shortcoming is more evident during the system description in chapter four and the example analysis of images in chapter five. Chapter four gives a brief description of how to remedy this situation.

A vision system must strongly couple the processes of segmentation and interpretation [Kohl, Hanson, Riseman]. The system cannot separate the two. The perceptual cycle (introduced in the previous chapter) requires close interaction between the two processes. Although not mentioning it specifically, several other researchers agree with this concept and are attempting to use it. Those researchers, however, are using it to improve the accuracy of the low level operators. They should apply this same idea in the overall approach.

The following paragraphs give an example of how to use these concepts in the analysis of an image. Figure 3.1 shows the top portion of the hierarchy used in the following discussion. Processing begins at the *world* node. Because of context, the computer vision system only has to distinguish between a *culture* area and a *landscape* area. This is a simpler, more solvable goal than "find all the objects in the image and label them from among 200 different possibilities."

The first step is data driven. The computer vision system invokes a low level operator that will return a sign of *culture* or *landscape*. The interpretation process drives the segmentation process. The data driven operation will not determine the final selection. It is only a first step. The operator does not have to be a single, general purpose operator. Because of selective attention, the computer vision system selects a simple operator that is appropriate to distinguish between *culture* and *landscape* areas.

The second step is goal driven. The computer vision system uses the quick impression given by the first step as a goal. Now expectations drive the analysis of the area under investigation. Suppose step one returned a guess of *culture*. The goal of the second step would be *culture* as opposed to *landscape*. The system uses appropriate operators that look for objects that are in *culture* areas and not in *landscape* areas. The segmentation process has altered the interpretation process which will now steer the segmentation process in a new direction.

Because of selective attention, the second step does not spend equal time processing all objects in the area. It only processes those objects that are appropriate to this very limited situation. The interpretation process focuses the segmentation process.

If the results of step two agree with those of step one, then label the area as *culture* and move down the hierarchy. This is because both bottom-up and top-down analysis agreed that the area is *culture*. If the results of step two disagree with those of step one, then repeat step two using *landscape* as the goal.

Suppose analysis labels the area *culture*. The next step must decide between *urban (built up areas)*, *transportation/navigation*, and *landmark/rural features*. Once again, context reduces the complexity of the task. Expectations produce reachable goals. Selective

```
world
 ├─ culture
 │   ├─ urban (built up areas)
 │   │   ├─ industry/utility
 │   │   ├─ commercial/residential
 │   │   └─ institutional/governmental
 │   ├─ transportation/navigation
 │   │   ├─ railroads
 │   │   ├─ roads
 │   │   ├─ aeronautical/aerospace
 │   │   ├─ naval/marine
 │   │   └─ associated transportation features
 │   └─ landmark/rural features
 │       ├─ communication/transportation
 │       ├─ storage
 │       ├─ agricultural
 │       ├─ recreational
 │       └─ miscellaneous
 └─ landscape
     ├─ hydrography
     │   ├─ water
     │   └─ snow/ice
     ├─ physiography
     │   ├─ exposed soils (surface composition)
     │   └─ landforms
     └─ phytography
         ├─ cropland
         ├─ rangeland
         ├─ woodland
         └─ wetland
```

Figure 3.1 - DMA Hierarchy (top portion)

attention focuses operators to only those objects of interest.

```
world
|
|__ culture
    |
    |__ urban (built up areas)
        |
        |__ commercial/residential
            |
            |__ 039 commercial buildings
            |
            |__ 040 apartments/hotel
            |
            |__ 041 house (single family dwelling)
            |
            |__ 042 mobile homes
            |
            |__ 043 barracks
            |
            |__ 044 display sign (large
                     billboard, highway sign)
```

Figure 3.2 - Sample Path Through Hierarchy

The analysis process continues down through the hierarchy until the analysis reaches either the bottom of the hierarchy or the desired level of analysis. (Note that the system user does not have to direct the system to locate the lowest level objects in an image. The analysis can stop at any desired level.) At the bottom of the hierarchy, the system locates and labels individual objects. For example, using the complete hierarchy given in appendix 2, figure 3.2 shows one path through the hierarchy. The objects the system will locate and label are all different types of buildings. Size and shape distinguish them. The system calls a low level operator that specializes in detecting buildings (possibly one of the excellent operators described in chapter two). Once again, context, expectations, and selective attention simplify the task. In this context the system only expects to find buildings and selective attention narrows the focus to a special operator. The system does not need a super operator capable of detecting any object. This situation is possible because of the hierarchy of labels.

3.2 - Knowledge Representation - Frames and Expectation Driven Processing

This section discusses knowledge representation in computer vision. In particular it examines frames and expectation driven processing. A major question in any computer vision system is how to represent the knowledge needed to analyze an image. Chapter two pointed out the inadequacy of the simple rule-based system. The thoughts on perceptual psychology in chapter two pointed to context, expectations, selective attention, and the perceptual cycle. The frame is a logical choice of knowledge representation for a system using these concepts.

Minsky [Minsky] originated the concept of frames. Many others including [Rich], [Gevarter], [Barr, Feigenbaum], [Rao, Jain], and [Neisser 1976] also described frames. Frames are complex data structures that contain information describing objects and relations that are appropriate to a given situation. Frames provide a structure or framework for expectations given the context of the situation. Frames focus or select attention on the things that should occur.

Frames are one of the corners of Neisser's perceptual cycle. The perceptual cycle was described earlier and figure 3.3 shows it again. Neisser describes them as anticipatory schemata. The term schemata describes a frame in a visual context. The schemata contains plans for perceptual action as well as readiness for a particular optical structure. The schemata is that part of the perceptual cycle that is internal to the observer. The observer's experience modifies the schemata. The schemata controls the activity of looking. The schemata determines what will be perceived.

Frames have several advantages over other knowledge representation schemes in computer vision. Among these advantages are that frames encourage expectation driven processing. Frames can express the hierarchical nature of images. Frames express

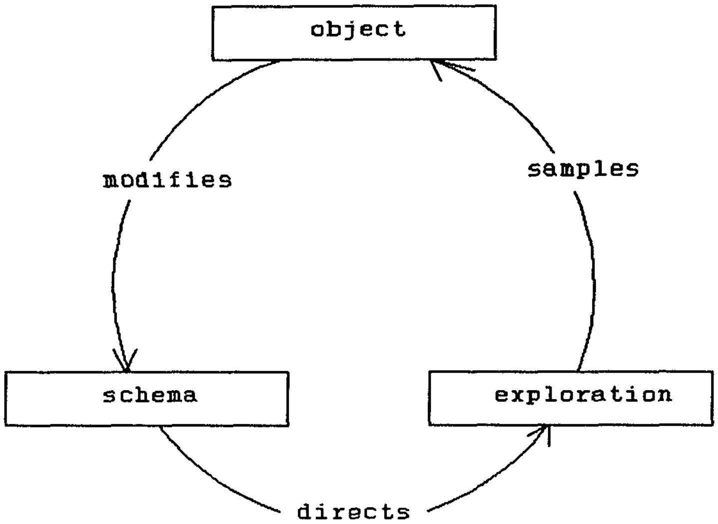

Figure 3.3 - The Perceptual Cycle

knowledge in a more explicit and modular manner than other representations. Frames are a good compromise between procedural and declarative forms of knowledge.

The primary mechanism of the frame is the slot. A slot is a blank field in the frame that fills when an expected object or property is discovered. Because of this, frames have earned the name slot and filler structures. Figure 3.4 shows an example frame for a restaurant. A person entering a restaurant for the first time expects to see each of these slots. The person fills the slots with the appropriate answers as he observes objects. In this respect, frames evolve from general to specific instantiations during the analysis process.

Restaurant Frame
  Type of: service industry
  Menu: plastic book, place mat, billboard, ...
  Serving place: covered table, picnic table, bar, ...
  Type of food: traditional, Italian, Chinese, ...

Figure 3.4 - Example Frame

Shown in figure 3.5 is the general structure of a frame used in the computer vision system described in chapter four. Frames encourage expectation driven processing. The is_a slot in the frame names what the system expects to find. The is_a becomes the goal of analysis. If analysis satisfies the goal, then the is_a is no longer a goal but a reality. If analysis does not satisfy the goal, then the frame is replaced by an alternate and the alternate is_a becomes the new goal.

```
is_a ...
is_part_of ...
level_in_tree ...
goal_of_analysis ...
intrinsic_characteristics ...
        distinguishing characteristics and
        assertions of belief
sub_node_names ...
sub_node_operator ...
```

Figure 3.5 - General Frame from TDBU System

Frames express well the hierarchical nature of images. You can describe images naturally by hierarchies. The DMA hierarchy shown in chapter three describes the image in general terms at the top levels. The description becomes more specific as you move down to lower levels. The is_part_of and level_in_tree slots connect the frame with the frames above it. The sub_node_names and sub_node_operator slots connect the frame with the frames below it. Using these slots, the system moves up and down through the hierarchy.

Frames express knowledge in a more explicit and modular manner. The frame keeps all of the information concerning an entity in one place. The knowledge in frames is readable and understandable. Rule-based systems spread the information over many separate rules. The rules are often difficult to read and understand. Modifying the rules is especially difficult because you must first find many different rules and then modify them consistently.

Frames are a good compromise between procedural and declarative forms of knowledge. The is_a, is_part_of, level_in_tree, goal_of_analysis, and sub_node_names slots are declarative. They declare facts. The distinguishing_characteristics and sub_node_operator slots are procedural. They invoke attached procedures, i.e. they call procedures that operate on the image.

3.3 - Belief Maintenance Systems - The Dempster-Shafer Theory of Evidence

This section discusses belief maintenance and the Dempster-Shafer theory of evidence. It describes how the Dempster-Shafer theory works and concludes that it is appropriate for computer vision.

A belief maintenance system must take information from different sources at different times, pool this information, and draw a reasonable conclusion. There are two questions in belief maintenance. Is a belief maintenance system needed? If so, what system should you use? The first question may seem out of place, but it is valid. Any number of texts and handbooks on artificial intelligence and expert systems exist that do not mention the subject. At the same time, volumes have been written about the subject.

Computer vision systems need a belief maintenance system because of the inherent uncertainty and inadequacy of the individual pieces of evidence. In many expert systems the reasoning process hinges on a single piece of evidence. These systems do not require a belief maintenance system. The labeling of areas in an aerial image, however, requires pooling individual pieces of evidence. The individual pieces of evidence can be uncertain, incomplete, incorrect, and often contradictory [Lowrance, Garvey], [Wesley], [Rao, Jain], and [Hink, Woods].

The computer vision system described in chapter four uses the Dempster-Shafer (D-S) theory of evidence described by [Shafer 1976], [Lowrance, Garvey], and [Gordon, Shortliffe 1984 and 1985]. Several hundred subroutines comprising several thousand lines of code were written to implement a form of the D-S theory of evidence as part of the computer vision system. This theory performs well, is not computationally complex, and has a sound theoretical basis.

A survey of belief maintenance systems was performed by [Goldberg, et. al.], [Goodenough, et. al.], and by [Cheng, Kashyap]. These surveys favored the D-S theory. Zadeh [Zadeh] also reviewed the D-S theory and concluded that it was appropriate for use in expert systems. Other researchers favoring the D-S theory include [Lee], [Lehrer, Reynolds, Griffith February 1987], [Lehrer, Reynolds, Griffith June 1987], [Lee, Shin], and [Stephanou, Lu].

In the D-S theory of evidence, the set of all hypotheses that describe a situation is called the frame of discernment. The letter Θ denotes the frame of discernment. The hypotheses in Θ must be mutually exclusive and exhaustive.

There are two properties of the D-S theory to note. First, the D-S theory allows one to assign belief not only to single hypotheses, but also to subsets of hypotheses. Second, one can represent ignorance by assigning belief to the union of all the basic hypotheses. The expert can be vague in the early stages of analysis. This is done by assigning belief to subsets of hypotheses. This enlarges the set of possible interpretations to 2^{Θ}. The expert can narrow his assertions about the problem later when more specific evidence is available.

The assignment of belief to ignorance allows an expert to delay judgment about a problem until he acquires adequate evidence. This mirrors the human tendency to procrastinate. It allows the expert to express doubt and wait until further evidence appears before becoming more specific in the reasoning process.

There are several methods to describe the D-S theory, but the easiest to understand is that used by [Gordon, Shortliffe 1984 and 1985]. This method uses actual examples as opposed to mathematical theory. Please refer to [Shafer 1976] and [Shafer 1985] for more theoretical discussions. Consider the situation where there are three hypotheses, A, B, and C. There would be eight subsets of hypotheses in the frame of discernment as shown in figure 3.6.

(A,B,C)

(A,B) (A,C) (B,C)

(A) (B) (C) (NULL)

Figure 3.6 - Example Frame of Discernment

The first subset (A,B,C) corresponds to the hypothesis A or B or C. Since this hypothesis includes all three basic hypotheses it distinguishes nothing. This is how the D-S theory represents ignorance. The first three hypotheses in the bottom row of the hierarchy, (A), (B), and (C) are the basic hypotheses. They are singletons. The final set in the bottom row (NULL) is the empty set and corresponds to the hypothesis known to be false. The belief in (NULL), therefore, must always be zero. The NULL set normalizes the combination of two assertions.

Assertions of belief are basic probability assignments (bpa). A bpa represents the impact of a piece of evidence. It is a generalization of the Bayesian probability density function. The bpa is more general because it can assign degrees of belief to all of the subsets in the frame of discernment - not just to the singletons. The degrees of belief must sum to 1.0. An example clearly demonstrates a bpa. Suppose there is a piece of evidence that supports hypothesis (A or B) and also supports hypothesis (A). The bpa (represented by m for measure of belief) for this piece of evidence might be:

$m(A,B) = 0.6 \quad m(A) = 0.3 \quad m(\Theta) = 0.1$

The quantity m(A,B) is the portion of total belief committed exactly to the subset (A or B). In the same manner m(A) and m(B) represent the portions of total belief committed exactly to (A) and exactly to (B). The sum m(A) + m(B) + m(A,B) represents the total portion of belief committed to (A or B) and is denoted by Bel(A,B). The quantity $m(\Theta)$ represents ignorance. In this case $m(\Theta)$ represents the subset (A or B or C).

Dempster's rule of combination provides the means to combine two bpa's. This allows the system to pool assertions from multiple pieces of evidence and draw a conclusion. The combination rule employs an intersection tableau. Given two bpa's shown in figure 3.7, an intersection tableau is constructed with the first bpa across the top and the second bpa down the left side as shown in figure 3.8.

$m1(A,B) = 0.8$
$m1(\Theta) = 0.2$

$m2(B) = 0.7$
$m2(C) = 0.2$
$m2(\Theta) = 0.1$

Figure 3.7 - Two Basic Probability Assignments

The subsets inside the tableau are the intersection of the subsets along the top and down the side. The intersection of (A,B) and (B) is (B). The value given to (B) in the upper left corner of the tableau is the product of the subsets (A,B) and (B). The other subsets and values inside the tableau are obtained in the same manner.

		m1	
		(A,B) 0.8	(THETA) 0.2
	(B) 0.7	(B) 0.56	(B) 0.14
m2	(C) 0.2	(NULL) 0.16	(C) 0.04
	(THETA) 0.1	(A,B) 0.08	(THETA) 0.02

Figure 3.8 - Intersection Tableau

Note the NULL set in the second row inside the tableau. The intersection of (C) and (A,B) is NULL. As mentioned earlier, the belief attributed to the NULL set must equal zero. This value will normalize the other beliefs. Let a value K equal the sum of all NULL sets in the tableau. To remove the belief attributed to NULL, you sum the other values in the tableau and divide them by 1 - K. This yields the result shown in figure 3.9. Notice how the combination of the two bpa's has narrowed the hypothesis set. The first bpa pointed to the subset (A,B). The second bpa pointed to the singleton (B). Combining the two bpa's narrowed the decision to (B).

The D-S theory has proven its value, but what about using it in the computer vision field? Does it fit into computer vision systems? In expert system terms, the computer vision system in this dissertation is an analysis system. The term analysis system contrasts with diagnostic system (diagnosing a problem such as a bad car engine) or advisory system (giving advice in a field such as a financial consultant). Other analysis systems could be in

K = 0.16

1 - K = 0.84

m1(A,B) + m2(A,B) = 0.08/0.84 = 0.96

m1(B) + m2(B) = (0.56 + 0.14)/0.84 = 0.832

m1(C) + m2(C) = 0.02/0.84 = 0.024

Figure 3.9 - Result of Intersection Tableau

geology (this is an unknown rock, analyze it and list its characteristics) or military (here are the physical specifications of a weapon, analyze them and list the capabilities and limitations of the weapon). The difference in the computer vision system is in the evidence. In computer vision, the computer must derive the evidence by itself using less than perfect operators. There is more ignorance (more belief attributed to Θ) here than in other fields. This means the decisions can be less conclusive. Therefore, the system requires more pieces of evidence to reach a conclusion with the same high degree of certainty.

A problem with the D-S theory involves the number of computations involved in the combination of two bpa's. The above example is a trivial case. There are three singletons and only 2^3 or 8 total hypotheses. If there are 100 singletons, then there are 2^{100} total hypotheses and the combination of two bpa's becomes intractable. Researchers have worked on reducing the number of computations when the hypotheses are in a hierarchy [Gordon, Shortliffe 1985], and [Shafer 1985]. The best way to explain the reduction in computations is with an example.

Shown in figure 3.10 is the top portion of the DMA hierarchy. It is important to realize that the label *culture* is the subset *(urban (built up areas)* or *transportation/navigation* or *landmark/rural areas*). The label *landscape* is the subset (*hydrography* or *physiography* or *phytography*). The label *world* is equal to Θ. When the hypotheses are in a hierarchy there are two major simplifications you can make. (1) The number of meaningful hypotheses is a small percentage of the total hypotheses. (2) The number of meaningful bpa's you can make is also a small percentage of the total possible and context limits them.

The hierarchy shown in figure 3.10 has six singletons - *urban (built up areas),*

```
world
 ├─ culture
 │   ├─ urban (built up areas)
 │   │   ├─ industry/utility
 │   │   ├─ commercial/residential
 │   │   └─ institutional/governmental
 │   ├─ transportation/navigation
 │   │   ├─ railroads
 │   │   ├─ roads
 │   │   ├─ aeronautical/aerospace
 │   │   ├─ naval/marine
 │   │   └─ associated transportation features
 │   └─ landmark/rural features
 │       ├─ communication/transportation
 │       ├─ storage
 │       ├─ agricultural
 │       ├─ recreational
 │       └─ miscellaneous
 └─ landscape
     ├─ hydrography
     │   ├─ water
     │   └─ snow/ice
     ├─ physiography
     │   ├─ exposed soils (surface composition)
     │   └─ landforms
     └─ phytography
         ├─ cropland
         ├─ rangeland
         ├─ woodland
         └─ wetland
```

Figure 3.10 - DMA Hierarchy (top portion)

transportation/navigation, landmark/rural areas, hydrography, physiography, and *phytography.* In the pure D-S theory there would be 2^6 or 64 subsets or hypotheses. Most of these subsets have no meaning in the labeling process so you can drop them. For example, the subset (*urban (built up areas)* or *hydrography* or *phytography*) has no meaning. There is no reason to make computations on its behalf. There are only 9 meaningful hypotheses in the hierarchy - 1 meaningful hypothesis for each node in the hierarchy.

The context and form of bpa's will be limited to small portions of the hierarchy. If the analysis is at a node and you assert a bpa about its sub-nodes, then that bpa will mention only one sub-node and its complement within the other sub-nodes. For example, if the analysis is at the *culture* node figure 3.11 shows a possible bpa. The m^(not *urban (built up areas)*) is not the true complement of *urban (built up areas)* - it is the simplified complement. The simplified complement of a node is the union of the other nodes with the same parent node. (The term simplified complement is original to this paper.) Figure 3.12 shows the difference between the true complement and the simplified complement. The true complement of *urban (built up areas)* has no meaning. The simplified complement, however, has meaning in the context of *culture.* It means "give belief to everything but *urban (built up areas)* under the context of *culture.*"

m(urban (built up areas) = 0.7

m^(not urban (built up areas)) = 0.2

m(THETA) = 0.1

Figure 3.11 - Example bpa

The above two simplifications greatly reduce the computations in the combination of two bpa's. The research of [Gordon, Shortliffe 1985] showed that using a knowledge of the hierarchy and modifying the D-S could reduce the computational load. Shafer, in [Shafer 1985] and [Shafer 1987], demonstrated that the unmodified D-S theory is usable with reduced computations given the special conditions mentioned above.

The computer vision system described in chapter four uses a form of the D-S theory to maintain beliefs in the analysis of aerial images. Since the system organizes knowledge in a hierarchy, it can use the simplifications mentioned earlier. This required creating a large set of custom subroutines to implement the D-S theory and the intersection tableau. In itself, this was a major research and programming effort.

If the knowledge base were not organized in a hierarchy, the size and complexity of the problem would be much greater. Taking the DMA hierarchy and using the labels one level up from the leaf nodes results in a flat list of area labels. The labels would include *commercial/residential, naval/marine, agricultural,* and *disposal.* Using the flat list of labels instead of the hierarchy, there would be 29 labels and no simplifications. In this case Θ is 29 and there are 2^{29} or 536,870,912 possible combinations. This explodes the computational complexity of the problem. The D-S theory would be unusable and its advantages would be lost.

4 - Implementing the New Approach

This chapter examines the TDBU (Top-Down Bottom-Up) computer vision system. TDBU is a computer vision system that analyzes multi-spectral, high resolution, aerial images of urban areas. This work began with [Harlow, et. al.] and continued in [Phillips].

A Kaypro 286i (IBM-PC AT compatible) currently hosts TDBU. The computer has a 40 Mbyte hard disk and 640K of memory. TDBU is written entirely in C and is currently using the Microsoft C compiler version 5.0. This is the fifth different compiler and fourth different machine used during the development of TDBU. It is a portable system.

Figure 4.1 shows the basic structure of this system. In general, the system uses the knowledge stored in frames to investigate areas in the image. The control mechanism calls low level routines in the image processing interface. They report the results back to the control mechanism which invokes the belief maintenance system. The belief maintenance system reports which frame is the most probable label for the area under investigation. The control mechanism either uses the most probable frame for further investigation or for labeling the area.

The control mechanism implements the cyclical nature of the analysis. It does this with a simple closed loop algorithm. This is the interpretation portion of the interpretation and segmentation interplay. It interprets results from the image processing interface and belief maintenance system and then redirects the segmentation. The control mechanism also guides the system down through the hierarchy. When analysis decides on an area label, the control mechanism moves the system down to the next level in the hierarchy.

The knowledge base is the hierarchy of frames. Figure 4.2 shows the top portion of the hierarchy. Section 3.2 discussed the advantages of using frames in knowledge representation in computer vision systems. The hierarchy brings context into the processing. Each node in the hierarchy contains implicit knowledge based on its position in the hierarchy. The frames imbed expectations. The slots in the frames contain what the system expects to find at each point in the analysis.

The belief maintenance system aids the control mechanism in drawing a reasonable conclusion. It pools evidence from a group of less than perfect operators and sends the result to the control mechanism. The system uses a form of the Dempster-Shafer theory of evidence. It comprises a large set of subroutines that implement the intersection tableau as described in section 3.3. These subroutines use and implement the computational simplifications available from using a hierarchy of hypotheses (labels or frames). They greatly limit the amount of computations needed to use the D-S theory and they were a

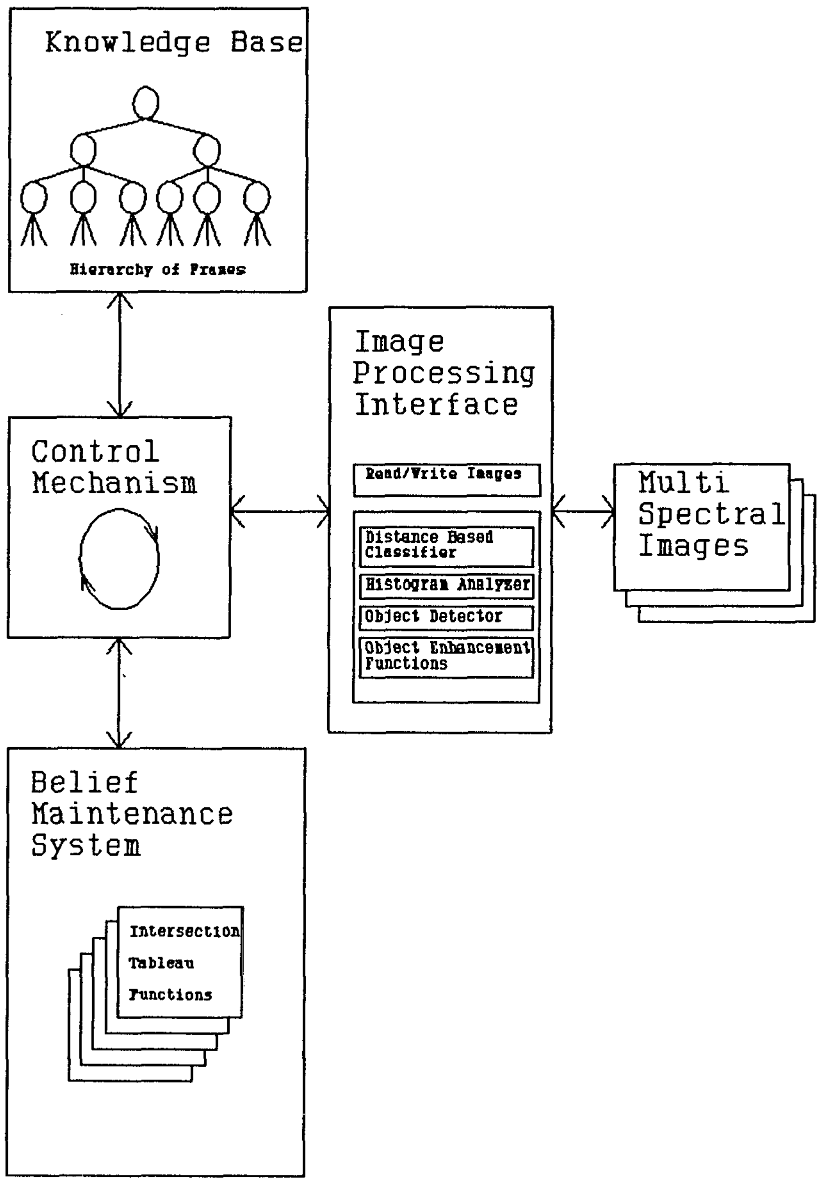

Figure 4.1 - Top-Down Bottom-Up System

major programming effort in themselves.

The image processing interface works on the image data. It is the segmentation portion of the segmentation and interpretation interplay. There are several different operators in the image processing interface. The control mechanism chooses the operators depending on the situation and the expectations given by the knowledge base. The relation between the knowledge base and the operators is very close. The operators are procedural attachments in the knowledge base. The knowledge base contains their names and directs the control mechanism to call them when appropriate. The code of the operators is not in the knowledge base, but their names are.

4.1 - The Images

The system analyzed two different types of images. The primary image is a three channel multi-spectral aerial image of an apartments area south of the Baton Rouge campus of Louisiana State University. The image was taken from an altitude of 1500 feet. The image is 511x512 pixels and has 256 gray levels. Each pixel covers an area 2.25 feet by 2.25 feet. You can identify easily buildings, parking lots, and carports from this altitude. The image is in a format tied to the ELAS [NASA] image processing system.

Appendix 1 contains photographs of the image. The photographs show how each spectral channel aids in detecting objects of interest.

Photograph 1 shows the green channel (.5 - .55 μm band). The bright white objects are carports. The roads appear as long bright objects. Buildings sometimes do not appear because they are almost the same gray level as the surrounding grass. Photograph 2 shows the red channel (.65 - .69 μm). The bright white objects are buildings. The long gray areas are roads and cement. Photograph 3 shows the thermal IR channel (8.5 - 13.0 μm). The dark gray objects are buildings. Photograph 6 shows the results of analysis.

To illustrate the power and portability of the system, a second type of image was analyzed. The basis of the second type of image was an 8"x10" color aerial photograph of the Atlantic Undersea Test and Evaluation Center (a U.S. Navy test facility) on Andros Island in the Bahamas. A Hewlett-Packard ScanJet Plus desktop scanner scanned several portions of the photograph into a single channel, 256 gray level image. The scanner produced an image in the TIFF (Tagged Image File Format) format. This was transformed to the ELAS format. Photograph 7 shows the first area of this image. This is a section of house trailers separated by a road. This image is 200x200 pixels. Photograph 8 shows another section of house trailers. This image is 100x200 pixels. Photograph 9 shows a third section of house trailers. This image is also 100x200 pixels. These three images (photographs 7, 8, and 9) possess vastly different spectral properties from the first image (photographs 1, 2, and 3).

4.2 - The Knowledge Base

A hierarchy of frames stores the knowledge base. As described in chapter three, frames are a basic slot and filler notation. They satisfy the three perceptual psychology properties of context, expectation, and selective attention and they are a cornerstone of the perceptual cycle.

You could use any number of hierarchies for the frames. This system uses a hierarchy of labels given by the Defense Mapping Agency (DMA) [Rusco]. The hierarchy was introduced in earlier chapters and some of its weakness were discussed briefly (see chapter six for a discussion of how this hierarchy might be modified in future work). There are several uses of the hierarchy that have not yet been carried forward from the discussions in [Harlow, et. al]. One such use is storing knowledge about the spatial relations among the low level objects. This would be knowledge such as "carports located next to buildings" and "sidewalks located next to roads." Figure 4.1.A illustrates these concepts. These are examples of some of the powerful knowledge imbedded into each position in the hierarchy that can aid and direct the low level operators. Some of the systems described in chapter 2 used knowledge similar to this. Those systems, however, expressed that knowledge as special rules in low level operators - not as part of the total system framework. This use of the hierarchy has not yet been implemented in the current system because it is essentially a feature used by low level operators. The slots for this information are in the frames and will be incorporated in the future (see chapter 6 for this discussion).

Appendix 2 lists the DMA hierarchy and figure 4.2 shows the top four levels. The hierarchy has at least one more level lower than figure 4.2 shows. The DMA hierarchy is quite extensive and has 247 nodes in the hierarchy and 207 leaf nodes or final labels.

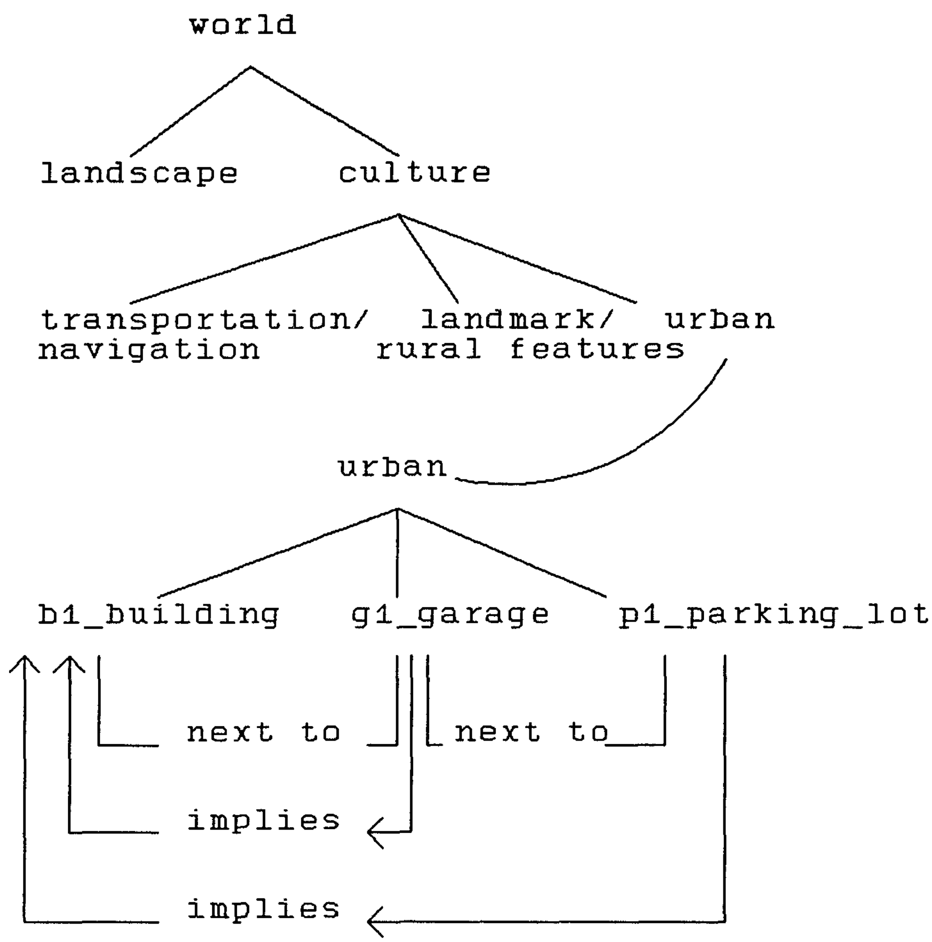

Figure 4.1.A - Adding a Vision Hierarchy to the DMA Hierarchy

Every non-leaf node in the hierarchy has a corresponding frame.

Each frame has two types of information. The first type of information is a list of intrinsic characteristics. Intrinsic characteristics distinguish a node from the other nodes in the hierarchy having the same parent. The second type of information lists the node's sub nodes. It also gives a data driven operator which will generate an initial belief vector about the sub nodes. In general, each frame is as shown in figure 4.3. Appendix 3 lists the frames.

```
is_a ...
is_part_of ...
level_in_tree ...
goal_of_analysis ...
intrinsic_characteristics ...
        distinguishing characteristics and
        assertions of belief
sub_node_names ...
sub_node_operator(s) ...
```

Figure 4.3 - Frames

The is_a slot gives the name of the frame. The is_part_of slot ties the frame to the node above it in the hierarchy. This demonstrates one of the major advantages that frames have over IF...THEN production rules. Rules in a production system do not express the hierarchical nature of scenes as well as frames.

The level_in_tree slot gives the distance down in the hierarchy from the root node. This helps to direct the control mechanism. When the level_in_tree = "root," the node does not have any intrinsic characteristics. This is because the root has no siblings from which it must be distinguished. When the level_in_tree = "leaf," then the processing on that area in the image has finished.

The goal_of_analysis slot describes the desired output from this stage of analysis. The goal is usually to give a single label to the area under investigation. When the processing reaches the bottom of the hierarchy the goal changes to labeling the individual objects in the area. The intrinsic characteristics list information that distinguishes the node from the other nodes in the hierarchy with the same parent node. Figure 4.4 show a

```
world
├─ culture
│   ├─ urban (built up areas)
│   │   ├─ industry/utility
│   │   ├─ commercial/residential
│   │   └─ institutional/governmental
│   ├─ transportation/navigation
│   │   ├─ railroads
│   │   ├─ roads
│   │   ├─ aeronautical/aerospace
│   │   ├─ naval/marine
│   │   └─ associated transportation features
│   └─ landmark/rural features
│       ├─ communication/transportation
│       ├─ storage
│       ├─ agricultural
│       ├─ recreational
│       └─ miscellaneous
└─ landscape
    ├─ hydrography
    │   ├─ water
    │   └─ snow/ice
    ├─ physiography
    │   ├─ exposed soils (surface composition)
    │   └─ landforms
    └─ phytography
        ├─ cropland
        ├─ rangeland
        ├─ woodland
        └─ wetland
```

Figure 4.2 - DMA Hierarchy (top portion)

possible form of these characteristics.

---

has b1_building
has p1_parking_lot
has 040 apartments/hotel
has 041 house (single family dwelling)

Figure 4.4 - Example Intrinsic Characteristics

The intrinsic characteristics cause the control mechanism to call the proper low level image processing operators. The characteristics listed here would cause the control mechanism to call the object detector operator (see section 4.5.2). The frames store assertions of belief with the intrinsic characteristics. An assertion of belief is in the form of a basic probability assignment (bpa) as described in chapter three.

The low level image processing operators are known as attached procedures. They are procedures that are connected to the data structure of the frame (recall this is one of the advantages of using frames for knowledge representation). It is easy to attach any procedure to the frame. This allows the system to incorporate operators developed by other researchers. The details of this are (1) take the operator written in C, (2) compile and link it to the system software, (3) put its name in the frame at the intrinsic characteristic slot, and (4) modify the control mechanism code to call it when its name appears in the frame. The most difficult step in the process is (1). Many operators produced in research efforts are not portable.

If the object detector finds an intrinsic characteristic, then the control mechanism asserts a bpa that is favorable to this frame. If the object detector does not find an intrinsic characteristic, then the control mechanism asserts a bpa that is unfavorable to this frame. For example, figure 4.5 gives a portion of the *culture* frame. The first intrinsic characteristic is "b1_building." If the object detector operator finds objects meeting the description of "b1_building," then the system would make the positive assertions. If the system did not find such objects, then the system would make the negative assertions. The positive assertion says that place=2, i.e. node 2 *culture*, receives 0.8 belief and place=1, i.e. node 1 *world* or Θ, receives 0.2. Please note the descriptions of the objects for which the object detector searches. This is another major knowledge base the system uses. Refer to section 4.5.2 for details. An expert photointerpreter creates the assertions in a purely subjective manner. He bases them on experience and judgment. The author created the assertions in this work.

The frame.is_a is:
    ->culture

characteristic[0] is ->b1_building

characteristic[0].positive_assertion[0].place=2
characteristic[0].positive_assertion[0].belief=0.8
characteristic[0].positive_assertion[1].place=1
characteristic[0].positive_assertion[1].place=0.2

characteristic[0].negative_assertion[0].place=3
characteristic[0].negative_assertion[0].belief=0.6
characteristic[0].negative_assertion[1].place=1
characteristic[0].negative_assertion[1].belief=0.4

Figure 4.5 - Portion of *Culture* Frame

The sub_node_names slots list the names of the nodes branching downwards from this node in the hierarchy. The sub_node_operator specifies a simple, data driven operator that will help to distinguish the sub nodes. This operator is some type of statistics based operator. TDBU uses the simple average of the gray levels for this operator (see section 4.5.2).

Figure 4.6 shows another portion of the *culture* frame. This portion lists the sub_node_operator and the sub_nodes. The operator is the simple average_of_pixels (see section 4.5.2). Each sub_node has a mean value associated with it. The average_of_pixels operator calculates the mean gray level of the area and then uses each sub node's mean value to calculate an assertion of belief. The system gives each sub node a measure of belief based on the distance between the mean gray level of the area and the sub node's mean value.

The frames are implemented using the database program PC-FILE (ButtonWare, Inc. Bellevue, Washington). PC-FILE is a simple yet capable database program that runs on IBM PC's and compatibles. The author uses PC-FILE to enter and modify the frames.

Portions of the TDBU system's code reads the database files into C language structures. A utility program uses some of the same code to read the files and print them in the form of appendix 3. Appendix 4 describes the implementation details of the database files.

The frame.is_a is:
->culture

The frame has 3 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->urban (built up areas)
Sub node[0].bel_element_number is 4
Sub node[0].mean is 90.0

Sub node[1].is_a is ->transportation/navigation
Sub node[1].bel_element_number is 5
Sub node[1].mean is 60.0

Sub node[2].is_a is ->landmark/rural features
Sub node[2].bel_element_number is 6
Sub node[2].mean is 40.0

Figure 4.6 - Portion of *Culture* Frame

4.3 - The Control Mechanism

The control mechanism directs the overall flow of the image analysis process. The perceptual psychology principles described in chapter two and the combination top-down and bottom-up processing described in chapter three form the foundation of the control algorithm. Figure 4.7 shows the basic control algorithm (the <- symbol indicates assignment).

There are two frame pointers used in the analysis. They are (1) the Top_Frame and (2) the Candidate_Frame. The Top_Frame is the frame that is the current label of the area. The Candidate_Frame is the frame that has the highest probability of being the new

```
1) Given a desired level of analysis

2) Top_Frame <- frame_named(World)

3) Candidate_Frame.is_a <- nothing

4)   While
             Top_Frame.level_in_tree != leaf
     OR
             Top_Frame.level_in_tree <= desired level
             of analysis
     Do steps 5) to 9)

     5)      If
                     first pass through While loop of step 4)
             Then
                     apply Top_Frame.sub_node_operator to initialize
                     the belief_vector

     6) Most_probable_label <- Maximum_of(belief_vector)

     7)      If
                     Most_probable_label != Candidate_Frame.is_a

     8)      Then
                     Candidate_Frame <- frame_named(Most_probable_label)
                     Investigate area using intrinsic characteristics
                     of Candidate_Frame
                     8A)    Loop over the number of intrinsic characteristics
                     8B)            object detector looks for object
                     8C)            set assertion of belief
                     8D)            alter belief_vector using Dempster's
                                    rule of combination
                     8E)    end loop

     9)      Else
                     Region_label <- Candidate_Frame.is_a
                     Top_Frame <- frame_named(Region_label)
                     Candidate_Frame.is_a <- nothing

10)  If Top_frame.level_in_tree = leaf
     Then
             10A)   find and label the lowest level objects
             10B)   move on to the next area to analyze
```

Figure 4.7 - Control Algorithm

label of the area. The Candidate_Frame is below the Top_Frame in the hierarchy and is a sub node of the Top_Frame.

Step 1) - The user specifies how far down into the tree he wants the analysis to proceed.

Step 2) - The analysis of an area begins at the top of the hierarchy. The frame *world* is the label of the area when the analysis begins.

Step 3) - The Candidate_Frame is the most likely sub node of the Top_Frame. Since analysis has not started, the Candidate_Frame is set to *nothing*.

Step 4) - This test checks the level of analysis. The system analyzes and labels the area on through the hierarchy until it reaches the desired level of analysis.

Step 5) - The first stage of investigation uses the data driven sub node operator. The goal of analysis is to label the area as one of the Top_Frame's sub nodes. The sub node operator uses known statistical properties of the sub nodes. It generates a basic probability assignment or assertion about the sub nodes (see section 4.5.2). This is the data driven first step of analysis described in chapter three. The simple operator is chosen using selective attention. It is not an all around general purpose operator. It is one chosen specifically for this situation because of its ability to give an estimate at a glance.

Step 6) - The system sets The Most_probable_label to the is_a slot of the node with the highest belief. The system determines the node with the highest belief by using the sub node operator in step 5).

Step 7) - The system performs the test for conclusion. If the Most_probable_label is the same as the Candidate_Frame.is_a, then proceed to step 9) to label the area.

If the analysis is on the first pass of the WHILE statement of step 4), the Candidate_Frame.is_a = *nothing* so the analysis will proceed to step 8).

If the analysis is on a subsequent pass of the WHILE statement of step 4), then detailed analysis of the area has taken place. If the Most_probable_label is the same before and after detailed analysis, then label the area with the Most_probable_label in step 9).

Step 8) - At this point, the system sets the Candidate_Frame equal to the Most_probable_label. The analysis now uses the intrinsic characteristics of Candidate_Frame to investigate the area. The results of the investigation will alter the belief_vector and control will go back to step 4).

This implements the goal driven second step of the analysis mentioned in chapter three. Expectations drive the analysis in this step. Using selective attention, the system does not spend equal time on all areas and objects in the area. It processes only a small percentage of objects in the area. The context of the frame and the frame's location in the hierarchy limits the number of processed objects.

This is the segmentation process. The interpretation of the first, simple segmentation of step 5) guides the process. That interpretation greatly simplifies this one. Steps 8A) through 8E) describe this process in detail. The Candidate_Frame has a given number of intrinsic characteristics. In step 8B), the object detector looks for the object listed in the intrinsic characteristic. The results of the object detector set an assertion of belief in step 8C) (see section 4.2). The system alters the belief_vector using Dempster's rule of combination in step 8D) (see section 4.4 for details). The process of investigate and alter the belief_vector repeats itself in step 8E).

Step 9) - Analysis at this level in the hierarchy ends. The system labels the area Candidate_Frame.is_a. The Top_Frame now becomes the Candidate_Frame and the Candidate_Frame.is_a again becomes *nothing*. Control goes back to step 4) and analysis continues at a lower level in the hierarchy.

Step 10) - At this point the system reaches the bottom of the hierarchy. At step 10A) the system tries to find and label the lowest level objects, i.e. the leaf nodes. The system uses the object detector and the object enhancement functions to do this. The system now departs from the usual top down traversal of the hierarchy. It uses all leaf nodes on this level of the hierarchy as possible objects. If the system did not do this, then it would detect only one type of object. If the system had labeled an area as *commercial/residential*, then it would only try to detect types of buildings. It would leave out streets, sidewalks, and other common objects found in residential areas.

Figure 4.8 shows all of the leaf nodes found on the same level as the leaf nodes of the *commercial/residential* frame. When the system reaches this level of analysis, it uses the object detector to look for each object shown in figure 4.8. It also uses the object enhancement functions to clean up the object detector's results (see section 4.5.2). Step 10A) ends the analysis of an area. Step 10B) resets the system, adjusts the coordinates of

039 commercial buildings
040 apartments/hotel
041 house (single family dwelling
042 mobile homes
043 barracks
044 display sign (large billboard, highway sign)

045 governmental administration building
046 military admin/operations building
047 capitol building
048 hospital
049 prison
050 palace
051 museum
052 observatory
053 church/tabernacle
054 mosque
055 cemetery building

056 single track railway
057 double track railway
058 multiple track railway
059 RR yard/siding
060 tramway/inclined railway
061 monorail
062 RR storage/repair building
063 RR terminal building
064 RR station/depot
065
066 roundhouse

067 multi lane, divided (grass median) highway
068 multi lane highway
069 primary road (dual lane, hard surface)
070 secondary road (dual lane, loose/dirt surface)
071 trail/track (one lane)
072 toll gates
073 cloverleaf/interchange
074 garage, service/repair facilities (landmark)

Figure 4.8 - All Leaf Nodes on the Same Level as *commercial/residential* Leaf Nodes

the area under analysis, and returns control back to step 2).

4.4 - The Belief Maintenance System

The belief maintenance system uses a form of the Dempster-Shafer (D-S) theory of evidence described in chapter three. The belief maintenance system comprises many subroutines that implement Dempster's rule of combination (the intersection tableau). This was a significant software effort in itself. The subroutines are tailored to the DMA hierarchy. These subroutines use and implement the advantages inherent when employing the D-S theory on a hierarchy of hypotheses (labels or frames). These advantages significantly reduce the amount of computations. As discussed in section 3.3, there are 247 nodes in the hierarchy so there are that many places for computations. If a hierarchy were not used, then there would be a flat list of 29 labels. This would mean 2^{29} or 536,870,912 computations whenever the system combined two simple assertions. The subroutines in the system are custom written for just this hierarchy.

The belief maintenance system uses one belief vector. There is an element in the belief vector for each non-leaf node in the hierarchy. Figure 4.6 shows that each node has its own belief element number. For example, *urban (built up areas)* is belief element number 4, *transportation/navigation* is belief element number 5, and *landmark/rural* features is belief element number 6.

The low level operators (see section 4.5.2) pass their results to the belief maintenance system in the form of assertions or basic probability assignments. Figure 4.5 shows how each intrinsic characteristic has positive and negative assertions. If the operator finds the cue for which it is searching, it passes the positive assertions to the belief maintenance system. If it does not find the cue, it passes the negative assertions to the belief maintenance system.

The results of the low level operators are pieces of evidence. Examples of evidence are "found concrete in the area" or "located several parking lots." As stated above, the operators state these pieces of evidence in the form of assertions. The pieces of evidence can be suspicious or incorrect depending on the ability of the low level operator. The pieces of evidence can also be contradictory. The histogram analyzer (see section 4.5.2) can return a statement "found concrete" which indicates man made objects. It can then return a statement "did not find any roofing materials" which indicates a natural area. This contradicts the first piece of evidence. The belief maintenance system takes the pieces of evidence, in the form of assertions, and combines them to draw a conclusion. The conclusion depends on the ability of the operators and the relative beliefs in the assertions.

The belief maintenance system uses Dempster's rule of combination to pool the assertions with the belief vector. This alters the contents of the belief vector. The belief maintenance system contains a large set of subroutines that implement the intersection tableau. There is one subroutine for each belief element number. Given the "place" shown in figure 4.5, the belief maintenance system calls the subroutine of the same number. That subroutine implements the intersection tableau to combine that place with all of the other places in the belief vector.

The belief maintenance system returns the altered belief vector to the control mechanism. After the object detector finishes looking for each intrinsic characteristic, the control mechanism uses the altered belief vector. The control mechanism either labels the area or it selects another frame to be the candidate frame and analysis continues.

4.5 - The Image Processing Interface

The image processing interface is a set of functions that operates directly on the image. There are two types of functions. There are functions used to read and write the images from disk and there are functions which operate on the image.

4.5.1 - Image Processing Read/Write Functions

These functions read and write images from and to disk. The images are in a format tied to the ELAS [NASA] image processing system. Special C functions were written to perform the read and write functions. This is because the author does not have an ELAS system on the IBM-PC compatible machine used in this research.

An ELAS image is divided into two parts. They are (1) the image header and (2) the image data. The image header is 1024 bytes long and contains information such as the size of the image, the number of spectral channels, and the x and y spot size of each pixel. The format stores the multi-channel image data row by row. The first row of the first channel is followed by the first row of the second channel and the first row of the third channel. Next is the second row of the first channel, the second row of the second channel, and the second row of the third channel and so on.

4.5.2 - Image Processing Operators

The system contains four types of operators. These are (1) Distance Based Operator, (2) Histogram Analyzer, (3) Object Detector, and (4) Object Enhancement Functions. These operators are basic and simple. They are neither as complex nor as capable as the operators described in chapter two. As stated in chapter one, the development of low level signal processing operators is not the objective of this dissertation. These operators are attached procedures. They are attached or linked to the intrinsic characteristic slots of the frame. It is easy to add new operators to the system. This is one of the advantages of this computer vision system. Operators produced by other researchers can be incorporated into the system.

In the following discussion, please note how the system employs different operators looking for different cues at different levels in the hierarchy. There are several advantages to this. One of the biggest advantages is the reduction in computational expense. There is no reason to use complex, computationally expensive operators to solve simple problems. At the highest level of analysis the system uses a histogram analyzer. This is a simple and quick operator. It determines the presence or absence of materials (concrete, roofing, etc.) by examining the histogram of the area. At the next level of analysis the system uses an object detector applied to a reduced resolution image. The reduced resolution image used 1 pixel to represent a 4x4 pixel area from the original image. This speeds up operation. At the next level down the system again uses an object detector applied to a reduced resolution image. The reduced resolution at this level uses 1 pixel to represent a 2x2 area in the original image. Again, this speeds up operation and still retains enough ability to solve the problem at hand.

Distance Based Operator

The distance based operator calculates the mean gray level of an area and uses this to initialize the belief vector. This performs the role of the sub_node_operator in the frames. The sub_node_operator generates an initial belief vector when the analysis of an area begins. The distance based operator measures the distance from the mean gray level of the area to the means of each sub node. It uses this distance to assign a measure of belief to each sub node. Figure 4.10 shows the algorithm for this operator.

Given: N sub nodes
mean = mean gray level of area

1. Calculate distance_i = | mean - mean of sub node_i | for i=1,N

2. Set distance_{MIN} = minimum of distance_i for i=1,N
i_{MIN} = i for $\text{distance}_i = \text{distance}_{MIN}$

3. Calculate denominator = Sum for i=1,N of $\text{distance}_{MIN}/\text{distance}_i$

4. Set m(THETA) = 0.1 for arbitrary ignorance
$m(i_{MIN})$ = 0.9/denominator

5. Calculate for i=1,N i != i_{MIN}
$m(i) = (\text{distance}_{MIN}/\text{distance}_i) * m(i_{MIN})$

Figure 4.10 - Sub Node Operator Algorithm

Figure 4.11 gives an example that demonstrates the operator. The sub node names and means come from figure 4.6. The result is that the *transportation/navigation* sub node receives the highest measure of belief. The operator bases the measure of belief given to each sub node on the distance from that sub node's mean to the mean of the area. Θ receives an arbitrary measure of belief.

Histogram Analyzer

The system uses the histogram analyzer to investigate the intrinsic characteristics of an area when the level of analysis is 2. At this level of analysis the system must differentiate between *culture* and *landscape*. This decision is basic enough that the system can use a histogram analyzer.

The histogram analyzer takes the histogram of an area, smoothes it, and then examines the peaks in the smoothed histogram. The frame tells the histogram analyzer what material should be present. The histogram analyzer looks up the spectral properties

Node is culture

sub $node_1$ is urban (built up areas)
mean of sub $node_1$=90

sub $node_2$ is transportation/navigation
mean of sub $node_2$=60

sub $node_3$ is landmark/rural features
mean of sub $node_3$=40

Given mean = 70

1. $distance_1 = | 70 - 90 | = 20$
   $distance_2 = | 70 - 60 | = 10$
   $distance_3 = | 70 - 40 | = 30$

2. $distance_{MIN} = 10 \quad i_{MIN}=2$

3. denominator = 10/20 + 10/10 + 10/30 = 11/6

4. m(THETA) = 0.1
   m(2) = 0.9 / (11/6) = 0.49

5. m(1) = 10/20 * 0.49 = 0.25
   m(3) = 10/30 * 0.49 = 0.16

Result: m(THETA) = 0.1 m(1) = 0.25
m(2) = 0.49 m(3) = 0.16

Figure 4.11 - Example Sub Node Operator Calculation

of that material (see the description of the spectral information in the next section). These spectral properties contain information that reveals what peaks should be in the histogram. If the desired peaks are in the histogram, then the material is present in the area of the image. The results of the histogram analyzer determine which assertions the belief maintenance system will make.

Note the *culture* and *landscape* frames at the beginning of Appendix 3. The characteristics portions of the frames hold a material name. The histogram analyzer takes this name, reads the spectral information for that material, and examines the histogram. The *culture* frame lists asphalt roofing in two spectral channels (the red and the thermal IR channels) and aluminum roofing as its characteristics. These materials have unique spectral properties that stand out in the histogram.

Object Detector

The most often employed operator is an object detector. Given spectral and spatial information, the operator locates and calculates the parameters of the objects. Figure 4.12 shows the basic algorithm.

1. Given:
    A. Area to analyze
    B. Spectral Properties
        1. Number of spectral channels for object
        2. Gray level thresholds
    C. Spatial Properties
        1. Limits on size of object
        2. Limits on height of object
        3. Limits on width of object
        4. Limits on width to height ration
        5. Limits on principle axis of object

2. Using properties B.1. and B.2. above, threshold the image into a 1 0 image.

3. Grow regions (see figure 4.13).

4. Compute the principle axis, height, and width of each region of step 3 (see figure 4.14 for principle axis algorithm).

5. Eliminate any region whose parameters fall outside the limits of C.1. through C.5. above.

6. Results:
    List each object giving the area, height, width, width to height ratio, and principle axis.
    Create an output 1 0 image showing the detected objects.

Figure 4.12 - Basic Object Detector Algorithm

Figure 4.13 shows the region growing algorithm of figure 4.12 step 3.

```
1. Segment the image (m x n).  Assume a picture function
     g(i,j) for i=1,m j=1,n
     g(i,j)    = 1 for object
               = 0 for background

2. set g_label=2  this is the label value

3.   for i=1 to m do
     begin scanning i^th row
          for j=1 to n do
          begin checking j^th element
               stack_empty = true
               if g(i,j) = 1 then
                    begin
                    label_and_check_neighbor(g(i,j),g_label)
                    end
               while stack_empty = false do
                    begin
                    pop on element (i,j) off the stack
                    label_and_check_neighbor(g(i,j),g_label)
                    end
               g_label = g_label + 1
          end of checking j^th element
     end of scanning i^th row

     procedure label_and_check_neighbor(g(i,j), g_label)
          begin
          g(r,e) = g_label
          check g(R,E) where R=r-1,r,r+1 and E=e-1,e,e+1
          if g(R,E) = 1 then
               begin
               push (R,E) onto the stack
               stack_empty = false
               end
     end procedure label_and_check_neighbor
```

The procedure label_and_check_neighbor also calculates the max and min x and y coordinates of each object found.

Figure 4.13 - Region Growing Algorithm

Figure 4.14 shows the principle axis calculation of figure 4.12 step 4. The principle axis algorithm is from [Castleman].

The formula for the principle axis Φ is:

$$\tan(2\Phi) = 2\mu_{11}/(\mu_{20} - \mu_{02})$$

where:

μ_{11} = Sum over x,y in object [(x - center_x)(y - center_y)]

μ_{20} = Sum over x,y in object [(x - center_x)(x - center_x)]

μ_{02} = Sum over x,y in object [(y - center_y)(y - center_y)]

center_x = (Sum of x over x,y in object)/(no. of points in object)

center_y = (Sum of y over x,y in object)/(no. of points in object)

Figure 4.14 - Principle Axis Algorithm

An object descriptor describes each object. Figure 4.15 shows the C structures used in the object descriptor. Figure 4.16 shows the object descriptor for the object *041 house (single family dwelling)*. The descriptors express dimensions in feet and angles in degrees.

```
struct spectral_structure{
        short channel;
        short low_threshold;
        short high_threshold;
};

struct spectral_signature{
        char  is_a[MAX_NAME_LENGTH];
        short num_of_channels;
        struct spectral_structure spectrum[SPECTRUM_LENGTH];
};

struct spatial_structure{
        char  is_a[MAX_NAME_LENGTH];
        char  material[MAX_NAME_LENGTH];
        long  min_area;
        long  max_area;
        long  min_width;
        long  max_width;
        long  min_height;
        long  max_height;
        float min_w_to_h_ratio;
        float max_w_to_h_ratio;
        float min_angle;
        float max_angle;
};

struct descriptor_structure{
        char  is_a[MAX_NAME_LENGTH];
        short num_of_channels;
        struct spectral_structure spectrum[SPECTRUM_LENGTH];
        long  min_area;
        long  max_area;
        long  min_width;
        long  max_width;
        long  min_height;
        long  max_height;
        float min_w_to_h_ratio;
        float max_w_to_h_ratio;
        float min_angle;
        float max_angle;
};
```

Figure 4.15 - C Structures for the Object Descriptor

```
strcpy(d.is_a,"041 house (single family dwelling)");
strcpy(d.material, "roofing");

d.min_area            = 56;
d.max_area            = 840;
d.min_width           = 7;
d.max_width           = 46;
d.min_height          = 11;
d.max_height          = 68;
d.min_w_to_h_ratio    = 0.3;
d.max_w_to_h_ratio    = 1.0;
d.min_angle           = -90.0;
d.max_angle           = 90.0;
```

Figure 4.16 - Object Descriptor for *041 house (single family dwelling)*

Figure 4.16 clearly shows the separation of spectral and spatial information in the object descriptor. The material slot ("roofing" for this example) contains the spectral information. The remaining slots (min_area down though max_angle) contain the spatial information. The descriptor expresses the spatial information in feet. Therefore, this information will not change from image to image (an image contains a header that allows the system to transform feet to pixels).

The only portion of the knowledge base that changes from image to image is the spectral information. This may change if you obtain different images from different scanners and under different atmospheric conditions. Nevertheless, the system keeps this changeable information in one, easy to find, and easy to modify location.

```
strcpy(s.is_a, "roofing");
s.num_of_channels              =   1;
s.spectrum[0].channel          =   2;
s.spectrum[0].low_threshold    = 185;
s.spectrum[0].high_threshold   = 255;
```

Figure 4.17 - Spectral Information for "roofing" Material

Figure 4.17 shows the spectral information for the "roofing" material for the primary image (photographs 1, 2, and 3). The system used different spectral information for the images of photographs 7, 8, and 9. Each new image required modifications to the spectral knowledge base. The spatial information did not change. Appendix 5 lists the object descriptors used in the system. Appendix 6 lists the spectral information used for the first image (photographs 1, 2, and 3).

The object detector can function at different resolutions. When the level of analysis is 3, the object detector uses a reduction in resolution of four, i.e. 1 pixel represents a 4x4 area of pixels in the full resolution image. Photograph 4 shows the same channel as photograph 2 with a reduction in resolution of four. When the level of analysis is 4, the object detector uses a reduction in resolution of two, i.e. 1 pixel represents a 2x2 area of pixels in the full resolution image. Photograph 5 shows the same channel as photograph 2 with a reduction in resolution of two.

Object Enhancement Functions

The object enhancement functions are a set of functions used to enhance the output of the object detector. The system uses them to improve the accuracy and presentability of the object detector output. The system uses different functions depending on the type of object the object detector processes. There are three uses for them. The uses are (1) with buildings, (2) with roads, and (3) with sidewalks.

The enhancement used with buildings fills them out to their actual edges. The object detector uses spectral qualities of the buildings to threshold and merge the pixels into a building. There are, however, several problems in this process. The biggest problem is the shadows on the rooftops of buildings. These shadows distort the spectral properties of the buildings and the object detector output often yields only half a building. The system attempts to compensate for this by using several object enhancement functions. The most important of these is an edge detector. The system uses an edge detector known as the Kirsch operator [Levine]. It is a compass gradient operator. It convolves eight different 3x3 masks with each pixel in the area. Figure 4.18 shows the eight convolution masks. This is a computationally expensive edge detector, but it is used sparingly and is worth the expense.

```
                      Directions
                        7  0  1
                        6  x  2
                        5  4  3

0 direction mask               4 direction mask
     5   5   5                      -3  -3  -3
    -3   0  -3                      -3   0  -3
    -3  -3  -3                       5   5   5

1 direction mask               5 direction mask
    -3   5   5                      -3  -3  -3
    -3   0   5                       5   0  -3
    -3  -3  -3                       5   5  -3

2 direction mask               6 direction mask
    -3  -3   5                       5  -3  -3
    -3   0   5                       5   0  -3
    -3  -3   5                       5  -3  -3

3 direction mask               7 direction mask
    -3  -3  -3                       5   5  -3
    -3   0   5                       5   0  -3
    -3   5   5                      -3  -3  -3
```

Figure 4.18 - Convolution Masks for Kirsch Edge Detector

The following figures lead through the process of enhancing the buildings detected by the object detector. These examples are from the lower right hand corner of the image. Figure 4.19 shows the raw output of the object detector (the remainder of the figures in this chapter are at the end of the chapter). The operator did detect the buildings, but the edges are rough. Figure 4.20 shows the output of the Kirsch edge detector. This is also rough. Note the edge running through the middle of one of the buildings. This is the dividing line of the roof caused by a shadow. There are edges in this figure that do not relate to buildings. The next step will eliminate these.

Figure 4.21 shows the result of overlaying the output of the object detector with that of the edge detector. This is still incomplete, but it smoothes out most of the edges from the object detector. It still has several edges produced by the edge detector that are not related to the buildings. Therefore, the next step is to remove these edges. An operator takes the pixels produced by the edge detector and merges them with the output of the object detector. If an edge pixel is adjacent to an object pixel, then the operator changes the edge pixel to an object pixel. The is an iterative operator that runs over and over until there are no more edge pixels to convert to object pixels. After this operation finishes, the operator eliminates the remaining edge pixels that were not converted to object pixels. Figure 4.22 shows the result of this. The buildings are almost complete.

The final enhancement operation removes the holes inside the buildings. This operator looks at zero pixels and counts the number of non-zero neighbor pixels. If this count is greater than a threshold value, the operator changes the zero pixel to a non-zero pixel. Figure 4.23 shows the final result. This is a definite improvement over figure 4.19.

The enhancement used with roads performs the operations used with buildings and it also separates the roads from sidewalks. Roads and sidewalks have the same spectral intensities. Therefore, the object detector tends to join the two objects. Figure 4.24 shows the output of the object detector. This shows sidewalks as roads. The first object enhancement function removes the sidewalks from the roads by rejecting parts of the detected object that are "thin." It does this by examining each non-zero pixel and the surrounding pixels. It counts the number of non-zero pixels in the 12 foot by 12 foot area. If this count is less than a threshold, then the pixel under examination is part of a thin object and it is set to zero. Figure 4.25 shows the result of this operator. The sidewalks are gone. The operator also thinned the road a small amount.

Next, the system uses the Kirsch edge detector mentioned earlier. Figure 4.26 shows the result of the edge detector. This shows edges of the road and also edges not associated with the road. Figure 4.27 shows the result of overlaying the output of the object detector with the edge detector. This fills out some of the road edges the previous operator thinned. Figure 4.27 also shows unwanted edges that are not related to the road.

The next step must remove them. Figure 4.28 shows the final result of road enhancement. The final operator merged the output of the edge detector with that of the object detector. This operator eliminated edge detector pixels that were not adjacent to object detector pixels. It is the same operator described earlier that operated on the buildings to produce figure 4.22. Figure 4.28 gives an accurate representation of the road.

The final usage of enhancement functions is with sidewalks. As stated earlier, the spectral properties of roads and sidewalks are the same. The object detector joins the two so the enhancement functions must separate them. Figure 4.29 shows the output of the object detector. This is the same as figure 4.24.

Eliminating roads from sidewalks is similar to eliminating sidewalks from roads. An object enhancement function rejects parts of the sidewalk and road object that are "wide." The function looks at each non-zero pixel and the surrounding pixels. It counts the number of non-zero pixels in the 12 foot by 12 foot neighborhood. If the number is greater than a threshold, then the pixel under investigation is part of a "thick" object and is set to zero. Figure 4.30 shows the results of this function. This is the final output for the sidewalks. This is an accurate depiction of the sidewalk.

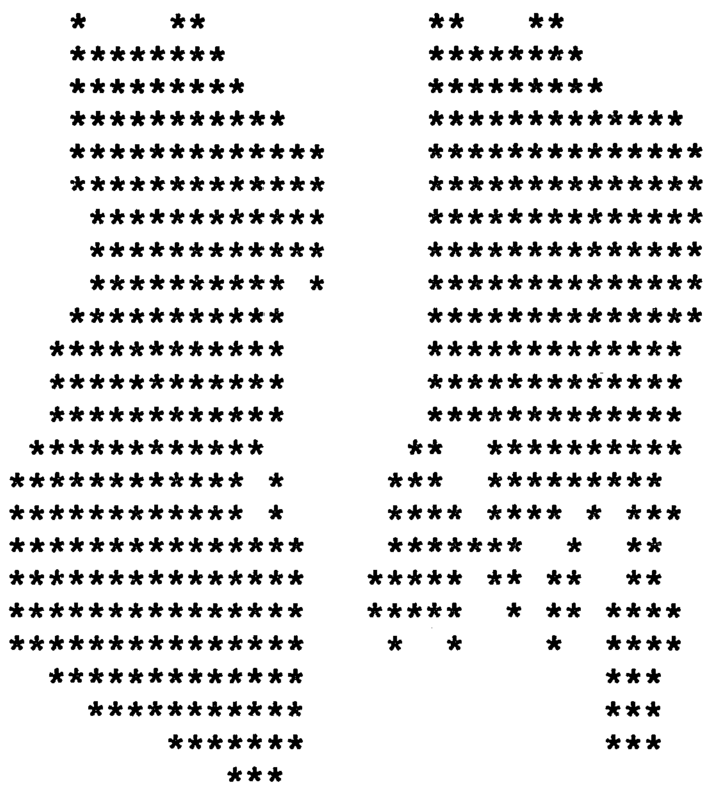

Figure 4.19 - Raw Output of the Object Detector (Building Example)

```
****
 ****                                  *
  ******                         *    **        **
   *****                              **
     ***                              **   *
     ***                              **     *
     ***  ***  ****         ***   **   *
    ****  **********        **********    *
    ****  **********        *********
-   ****  **     ****       ***     *
***  ***  **      *******   ***           *
**** **   **        *****   ***          **
*** ***   **           **   **           **
  * ***   **           **   **           **
  * ***   **           **   **           **
  * ***  ***           **   **           **
  * ***  ***           **   **           **
  * **   ***           **   **  *        *
  ****   **   *        **   **  **       **
- ****   **   ******   **   **   *       **  *
******  ****  ******   **  ***   *   *   **
****** ********   **   *   ***   *       **
 *****  ***       **  **   ***           **
        **        ** ***  ****              ***
*      ****       ******  ** *         ***    *
       ******      *****  *****        ***
       ********     ****  ******  *   ****
        *********** ***      *** ***  ***
          *************       * ***** ***
-             *********             *****
                  *****             *****
                   ***
```

Figure 4.20 - Output of Kirsch Edge Detector (Building Example)

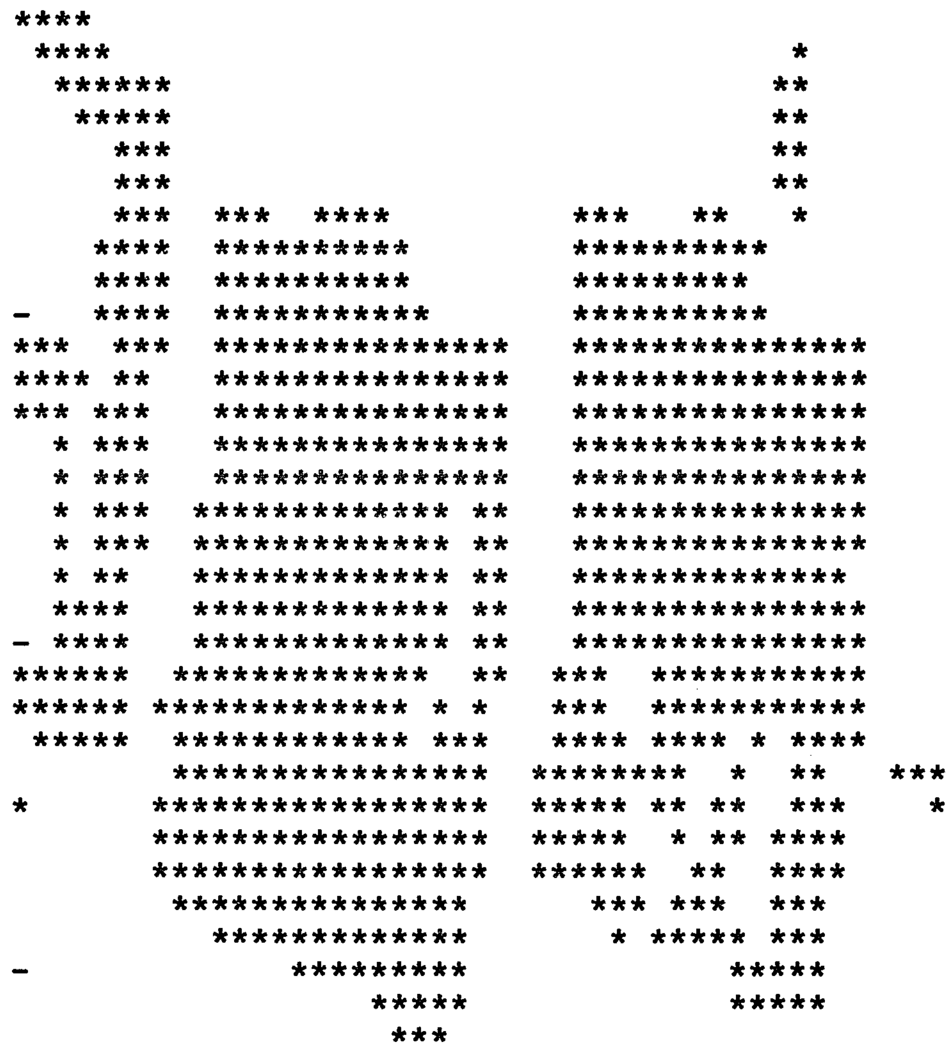

Figure 4.21 - Overlaying the Object Detector and Edge Detector Outputs (Building Example)

```
   ***  ****         ***   **
   **********        **********
   **********        *********
   ***********       **********
   ***************   ***************
   ***************   ***************
   ***************   ***************
   ***************   ***************
   ***************   ***************
  ************* **   ***************
  ************* **   ***************
  ************* **   **************
  ************* **   ***************
  ************* **   ***************
 *************  **  ***  ***********
************* * *   ***  ***********
 ************ ***   **** **** * ****
 ****************  ********  *  **
*****************  ***** ** **  ***
*****************  *****  * ** ****
*****************  ******  **  ****
 ***************      *** ***  ***
   *************       * ***** ***
       *********             *****
           *****             *****
            ***
```

Figure 4.22 - Result of Removing Non-Object Pixels (Building Example)

```
   ***  ****         ***   **
   **********        **********
   **********        *********
   ***********       **********
   ***************   ***************
   ***************   ***************
   ***************   ***************
   ***************   ***************
   ***************   ***************
  ****************   ***************
  ****************   ***************
  ****************   **************
  ****************   ***************
  ****************   ***************
 *****************  ****************
*****************   ****************
 ****************   ****************
 ****************  ***************
*****************  ****************
*****************  ****************
*****************  ****************
 ***************      ************
   *************       ***********
       *********             *****
           *****             *****
            ***
```

Figure 4.23 - Final Result of Building Object Enhancement (Building Example)

Figure 4.24 - Raw Output of the Object Detector (Road Example)

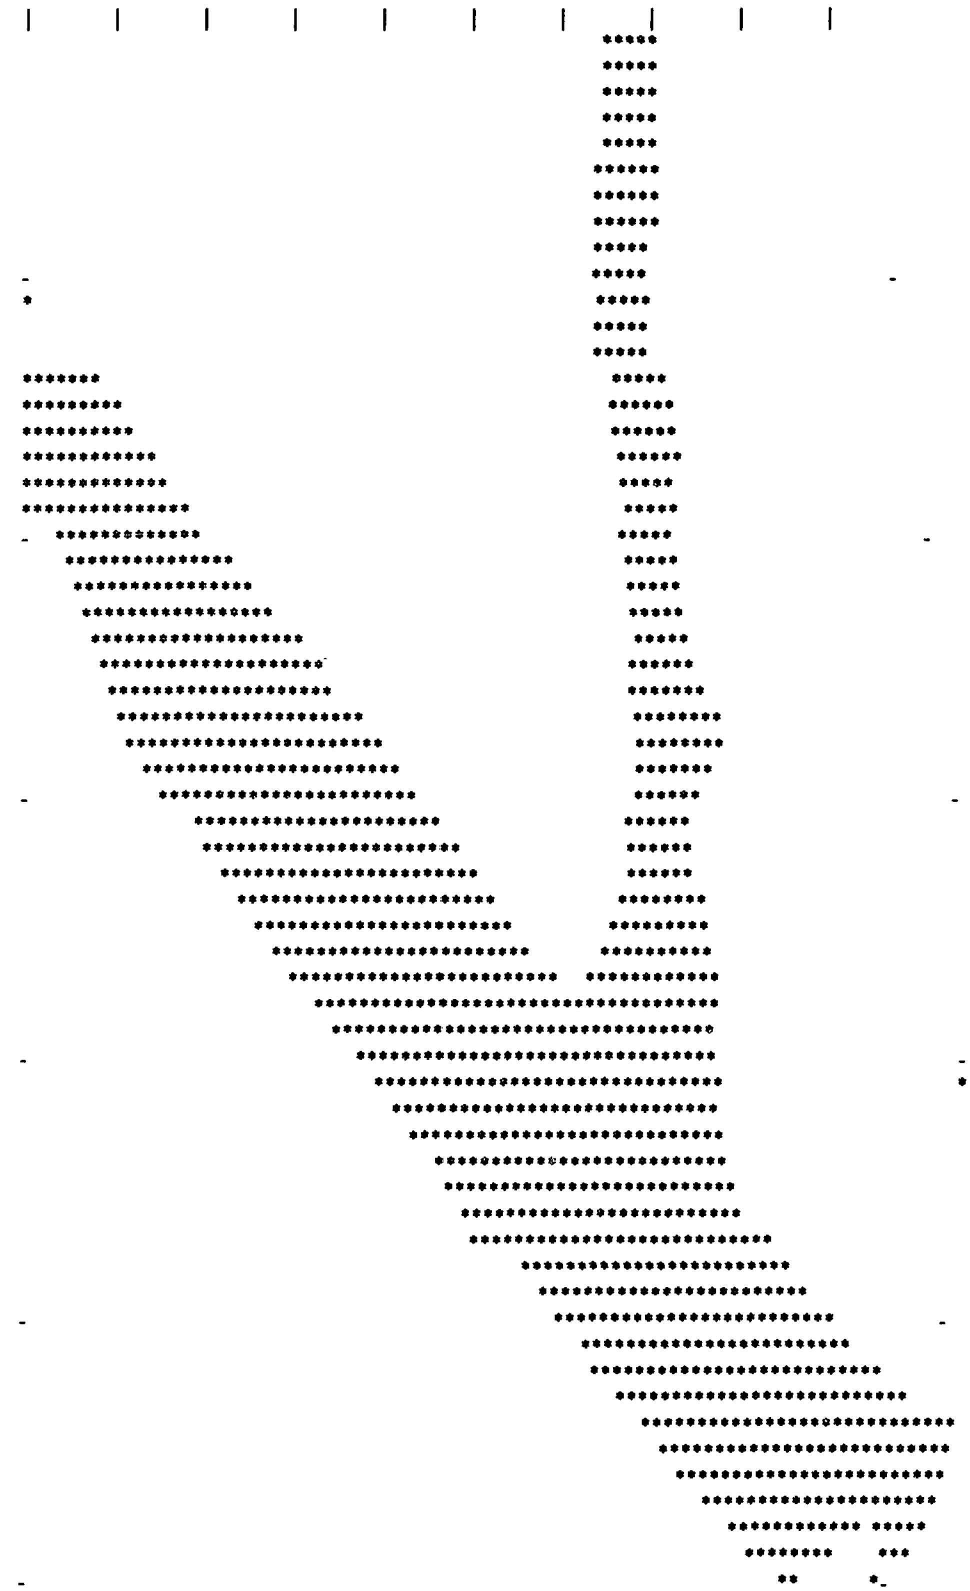

Figure 4.25 - Sidewalks Removed from Roads (Road Example)

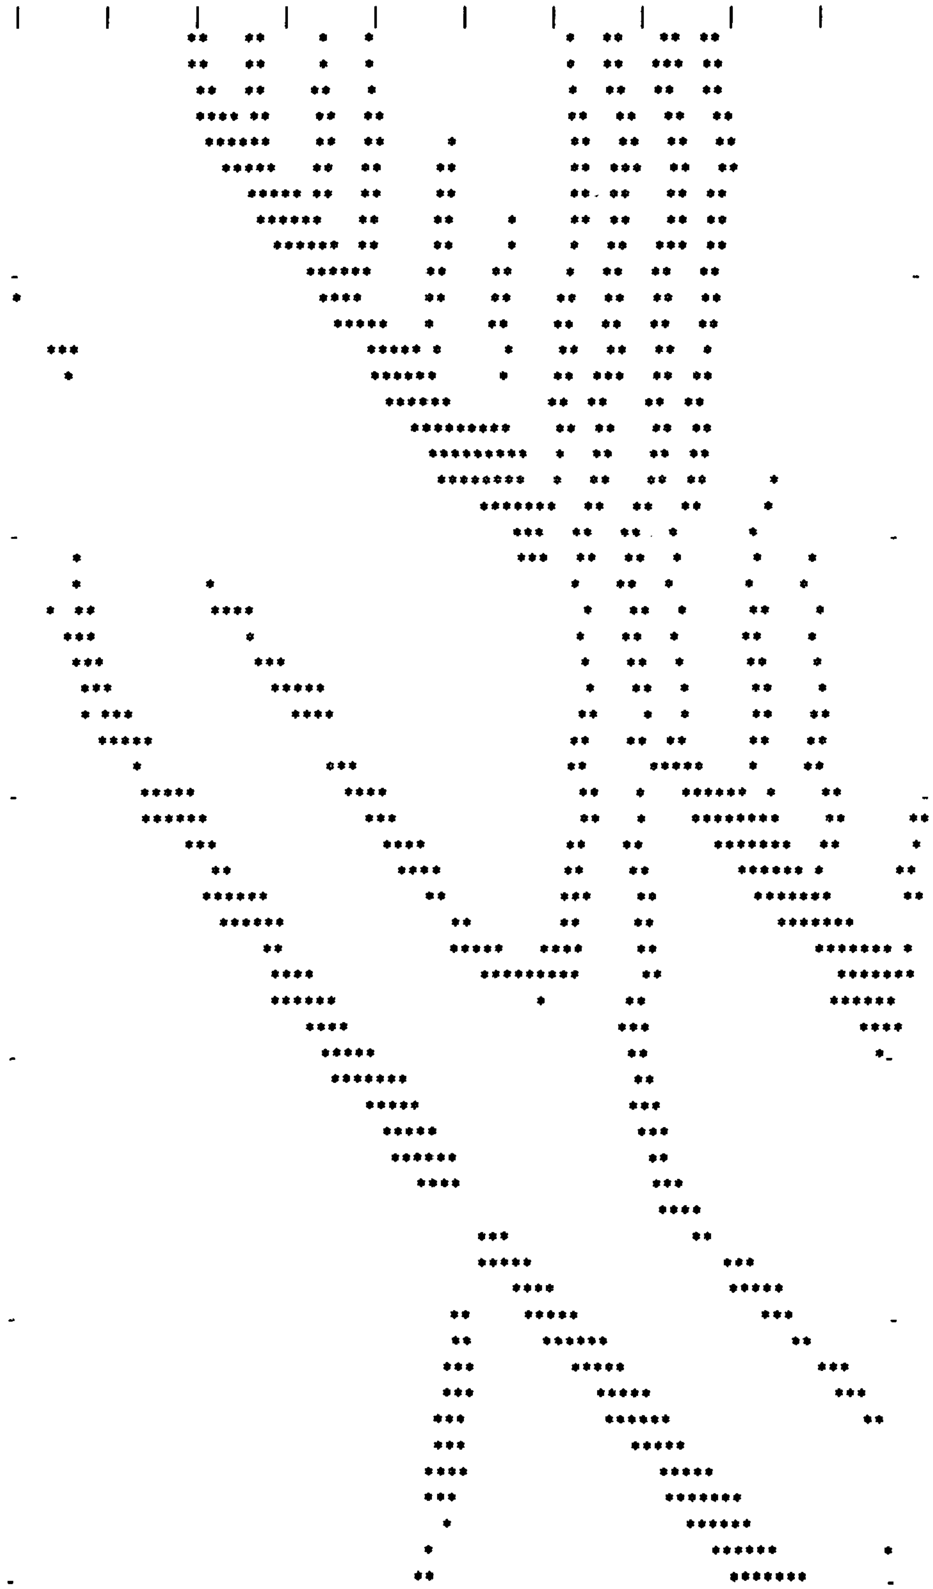

Figure 4.26 - Output of Kirsch Edge Detector (Road Example)

Figure 4.27 - Overlaying the Object Detector and Edge Detector Outputs (Road Example)

Figure 4.28 - Final Result After Removing Non-Object Pixels (Road Example)

Figure 4.29 - Raw Output of the Object Detector (Sidewalk Example)

Figure 4.30 - Final Result After Removing Non-Object Pixels (Sidewalk Example)

5 - Discussion of Approach

This chapter takes a close look at the approach proposed in this dissertation. Two vastly different types of images were examined successfully. This demonstrates the approach and system can analyze aerial photographs. This chapter closes with a section discussing the advantages this approach has over other approaches and systems. That section points out several unique abilities of this system.

5.1 - Examples of Image Analysis

The analysis of an image proceeds in a raster scan of 100x100 areas. The first area analyzed is in the upper left corner of the image. The next area is to the right of the first and so on. Figure 5.1 shows this. This figure shows how the system divides a 500x500 pixel image into 100x100 pixel areas. It also shows them numbered in order of analysis. The system analyzes the first area completely (it works its way down through the hierarchy and then finds the low level objects) before moving on to the second area and so on until it finishes the entire image. Analysis only works on 100x100 pixels at a time. This is a limitation of the compiler, personal computer technology, and operating system. The TDBU software uses an array of short type (8 bits per pixel) for the images. Because of the 64K byte limit on a single item of data, it is not practical to have arrays larger than 100x100. Therefore, the system restricts image analysis to 100x100 areas. The limitation had a major influence on the current implementation of the system. The 100x100 areas hold only a single type of area, i.e. they hold only an apartments area or only a wooded area but not both. If the limitation were not present, the system would analyze the entire 512x512 image at once. This would mean there would be more than one type of urban area in the image being analyzed. The basic structure of the system would not change. The implementation details, however, would change drastically. The system would have to look at cues from the entire image and then work to partition the image into naturally bounded areas, i.e. separate the housing areas from the wooded areas. This would be a challenging but welcome problem and will be addressed in the future.

Photograph 6 shows the final results of the analysis of the images shown in photographs 1, 2, and 3. Photographs 10, 11, and 12 show the results of analysis of the images in photographs 7, 8, and 9. These are the result of the combination top-down and

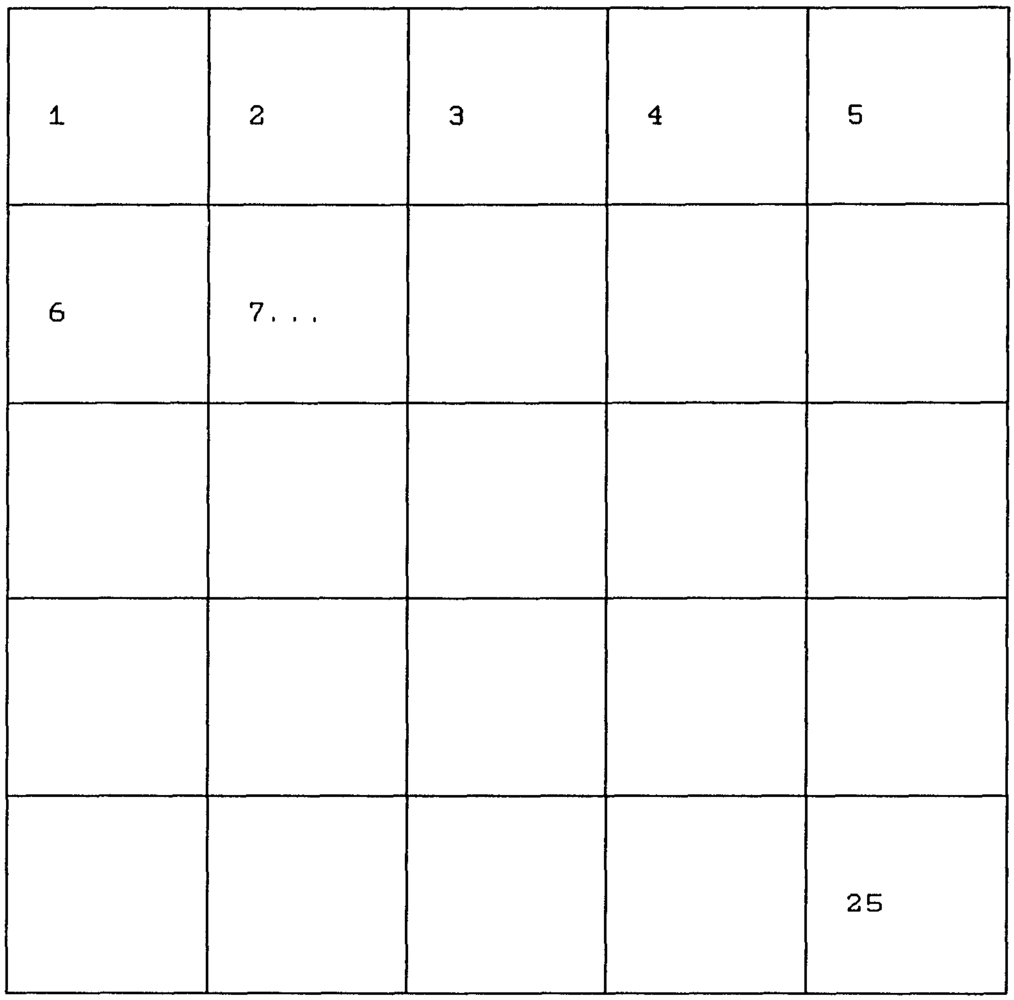

Figure 5.1 - Image Analysis Proceeds in a Raster Scan of 100x100 Pixel Areas

bottom-up processing approach applied throughout the levels of the DMA hierarchy. They are deceptive in that they appear to be the result of a simple object detector. They are the result of the system working its way down through the hierarchy until it reaches either the *commercial/residential* or *woodland* nodes. The object detector locates the final low level objects. The system could have stopped analysis of the areas in the image at any desired level. If the user desired analysis to the first level, then the system would have labeled each area *culture* or *landscape* as shown in figure 5.2. If the user desired analysis to the second level, then the system would have labeled each area *urban* or *phytography* as shown in figure 5.3. If the user desired analysis to the third level, then the system would have labeled each area *commercial/residential* or *woodland* as shown in figure 5.4.

culture	culture	culture	culture	culture
culture	culture	culture	culture	culture
land-scape	culture	culture	culture	culture
land-scape	culture	land-scape	culture	culture
land-scape	culture	land-scape	land-scape	culture

Figure 5.2 - Image Analysis to the First Level of the Hierarchy

urban	urban	urban	urban	urban
urban	urban	urban	urban	urban
phyto-graphy	urban	urban	urban	urban
phyto-graphy	urban	phyto-graphy	urban	urban
phyto-graphy	urban	phyto-graphy	phyto-graphy	urban

Figure 5.3 - Image Analysis to the Second Level of the Hierarchy

commercial/ residential	commercial/ residential	commercial/ residential	commercial/ residential	commercial/ residential
commercial/ residential	commercial/ residential	commercial/ residential	commercial/ residential	commercial/ residential
woodland	commercial/ residential	commercial/ residential	commercial/ residential	commercial/ residential
woodland	commercial/ residential	woodland	commercial/ residential	commercial/ residential
woodland	commercial/ residential	woodland	woodland	commercial/ residential

Figure 5.4 - Image Analysis to the Third Level of the Hierarchy

In photograph 6, the small blocks are buildings. The system correctly labeled them as *041 house (single family dwelling)*. The larger rectangular blocks are larger buildings. The system correctly labeled them as *040 apartments/motel*. The large areas are woods and the system correctly labeled them as *200 deciduous*. The roads are *069 primary road (dual lane, hard surface)* and the sidewalks are *097 footpath/trail*.

The residential area in the lower right hand corner of photographs 1, 2, and 3 will serve as an example of the analysis process. The first level of analysis labels the area as *culture*. This is because data-driven and goal-driven analysis finds man made materials. These man made materials differentiate between *culture* and *landscape*. Figure 5.5 shows this situation. The lower right hand corner of figure 5.5 gives a sketch showing houses, carports, sidewalks, and streets.

world

landscape culture

current label

candidates

Intrinsic characteristics of culture
Histogram Analyzer looks for materials
roofing, roofing2, aluminum

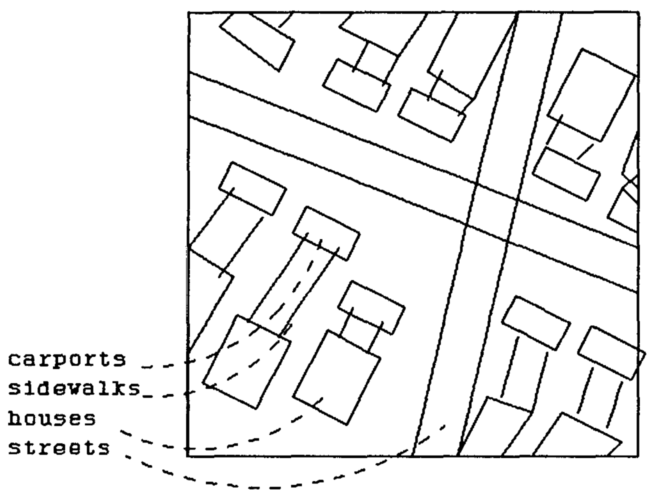

Figure 5.5 - First Level of Analysis

At the beginning of analysis, the Top_Frame or label was *world* and the choice for the next label was between *culture* and *landscape*. The sub node operator pointed to *culture* as the Candidate_frame. Therefore, the goal-driven investigation of the area used the intrinsic characteristics of *culture*. Figure 5.6 shows the intrinsic characteristics of *culture*. These are three man made materials. The histogram analyzer (described in section 4.5.2) takes these materials, looks up their spectral properties, and looks for the peaks of those properties in the smoothed histogram. The histogram analyzer detected these peaks in the smoothed histogram of the area. The presence of these peaks altered the belief vector in favor of *culture*. The system labeled the area *culture*.

```
The frame.is_a is:
   ->culture
The frame.is_part_of is:
   ->world
The frame.goal_of_analysis is:
   ->region label
The frame.level_in_tree is ->2


characteristic[0] is ->roofing
characteristic[0] function is ->histogram

characteristic[1] is ->2roofing
characteristic[1] function is ->histogram

characteristic[2] is ->aluminum
characteristic[2] function is ->histogram
```

Figure 5.6 - Intrinsic Characteristics of the *culture* Frame

The context makes the decision process simple. Expectations direct the operators to look for specific properties. Selective attention limits the processing to only the necessary parts of the area.

The first operator, an initial "segmentation," gave an interpretation. This interpretation directed the second segmentation. The second segmentation gave a second interpretation. This, when combined with the first interpretation, gave a final interpretation. The segmentation and interpretation processes worked together step by step through the analysis. This is top-down and bottom-up processing.

The next level of analysis labels the area *urban (built up areas)*. Figure 5.7 shows this situation. At the start of analysis the Top_frame was *culture* and the choices were *urban (built up areas), transportation/navigation*, and *landmark/rural features*. The sub node operator attributes more belief to *urban (built up areas)* than to *transportation/navigation* and *landmark/rural features*. The Candidate_frame becomes *urban (built up areas)* and goal driven analysis begins with the intrinsic characteristics of *urban (built up areas)*. Figure 5.8 shows the intrinsic characteristics of *urban (built up areas)*. These are simple, generic objects found in an urban area and neither in a transportation network nor a rural area. Since *urban (built up areas)* is on level three of the hierarchy, the object detector works on a reduced resolution image. The system reduces the resolution by a factor of four, i.e. the system represents a 4x4 pixel area with 1 pixel.

world

landscape culture current label

candidates

transportation/ navigation landmark/ rural features urban

Intrinsic Characteristics of urban
(object detector reduced resolution
by a factor of 4)
b1_building p1_parking_lot g1_garage

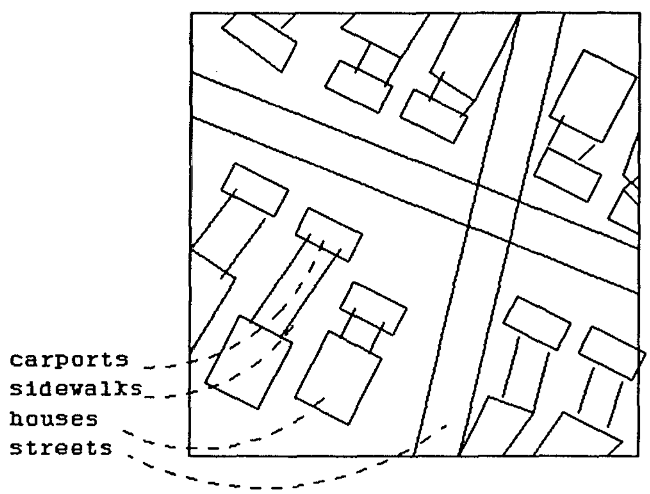

Figure 5.7 - Second Level of Analysis

```
The frame.is_a is:
   ->urban (built up areas)
The frame.is_part_of is:
   ->culture
The frame.goal_of_analysis is:
   ->region label
The frame.level_in_tree is ->3

characteristic[0] is ->b1_building
characteristic[1] is ->p1_parking_lot
characteristic[2] is ->g1_garage
```

Figure 5.8 - Intrinsic Characteristics of the *urban (built up areas)* Frame

The object detector located the objects. Appendix 7 section 1 lists the output of the object detector. The presence of the objects altered the belief vector in favor of *urban (built up areas)*. The system labeled the area *urban (built up areas)*.

The next level of analysis labels the area *commercial/residential.* Figure 5.9 shows this situation. At the start of the analysis the Top_Frame was *urban (built up areas)* and the choices were *industry/utility, commercial/residential,* and *institutional/governmental.* The sub node operator attributes more belief to *commercial/residential.* The Candidate_frame becomes *commercial/residential* and the system investigates the area using the intrinsic characteristics of *commercial/residential.*

world

landscape culture

transportation/ navigation

landmark/ rural features

urban

current label

candidates

industry/ utility

institutional/ governmental

commercial/ residential

Intrinsic Characteristics of commercial/ residential (object detector reduced resolution by a factor of 2)

040 apartments
hotel

041 house
single family
dwelling

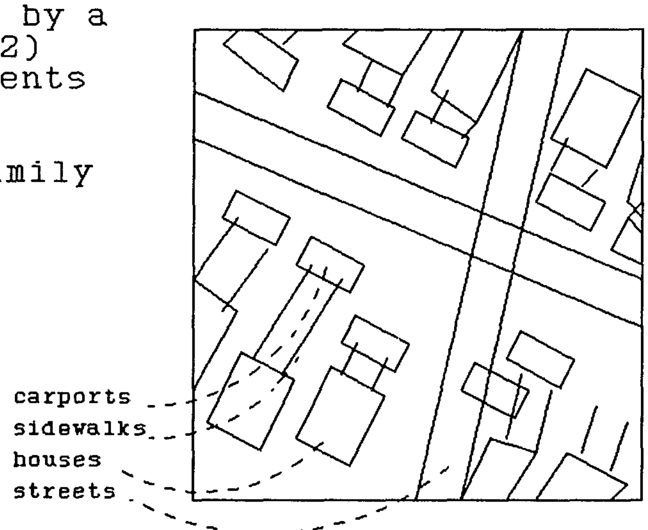

Figure 5.9 - Third Level of Analysis

Figure 5.10 shows the intrinsic characteristics of *industry/utility, commercial/residential*, and *institutional/governmental.* Note that *industry/utility* has sub nodes below it in the hierarchy, but *commercial/residential* and *institutional/governmental* do not. These two, *commercial/residential* and *institutional/governmental*, are at the bottom of the hierarchy (level in tree is 99) and their sub nodes are individual objects.

The frame.is_a is:
->*industry/utility*
The frame.is_part_of is:
->*urban (built up areas)*
The frame.goal_of_analysis is:
->region label
The frame.level_in_tree is ->4

characteristic[0] is ->b4_building
characteristic[1] is ->d1_disposal

The frame.is_a is:
->*commercial/residential*
The frame.is_part_of is:
->*urban (built up areas)*
The frame.goal_of_analysis is:
->objects label
The frame.level_in_tree is ->99

characteristic[0] is ->040 apartments/hotel
characteristic[1] is ->041 house (single family dwelling)

The frame.is_a is:
->*institutional/governmental*
The frame.is_part_of is:
->*urban (built up areas)*
The frame.goal_of_analysis is:
->region label
The frame.level_in_tree is ->99

characteristic[0] is ->045 governmental administration building
characteristic[1] is ->046 military admin/operations building

Figure 5.10 - Intrinsic Characteristics of Frames *industry/utility, commercial/residential*, and *institutional/governmental*

Using the intrinsic characteristics of *commercial/residential*, the object detector located several instances of *041 house (single family dwelling)*. Appendix 7 section 2 lists the output of the object detector. This altered the belief vector in favor of *commercial/residential*. The system labeled the area *commercial/residential*.

The final level of analysis finds and labels the objects or leaf nodes in the hierarchy. Figure 5.11 shows this situation. As discussed earlier, at this point the system departs from the top down traversal of the hierarchy. The system used the leaf nodes of *commercial/residential* and the leaf nodes of the other frames on the same level of the hierarchy. Figure 5.12 shows these objects. In this part of the image, the system found *041 house (single family dwelling)*, *069 primary road*, and *097 footpath/trail*. Appendix 7 section 3 lists the output of the object detector. At this final level, the object detector also inserts the objects into the output image. Photograph 6 shows this output image. This explains the appearance of photograph 6. It is the final output of the object detector at the lowest level of analysis.

world

landscape culture

transportation/ navigation

landmark/ rural features

urban

current label

industry/ utility

institutional/ governmental

commercial/ residential

Candidate Objects

040 apartments/ hotel
041 house single family dwelling
069 primary road
097 footpath/ trail

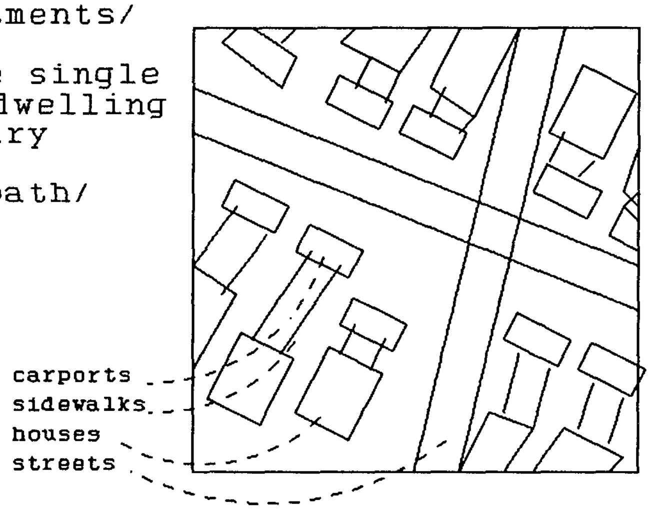

Figure 5.11 - Fourth Level of Analysis

The system has the ability to move up and down the branches of the hierarchy. The above example was a simple downward traversal of the hierarchy. There are, however, cases where the labeling process is not so simple. Suppose the analysis was at the *culture* node. If the first step data-driven operator pointed to *transportation/navigation*, then it would become the candidate node. Analysis would move down to *transportation/navigation*. The system would use the intrinsic characteristics of *transportation/navigation* to investigate the area closely. The close investigation could produce evidence unfavorable to *transportation/navigation*. In that case the analysis would go back up to *culture* and then down to another sub node such as *urban (built up areas)*. The system would then investigate the area using the intrinsic characteristics of *urban (built up areas)*. Depending on the evidence returned at this point, the system could decide to label the area *urban (built up areas)* or it could move back up to *culture* and then down to another sub node.

It is important to realize that the system also labels areas of woods and finds the trees. The large white sections of photograph 6 are wooded areas. The system correctly analyzed these areas as *woodland* and then labeled the patches of trees. The reader should not be mislead by the title of this dissertation. "Urban" areas are not limited to just housing areas. A city is an urban area. A city usually has parks, woods, factories, and transportation networks as well as residential areas. Photographs 1, 2, and 3 show an urban area that has single family homes, apartments, and woods. The system correctly finds the *commercial/residential* areas and the *woodland* areas.

The areas in the lower left hand corner of photographs 1, 2, and 3 serve as an example of labeling a *woodland* area. Figure 5.13 shows the hierarchy traversal for a wooded area. The first level of analysis labels the area *landscape*. The data-driven and goal-driven analysis could not find any man made materials. At the beginning of analysis, the Top_Frame was *world* and the choice was between *culture* and *landscape*. The first step operator pointed to *landscape*. Closer investigation, at this level the histogram analyzer, agreed with this and the system labeled the area *landscape*.

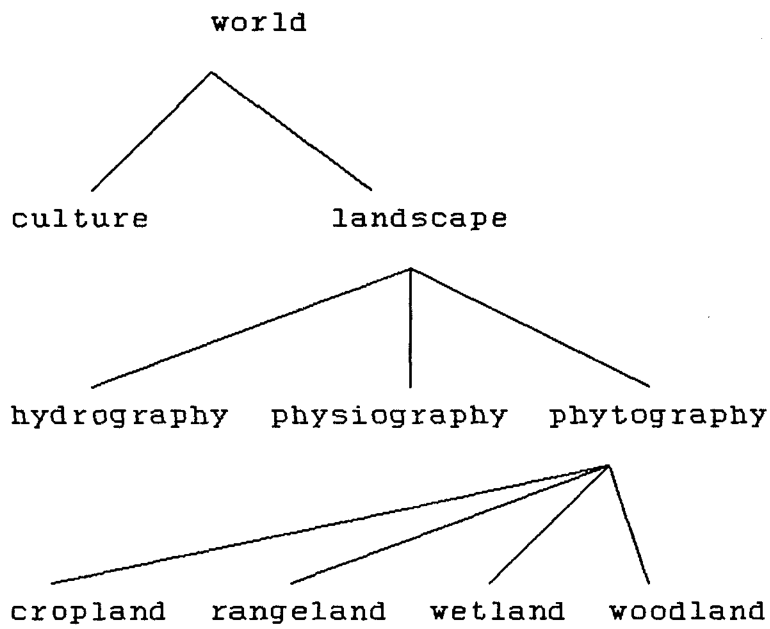

Candidate Objects of woodland
200 decidous
201 caniferous
202 mixed (decidous and caniferous)
203 mangrove
204 nipa palm

Figure 5.13 - Hierarchy Traversal of Wooded Area

The next level of analysis labeled the area *phytography*. At the start of analysis at this level, the Top_Frame was *landscape* and the choices were *hydrography, physiography*, and *phytography*. The first data-driven operator pointed to *phytography*. The careful, goal-driven analysis found wooded type objects (various foliage). The combination of these processes caused the system to label the area *phytography*.

The next level of analysis labeled the area *woodland*. At the start of analysis at this level, the Top_Frame was *phytography* and the choices were *cropland, rangeland, woodland*, and *wetland*. Again the analysis began with a first, data-driven operator and ended with careful investigation by a goal-driven object detector. The result was the system labeled the area *woodland*.

Now the system is at the bottom of the label hierarchy. The final step is to have the object detector locate the individual objects pertaining to this type of area. The *woodland* node of the hierarchy is node 29 in appendix 2. The individual objects below it are different types of trees. The system directs the object detector to look for these. The object detector found a large area of deciduous trees. This is the final output of the system for the *woodland* areas.

The system also analyzed the images shown in photographs 7, 8, and 9. These images came from a color aerial photograph. A Hewlett-Packard ScanJet Plus desktop scanner produced the separate images by scanning separate areas of the original photograph. The scanner produced 256 gray levels. The spectral properties of the different areas were vastly different from the 3 channel image described earlier and slightly different from each other. Because of this, the spectral part of the knowledge base was changed. That, however, was the only change required in the entire system.

Photograph 10 shows the results of analysis of photograph 7 (note, in photographs 10, 11, and 12 the black objects are *041 house (single family dwelling)*, the dark gray objects are *097 footpath/trail*, and the light gray objects are *069 primary road*). The system correctly labeled the areas as *commercial/residential*. The object detector labeled the homes, sidewalks, and roads.

Photograph 11 shows the results of analysis of photograph 8. The system correctly labeled the areas as *commercial/residential*. The object detector labeled the homes, sidewalks, and roads. The object detector incorrectly labeled four homes as sidewalks. This shows the weakness of the object detector currently used by the system. Again, developing accurate low level operators is not the purpose of this dissertation.

Photograph 12 shows the results of analysis of photograph 9. The system correctly labeled the areas as *commercial/residential*. The object detector labeled the homes, sidewalks, and roads.

039 commercial buildings
040 apartments/hotel
041 house (single family dwelling
042 mobile homes
043 barracks
044 display sign (large billboard, highway sign)

045 governmental administration building
046 military admin/operations building
047 capitol building
048 hospital
049 prison
050 palace
051 museum
052 observatory
053 church/tabernacle
054 mosque
055 cemetery building

056 single track railway
057 double track railway
058 multiple track railway
059 RR yard/siding
060 tramway/inclined railway
061 monorail
062 RR storage/repair building
063 RR terminal building
064 RR station/depot
065
066 roundhouse

067 multi lane, divided (grass median) highway
068 multi lane highway
069 primary road (dual lane, hard surface)
070 secondary road (dual lane, loose/dirt surface)
071 trail/track (one lane)
072 toll gates
073 cloverleaf/interchange
074 garage, service/repair facilities (landmark)

Figure 5.12 - Sub Nodes of the *commercial/residential* frame

5.2 - Advantages of Approach

A key question in the development of this approach and system is "given an aerial photograph of a known urban area, what is the difference between this approach and system and other systems?" The following pages discuss the advantages that this approach and system has over other systems. These were mentioned briefly in chapter 1.

A major advantage is that each point in the hierarchy has implicit knowledge imbedded in it because of the context and framework of the hierarchy. The amount and power of this knowledge cannot be underestimated. This knowledge frees the low level operators from worrying with unnecessary details. For example, suppose the system was at the *urban (built up areas)* node (refer to figure 5.14) and it was looking for linear features. The linear features here would be sidewalks and roads. The linear feature detector would not consider things such as streams, runways, or rows of crops. This would free the detector from large amounts of details and complications. In simple terms, if the linear feature detector was a rule-based operator employing 10,000 rules, then it could eliminate 9000 of those rules because of the implicit knowledge imbedded in the position in the hierarchy.

Some of the major advantages to this approach are in the area of reducing the computational complexity. This is to be expected since a key to the approach is the hierarchy of labels. In general, hierarchies reduce the complexity of problems so an approach based on a high level hierarchy would have reduced complexity.

First, the system uses different operators at different levels of analysis. At the top level of analysis, the system uses a histogram analyzer. This is a simple and quick operator yet it is sufficient for the task at that level. At the next level of analysis the system uses an object detector working on a reduced resolution image. At this level each pixel represents a 4x4 pixel area in the original image. Therefore, the system reduces the size of the problem by one fourth. At the next level of analysis the system uses an object detector working on a reduced resolution image. At this level each pixel represents a 2x2 pixel area in the original image. Therefore, the system reduces the size of the problem by one half.

A major reduction in complexity comes in the way in which the system uses the

Dempster-Shafer (D-S) theory of evidence. The D-S theory has a number of advantages as described in section 3.3. There are, however, computational problems with the D-S theory. Simplifications exist if the hypotheses are in a hierarchy. This approach uses the hierarchy of labels (hypotheses) and takes advantage of these simplifications. Therefore, the number of meaningful hypotheses is 247 (the number of nodes in the hierarchy). If the hierarchy were not used, a flat set of final hypotheses would be required. There would be 29 singleton hypotheses and 2^{29} or 536,870,912 meaningful hypotheses. The combination of two simple assertions would require this many floating point multiplication, addition, and division operations. This reduction in complexity simply comes from the use of the hierarchy.

Another reduction in complexity comes from the object models used by the object detector. In the higher levels of the hierarchy the system uses generic object models. At the lowest level of analysis the system uses specific objects such as *040 house (single family dwelling)* or *041 apartments/motel*. In the higher levels the system uses generic objects such as b1_building, b2_building, or l1_lawn. These are simple, flexible models that do not require a high degree of accuracy or complexity from the operators.

An advantage to this approach is it can label all types of urban areas - not just one. It can label areas such as woods, residential, airports, and transportation networks. The system demonstrated this in chapter five. It successfully analyzed a multi-spectral image containing three types of urban areas - a single family home area, an apartments area, and a wooded area. It also successfully analyzed three other images containing a fourth type of urban area - a house trailer park. The current implementation, however, cannot analyze all types of areas given in the DMA hierarchy. The limiting factor is the availability of images with which to experiment. Other research systems could make the same claim. If they were given an image, they could write new code and adapt their system to the new image. There is a major difference in this point. The system in this dissertation has a framework for the entire world (at least as seen by the Defense Mapping Agency). The hierarchy of frames and labels is all inclusive. The frames for icebergs and railroad yards haven't been created, but places are available for them. Other research systems are built to label the objects in a known image type. They would need new operators and models for each particular image.

The next two unique abilities of the approach concern working on real world problems involving hundreds and thousands of images. These abilities are (1) stopping analysis at any level in the hierarchy and (2) finding target type images and eliminating non-target images.

This approach can stop its analysis at any level in the DMA hierarchy. The system can label an area with any intermediate node. For example, if the desired level of analysis is two, the system would label an area *culture* or *landscape* (refer to figure 5.14). If the desired level of analysis is three, the system would label the area *urban (built up areas), transportation/navigation, landmark/rural features, hydrography, physiography,* or *phytography.*

This ability is crucial for automated aerial photograph analysis systems. Suppose you were given 10,000 images on a magnetic tape and were required to label them to level 3 (listed above). This problem requires neither a set of complex, state of the art operators nor a costly, expert, human photointerpreter. The approach is uniquely suited for this practical, real world task. It can quickly work its way down to this level of the hierarchy using the simple histogram analyzer and the object detector employing a reduced resolution.

This approach can quickly locate target type images and eliminate all non-target images. For example, suppose a special interest group desires to find the *aeronautical/aerospace* images from a group of 10,000 images (again refer to figure 5.14). Using this approach, the system would start at the *world* node and work its way down through the hierarchy in each image. If the system branches down the *landscape* side of the hierarchy, then it would eliminate that image at this point. Further processing on that image is not necessary. If the image was classified as *culture*, then the processing would continue. If the next classification was *urban (built up areas)* or *landmark/rural* features, then the system would stop processing the image at that point. The calculations performed to this point are quick and simple. The effort expended to eliminate non-target images is small. Using these techniques, the system could quickly find the target images and eliminate the non-target ones. It could work through thousands of images in short order.

Other systems cannot perform this type of task this quickly and cheaply. If a system was designed to detect runways, airplanes, and buildings in an airport, it would have a long, difficult time eliminating non-airport images. If it was given a forest image, it would work for hours on the hundreds of confusing line segments found in the forest image. It would

```
world

 ├─ culture
 │
 │   ├─ urban (built up areas)
 │   │
 │   │   ├─ industry/utility
 │   │   ├─ commercial/residential
 │   │   └─ institutional/governmental
 │   │
 │   ├─ transportation/navigation
 │   │
 │   │   ├─ railroads
 │   │   ├─ roads
 │   │   ├─ aeronautical/aerospace
 │   │   ├─ naval/marine
 │   │   └─ associated transportation features
 │   │
 │   └─ landmark/rural features
 │
 │       ├─ communication/transportation
 │       ├─ storage
 │       ├─ agricultural
 │       ├─ recreational
 │       └─ miscellaneous
 │
 └─ landscape

     ├─ hydrography
     │
     │   ├─ water
     │   └─ snow/ice
     │
     ├─ physiography
     │
     │   ├─ exposed soils (surface composition)
     │   └─ landforms
     │
     └─ phytography

         ├─ cropland
         ├─ rangeland
         ├─ woodland
         └─ wetland
```

Figure 5.14 - DMA Hierarchy (top portion)

have such problems because researchers tailor the operators in airport analysis systems to the regular, straight, connecting lines in airport images.

Another advantage of this system is the ease of adding operators from other research efforts. All of the operators are linked to the system as attached procedures. They are attached to the frame that needs them. Each frame at each level of the hierarchy has different goals and needs. They can each have a different operator attached to meet these goals and needs. The *commercial/residential* frame needs operators that can find houses, sidewalks, yards, and streets. The *woodland* frame needs texture operators to indicate the texture of a wooded area.

Since the operators are attached procedures, it is easy to attach or remove them from the system. Section 4.2 discussed this process. The major problem in the process is obtaining research operators in the form of portable, usable C language functions. Most researchers do not develop their operators in this manner.

Another advantage to the approach is that it strongly couples the segmentation and interpretation processes in the high level analysis. Segmentation and interpretation cannot be separated. The current state of the art in segmentation operators cannot segment the image data well enough for interpretation systems to work with their results. The interpretation process must guide and direct the segmentation operators, use their results, and then direct them again. The system in this dissertation uses such an idea. Its basis is the perceptual cycle.

A major advantage to this system is it is portable. It can move easily from image type to image type and from computer system to computer system. It can move from image type to image type because the knowledge is located in a separate, easy to modify portion of the knowledge base. The spectral information is the relative reflectance of the materials in the images. This information may change from image type to image type. It changes because of different types of scanners and atmospheric conditions. The two vastly different types of images analyzed in chapter five demonstrated how easy it is to move the system from one image type to another.

The system is written in the C language. This is a common, inexpensive language that is standard in most computing environments. Many research systems are written in custom languages or using rare, hard to find, and expensive artificial intelligence dialects.

The system is not tied to a host image processing software system. Therefore, it can be ported to any computer system with a C compiler.

This final advantage may be the most important because it makes the system something that the vast majority of research systems are not - practical. The system was created in a practical environment with real economic and computer constraints. The author was forced to move the system among computers, compilers, and operating systems. These moves made the author consider problems and constraints that usually do not enter into a dissertation project.

6 - Conclusions

This dissertation presented a new approach to analyzing aerial photographs and described a system created around this approach. There are several different problems in the analysis of aerial photographs. These problems include ambiguities, guiding low level operators, processing only the salient features, linking operators together to help themselves and other operators, explaining the results of less than perfect operators, and tying the system together with an overall theme. Ambiguities in aerial photograph interpretation systems relate to the explosion of possible objects in possible scenes. The system must select and guide low level operators. The system must apply them to the image only if they are appropriate for the given situation. The operators must process certain salient features that contain the minimum information required to recognize something. The system must link the different operators together to feed information to each other and direct one another. The low level operators need to work with the understanding portion of the system to improve the accuracy of their results.

This dissertation presented three basic ideas as the foundation for a new analysis approach. The ideas are (1) the use of selective attention, expectations, context, and the perceptual cycle, (2) analyzing the image through a hierarchy of increasingly specific labels, and (3) the interplay between the segmentation and interpretation processes. These concepts have been expressed in many ways in the perceptual psychology, artificial intelligence, and computer vision literature. Nevertheless, they have never been tied together as the basis for an approach to analyzing aerial photographs.

Humans use selective attention. We direct our attention to only those items which interest us and ignore all else. Expectations play an important role in perception because they guide our attention. Context works closely with expectations in driving human perception. The context of a situation limits the number of alternatives and, thereby, reduces the difficulty in decision making. Humans perceive things using the perceptual cycle. This is a simple feedback system. The point is that knowledge of the world directs our exploration. The findings modify the knowledge which then directs the exploration in a modified manner.

A key to the overall approach is the hierarchy of labels. The hierarchy of labels allows the system to work gradually to the point where it knows the type of scene in the image. At this point, it is ready to identify the pertinent objects. Each node in the hierarchy has only a few possible choices. This greatly reduces the complexity of the problem. It also reduces the required capabilities of the system.

These concepts drove the creation of a computer vision system. The computer vision system created using this approach is simple and modular. It has four basic parts: (1) the knowledge base, (2) the belief maintenance system, (3) the control mechanism, and (4) the image processing interface. The control mechanism implements the cycle and the interpretation part of the interpretation and segmentation interplay. The knowledge base is the hierarchy and also imbeds the ideas of expectations, selective attention, and context. The belief maintenance system works with the control mechanism to draw conclusions from the analysis. The image processing interface is the segmentation portion of the segmentation and interpretation interplay.

The approach described in this dissertation and the system created herein have several advantages over other systems. A major advantage is the implicit knowledge imbedded in each point in the hierarchy. This frees low level operators from trying to solve problems that are more difficult than exist. Among other advantages is the reduction in the computational complexity of several aspects of the system. The system is based on a hierarchy of labels. Hierarchies inherently reduce the complexity of problems. Therefore, it is no surprise that the hierarchy brings several advantages. The system uses different operators at different levels of analysis. The operators at the higher levels of analysis are simple and quick yet effective. The Dempster-Shafer theory of evidence receives several simplifications from the use of a hierarchy that reduces the computations. In this case the reduction is from approximately 500,000,000 to 247.

Before closing, a few words must be directed towards the DMA hierarchy. This hierarchy was chosen for this project back in 1986. Several weaknesses surfaced as the project progressed. The greatest weakness is at the bottom of the hierarchy. Suppose the system labels an image as *commercial/residential* (see node 11 in Appendix 2). You would expect to find objects such as roads, yards, trees, homes, offices, sidewalks, etc. in this area and spatial relationships among the objects. These would help describe the objects and help the low level operators find them. You would also expect to find these objects under the *commercial/residential* node in the hierarchy. This, however, is not the case. As was pointed out in chapter four, the system must break from the DMA and reach out across the hierarchy at this point. This introduced an unwanted inconsistency into the system.

As stated at the beginning, analyzing aerial photographs is beyond the scope of a single dissertation. There is much work yet to be done. The future work concerns more images, better computing resources, operators from other researchers, work on the belief maintenance system, modifications to the DMA hierarchy, and the inclusion of spatial relationships in the frames.

The system needs low level operators from other researchers. One of the advantages of this system is that it can accept new operators easily - if they are portable operators. If the operators are not portable, then this author must work on them either as sent by other researchers or by writing them from descriptions in the literature. Creating low level operators was not the goal of the dissertation. The system's current operators are not the best. Incorporating other operators can only improve the system and also prove the flexibility of this system.

The belief maintenance system needs expanding. The current system implements the D-S theory only for those portions of the hierarchy used in the example images. Code must be written for the remainder of the hierarchy. That would be a sizeable programming effort (over 10,000 lines of code). Another possibility is to revise completely the manner in which the D-S theory is implemented. This would be interesting research into the implementation details of the D-S theory.

The DMA hierarchy needs modifications. This dissertation pointed out several weaknesses of the hierarchy with the main weakness at the bottom of the hierarchy. Modifications would include tying the low level objects in the DMA hierarchy to more than one parent. This would allow searching for the objects regardless of the final label for an area under investigation. This would remove the inconsistency in the system's analysis algorithm. A major modification would be the inclusion of spatial relationships among the low level objects. This is part of the implicit knowledge imbedded in each node in the hierarchy. This concept was used in [Harlow, et. al.], but has not yet been incorporated into the current system. This would help direct the low level operators and would prevent the omission of fuzzy or occluded low level objects in the final results.

Bibliography

[1] Amadasun, M., R. A. King, "Low-Level Segmentation of Multispectral Images Via Agglomerative Clustering of Uniform Neighborhoods," Pattern Recognition, Vol. 21, No. 3, 1988, pp. 261-268.

[2] Amat, Jean-Louis, Jean-Pierre Rogala, Frederic Momal, "Design of a Vision System in Remote Sensing: Modeling Image Processing Sequences," Eighth International Conference on Pattern Recognition, IEEE Computer Society Press, October 1986, pp. 1244-1246.

[3] Arkin, Ronald C., Edward Riseman, Allen Hanson, "Visual Strategies for Mobile Robot Navigation," Proceedings of the IEEE Computer Society Workshop on Computer Vision, Computer Society Press of the IEEE, December 1987, pp. 176-181.

[4] Aviad, Z., P. D. Carnine Jr., "Road Finding for Road Network Extraction," Proceedings CVPR 88 The Computer Society Conference on Computer Vision and Pattern Recognition, The Computer Society Press of the IEEE, June 1988, pp. 814-819.

[5] Ballard, Dana H., Altan Ozcandarli, "Eye Fixation and Early Vision: Kinetic Depth," Second International Conference on Computer Vision, IEEE Computer Society Press, December 1988, pp. 524-531.

[6] Barr, Avron, Edward A. Feigenbaum, "The Handbook of Artificial Intelligence, Volume 1," William Kaufman Inc., 1981.,

[7] Barr, Avron, Edward A. Feigenbaum, "Frames and Scripts," pp. 216-222 of "The Handbook of Artificial Intelligence."

[8] Barrow, Harry G., Jay M. Tannenbaum, "Computational Vision," Proceedings of the IEEE, Vol. 69, No. 5, May 1981.

[9] Barry, Michele, David Cyrluk, Deepak Kapur, Joseph Mundy, Van-Duc Nguyen, "A Multi-Level Geometric Reasoning System for Vision," Artificial Intelligence, Vol. 37, 1988, pp. 291-332.

[10] Basu, S., Image Segmentation by Semantic Method," Pattern Recognition, Vol. 20, No. 5, 1987, pp. 497-511.

[11] Bennet, Bruce M., Donald D. Hoffman, Chetan Prakash, "Perception and Computation," First International Conference on Computer Vision, IEEE Computer Society Press, June 1987, pp. 356-364.

[12] Besl, Paul J., Ramesh C. Jain, "Segmentation Through Variable-Order Surface Fitting," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 10, No. 2, March 1988, pp. 167-192.

[13] Bhanu, Bir, Bahram A. Parvin, "Segmentation of Natural Scenes," Pattern Recognition, Vol. 20, No. 5, 1987, pp. 487-496.

[14] Biederman, Irving, "Human Image Understanding: Recent Research and a Theory," Computer Vision, Graphics, and Image Processing, Vol. 32, October 1985, pp. 29-73.

[15] Binford, Thomas O., "Survey of Model-Based Image Analysis Systems," The International Journal of Robotics Research, Vol. 1, No. 1, Spring 1982, pp. 18-64.

[16] Bobrow, D.G., A. Collins, eds., "Representation and Understanding," Academic Press, 1975.

[17] Brown, Christopher, Editor, "Advances in Computer Vision, Volume 1," Lawrence Erlbaum Associates, Publishers, 1988.

[18] Browse, Roger A., Marion G. Rodrigues, "Propagation of Interpretations Based on Graded Resolution Input," First International Conference on Computer Vision, IEEE Computer Society Press, June 1987, pp. 405-410.

[19] Bruner, Jerome S., A. Leigh Minturn, "Perceptual Identification and Perceptual Organization," The Journal of General Psychology, Vol. 53, 1955, pp. 21-28.

[20] Burt, Peter J., "Smart Sensing within a Pyramid Vision Machine," Proceedings of the IEEE, Vol. 76, No. 8, August 1988, pp. 1006-1015.

[21] Castleman, Kenneth R., "Digital Image Processing," Prentice-Hall, 1979.

[22] Catros, J.Y., D. Mischler, "An Artificial Intelligence Approach for Medical Picture Analysis," Pattern Recognition Letters, Vol. 8, No. 2, September 1988, pp. 123-130.

[23] Chassery, Jean-Marc, Catherine Garbay, "Expert Systems, Image Processing and Image Interpretation," Eighth International Conference on Pattern Recognition, IEEE Computer Society Press, October 1986, pp. 175-177.

[24] Chen, Chu-Xin, "An Approach to Object Detection in Mulltispectral Images," Master's Thesis, Louisiana State University, May 1986.

[25] Cheng, Yizang, Rangasami L. Kashyap, "Study of the Different Methods for Combining Evidence," SPIE Vol. 635, Applications of Artificial Intelligence III, 1986, pp. 384-393.

[26] Chien, C.H., J.K. Aggarwal, "Shape Recognition from Single Silhouettes," First International Conference on Computer Vision, IEEE Computer Society Press, June 1987, pp. 481-490.

[27] Clark, James J., Nicola J. Ferrier, "Modal Control of an Attentive Vision System," Second International Conference on Computer Vision, IEEE Computer Society Press, December 1988, pp. 514-523.

[28] Conners, RIchard W., Mohan M. Trivedi, Charles A. Harlow, "Segmentation of a High-Resolution Urban Scene Using Texture Operators," Computer Vision, Graphics, and Image Processing, Vol. 25, No. 3, March 1984, pp. 273-310.

[29] Daily, Michael J., John G. Harris, Kurt Reiser, "An Operational Perception System for Cross-Country Navigation," Proceedings CVPR 88 The Computer Society Conference on Computer Vision and Pattern Recognition, The Computer Society Press of the IEEE, June 1988, pp. 794-802.

[30] Dainis, George A., III, Matthew O. Ward, "Rule-Based Inspection of Leadframes," Proceedings CVPR 88 The Computer Society Conference on Computer Vision and Pattern Recognition, The Computer Society Press of the IEEE, June 1988, pp. 580-585.

[31] Darwish, Ahmed M., Anil K. Jain, "A Rule-Based System for Automated Industrial Inspection," SPIE Vol. 635, Applications of Artificial Intelligence III, 1986, pp. 172-181.

[32] Darwish, Ahmed M., Anil K. Jain, "A Rule-Based Approach for Visual Pattern Inspection," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 10, No. 1, January 1988, pp. 56-68.

[33] Dickinson, Sven J., Larry S. Davis, "An Expert Vision System for Autonomous Land Vehicle Road Following," Proceedings CVPR 88 The Computer Society Conference on Computer Vision and Pattern Recognition, The Computer Society Press of the IEEE, June 1988, pp. 826-831.

[34] Dimanzo, Mauro, Giovanni Adorni, Fausto Giunchiglia, "Reasoning About Scene

Descriptions," Proceeding of the IEEE, Vol. 74, No. 7, July 1986, pp. 1013-1025.

[35] Doherty, Mark F., Carolyn M. Bjorklund, Mark T. Noga, "Split-Merge-Merge: An Enhanced Segmentation Capability," Proceedings Computer Vision and Pattern Recognition, IEEE Computer Society Press, June 1986, pp. 325-330.

[36] Doherty, Mark F., Carolyn M. Bjorklund, Connie Y. Wang, Mark T. Noga, Robert H. Laprade, "Improved Cartographic Classification Via Expert Systems," SPIE Vol. 635, Applications of Artificial Intelligence III, 1986, pp. 35- 39.

[37] Dreman, David, "The Not-So-Expert Expert," chapter 5 of "The New Contrarian Investment Strategy," Random House, 1982.

[38] Duda, Richard O., Peter E. Hart, Nils J. Nilsson, George L. Sutherland, "Semantic Network Representations in Rule-Based Inference Systems," pp.203-221 of "Pattern-Directed Inference Systems," D.A Waterman, F. Hayes-Roth, eds.

[39] Draper, Bruce A., John Brolio, Robert T. Collins, Allen R. Hanson, Edward M. Riseman, Proceedings CVPR 88 The Computer Society Conference on Computer Vision and Pattern Recognition, The Computer Society Press of the IEEE, June 1988, pp. 129-135.

[40] Eshera, M.A., King-Sun Fu, "An Image Understanding System Using Attributed Symbolic Representation and Inexact Graph Matching," IEEE Transactions on PAMI, Vol. PAMI-8, No. 5, September 1986, pp. 604-618.

[41] Fan, T. J., G. Medioni, R. Nevatia, "Recognizing 3-D Objects Using Surface Descriptions," Second International Conference on Computer Vision, IEEE Computer Society Press, December 1988, pp. 474-481.

[42] Faugeras, Oliver D., Marc Berthod, "Improving Consistency and Reducing Ambiguity in Stochastic Labeling: An Optimization Approach," IEEE Transactions on PAMI, Vol. PAMI-3, No. 4, July 1981, pp. 412-424.

[43] Faugeras, Oliver D., Keith E. Price, "Semantic Description of Aerial Images Using Stochastic Labeling," IEEE Transactions on PAMI, Vol. PAMI-3, No. 6, November 1981, pp. 633-643.

[44] Ferrante, Richard D., Mark J. Carlotto, Jean-Michel Pomarede, Paul M. Baim, "Multi-Spectral Image Analysis System," pp. 357-363 of First Conference on Artificial Intelligence Applications, August 1984.

[45] Freeman, Herbert, Editor, "Machine Vision Algorithms, Architectures, and Systems," Academic Press, 1988.

[46] Fu, K.S., Rosenfeld, Azriel, "Pattern Recognition and Computer Vision," Computer, Vol. 17, No. 10, October 1984, pp. 274- 282.

[47] Fua, Pascal, Andrew S. Hanson, "Locating Cultural Regions in Aerial Imagery using Geometric Clues," Proceedings of DARPA Workshop on Image Understanding, Miami Beach, December 1985.

[48] Fua, Pascal, Andrew S. Hanson, "Using Generic Geometric Models for Intelligent Shape Extraction," Proceedings: Image Understanding Workshop, February 1987, Morgan Kaufmann Publishers, pp. 227-233.

[49] Georgeff, Michael, Umberto Bonollo, "Procedural Expert Systems," SRI International Technical Note 314, December 9, 1983.

[50] Gevarter, William B., "Artificial Intelligence Expert Systems Computer Vision and Natural Language Processing," Noyes Publications, 1984.

[51] Gilmore, John F., Melinda M. Fox, Alicia L. Stevenson, Marla X. Rabin, "TESS - The Tactical Expert System," SPIE Vol. 635, Applications of Artificial Intelligence III, 1986, pp. 143- 150.

[52] Goldberg, Morris, David G. Goodenough, Mayer Alvo, Gerald M. Karam, "A Hierarchical Expert System for Updating Forestry Maps with LANDSAT Data," Proceedings of the IEEE, Vol. 73, No. 6, June 1985, pp. 1054-1063.

[53] Goldstein, E. Bruce, "Sensation and Perception," 2nd edition, Wadsworth Publishing, 1984.

[54] Gonzales, Rafael C., Paul Wintz, "Digital Image Processing," Addison Wesley, 1977.

[55] Goodenough, David G., Morris Goldberg, Gordon Plunkett, John Zelek, "An Expert System for Remote Sensing," IEEE Transactions on Geoscience and Remote Sensing, Vol. GE-25, No. 3, May 1987,, pp. 349-359.

[56] Gordon, Jean, Edward H. Shortliffe, "The Dempster-Shafer, Theory of Evidence," pp. 272-292 of Shortliffe, Edward H., Bruce G. Buchanan, eds., "Rule Based Expert Systems," Addison Wesley Publishing Company, 1984.

[57] Gordon, Jean, Edward H. Shortliffe, "A Method for Managing Evidential Reasoning in a Hierarchical Hypothesis Space," Artificial Intelligence, Vol. 26, No. 3, July 1985,

pp. 323-357.

[58] Han, Joon H., Richard A. Volz, "Region Grouping from a Range Image," Proceedings CVPR 88 The Computer Society Conference on Computer Vision and Pattern Recognition, The Computer Society Press of the IEEE, June 1988, pp. 241-248.

[59] Haralick, Robert M., "Computer Vision Theory: The Lack Thereof," Computer Vision, Graphics, and Image Processing, Vol. 36, 1986, pp. 372-386.

[60] Haralick, Robert M., James S. J. Lee, "Context Dependent Edge Detection," Proceedings CVPR 88 The Computer Society Conference on Computer Vision and Pattern Recognition, The Computer Society Press of the IEEE, June 1988, pp. 223-228.

[61] Harlow, Charles A., Mohan M. Trivedi, Richard W. Conners, Dwayne Phillips, "Scene Analysis of High Resolution Aerial Scenes," Optical Engineering, Vol. 25, No. 3, March 1986, pp. 347- 355.

[62] Harwood, David, Susan Chang, Larry S. Davis, "Interpreting Aerial Photographs by Segmentation and Search," Proceedings: Image Understanding Workshop, February 1987, Morgan Kaufmann Publishers, pp. 507-520.

[63] Havens, William, Alan Mackworth, "Representing Knowledge of the Visual World," Computer, October 1983, pp. 90-96.

[64] Hayes-Roth, Barbara, "A Blackboard Architecture for Control," Artificial Intelligence, Vol. 26, No. 3, July 1985, pp. 251-321.

[65] Hayes-Roth, Frederick, "The Role of Partial and Best Matches in Knowledge Systems," pp. 557-574 of "Pattern Directed Inference Systems".

[66] Hayes-Roth, Frederick, "Knowledge-Based Expert Systems," Computer, Vol. 17, No. 10, October 1984, pp. 263-273.

[67] Hayes-Roth, Frederick, Donald A. Waterman, Douglas B. Lenat, eds. "Building Expert Systems," Addison-Wesley Publishing Co., 1983.

[68] Hink, Robert F., David L. Woods, "How Humans Process Uncertain Knowledge: An Introduction for Knowledge Engineers," AI Magazine, Vol. 8, No. 3, Fall 1987, pp. 41-52.

[69] Hochberg, Julian, "Machines Should Not See as People Do, but Must Know How People See," Computer Vision, Graphics and Image Processing, Vol. 37, 1987, pp.

221-237.

[70] Horaud, Radu, Michael Brady, "On the Geometric Interpretation of Image Contours," Artificial Intelligence, Vol. 37, 1988, pp. 333-353.

[71] Huertas, A., W. Cole, R. Nevatia, "Detecting Runways in Aerial Images," Proceedings: Image Understanding Workshop, February 1987, Morgan Kaufmann Publishers, pp. 272-297.

[72] Huertas, A., R. Nevatia, "Detecting Buildings in Aerial Images," Computer Vision, Graphics, and Image Processing, Vol. 41, 1988, pp. 131-152.

[73] Hummel, Robert A., Michael S. Landy, "A Statistical Viewpoint on the Theory of Evidence," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 10, No. 2, March 1988, pp. 235-247.

[74] Kanade, Takeo, "Model Representations and Control Structures in Image Understanding," Proceedings of the 5th IJCAI, pp. 1074-1082.

[75] Kender, John R., "Vision Expert Systems Demand Challenging Expert Interactions," part of "Dialog," Computer Vision, Graphics, and Image Processing, Vol. 34, April 1986, pp.99-117.

[76] Kohl, Charles A., Allen R. Hanson, Edward M. Riseman, "Goal-Directed Control of Low-Level Processes for Image Interpretation," Proceedings: Image Understanding Workshop, February 1987, Morgan Kaufmann Publishers, pp. 538-551.

[77] Koons, D. B., B. H. McCormack, "A Model of Visual Knowledge Representation," First International Conference on Computer Vision, IEEE Computer Society Press, June 1987, pp. 365-372.

[78] Kovalevsky, V.A., "Dialog on "Expert" Vision Systems: Comments," part of "Dialog," Computer Vision, Graphics, and Image Processing, Vol. 34, April 1986, pp.99-117.

[79] Kuan, Darwin, Hormoz Shariat, Kalyan Dutta, Pat Ransil, "A Constraint-Based System for Interpretation of Aerial Imagery," Second International Conference on Computer Vision, IEEE Computer Society Press, December 1988, pp. 601-609.

[80] Lamdan, Yehezkel, Haim J. Wolfson, "Geometric Hashing: A General and Efficient Model-Based Recognition Scheme," Second International Conference on Computer Vision, IEEE Computer Society Press, December 1988, pp. 238-249.

[81] Laprade, Robert H., "Split-and-Merge Segmentation of Aerial Photographs," Computer

Vision, Graphics, and Image Processing, Vol. 44, No. 1, October 1988, pp. 77-86.

[82] Lee, Chia-Huang, "A Comparison of Two Evidential Reasoning Scenes," Artificial Intelligence, Vol. 35, No. 1, May 1988, pp. 127-134.

[83] Lee, Sunggu, Kang G. Shin, "Uncertain Inference Using Belief Functions," Third Conference on Artificial Intelligence Applications, IEEE Computer Society Press, February 1987, pp. 238-243.

[84] Lehrer, Nancy Bonar, George Reynolds, Joey Griffith, "Initial Hypothesis Formation in Image Understanding," Proceedings: Image Understanding Workshop, February 1987, Morgan Kaufmann Publishers, pp. 521-537.

[85] Lehrer, Nancy Bonar, George Reynolds, Joey Griffith, "A Method for Initial Hypothesis Formation in Image Understanding," First International Conference on Computer Vision, IEEE Computer Society Press, June 1987, pp. 578-585.

[86] Levesque, Hector J., "Making Believers Out of Computers," Artificial Intelligence, Vol. 30, No. 1, October 1986, pp. 81-108.

[87] Levine, Martin D., "Vision in Man and Machine," McGraw-Hill, 1985.

[88] Levine, Martin D., Samir I.Shaheen, "A Modular Computer Vision System for Picture Segmentation and Interpretation," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. PAMI-3, No. 5, September 1981, pp. 540-556.

[89] Levine, Martin D., Ahmed M. Nazif, "Low Level Image Segmentation: An Expert System," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. PAMI-6, No. 5, September 1984, pp. 555-577.

[90] Levine, Martin D., Ahmed M. Nazif, "Rule-Based Image Segmentation: A Dynamic Control Strategy Approach," Computer Vision, Graphics, and Image Processing, Vol. 32, No. 1, October 1985, pp. 104-126.

[91] Li, Hungwen, John R. Kender, "Scanning the Issue: Introduction to a Special Issue on Computer Vision," Proceedings of the IEEE, Vol. 76, No. 8, August 1988, pp. 859-862.

[92] Lindsay, Peter H., Donald A. Norman, "Human Information Processing," 2nd edition, Academic Press, 1977. Lowrance, John D., Thomas D. Garvey, "Evidential Reasoning: An Implementation for Multisensor Integration," SRI International, Technical Note 307, December 1983.

[93] Lowe, David G., "Perceptual Organization and Visual Recognition," Kluwer Academic Publishers, 1985.

[94] Lowe, David G., "Three-Dimensional Object Recognition from Single Two-Dimensional Images," Artificial Intelligence, Vol. 31, 1987, pp. 355-395.

[95] Mantas, J., "Methodologies in Pattern Recognition and Image Analysis - A Brief Survey," Pattern Recognition, Vol. 20, No. 1, 1987, pp. 1-6.

[96] Matsuyama, Takashi, "Knowledge Organization and Control Structure in Image Understanding," 7th ICPR, pp. 1118-1127.

[97] Matsuyama, Takashi, "Knowledge-Based Aerial Image Understanding Systems and Expert Systems for Image Processing," IEEE Transactions on Geoscience and Remote Sensing, Vol. GE-25, No. 3, May 1987,, pp. 305-316.

[98] McKeown, David M., "The Role of Artificial Intelligence in the Integration of Remotely Sensed Data with Geographic Information Systems," IEEE Transactions on Geoscience and Remote Sensing, Vol. GE-25, No. 3, May 1987,, pp. 330-348.

[99] McKeown, David M., Wilson A. Harvey Jr., John McDermott, "Rule-Based Interpretation of Aerial Imagery," IEEE Transactions on PAMI, Vol. PAMI-7, No. 5, September 1985, pp. 570-585.

[100] McKeown, David M., Wilson A. Harvey, "Automatic Knowledge Acquisition for Aerial Image Interpretation," Proceedings: Image Understanding Workshop, February 1987, Morgan Kaufmann Publishers, pp. 205-221.

[101] McKeown, David M., Jerry L. Denlinger, "Cooperative Methods for Road Tracking in Aerial Imagery," Proceedings CVPR 88 The Computer Society Conference on Computer Vision and Pattern Recognition, The Computer Society Press of the IEEE, June 1988, pp. 662-672.

[102] Meisels, Amnon, Samuel Bergman, "Finding Objects on Aerial Photographs: A Rule-Based Low Level System," Proceedings CVPR 88 The Computer Society Conference on Computer Vision and Pattern Recognition, The Computer Society Press of the IEEE, June 1988, pp. 118-122.

[103] Menhardt, Wido, Karl-Heinrich Schmidt, "Computer Vision on Magnetic Resonance Images," Pattern Recognition Letters, Vol. 8, No. 2, September 1988, pp. 73-85.

[104] Miller, George A., "The Magical Number Seven, Plus or Minus Two: Some Limits on

Our Own Capacity for Processing Information," The Psychological Review, Vol. 63, No. 2, March 1956, pp. 81-97.

[105] Minsky, M., "A Framework for Representing Knowledge," pp. 211-277 of "The Psychology of Computer Vision," P. Winston editor, McGraw Hill, 1975.

[106] Mohan, Rakesh, Ramakant Nevatia, "Perceptual Grouping with Applications to 3D Shape Extraction," Proceedings of the IEEE Computer Society Workshop on Computer Vision, Computer Society Press of the IEEE, December 1987, pp. 158-163.

[107] Murthy, H. Krishna, M. Narasimha Murty, "A Conceptual Clustering Scheme for Frame-Based Knowledge Organization," SPIE Vol. 635, Applications of Artificial Intelligence III, 1986, pp. 408-414.

[108] Musavi, M. T., M. V. Shirvaikar, E. Ramanathan, A. R. Nekovei, "A Vision Based Method to Automate Map Processing," Pattern Recognition, Vol. 21, No. 4, 1988, pp. 319-326.

[108] Nagao, Makoto, "Control Strategies in Pattern Analysis," Pattern Recognition, Vol. 17, No. 1, pp. 45-56.

[109] Nagao, Makota, Takashi Matsuyama, "A Structural Analysis of Complex Aerial Photographs," Plenum Press, 1980.

[110] Nandhakumar, N., J.K. Aggarwal, "The Artificial Intelligence Approach to Pattern Recognition - A Perspective and an Overview," Pattern Recognition, Vol. 18, No. 6,1985, pp. 383-389.

[111] Nandhakumar, N., J.K. Aggarwal, "Integrating Information from Thermal and Visual Images for Scene Analysis," SPIE Vol. 635, Applications of Artificial Intelligence III, 1986, pp. 132-142.

[112] Nandhakumar, N., J.K. Aggarwal, "Integrated Analysis of Thermal and Visual Images for Scene Interpretation ," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 10, No. 4, July 1988, pp. 469-481.

[113] National Aeronautics and Space Administration, National Space Technology Laboratories, Earth Resources Laboratory, "ELAS -Earth Resources Laboratory Applications Software ... A Geobased Information System," March 1980.

[114] Neisser, Ulrich, "Cognitive Psychology," Appleton-Century-Crofts, New York, 1967.

[115] Neisser, Ulrich, "Cognition and Reality Principles and Implications of Cognitive Psychology," W. H. Freeman and Company, San Francisco, California, 1976.

[116] Niblack, Wayne, Dragutin Petkovic, Dick Damian, "Experiments and Evaluations of Rule Based Methods in Image Analysis," Proceedings CVPR 88 The Computer Society Conference on Computer Vision and Pattern Recognition, The Computer Society Press of the IEEE, June 1988, pp. 123-128.

[117] Nicolin, Bernhard, Richard Gabler, "A Knowledge-Based System for the Analysis of Aerial Images," IEEE Transactions on Geoscience and Remote Sensing, Vol. GE-25, No. 3, May 1987,, pp. 317-329.

[118] Nieman, H., "A Homogeneous Architecture for Knowledge Based Image Understanding Systems," The Second Conference on Artificial Intelligence Applications, IEEE Press, December 1985, pp. 88-93.

[119] Nii, Penny, "Blackboard Systems: The Blackboard Model of Problem Solving and the Evolution Blackboard Architectures, Part 1", AI Magazine, Summer 1986, pp. 38-53.

[120] Nii, Penny, "Blackboard Systems: The Blackboard Model of Problem Solving and the Evolution Blackboard Architectures, Part 2", AI Magazine, August 1986, pp. 82-106.

[121] Nilsson, Nils J., "Probabalistic Logic," SRI International, Technical Note 321, February 6, 1984.

[122] Numao, Masayuki, Mitsura Ishizuka, "A Frame-Like Knowledge Representation System for Computer Vision," pp. 1128-1130 of 7th International Conference on Pattern Recognition.

[123] Osbourn, G. C., A New Approach to Machine-Based Perception of Monocular Images," Proceedings CVPR 88 The Computer Society Conference on Computer Vision and Pattern Recognition, The Computer Society Press of the IEEE, June 1988, pp. 864-870.

[124] Parvin, B., G. Medioni, "Adaptive Multiscale Feature Extraction from Range Data," Computer Vision, Graphics, and Image Processing, Vol. 45, No. 3, March 1989, pp. 346-356.

[125] Pavlidis, T. "A Critical Survey of Image Analysis Methods," Eighth International Conference on Pattern Recognition, IEEE Computer Society Press, October 1986, pp. 502-511.

[126] Pearl, Judea, "On Evidential Reasoning in a Hierarchy of Hypotheses," Artificial Intelligence, Vol. 28, No. 1, February 1986, pp. 9-15.

[127] Pentland, Alex P., "Perceptual Organization and the Representation of Natural Form," Artificial Intelligence, Vol. 28, 1986, pp. 293-331.

[128] Pentland, Alex P., "Part Models," Proceedings Computer Vision and Pattern Recognition, IEEE Computer Society Press, June 1986, pp. 242-249.

[129] Pentland, Alex P., Editor, "From Pixels to Predicates," Ablex Publishing Corp., 1986.

[130] Pentland, Alex P., "Recognition by Parts," First International Conference on Computer Vision, IEEE Computer Society Press, June 1987, pp. 612-620.

[131] Pentland, Alex P., "Shape Information from Shading: A Theory about Human Perception," Second International Conference on Computer Vision, IEEE Computer Society Press, December 1988, pp. 404-413.

[132] Perkins, W.A., T.J. Laffey, T.A. Nguyen, "Rule-Based Interpreting of Aerial Photographs Using LES,"Proceedings Applications of Artificial Intelligence, 1985, Arlington, Va., pp. 138-146.

[133] Phillips, Dwayne, "Using a Top-Down and Bottom-Up Strategy to Analyze High Resolution Aerial Photographs of Urban Areas," Proceedings: First International Conference on Industrial & Engineering Applications of Artificial Intelligence & Expert Systems," June 1988.

[134] Preperata, Franco P., Sylvian R. Rays, "An Approach to Artificial Nonsymbolic Cognition," Information Sciences, Vol. 4, 1972, pp. 65-86.

[135] Price, Keith, "Anything You Can Do, I Can Do Better (No You Can't)...,"Computer Vision, Graphics, and Image Processing, Vol. 36, 1986, pp. 387-391.

[136] Provan, Gregory M., "The Vision Constraint Recognition System: Analyzing the Role of Reasoning in High Level Vision," Proceedings of the IEEE Computer Society Workshop on Computer Vision, Computer Society Press of the IEEE, December 1987, pp. 170-175.

[137] Provan, Gregory M., "Model-Based Object Recognition: A Truth Maintenance Approach," Proceedings The Fourth Conference on Artificial Intelligence Applications, Computer Society Press of the IEEE, March 1988, pp. 230-235.

[138] Rao, A. Ravishankar, Ramesh Jain, "Knowledge Representation and Control in

Computer Vision Systems," IEEE Expert, Vol. 3, No. 1, Spring 1988, pp. 64-79.

[139] Rearick, T.C., "Real-Time Image Understanding," SPIE Vol. 635, Applications of Artificial Intelligence III, 1986, pp. 166- 171.

[140] Reynolds, George, J. Ross Beveridge, "Searching for Geometric Structures in Images of Natural Scenes," Proceedings: Image Understanding Workshop, February 1987, Morgan Kaufmann Publishers, pp. 257-271.

[141] Rich, Elaine, "Artificial Intelligence,"McGraw Hill, 1983.

[142] Rosenfeld, Azriel, "Picture Processing 1980," Computer Vision, Graphics, and Image Processing, Vol. 16, No. 1, May 1981, pp. 52-89.

[143] Rosenfeld, Azriel, "Picture Processing 1981," Computer Vision, Graphics, and Image Processing, Vol. 19, No. 1, May 1982, pp. 35-75.

[144] Rosenfeld, Azriel, "Picture Processing 1982," Computer Vision, Graphics, and Image Processing, Vol. 22, No. 3, June 1983, pp. 339-387.

[145] Rosenfeld, Azriel, "Picture Processing 1983," Computer Vision, Graphics, and Image Processing, Vol. 23, No. 3, June 1984, pp. 347-393.

[146] Rosenfeld, Azriel, "Picture Processing 1984," Computer Vision, Graphics, and Image Processing, Vol. 30, No. 2, May 1985, pp. 189-242.

[147] Rosenfeld, Azriel, "Picture Processing 1985," Computer Vision, Graphics, and Image Processing, Vol. 34, No. 2, May 1986, pp. 204-251.

[148] Rosenfeld, Azriel, "Image Analysis and Computer Vision: 1987," Computer Vision, Graphics, and Image Processing, Vol. 42, 1988, pp. 234-293.

[149] Rosenfeld, Azriel, "Image Analysis and Computer Vision: 1988," Computer Vision, Graphics, and Image Processing, Vol. 46, No. 2, May 1989, pp. 196-264.

[150] Rosenfeld, Azriel, "Computer Vision: Basic Principles," Proceedings of the IEEE, Vol. 76, No. 8, August 1988, pp. 863-868.

[151] Rosenfeld, Azriel, "Image Analysis: Problems, Progress and Prospects," Pattern Recognition, Vol. 17, No. 1, pp. 3-12.

[152] Rosenfeld, Azriel, "Expert Vision Systems: Some Issues," part of "Dialog," Computer Vision, Graphics, and Image Processing, Vol. 34, April 1986, pp.99-117.

[153] Rosenfeld, Azriel, Avinash C. Kak, "Digital Picture Processing," Academic Press, 1976.

[154] Rusco, Dan W., "Specification of MC&G Tangible Feature for Pilot Digital

Operations Experimentation," DMAAC/ACT, August 1978, revised July 1979.

[155] Sakai, T., T. Kanade, Y.I. Ohta, "Model Based Interpretation of Outdoor Scene," Proceedings IJCPR-3, pp.581-585.

[156] Sagerer, G., "Automatic Interpretation of Medical Image Sequences," Pattern Recognition Letters, Vol. 8, No. 2, September 1988, pp. 87-102.

[157] Schowengerdt, Robert A., "Techniques for Image Processing and Classification in Remote Sensing," Academic Press, 1983.

[158] Sha'ashua, Amnon, Shimon Ullman, "Structural Saliency: The Detection of Globally Salient Structures Using a Locally Connected Network," Second International Conference on Computer Vision, IEEE Computer Society Press, December 1988, pp. 321-327.

[159] Shacter, Ross D., David E. Heckerman, "Thinking Backward for Knowledge Acquisition," AI Magazine, Vol. 8, No. 3, Fall 1987, pp. 55-61.

[160] Shafer, Glen, "A Mathematical Theory of Evidence," Princeton University Press, 1976.

[161] Shafer, Glen, "Hierarchical Evidence," The Second Conference on Artificial Intelligence Applications, IEEE Press, December 1985, pp. 16-21.

[162] Shafer, Glen, Roger Logan, "Implementing Dempster's Rule for Hierarchical Evidence," Artificial Intelligence, Vol. 33, No. 3, November 1987, pp. 271-298.

[163] Shapiro, Linda G., "The Role of AI in Computer Vision," The Second Conference on Artificial Intelligence Applications, IEEE Press, December 1985, pp. 76-81.

[164] Shneier, Michael O., Ronald Lumia, Ernest W. Kent, "Model-Based Strategies for High-Level Robot Vision," Computer Vision, Graphics and Image Processing, Vol. 33, March 1986, pp. 293-306.

[165] Shortliffe, Edward H., Bruce G. Buchanan, eds., "Rule Based Expert Systems," Addison Wesley Publishing Company, 1984.

[166] Shvaytser, Haim, "Learnable and Non-Learnable Visual Concepts," Second International Conference on Computer Vision, IEEE Computer Society Press, December 1988, pp. 264-268.

[167] Silverberg, Teresa M., "Infrared Image Interpretation Using Spatial and Temporal Knowledge," Proceedings of the IEEE Computer Society Workshop on Computer Vision, Computer Society Press of the IEEE, December 1987, pp. 264-267.

[168] Sleigh, A.C., P.K. Bailey, "DIPOD: An Image Understanding Development System and Implementation System," Pattern Recognition Letters, Vol. 6, No. 2, July 1987, pp. 101-106.

[169] Smith, R.G., G.M.E. Lafue, E. Schoen, S.C. Vestal, "Declarative Task Description as a User-Interface Structuring Mechanism," Computer, September 1984, pp. 29-38.

[170] Smyrniotis, Chuck, Kaylon Dutta, "A Knowledge-Based System for Recognizing Man-Made Objects in Aerial Images," Proceedings CVPR 88 The Computer Society Conference on Computer Vision and Pattern Recognition, The Computer Society Press of the IEEE, June 1988, pp. 111-117.

[171] Stansfield, Sharon A., "ANGY: A Rule-Based Expert System for Automatic Segmentation of Coronary Vessels from Digital Subtracted Angiograms," IEEE Transactions on PAMI, Vol. PAMI-8, No. 2, March 1986, pp. 188-199.

[172] Stephanou, Harry F., Shin-Yee Lu, "Measuring Consensus Effectiveness by a Generalized Entropy Criterion," IEEE Transactions of Pattern Analysis and Machine Intelligence, Vol. 10, No. 4, July 1988, pp. 544-554.

[173] Thompson, William B., "Comments On: "Expert" Vision Systems: Some Issues," part of "Dialog," Computer Vision, Graphics, and Image Processing, Vol. 34, April 1986, pp.99-117.

[174] Thompson, William B., Rand P. Whillock, "Occlusion-Sensitive Matching," Second International Conference on Computer Vision, IEEE Computer Society Press, December 1988, pp. 285-289.

[175] Thorpe, Charles, Martial H. Hebert, Takeo Kanade, Steven A. Shafer, "Vision and Navigation for the Carnegie-Mellon Navlab," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 10, No. 3, May 1988, pp. 362-373.

[176] Tou, J. T., R. C. Gonzalez, "Pattern Recognition Principles," Addison-Wesley Publishing Company, 1974.

[177] Turk, Matthew A., David G. Mergenthaler, Keith D. Gremban, Martin Marra, IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 10, No. 3, May 1988, pp. 342-361.

[178] Tynor, Stephen D., Chi-Cheung Tsang, Kurt Kingher, John F. Gilmore, "VEST - The Visual Expert System Testbed," SPIE Vol. 635, Applications of Artificial Intelligence

III, 1986, pp. 207-215.

[179] Uhr, Leonard, "Workshop on Goal-Directed "Expert" Vision Systems: My Comments," part of "Dialog," Computer Vision, Graphics, and Image Processing, Vol. 34, April 1986, pp.99-117.

[180] Walker, Ellen Lowenfeld, Martin Herman, "Geometric Reasoning for Constructing 3D Scene Descriptions from Images," Artificial Intelligence, Vol. 37, 1988, pp. 275-290.

[181] Wallace, Andrew M., "A Comparison of Approaches to High-Level Image Interpretation," Pattern Recognition, Vol. 21, No. 3, 1988, pp. 241-259.

[182] Walters, Deborah, "Selection and Use of Image Features for Segmentation of Boundary Images," Proceedings Computer Vision and Pattern Recognition, IEEE Computer Society Press, June 1986, pp. 319-324.

[183] Walters, Deborah, "Selection of Image Primitives for General Purpose Visual Processing," Computer Vision, Graphics and Image Processing, Vol. 37, 1987, pp. 261-298.

[184] Walters, Deborah, Ganapathy Krishnan, "Color Separation Using General Purpose Computer Vision Algorithms," Third Conference on Artificial Intelligence Applications, IEEE Computer Society Press, February 1987, pp. 188-193.

[185] Waterman, D.A., Frederick Hayes-Roth, Editors, "Pattern-Directed Inference Systems," Academic Press, 1978.

[186] Weiss, Richard, Michael Boldt, "Geometric Grouping Applied to Straight Lines," Proceedings Computer Vision and Pattern Recognition, IEEE Computer Society Press, June 1986, pp. 489-495.

[187] Weiss, Sholom, Casmir A. Kulikowski, "A Practical Guide to Designing Expert Systems," Rowman and Allenheld Publishers, 1984.

[188] Wesley, Leonard P., "Evidential Knowledge-Based Computer Vision," Optical Engineering, Vol. 25, No. 3, March 1986, pp. 363- 379.

[189] Winograd, Terry, "Frame Representations and the Declarative /Procedural Controversy," from [Bobrow,Collins].

[190] Woods, William A., "What's Important About Knowledge Representation?," Computer, Vol. 16, No. 10, October 1983, pp. 22-27.

[191] Yen, John, "A Reasoning Model Based on an Extended Dempster-Shafer Theory,"

Proceedings AAAI-86, 5th National Conference on Artificial Intelligence, August 1986, Morgan Kaufman Publishers, pp. 125-131.

[192] Zadeh, Lofti A., "A Simple View of the Dempster-Shafer Theory of Evidence and it's Implication for the Rule of Combination," AI Magazine, Summer 1986, pp. 85-90.

[193] Zhu, Miao-Liang, Pen-Shu Yeh, "Automatic Road Detection on Aerial Photographs," Proceedings Computer Vision and Pattern Recognition, IEEE Computer Society Press, June 1986, pp. 34-40.

[194] Zucker, Steven W., "Production Systems with Feedback," pp. 539-555 of "Pattern Directed Inference Systems."

APPENDICES

APPENDIX 1 - Photographs

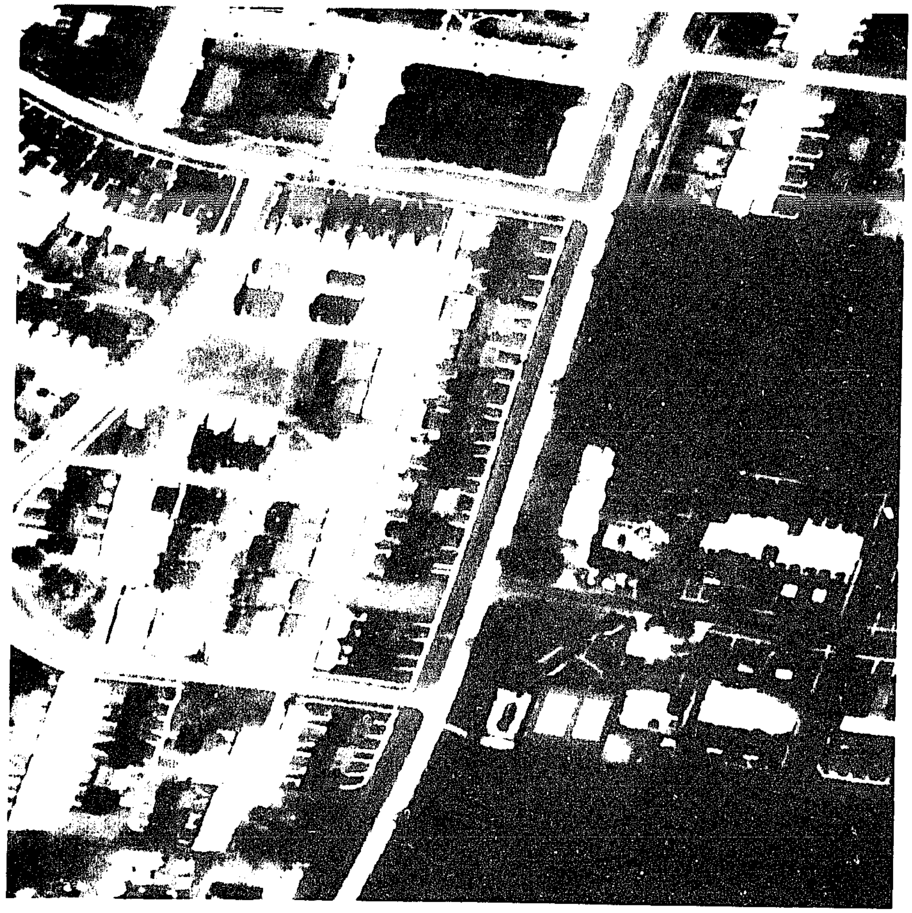

Photograph 1 - Green Channel .5 - .55 μm Band

Photograph 2 - Red Channel .65 - .69 μm Band

Photograph 3 - Thermal IR Channel 8.5 - 13.0 μm Band

Photograph 4 - Same Image as Photograph 2 with a Reduction in Resolution of 4

Photograph 5 - Same Image as Photograph 2 with a Reduction in Resolution of 2

Photograph 6 - Results of Analyzing Photographs 1, 2, and 3

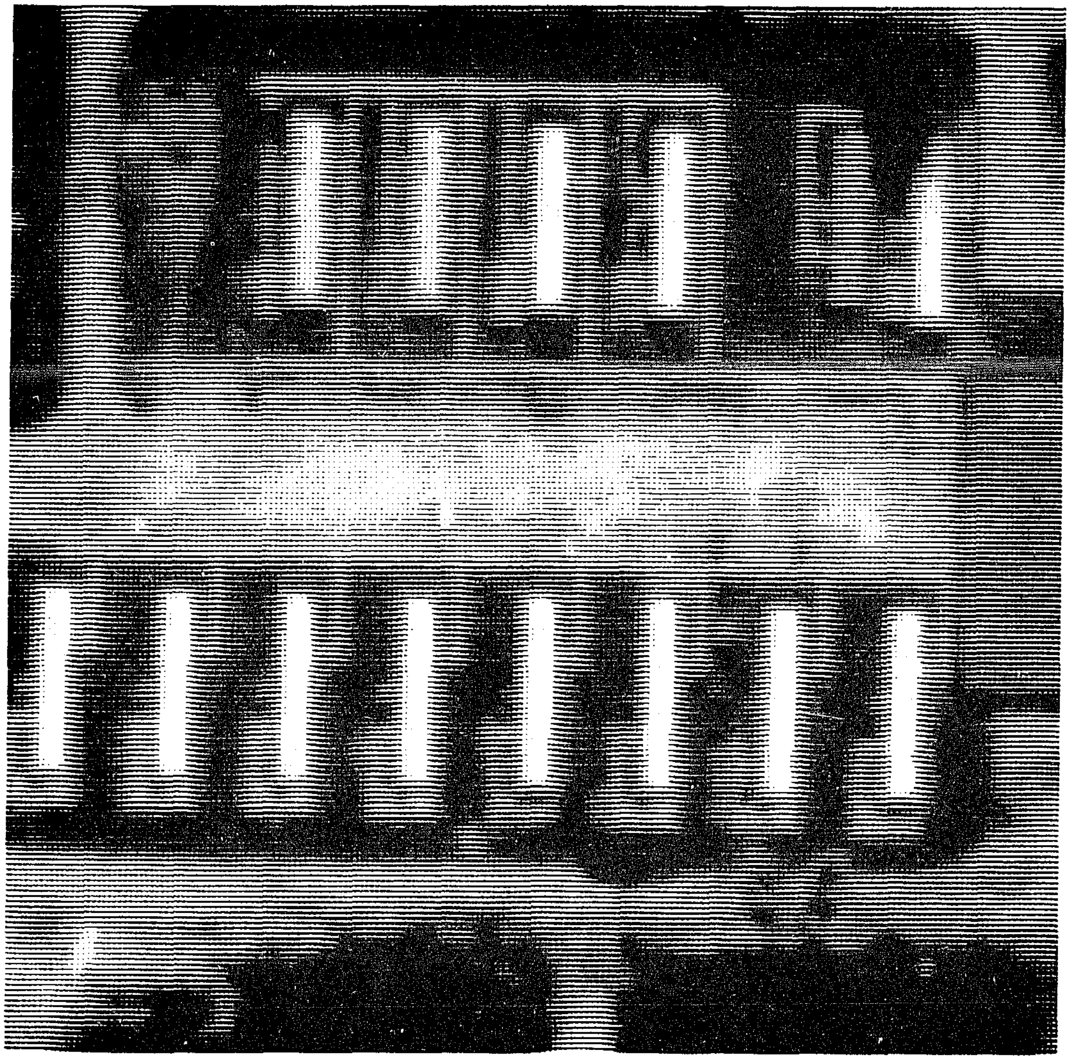

Photograph 7 - Second Example Image

Photograph 8 - Third Example Image

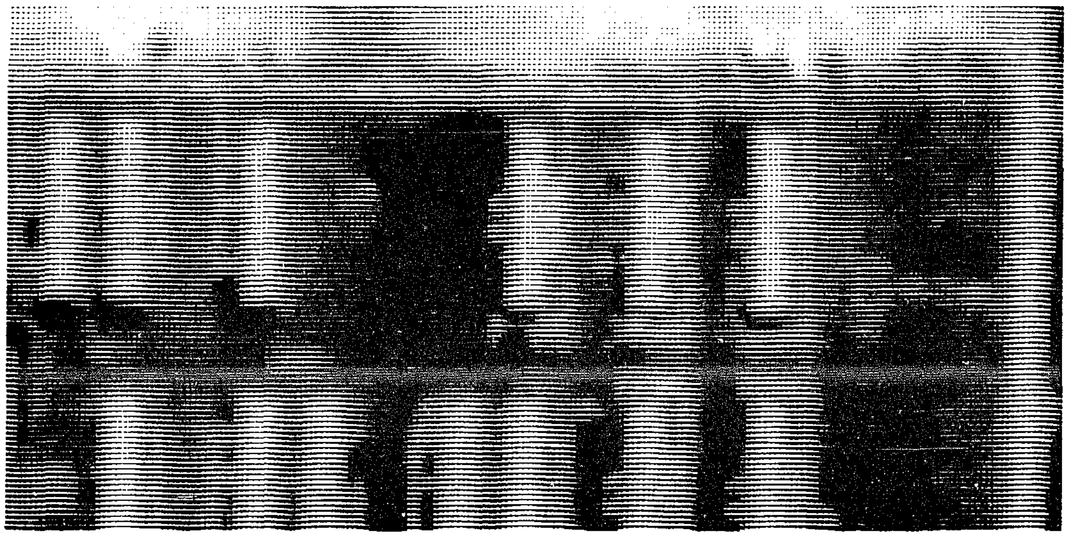

Photograph 9 - Fourth Example Image

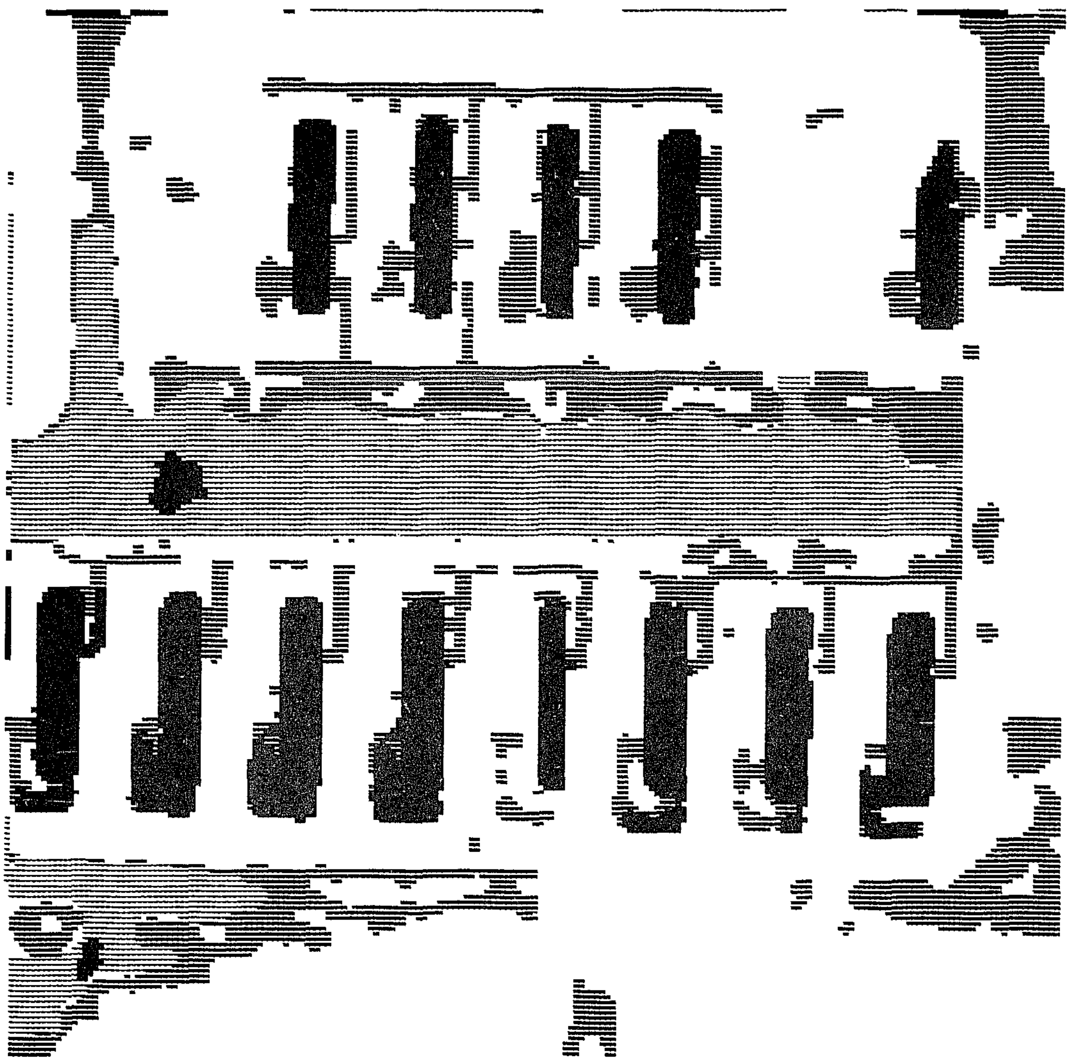

Photograph 10 - Results of Analyzing Photograph 7

Legend:
Black = 041 house (single family dwelling)
Dark Gray = 097 footpath/trail
Light Gray = 069 primary road
White = background

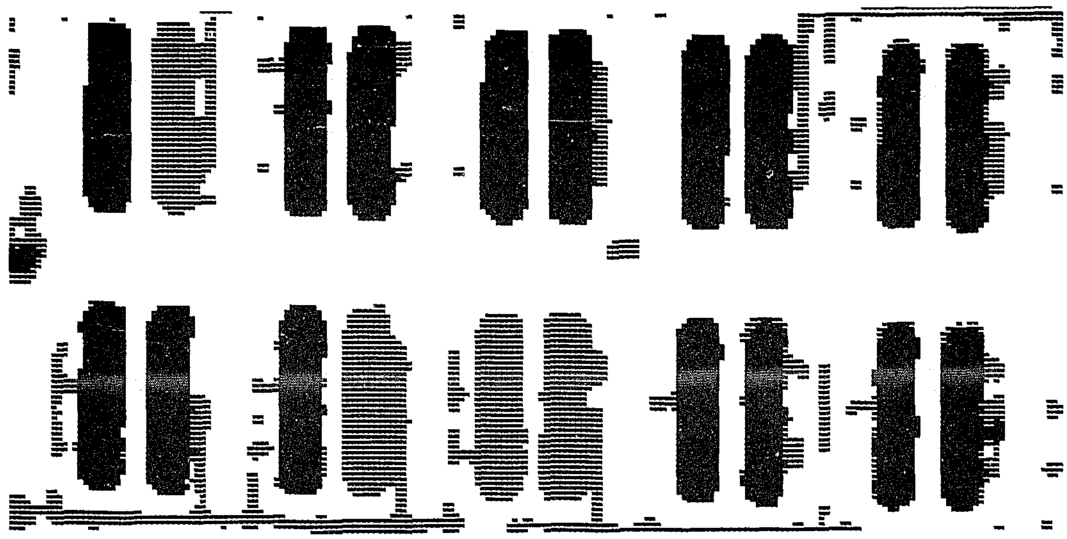

Photograph 11 - Results of Analyzing Photograph 8

Legend:
Black = 041 house (single family dwelling)
Dark Gray = 097 footpath/trail
Light Gray = 069 primary road
White = background

Photograph 12 - Results of Analyzing Photograph 9

Legend:
Black = 041 house (single family dwelling)
Dark Gray = 097 footpath/trail
Light Gray = 069 primary road
White = background

APPENDIX 2 - The Defense Mapping Agency Hierarchy

Node Title

1 world - level 1

culture - see node 2

landscape - see node 3

----

2 culture - level 2

part of world - see node 1

urban (built up areas) - see node 4

transportation/navigation - see node 5

landmark/rural features - see node 6

----

3 landscape - level 2

part of world - see node 1

hydrography - see node 7

physiography - see node 8

phytography - see node 9

4 urban (built up areas) - level 3

part of culture - see node 2

industry/utility - see node 10

commercial/residential - see node 11

institutional/governmental - see node 12

----

5 transportation/navigation - level 3

part of culture - see node 2

railroads - see node 13

roads - see node 14

aeronautical/aerospace - see node 15

naval/marine - see node 16

associated transportation features - see node 17

----

6 landmark/rural features - level 3

part of culture - see node 2

communication/transmission - see node 18

storage - see node 19

agricultural - see node 20

recreational - see node 21

miscellaneous - see node 22

7 hydrography - level 3

part of landscape - see node 3

water - see node 23

snow/ice - see node 23

----

8 physiography - level 3

part of landscape - see node 3

exposed soils (surface composition) - see node 25

landforms - see node 26

----

9 phytography - level 3

part of landscape - see node 3

cropland - see node 27

rangeland - see node 28

woodland - see node 29

wetland - see node 30

10 industry/utility - level 4

part of urban (built up) areas - see node 4

extraction - see node 31

processing - see node 32

power generation - see node 33

fabrication industry - see node 34

disposal - see node 35

associated industrial structures - see node 36

----

11 commercial/residential - level 4

part of urban (built up) areas - see node 4

039 commercial buildings

040 apartments/hotel

041 house (single family dwelling

042 mobile homes

043 barracks

044 display sign (large billboard, highway sign)

12 institutional/governmental - level 4

part of urban (built up) areas - see node 4

045 governmental administration building

046 military admin/operations building

047 capitol building

048 hospital

049 prison

050 palace

051 museum

052 observatory

053 church/tabernacle

054 mosque

055 cemetary building

13 railroads - level 4

part of transportation/navigation - see node 5

056 single track railway

057 double track railway

058 multiple track railway

059 RR yard/siding

060 tramway/inclined railway

061 monorail

062 RR storage/repair building

063 RR terminal building

064 RR station/depot

065

066 roundhouse

----

14 roads - level 4

part of transportation/navigation - see node 5

067 multi lane, divided (grass median) highway

068 multi lane highway

069 primary road (dual lane, hard surface)

070 secondary road (dual lane, loose/dirt surface)

071 trail/track (one lane)

072 toll gates

073 cloverleaf/interchange

074 garage, service/repair facilities (landmark)

15 aeronautical/aerospace - level 4

part of transportation/navigation - see node 5

075 runway/taxiway

076 aircraft parking area/apron

077 airport/airbase control tower

078 hangar

079 terminal/base operations building

080 aerospace assembly building

081 missile launch pad/gantry facility

082 engine test cell

083 wind tunnel

084 hellioport

085 seaplane base

16 naval/marine - level 4

part of transportation/navigation - see node 5

086 breakwater/jetty

087 wharf/pier/quay

088 dam locks

089 canal locks

090 sea wall

091 ramp/slip/ferry landing

092 dock/dry-dock

093 light ship

094 light house

095 off-shore loading facility

096 exposed wreck

17 associated transportation features - level 4

part of transportation/navigation - see node 5

097 footpath/trail

098 tunnel

099 underpass

100 ferry

101 aerial cableway/skilift

102 bridge

18 communication/transmission - level 4

part of landmark/rural features - see node 6

electrical/electronic - see node 37

fluid conduits - see node 38

----

19 storage - level 4

part of landmark/rural features - see node 6 tanks - see node 39

closed storage - see node 40

open storage - see node 41

----

20 agricultural - level 4

part of landmark/rural features - see node 6

128 farm buildings (house/shed)

129 barn

130 greenhouse

131 windmill-truss

132 windmill-solid

133 feedlot/stockyard/feeding pen

134 circular irrigation system

21 recreational - level 4

part of landmark/rural features - see node 6

135 racetrack

136 stadium

137 grandstand

138 athletic field

139 ampitheater

140 drive-in theater screen

141 fairground

142 campground/campsite

143 amusement park

144 roller coaster

145 ferris wheel

146 artificial mountain

22 miscellaneous landmarks - level 4

part of landmark/rural features - see node 6

147 ruins

148 fort

149 observation/lookout tower

150 watermill/gristmill

151 wall

152 fence

153 monument/oblisk

154 arch

155 pyramid

156 castle

157 dam

158 cemetery

159 fish pond/hatchery

160 sewage disposal pools

161 filtration/aeration beds

162 salt pan/evaporators

163 sluice gate

23 water - level 4

part of hydrography - see node 7

164 sea/ocean (sea state)

165 lake/pond/reservoir

166 river/stream

167 canal irrigation ditch

168 waterfall

169 rapids

----

24 snow/ice - level 4

part of hydorgraphy - see node 7

170 perennial (permanent) snowfield

171 perennial ice (glacier/ice cap)

172 glacial maraine

173 seasonal ice pack (limits)

174 polar ice pack (permanent)

25 exposed soils (surface composition) - level 4

part of physiography - see node 8

175 dry land (bare/barren soil/non-cultivated)

176 open cultivated ground

177 desert sand

178 sand dunes

179 exposed smooth (solid) rock

180 boulder field/lava

181 rock, rough area/region

182 dry lake/salt flat

183 mud/tidal flat

184 wet sand (beach/sand bar)

----

26 landforms - level 4

part of physiography - see node 8

185 levee/embankment/fill

186 cut

187 cliff/bluff/escarpment

188 reef shoals/rocks (in water)

189 terrace

27 cropland - level 4

part of phytography - see node 9

190 orchard/plantation

191 shelterbelt/hedgerow

192 nursery/grove

193 vineyard

194 crop (cultivated)

195 cranberry bog

196 rice paddy

----

28 rangeland - level 4

part of phytography - see node 9

197 herbaceous (grassland)

198 shrub/brush

199 mixed (shrub/brush and grass)

----

29 woodland - level 4

part of phytography - see node 9

200 decidous

201 caniferous

202 mixed (decidous and caniferous)

203 mangrove

204 nipa palm

30 wetland - level 4

part of phytography - see node 9

205 swamp (trees, brush/shrubs in water)

206 marsh (grass, cat tails, etc in water)

207 peat bog/cuttings

----

31 extraction - level 5

part of industry/utility - see node 10

001 quarry

002 sand/gravel/clay pit

003 gas/oil well

004 gas/oil rig with derrick

005 gas/oil rig offshore platform

006 gas/oil rig offshore platform with derrick

007 mine shaft superstructure

008 open-pit/strip mine

32 processing - level 5

part of industry/utility - see node 10

009 chemical plant

010 metal processing plant

011 sewage treatment plant

012 evaporative mining

013 coke plant

014 blast furnace

015 refinery

016 catalytic cracker

017 flare pipe

----

33 power generation - level 5

part of industry/utility - see node 10

018 hydroelectric power plant

019 thermal power plant

020 transformer yard

021 substation

022 windmotor

023 solar electric panels

024 solar heat panels

----

34 fabrication industry - level 5

part of industry/utility - see node 10

025 building

35 disposal - level 5

part of industry/utility - see node 10

026 oil sump/sludge pit

027 scrap yard

028 metal ore slag dump

029 tailings/mine dump

030 tailings pond

----

36 associated industrial structures - level 5

part of industry/utility - see node 10

031 buildings

032 smoke stack

033 conveyor

034 bridge crane

035 rotating crane

036 cooling tower

037 hopper

038 dredge/power shovel/dragline

----

37 electrical/electronic - level 5

part of communication/transmission - see node 18

103 microwave communication tower

104 radio/tv antenna tower/mast

105 telephone/telegraph lines

106 power transmission line

107 relay station/communication building

38 fluid conduits - level 5

part of communication/transmission - see node 18

108 pipeline (landmark)

109 penstock/flume

110 aqueduct

111 pumping station

----

39 tanks - level 5

part of storage - see node 19

112 tank

113 telescoping gasholder (gasometer)

----

40 closed storage - level 5

part of storage - see node 19

114 water tower building

115 ordinance storage bunker/mounds

116 grain elevator

117 grain bin

118 upright silo

119 warehouse

120 depot

41 open storage - level 5

part of storage - see node 19

121 trench silo

122 mineral pile

123 oil storage pit

124 vehicle storage/motor pool

125 vehicle parking area

126 aircraft storage area

127 ship storage area

APPENDIX 3 - Frames

***world* Frame**

The frame.is_a is:
->world
The frame.is_part_of is:
->nothing
The frame.goal_of_analysis is:
->region label
The frame.level_in_tree is ->1

This frame has 2 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->culture

Sub node[0].bel_element_number is 2
Sub node[0].mean is 88.000000

Sub node[1].is_a is ->landscape

Sub node[1].bel_element_number is 3
Sub node[1].mean is 44.000000

***culture* Frame**

The frame.is_a is:
->culture
The frame.is_part_of is:
->world
The frame.goal_of_analysis is:
->region label
The frame.level_in_tree is ->2

This frame has 3 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->urban (built up areas)

Sub node[0].bel_element_number is 4
Sub node[0].mean is 90.000000

Sub node[1].is_a is ->transportation/navigation

Sub node[1].bel_element_number is 5
Sub node[1].mean is 60.000000

Sub node[2].is_a is ->landmark/rural features

Sub node[2].bel_element_number is 6
Sub node[2].mean is 40.000000

characteristic[0] is ->roofing

characteristic[0] function is ->histogram

characteristic[0].positive_assertion[0].place=2
characteristic[0].positive_assertion[0].belief=0.800000
characteristic[0].positive_assertion[1].place=1
characteristic[0].positive_assertion[1].belief=0.200000

characteristic[0].negative_assertion[0].place=3
characteristic[0].negative_assertion[0].belief=0.600000
characteristic[0].negative_assertion[1].place=1
characteristic[0].negative_assertion[1].belief=0.400000

characteristic[1] is ->2roofing

characteristic[1] function is ->histogram

characteristic[1].positive_assertion[0].place=2
characteristic[1].positive_assertion[0].belief=0.800000
characteristic[1].positive_assertion[1].place=1
characteristic[1].positive_assertion[1].belief=0.200000

characteristic[1].negative_assertion[0].place=3
characteristic[1].negative_assertion[0].belief=0.600000
characteristic[1].negative_assertion[1].place=1
characteristic[1].negative_assertion[1].belief=0.400000

characteristic[2] is ->aluminum

characteristic[2] function is ->histogram

characteristic[2].positive_assertion[0].place=2

characteristic[2].positive_assertion[0].belief=0.800000
characteristic[2].positive_assertion[1].place=1
characteristic[2].positive_assertion[1].belief=0.200000

characteristic[2].negative_assertion[0].place=3
characteristic[2].negative_assertion[0].belief=0.600000
characteristic[2].negative_assertion[1].place=1
characteristic[2].negative_assertion[1].belief=0.400000

***landscape* Frame**

The frame.is_a is:
->landscape
The frame.is_part_of is:
->world
The frame.goal_of_analysis is:
->region label
The frame.level_in_tree is ->2

This frame has 3 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->hydrography

Sub node[0].bel_element_number is 7
Sub node[0].mean is 180.000000

Sub node[1].is_a is ->physiography

Sub node[1].bel_element_number is 8
Sub node[1].mean is 100.000000

Sub node[2].is_a is ->phytography

Sub node[2].bel_element_number is 9
Sub node[2].mean is 50.000000

characteristic[0] is ->roofing

characteristic[0] function is ->histogram

characteristic[0].positive_assertion[0].place=2
characteristic[0].positive_assertion[0].belief=0.700000
characteristic[0].positive_assertion[1].place=1

characteristic[0].positive_assertion[1].belief=0.300000

characteristic[0].negative_assertion[0].place=3
characteristic[0].negative_assertion[0].belief=0.800000
characteristic[0].negative_assertion[1].place=1
characteristic[0].negative_assertion[1].belief=0.200000

characteristic[1] is ->aluminum

characteristic[1] function is ->histogram

characteristic[1].positive_assertion[0].place=2
characteristic[1].positive_assertion[0].belief=0.700000
characteristic[1].positive_assertion[1].place=1
characteristic[1].positive_assertion[1].belief=0.300000

characteristic[1].negative_assertion[0].place=3
characteristic[1].negative_assertion[0].belief=0.800000
characteristic[1].negative_assertion[1].place=1
characteristic[1].negative_assertion[1].belief=0.200000

characteristic[2] is ->concrete

characteristic[2] function is ->histogram

characteristic[2].positive_assertion[0].place=2
characteristic[2].positive_assertion[0].belief=0.700000
characteristic[2].positive_assertion[1].place=1
characteristic[2].positive_assertion[1].belief=0.300000

characteristic[2].negative_assertion[0].place=3
characteristic[2].negative_assertion[0].belief=0.800000
characteristic[2].negative_assertion[1].place=1
characteristic[2].negative_assertion[1].belief=0.200000

urban (built up areas) **Frame**
The frame.is_a is:
->urban (built up areas)
The frame.is_part_of is:
->culture

The frame.goal_of_analysis is:
->region label
The frame.level_in_tree is ->3

This frame has 3 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->industry/utility

Sub node[0].bel_element_number is 10
Sub node[0].mean is 80.000000

Sub node[1].is_a is ->commercial/residential

Sub node[1].bel_element_number is 11
Sub node[1].mean is 85.000000

Sub node[2].is_a is ->institutional/governmental

Sub node[2].bel_element_number is 12
Sub node[2].mean is 90.000000

characteristic[0] is ->b1_building

characteristic[0] function is ->grow

characteristic[0].positive_assertion[0].place=4
characteristic[0].positive_assertion[0].belief=0.800000
characteristic[0].positive_assertion[1].place=1
characteristic[0].positive_assertion[1].belief=0.200000

characteristic[0].negative_assertion[0].place=248
characteristic[0].negative_assertion[0].belief=0.600000
characteristic[0].negative_assertion[1].place=1
characteristic[0].negative_assertion[1].belief=0.400000

characteristic[1] is ->p1_parking_lot

characteristic[1] function is ->grow

characteristic[1].positive_assertion[0].place=4

characteristic[1].positive_assertion[0].belief=0.800000
characteristic[1].positive_assertion[1].place=1
characteristic[1].positive_assertion[1].belief=0.200000

characteristic[1].negative_assertion[0].place=248
characteristic[1].negative_assertion[0].belief=0.600000
characteristic[1].negative_assertion[1].place=1
characteristic[1].negative_assertion[1].belief=0.400000

characteristic[2] is ->g1_garage

characteristic[2] function is ->grow

characteristic[2].positive_assertion[0].place=4
characteristic[2].positive_assertion[0].belief=0.800000
characteristic[2].positive_assertion[1].place=1
characteristic[2].positive_assertion[1].belief=0.200000

characteristic[2].negative_assertion[0].place=248
characteristic[2].negative_assertion[0].belief=0.700000
characteristic[2].negative_assertion[1].place=1
characteristic[2].negative_assertion[1].belief=0.300000

***transportation/navigation* Frame**

The frame.is_a is:
->transportation/navigation
The frame.is_part_of is:
->culture
The frame.goal_of_analysis is:
->region label
The frame.level_in_tree is ->3

This frame has 5 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->railroads

Sub node[0].bel_element_number is 13
Sub node[0].mean is 22.000000

Sub node[1].is_a is ->roads

Sub node[1].bel_element_number is 14

Sub node[1].mean is 22.200001

Sub node[2].is_a is ->aeronautical/aerospace

Sub node[2].bel_element_number is 15
Sub node[2].mean is 22.200001

Sub node[3].is_a is ->naval/marine

Sub node[3].bel_element_number is 16
Sub node[3].mean is 22.200001

Sub node[4].is_a is ->associated transportation features

Sub node[4].bel_element_number is 17
Sub node[4].mean is 22.200001

characteristic[0] is ->r1_runway

characteristic[0] function is ->grow

characteristic[0].positive_assertion[0].place=1
characteristic[0].positive_assertion[0].belief=0.200000
characteristic[0].positive_assertion[1].place=5
characteristic[0].positive_assertion[1].belief=0.800000

characteristic[0].negative_assertion[0].place=1
characteristic[0].negative_assertion[0].belief=0.300000
characteristic[0].negative_assertion[1].place=249
characteristic[0].negative_assertion[1].belief=0.700000

characteristic[1] is ->h1_hanger

characteristic[1] function is ->grow

characteristic[1].positive_assertion[0].place=1
characteristic[1].positive_assertion[0].belief=0.200000
characteristic[1].positive_assertion[1].place=5
characteristic[1].positive_assertion[1].belief=0.800000

characteristic[1].negative_assertion[0].place=1
characteristic[1].negative_assertion[0].belief=0.300000

characteristic[1].negative_assertion[1].place=249
characteristic[1].negative_assertion[1].belief=0.700000

***landmark/rural features* Frame**

The frame.is_a is:
->landmark/rural features
The frame.is_part_of is:
->culture
The frame.goal_of_analysis is:
->region label
The frame.level_in_tree is ->3

This frame has 5 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->communication/transmission

Sub node[0].bel_element_number is 18
Sub node[0].mean is 22.200001

Sub node[1].is_a is ->storage

Sub node[1].bel_element_number is 19
Sub node[1].mean is 22.200001

Sub node[2].is_a is ->agricultural

Sub node[2].bel_element_number is 20
Sub node[2].mean is 22.200001

Sub node[3].is_a is ->recreational

Sub node[3].bel_element_number is 21
Sub node[3].mean is 22.200001

Sub node[4].is_a is ->miscellaneous

Sub node[4].bel_element_number is 22
Sub node[4].mean is 22.200001

characteristic[0] is ->s1_storage

characteristic[0] function is ->grow

characteristic[0].positive_assertion[0].place=1
characteristic[0].positive_assertion[0].belief=0.200000
characteristic[0].positive_assertion[1].place=6
characteristic[0].positive_assertion[1].belief=0.800000

characteristic[0].negative_assertion[0].place=1
characteristic[0].negative_assertion[0].belief=0.200000
characteristic[0].negative_assertion[1].place=250
characteristic[0].negative_assertion[1].belief=0.800000

characteristic[1] is ->s1_stockyard

characteristic[1] function is ->grow

characteristic[1].positive_assertion[0].place=1
characteristic[1].positive_assertion[0].belief=0.200000
characteristic[1].positive_assertion[1].place=6
characteristic[1].positive_assertion[1].belief=0.800000

characteristic[1].negative_assertion[0].place=1
characteristic[1].negative_assertion[0].belief=0.200000
characteristic[1].negative_assertion[1].place=250
characteristic[1].negative_assertion[1].belief=0.800000

hydrography **Frame**

The frame.is_a is:
->hydrography
The frame.is_part_of is:
->landscape
The frame.goal_of_analysis is:
->region label
The frame.level_in_tree is ->3

This frame has 2 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->water

Sub node[0].bel_element_number is 23
Sub node[0].mean is 22.000000

Sub node[1].is_a is ->snow/ice

Sub node[1].bel_element_number is 24
Sub node[1].mean is 222.000000

characteristic[0] is ->s1_snow

characteristic[0] function is ->grow

characteristic[0].positive_assertion[0].place=1
characteristic[0].positive_assertion[0].belief=0.200000
characteristic[0].positive_assertion[1].place=7
characteristic[0].positive_assertion[1].belief=0.800000

characteristic[0].negative_assertion[0].place=1
characteristic[0].negative_assertion[0].belief=0.400000
characteristic[0].negative_assertion[1].place=251
characteristic[0].negative_assertion[1].belief=0.600000

characteristic[1] is ->w1_water

characteristic[1] function is ->grow

characteristic[1].positive_assertion[0].place=1
characteristic[1].positive_assertion[0].belief=0.200000
characteristic[1].positive_assertion[1].place=7
characteristic[1].positive_assertion[1].belief=0.800000

characteristic[1].negative_assertion[0].place=1
characteristic[1].negative_assertion[0].belief=0.400000
characteristic[1].negative_assertion[1].place=251
characteristic[1].negative_assertion[1].belief=0.600000

***physiography* Frame**
The frame.is_a is:
->physiography
The frame.is_part_of is:
->landscape
The frame.goal_of_analysis is:
->region label

The frame.level_in_tree is ->3

This frame has 2 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->exposed soils (surface composition)

Sub node[0].bel_element_number is 25
Sub node[0].mean is 22.000000

Sub node[1].is_a is ->landforms

Sub node[1].bel_element_number is 26
Sub node[1].mean is 33.000000

characteristic[0] is ->s1_soil

characteristic[0] function is ->grow

characteristic[0].positive_assertion[0].place=1
characteristic[0].positive_assertion[0].belief=0.800000
characteristic[0].positive_assertion[1].place=8
characteristic[0].positive_assertion[1].belief=0.000000

characteristic[0].negative_assertion[0].place=1
characteristic[0].negative_assertion[0].belief=0.400000
characteristic[0].negative_assertion[1].place=252
characteristic[0].negative_assertion[1].belief=0.600000

***phytography* Frame**
The frame.is_a is:
->phytography
The frame.is_part_of is:
->landscape
The frame.goal_of_analysis is:
->region label
The frame.level_in_tree is ->3

This frame has 4 sub nodes
The sub node operator is ->average of pixels

```
Sub node[0].is_a is ->cropland

Sub node[0].bel_element_number is 27
Sub node[0].mean is 80.000000

Sub node[1].is_a is ->rangeland

Sub node[1].bel_element_number is 28
Sub node[1].mean is 85.000000

Sub node[2].is_a is ->woodland

Sub node[2].bel_element_number is 29
Sub node[2].mean is 60.000000

Sub node[3].is_a is ->wetland

Sub node[3].bel_element_number is 30
Sub node[3].mean is 95.000000


characteristic[0] is ->t1_tree

characteristic[0] function is ->grow


characteristic[0].positive_assertion[0].place=1
characteristic[0].positive_assertion[0].belief=0.200000
characteristic[0].positive_assertion[1].place=9
characteristic[0].positive_assertion[1].belief=0.800000

characteristic[0].negative_assertion[0].place=1
characteristic[0].negative_assertion[0].belief=0.400000
characteristic[0].negative_assertion[1].place=253
characteristic[0].negative_assertion[1].belief=0.600000


characteristic[1] is ->l1_lawn

characteristic[1] function is ->grow


characteristic[1].positive_assertion[0].place=1
characteristic[1].positive_assertion[0].belief=0.200000
characteristic[1].positive_assertion[1].place=9
characteristic[1].positive_assertion[1].belief=0.800000
```

characteristic[1].negative_assertion[0].place=1
characteristic[1].negative_assertion[0].belief=0.400000
characteristic[1].negative_assertion[1].place=253
characteristic[1].negative_assertion[1].belief=0.600000

***industry/utility* Frame**

The frame.is_a is:
->industry/utility
The frame.is_part_of is:
->urban (built up areas)
The frame.goal_of_analysis is:
->region label
The frame.level_in_tree is ->4

This frame has 6 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->extraction

Sub node[0].bel_element_number is 31
Sub node[0].mean is 22.000000

Sub node[1].is_a is ->processing

Sub node[1].bel_element_number is 32
Sub node[1].mean is 32.000000

Sub node[2].is_a is ->power generation

Sub node[2].bel_element_number is 33
Sub node[2].mean is 33.000000

Sub node[3].is_a is ->fabrication industry

Sub node[3].bel_element_number is 34
Sub node[3].mean is 34.000000

Sub node[4].is_a is ->disposal

Sub node[4].bel_element_number is 35
Sub node[4].mean is 35.000000

Sub node[5].is_a is ->associated industrial structures

Sub node[5].bel_element_number is 36
Sub node[5].mean is 36.000000

characteristic[0] is ->b4_building

characteristic[0] function is ->grow

characteristic[0].positive_assertion[0].place=1
characteristic[0].positive_assertion[0].belief=0.200000
characteristic[0].positive_assertion[1].place=10
characteristic[0].positive_assertion[1].belief=0.800000

characteristic[0].negative_assertion[0].place=1
characteristic[0].negative_assertion[0].belief=0.300000
characteristic[0].negative_assertion[1].place=254
characteristic[0].negative_assertion[1].belief=0.700000

characteristic[1] is ->d1_disposal

characteristic[1] function is ->grow

characteristic[1].positive_assertion[0].place=1
characteristic[1].positive_assertion[0].belief=0.200000
characteristic[1].positive_assertion[1].place=10
characteristic[1].positive_assertion[1].belief=0.800000

characteristic[1].negative_assertion[0].place=1
characteristic[1].negative_assertion[0].belief=0.300000
characteristic[1].negative_assertion[1].place=254
characteristic[1].negative_assertion[1].belief=0.700000

***commercial/residential* Frame**

The frame.is_a is:
->commercial/residential
The frame.is_part_of is:
->urban (built up areas)
The frame.goal_of_analysis is:
->objects label
The frame.level_in_tree is ->99

This frame has 6 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->039 commercial buildings

Sub node[0].bel_element_number is 37
Sub node[0].mean is 37.000000

Sub node[1].is_a is ->040 apartments/hotel

Sub node[1].bel_element_number is 38
Sub node[1].mean is 38.000000

Sub node[2].is_a is ->041 house (single family dwelling)

Sub node[2].bel_element_number is 39
Sub node[2].mean is 39.000000

Sub node[3].is_a is ->042 mobile homes

Sub node[3].bel_element_number is 40
Sub node[3].mean is 40.000000

Sub node[4].is_a is ->043 barracks

Sub node[4].bel_element_number is 41
Sub node[4].mean is 41.000000

Sub node[5].is_a is ->044 display sign (large billboard, high)

Sub node[5].bel_element_number is 42
Sub node[5].mean is 42.000000

characteristic[0] is ->040 apartments/hotel

characteristic[0] function is ->grow

characteristic[0].positive_assertion[0].place=1
characteristic[0].positive_assertion[0].belief=0.200000
characteristic[0].positive_assertion[1].place=11
characteristic[0].positive_assertion[1].belief=0.800000

characteristic[0].negative_assertion[0].place=1

characteristic[0].negative_assertion[0].belief=0.300000
characteristic[0].negative_assertion[1].place=256
characteristic[0].negative_assertion[1].belief=0.700000

characteristic[1] is ->041 house (single family dwelling)

characteristic[1] function is ->grow

characteristic[1].positive_assertion[0].place=1
characteristic[1].positive_assertion[0].belief=0.200000
characteristic[1].positive_assertion[1].place=11
characteristic[1].positive_assertion[1].belief=0.800000

characteristic[1].negative_assertion[0].place=1
characteristic[1].negative_assertion[0].belief=0.300000
characteristic[1].negative_assertion[1].place=256
characteristic[1].negative_assertion[1].belief=0.700000

institutional/governmental **Frame**

The frame.is_a is:
->institutional/governmental
The frame.is_part_of is:
->urban (built up areas)
The frame.goal_of_analysis is:
->region label
The frame.level_in_tree is ->4

This frame has 11 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->045 governmental administration building

Sub node[0].bel_element_number is 43
Sub node[0].mean is 43.000000

Sub node[1].is_a is ->046 military admin/operations building

Sub node[1].bel_element_number is 44
Sub node[1].mean is 44.000000

Sub node[2].is_a is ->047 capitol building

Sub node[2].bel_element_number is 45
Sub node[2].mean is 45.000000

Sub node[3].is_a is ->048 hospital

Sub node[3].bel_element_number is 46
Sub node[3].mean is 46.000000

Sub node[4].is_a is ->049 prison

Sub node[4].bel_element_number is 47
Sub node[4].mean is 47.000000

Sub node[5].is_a is ->050 palace

Sub node[5].bel_element_number is 48
Sub node[5].mean is 48.000000

Sub node[6].is_a is ->051 museum

Sub node[6].bel_element_number is 49
Sub node[6].mean is 49.000000

Sub node[7].is_a is ->052 observatory

Sub node[7].bel_element_number is 50
Sub node[7].mean is 50.000000

Sub node[8].is_a is ->053 church/tabernacle

Sub node[8].bel_element_number is 51
Sub node[8].mean is 51.000000

Sub node[9].is_a is ->054 mosque

Sub node[9].bel_element_number is 52
Sub node[9].mean is 52.000000

Sub node[10].is_a is ->055 cemetary building

Sub node[10].bel_element_number is 53
Sub node[10].mean is 53.000000

characteristic[0] is ->b4_building

characteristic[0] function is ->grow

characteristic[0].positive_assertion[0].place=1
characteristic[0].positive_assertion[0].belief=0.200000
characteristic[0].positive_assertion[1].place=12
characteristic[0].positive_assertion[1].belief=0.800000

characteristic[0].negative_assertion[0].place=1
characteristic[0].negative_assertion[0].belief=0.300000
characteristic[0].negative_assertion[1].place=255
characteristic[0].negative_assertion[1].belief=0.700000

characteristic[1] is ->p3_parking_lot

characteristic[1] function is ->grow

characteristic[1].positive_assertion[0].place=1
characteristic[1].positive_assertion[0].belief=0.200000
characteristic[1].positive_assertion[1].place=12
characteristic[1].positive_assertion[1].belief=0.800000

characteristic[1].negative_assertion[0].place=1
characteristic[1].negative_assertion[0].belief=0.300000
characteristic[1].negative_assertion[1].place=255
characteristic[1].negative_assertion[1].belief=0.700000

***cropland* Frame**

The frame.is_a is:
->cropland
The frame.is_part_of is:
->phytography
The frame.goal_of_analysis is:
->object label
The frame.level_in_tree is ->99

This frame has 7 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->190 orchard/plantation

Sub node[0].bel_element_number is 167
Sub node[0].mean is 167.000000

Sub node[1].is_a is ->191 shelterbelt/hedgerow

Sub node[1].bel_element_number is 168
Sub node[1].mean is 168.000000

Sub node[2].is_a is ->192 nursery/grove

Sub node[2].bel_element_number is 169
Sub node[2].mean is 169.000000

Sub node[3].is_a is ->193 vineyard

Sub node[3].bel_element_number is 170
Sub node[3].mean is 170.000000

Sub node[4].is_a is ->194 crop (cultivated)

Sub node[4].bel_element_number is 171
Sub node[4].mean is 171.000000

Sub node[5].is_a is ->195 cranberry bog

Sub node[5].bel_element_number is 172
Sub node[5].mean is 172.000000

Sub node[6].is_a is ->196 rice paddy

Sub node[6].bel_element_number is 173
Sub node[6].mean is 173.000000

characteristic[0] is ->c1_crops

characteristic[0] function is ->grow

characteristic[0].positive_assertion[0].place=1
characteristic[0].positive_assertion[0].belief=0.200000
characteristic[0].positive_assertion[1].place=27
characteristic[0].positive_assertion[1].belief=0.800000

characteristic[0].negative_assertion[0].place=1
characteristic[0].negative_assertion[0].belief=0.100000
characteristic[0].negative_assertion[1].place=267

characteristic[0].negative_assertion[1].belief=0.900000

rangeland **Frame**

The frame.is_a is:
->rangeland
The frame.is_part_of is:
->phytography
The frame.goal_of_analysis is:
->object label
The frame.level_in_tree is ->99

This frame has 3 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->197 herbaceous (grassland)

Sub node[0].bel_element_number is 174
Sub node[0].mean is 50.000000

Sub node[1].is_a is ->198 shrub/brush

Sub node[1].bel_element_number is 175
Sub node[1].mean is 40.000000

Sub node[2].is_a is ->199 mixed (shrub/brush and grass)

Sub node[2].bel_element_number is 176
Sub node[2].mean is 45.000000

characteristic[0] is ->l1_lawn

characteristic[0] function is ->grow

characteristic[0].positive_assertion[0].place=1
characteristic[0].positive_assertion[0].belief=0.200000
characteristic[0].positive_assertion[1].place=28
characteristic[0].positive_assertion[1].belief=0.800000

characteristic[0].negative_assertion[0].place=1
characteristic[0].negative_assertion[0].belief=0.300000
characteristic[0].negative_assertion[1].place=268
characteristic[0].negative_assertion[1].belief=0.700000

woodland **Frame**

The frame.is_a is:
->woodland
The frame.is_part_of is:
->phytography
The frame.goal_of_analysis is:
->object label
The frame.level_in_tree is ->99

This frame has 5 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->200 decidous

Sub node[0].bel_element_number is 177
Sub node[0].mean is 45.000000

Sub node[1].is_a is ->201 caniferous

Sub node[1].bel_element_number is 178
Sub node[1].mean is 50.000000

Sub node[2].is_a is ->202 mixed (decidous and caniferous)

Sub node[2].bel_element_number is 179
Sub node[2].mean is 55.000000

Sub node[3].is_a is ->230 mangrove

Sub node[3].bel_element_number is 180
Sub node[3].mean is 60.000000

Sub node[4].is_a is ->204 nipa palm

Sub node[4].bel_element_number is 181
Sub node[4].mean is 65.000000

characteristic[0] is ->200 decidous

characteristic[0] function is ->grow

characteristic[0].positive_assertion[0].place=1
characteristic[0].positive_assertion[0].belief=0.200000
characteristic[0].positive_assertion[1].place=29
characteristic[0].positive_assertion[1].belief=0.800000

characteristic[0].negative_assertion[0].place=1
characteristic[0].negative_assertion[0].belief=0.400000
characteristic[0].negative_assertion[1].place=269
characteristic[0].negative_assertion[1].belief=0.600000

***wetland* Frame**

The frame.is_a is:
->wetland
The frame.is_part_of is:
->phytography
The frame.goal_of_analysis is:
->object label
The frame.level_in_tree is ->99

This frame has 3 sub nodes
The sub node operator is ->average of pixels

Sub node[0].is_a is ->205 swamp (trees, brush/shrubs in water)

Sub node[0].bel_element_number is 182
Sub node[0].mean is 182.000000

Sub node[1].is_a is ->206 marsh (grass, cat tails, etc in wate

Sub node[1].bel_element_number is 183
Sub node[1].mean is 183.000000

Sub node[2].is_a is ->207 peat bog/cuttings

Sub node[2].bel_element_number is 184
Sub node[2].mean is 184.000000

characteristic[0] is ->w1_water

characteristic[0] function is ->grow

characteristic[0].positive_assertion[0].place=1

characteristic[0].positive_assertion[0].belief=0.200000
characteristic[0].positive_assertion[1].place=30
characteristic[0].positive_assertion[1].belief=0.000000

characteristic[0].negative_assertion[0].place=1
characteristic[0].negative_assertion[0].belief=0.400000
characteristic[0].negative_assertion[1].place=270
characteristic[0].negative_assertion[1].belief=0.600000

APPENDIX 4 - PC-FILE Database Implementation of Frames

In the TDBU system, the frame is a C data structure. Figure A4.1 shows the structure.

In the PC-FILE database program there is a limit to the number of fields in a data record. The frame data structure has far too many fields for a PC-FILE data record. To work around this limitation, TDBU uses multiple files. The following paragraphs describe the multiple files and how they fit together to form the frames.

There are three types of files; the *frames* file; the *sub nodes* file; and the *characteristics* file.

The principle file is frames.dta. It is the only file using the frames pattern. This file has one data record for each frame entered into the database. Each data record gives the frame and introduction. It also gives the names of the sub nodes file and the characteristics file for that frame. The system uses these file names to read that information.

Given in figure A4.2 is the data record for the world frame.

```
file frames.dta

is_a              [world]
is_part_of        [nothing]
level_in_tree     [1]
goal_of_anal[region label]
no_of_chars [0]
chars_file        [c:\file\data\nochar.dta]
no_of_nodes       [2]
op_name           [average of pixels]
subs_file         [c:\file|data\d1sub.dta]
```

Figure A4.2 - Data Record for the *world* Frame

Every frame uses the subs pattern. The subs pattern gives information regarding the sub nodes of a frame. Shown in figure A4.3 is the first entry in the sub nodes file for

```
struct assertion_type{
   short place;
   float belief_value;
};


struct structural_characteristics{
        char   is_a[MAX_NAME_LENGTH];
        short  space;
        char   funct_name[MAX_NAME_LENGTH];
        struct assertion_type positive_assertion[ASSERTION_LENGTH];
        struct assertion_type negative_assertion[ASSERTION_LENGTH];
        struct assertion_type blind_pos_assertion[ASSERTION_LENGTH];
        struct assertion_type blind_neg_assertion[ASSERTION_LENGTH];
        short  result;
        short  other_objects[OTHERS_LENGTH];
};

struct class{
        char  is_a[MAX_NAME_LENGTH];
        short bel_element_number;
        float mean;
};


struct frame{
        char   is_a[MAX_NAME_LENGTH];
        char   is_part_of[MAX_NAME_LENGTH];
        short  level_in_tree;
        char   goal_of_analysis[MAX_NAME_LENGTH];
        short  number_of_structural_characteristics;
        struct structural_characteristics characteristic[CHARACTERISTIC_LENGTH];
        short  count_of_sub_nodes;
        char   op_name[MAX_NAME_LENGTH];
        struct class sub_class[SUB_NODE_LENGTH];
};
```

Figure A4.1 - C Structure of a Frame

world. It shows that the first sub node is *culture*. It gives the information needed by the sub node operator. Since *world* has two sub nodes, there will be two data records in its sub nodes file.

---

file d1subs.dta

Record number 1

is_a	[culture]
bel_el_no	[2]
mean	[88.0]

Figure A4.3 - First Entry of Sub Nodes File for Frame *world*

Every frame uses the chars pattern. It lists the intrinsic characteristics of a frame. The system will detect the object named here and will make the assertions listed here. If the system detects the object, then it will make the positive assertion. If the system does not detect the object, then it will make the negative assertion.

Shown in figure A4.4 is the first data record from the d2chars.dta file. This shows that b1_building is the object of interest. The presence of a building is the first characteristic of the frame *culture*.

file d2chars.dta

Record number 1

is_a [b1_building]
funct_name [grow]

pos_1_place [2]
pos_1_belief[0.8]
pos_2_place [1]
pos_2_belief[0.2]

neg_1_place [3]
neg_1_belief[0.6]
neg_2_place [1]
neg_2_belief[0.4]

Figure A4.4 - First Data Record of d2chars.dta File

APPENDIX 5 - Object Descriptors

```
strcpy(d.is_a,"040 apartments/hotel");
strcpy(d.material, "roofing");
d.min_area        = 844;
d.max_area        = 2025;
d.min_width       =  3;
d.max_width       = 46;
d.min_height      = 46;
d.max_height      = 146;
d.min_w_to_h_ratio = 0.02;
d.max_w_to_h_ratio = 1.0;
d.min_angle        = -90.0;
d.max_angle        = 90.0;
```

```
strcpy(d.is_a,"041 house (single family dwelling)");
strcpy(d.material, "roofing");
d.min_area        = 56;
d.max_area        = 840;
d.min_width       = 7;
d.max_width       = 46;
d.min_height      = 11;
d.max_height      = 68;
d.min_w_to_h_ratio = 0.3;
d.max_w_to_h_ratio = 1.0;
d.min_angle        = -90.0;
d.max_angle        = 90.0;
```

```
strcpy(d.is_a,"069 primary road");
strcpy(d.material, "concrete");
d.min_area        = 400;
d.max_area        = 40000;
d.min_width       =  0;
d.max_width       = 203;
d.min_height      =  0;
d.max_height      = 203;
d.min_w_to_h_ratio = 0.0;
d.max_w_to_h_ratio = 1.0;
d.min_angle        = -90.0;
d.max_angle        = 90.0;
```

```
strcpy(d.is_a,"097 footpath/trail");
strcpy(d.material, "concrete");
d.min_area        =   5;
d.max_area        = 40000;
d.min_width       =   0;
d.max_width       =  82;
d.min_height      =   0;
d.max_height      =  82;
d.min_w_to_h_ratio = 0.0;
d.max_w_to_h_ratio = 1.0;
d.min_angle         = -90.0;
d.max_angle         = 90.0;


strcpy(d.is_a,"200 decidous");
strcpy(d.material, "leaves");
d.min_area        =   45;
d.max_area        = 18225;
d.min_width        =    -1;
d.max_width       =  248;
d.min_height       =    -1;
d.max_height      =  248;
d.min_w_to_h_ratio =   0.0;
d.max_w_to_h_ratio =   1.0;
d.min_angle         = -90.0;
d.max_angle         =  90.0;


strcpy(d.is_a,"b1_building");
strcpy(d.material, "roofing");
d.min_area        =  113;
d.max_area        = 1350;
d.min_width       =    0;
d.max_width       =  101;
d.min_height      =    0;
d.max_height      =  101;
d.min_w_to_h_ratio =  0.0;
d.max_w_to_h_ratio =  1.0;
d.min_angle         = -90.0;
d.max_angle       =  90.0;
```

```
strcpy(d.is_a,"b2_building");
strcpy(d.material, "roofing");
d.min_area        =  225;
d.max_area        =  720;
d.min_width       =   23;
d.max_width       =   45;
d.min_height      =   34;
d.max_height      =   68;
d.min_w_to_h_ratio =  0.01;
d.max_w_to_h_ratio =   1.0;
d.min_angle        = -90.0;
d.max_angle       =  90.0;


strcpy(d.is_a,"b3_building");
strcpy(d.material, "roofing");
d.min_area        =   720;
d.max_area        =  1125;
d.min_width       =    27;
d.max_width       =    40;
d.min_height      =    56;
d.max_height      =   101;
d.min_w_to_h_ratio =  0.01;
d.max_w_to_h_ratio =   1.0;
d.min_angle        = -90.0;
d.max_angle       =  90.0;


strcpy(d.is_a,"b4_building");
strcpy(d.material, "roofing");
d.min_area        = 8100;
d.max_area        = 14400;
d.min_width       = 135;
d.max_width       = 180;
d.min_height      = 135;
d.max_height      = 180;
d.min_w_to_h_ratio = 0.01;
d.max_w_to_h_ratio = 1.0;
d.min_angle        = 0.0;
d.max_angle        = 90.0;
```

```
strcpy(d.is_a,"c1_crops");
strcpy(d.material, "crops");
d.min_area          = 225;
d.max_area          = 18225;
d.min_width         = 23;
d.max_width         = 203;
d.min_height        = 23;
d.max_height        = 203;
d.min_w_to_h_ratio = 0.01;
d.max_w_to_h_ratio = 1.0;
d.min_angle          = -90.0;
d.max_angle          = 90.0;


strcpy(d.is_a,"d1_disposal");
strcpy(d.material, "waste");
d.min_area          = 5625;
d.max_area          = 11025;
d.min_width         = 113;
d.max_width         = 158;
d.min_height        = 113;
d.max_height        = 158;
d.min_w_to_h_ratio = 0.01;
d.max_w_to_h_ratio = 0.2;
d.min_angle          = -90.0;
d.max_angle          =  90.0;


strcpy(d.is_a,"r1_runway");
strcpy(d.material, "concrete");
d.min_area          =  788;
d.max_area          = 4050;
d.min_width         = 23;
d.max_width         = 45;
d.min_height        = 56;
d.max_height        = 203;
d.min_w_to_h_ratio = 0.01;
d.max_w_to_h_ratio = 0.2;
d.min_angle          = -90.0;
d.max_angle          =  90.0;


strcpy(d.is_a,"h1_hanger");
strcpy(d.material, "roofing");
```

```
d.min_area        = 8100;
d.max_area        = 14400;
d.min_width       = 135;
d.max_width       = 180;
d.min_height      = 135;
d.max_height      = 180;
d.min_w_to_h_ratio = 0.5;
d.max_w_to_h_ratio = 1.0;
d.min_angle        = -90.0;
d.max_angle        =  90.0;


strcpy(d.is_a,"d1_driveway");
strcpy(d.material, "concrete");
d.min_area        = 1125;
d.max_area        = 18225;
d.min_width       = 23;
d.max_width       = 203;
d.min_height      = 23;
d.max_height      = 203;
d.min_w_to_h_ratio = 0.01;
d.max_w_to_h_ratio = 1.0;
d.min_angle        = -90.0;
d.max_angle        = 90.0;


strcpy(d.is_a,"g1_garage");
strcpy(d.material, "aluminum");
d.min_area        = 90;
d.max_area        = 518;
d.min_width        = 0;
d.max_width       = 23;
d.min_height       = 0;
d.max_height      = 45;
d.min_w_to_h_ratio = 0.0;
d.max_w_to_h_ratio = 1.0;
d.min_angle        = -90.0;
d.max_angle        = 90.0;


strcpy(d.is_a,"l1_lawn");
strcpy(d.material, "grass");
```

```
d.min_area        =    56;
d.max_area        = 18225;
d.min_width       =   11;
d.max_width       = 203;
d.min_height      =   11;
d.max_height      = 203;
d.min_w_to_h_ratio = 0.01;
d.max_w_to_h_ratio = 1.0;
d.min_angle          = -90.0;
d.max_angle          = 90.0;



strcpy(d.is_a,"p1_parking_lot");
strcpy(d.material, "concrete");
d.min_area         = 23;
d.max_area         = 18225;
d.min_width          = 0;
d.max_width        = 203;
d.min_height         = 0;
d.max_height       = 203;
d.min_w_to_h_ratio = 0.0;
d.max_w_to_h_ratio = 1.0;
d.min_angle          = -90.0;
d.max_angle          = 90.0;



strcpy(d.is_a,"p2_parking_lot");
strcpy(d.material, "concrete");
d.min_area         = 3600;
d.max_area         = 4300;
d.min_width        = 68;
d.max_width        = 135;
d.min_height       = 135;
d.max_height       = 225;
d.min_w_to_h_ratio = 0.01;
d.max_w_to_h_ratio = 1.0;
d.min_angle          = -90.0;
d.max_angle          = 90.0;


strcpy(d.is_a,"p3_parking_lot");
strcpy(d.material, "concrete");
```

```
d.min_area        = 3600;
d.max_area        = 8100;
d.min_width       = 90;
d.max_width       = 135;
d.min_height      = 90;
d.max_height      = 135;
d.min_w_to_h_ratio = 0.01;
d.max_w_to_h_ratio = 1.0;
d.min_angle        = -90.0;
d.max_angle        = 90.0;


strcpy(d.is_a,"s1_storage");
strcpy(d.material, "aluminum");
d.min_area        = 8100;
d.max_area        = 18225;
d.min_width       = 135;
d.max_width       = 203;
d.min_height      = 135;
d.max_height      = 203;
d.min_w_to_h_ratio = 0.0;
d.max_w_to_h_ratio = 1.0;
d.min_angle        = -90.0;
d.max_angle        = 90.0;


strcpy(d.is_a,"s1_stockyard");
strcpy(d.material, "dirt");
d.min_area        = 5625;
d.max_area        = 11025;
d.min_width       = 113;
d.max_width       = 158;
d.min_height      = 113;
d.max_height      = 158;
d.min_w_to_h_ratio = 0.5;
d.max_w_to_h_ratio = 1.0;
d.min_angle        = -90.0;
d.max_angle        = 90.0;


strcpy(d.is_a,"s1_soil");
strcpy(d.material, "dirt");
d.min_area        = 8100;
d.max_area        = 18225;
```

```
d.min_width       = 23;
d.max_width       = 203;
d.min_height      = 23;
d.max_height      = 203;
d.min_w_to_h_ratio = 0.0;
d.max_w_to_h_ratio = 1.0;
d.min_angle        = -90.0;
d.max_angle        = 90.0;


strcpy(d.is_a,"t1_tree");
strcpy(d.material, "leaves");
d.min_area        =    45;
d.max_area        = 18225;
d.min_width       = -3;
d.max_width       = 248;
d.min_height      = -3;
d.max_height      = 248;
d.min_w_to_h_ratio = 0.0;
d.max_w_to_h_ratio = 1.0;
d.min_angle        = -90.0;
d.max_angle        = 90.0;


strcpy(d.is_a,"w1_water");
strcpy(d.material, "water");
d.min_area        = 8100;
d.max_area        = 18225;
d.min_width       = 23;
d.max_width       = 203;
d.min_height      = 23;
d.max_height      = 203;
d.min_w_to_h_ratio = 0.0;
d.max_w_to_h_ratio = 1.0;
d.min_angle        = -90.0;
d.max_angle        = 90.0;


strcpy(d.is_a,"s1_snow");
strcpy(d.material, "snow");
d.min_area        = 8100;
d.max_area        = 18225;
d.min_width       = 23;
```

```
d.max_width        = 203;
d.min_height       = 23;
d.max_height       = 203;
d.min_w_to_h_ratio = 0.0;
d.max_w_to_h_ratio = 1.0;
d.min_angle        = -90.0;
d.max_angle        = 90.0;
```

APPENDIX 6 - Spectral Information

```
strcpy(s.is_a, "roofing");
s.num_of_channels           =   1;
s.spectrum[0].channel       =   2;
s.spectrum[0].low_threshold   = 185;
s.spectrum[0].high_threshold  = 255;


strcpy(s.is_a, "2roofing");
s.num_of_channels           =   1;
s.spectrum[0].channel       =   3;
s.spectrum[0].low_threshold   =  30;
s.spectrum[0].high_threshold  =  50;


strcpy(s.is_a, "leaves");
s.num_of_channels           =   1;
s.spectrum[0].channel       =   1;
s.spectrum[0].low_threshold   =   0;
s.spectrum[0].high_threshold  =  45;


strcpy(s.is_a, "crops");
s.num_of_channels           =   2;
s.spectrum[0].channel       =   1;
s.spectrum[0].low_threshold   =  30;
s.spectrum[0].high_threshold  =  50;
s.spectrum[1].channel       =   2;
s.spectrum[1].low_threshold   =  75;
s.spectrum[1].high_threshold  =  95;


strcpy(s.is_a, "waste");
s.num_of_channels           =   1;
s.spectrum[0].channel       =   1;
s.spectrum[0].low_threshold   =  75;
s.spectrum[0].high_threshold  = 150;


s.num_of_channels           =   1;
s.spectrum[0].channel       =   1;
s.spectrum[0].low_threshold   =  75;
s.spectrum[0].high_threshold  = 175;
```

```
strcpy(s.is_a, "aluminum");
s.num_of_channels              =   1;
s.spectrum[0].channel          =   1;
s.spectrum[0].low_threshold    = 190;
s.spectrum[0].high_threshold   = 255;


strcpy(s.is_a, "gras");
s.num_of_channels              =   2;
s.spectrum[0].channel          =   1;
s.spectrum[0].low_threshold    =  30;
s.spectrum[0].high_threshold   =  70;
s.spectrum[1].channel          =   2;
s.spectrum[1].low_threshold    =  65;
s.spectrum[1].high_threshold   = 200;


strcpy(s.is_a, "dirt");
s.num_of_channels              =   1;
s.spectrum[0].channel          =   1;
s.spectrum[0].low_threshold    =  45;
s.spectrum[0].high_threshold   =  65;


strcpy(s.is_a, "water");
s.num_of_channels              =   1;
s.spectrum[0].channel          =   2;
s.spectrum[0].low_threshold    =  70;
s.spectrum[0].high_threshold   =  80;


strcpy(s.is_a, "snow");
s.num_of_channels              =   1;
s.spectrum[0].channel          =   1;
s.spectrum[0].low_threshold    = 225;
s.spectrum[0].high_threshold   = 255;
```

APPENDIX 7 - Output of the Object Detector

Section 1

REPRES> The descriptor was:
Name--b1_building
Number of channels--1
Thresholds 185 255
Areas 112 1350 (feet)
Widths 0 99 (feet)
Heights 0 99 (feet)
Ratios 0.000000 1.000000
Angles -90.000000 90.000000

b1_building - Object number 1
Center at x=1005 y=917
Area=470 (feet)
Height=40 Width=18 (feet)
Width to Height Ratio=0.444444 Angle=81.907654

b1_building - Object number 2
Center at x=1111 y=947
Area=758 (feet)
Height=45 Width=22 (feet)
Width to Height Ratio=0.500000 Angle=9.771790

b1_building - Object number 3
Center at x=935 y=1064
Area=1010 (feet)
Height=54 Width=36 (feet)
Width to Height Ratio=0.666667 Angle=43.494610

b1_building - Object number 4
Center at x=992 y=1070
Area=578 (feet)
Height=36 Width=27 (feet)
Width to Height Ratio=0.750000 Angle=6.920860

b1_building - Object number 5
Center at x=1100 y=1115
Area=938 (feet)
Height=63 Width=36 (feet)
Width to Height Ratio=0.571429 Angle=85.148834

REPRES> The descriptor was:
Name--p1_parking_lot
Number of channels--1
Thresholds 75 155
Areas 22 18225 (feet)
Widths 0 202 (feet)
Heights 0 202 (feet)
Ratios 0.000000 1.000000
Angles -90.000000 90.000000

p1_parking_lot - Object number 1
Center at x=922 y=913
Area=38 (feet)
Height=4 Width=4 (feet)
Width to Height Ratio=1.000000 Angle=0.000000

p1_parking_lot - Object number 2
Center at x=958 y=913
Area=38 (feet)
Height=4 Width=4 (feet)
Width to Height Ratio=1.000000 Angle=0.000000

p1_parking_lot - Object number 3
Center at x=1032 y=1005
Area=4898 (feet)
Height=58 Width=0 (feet)
Width to Height Ratio=0.000000 Angle=-38.920456

p1_parking_lot - Object number 4
Center at x=976 y=931
Area=110 (feet)
Height=22 Width=4 (feet)
Width to Height Ratio=0.200000 Angle=1.341988

p1_parking_lot - Object number 5
Center at x=994 y=931
Area=38 (feet)
Height=4 Width=4 (feet)
Width to Height Ratio=1.000000 Angle=0.000000

p1_parking_lot - Object number 6
Center at x=1012 y=940
Area=38 (feet)
Height=4 Width=4 (feet)
Width to Height Ratio=1.000000 Angle=0.000000

p1_parking_lot - Object number 7

Center at x=1093 y=967
Area=38 (feet)
Height=4 Width=4 (feet)
Width to Height Ratio=1.000000 Angle=0.000000

p1_parking_lot - Object number 8
Center at x=1111 y=976
Area=38 (feet)
Height=4 Width=4 (feet)
Width to Height Ratio=1.000000 Angle=0.000000

p1_parking_lot - Object number 9
Center at x=1093 y=985
Area=38 (feet)
Height=4 Width=4 (feet)
Width to Height Ratio=1.000000 Angle=0.000000

p1_parking_lot - Object number 10
Center at x=1129 y=985
Area=38 (feet)
Height=4 Width=4 (feet)
Width to Height Ratio=1.000000 Angle=0.000000

p1_parking_lot - Object number 11
Center at x=940 y=1007
Area=74 (feet)
Height=13 Width=4 (feet)
Width to Height Ratio=0.333333 Angle=3.562767

p1_parking_lot - Object number 12
Center at x=967 y=1016
Area=146 (feet)
Height=4 Width=4 (feet)
Width to Height Ratio=1.000000 Angle=-58.045357

p1_parking_lot - Object number 13
Center at x=922 y=1021
Area=38 (feet)
Height=4 Width=4 (feet)
Width to Height Ratio=1.000000 Angle=0.000000

p1_parking_lot - Object number 14
Center at x=1001 y=1039
Area=110 (feet)
Height=22 Width=13 (feet)
Width to Height Ratio=0.600000 Angle=27.536560

p1_parking_lot - Object number 15
Center at x=1084 y=1075
Area=38 (feet)
Height=4 Width=4 (feet)
Width to Height Ratio=1.000000 Angle=0.000000

p1_parking_lot - Object number 16
Center at x=1120 y=1079
Area=74 (feet)
Height=13 Width=4 (feet)
Width to Height Ratio=0.333333 Angle=3.562767

REPRES> The descriptor was:
Name--g1_garage
Number of channels--1
Thresholds 190 255
Areas 90 517 (feet)
Widths 0 22 (feet)
Heights 0 45 (feet)
Ratios 0.000000 1.000000
Angles -90.000000 90.000000

g1_garage - Object number 1
Center at x=915 y=929
Area=110 (feet)
Height=18 Width=13 (feet)
Width to Height Ratio=0.750000 Angle=90.000000

g1_garage - Object number 2
Center at x=971 y=949
Area=146 (feet)
Height=13 Width=9 (feet)
Width to Height Ratio=0.666667 Angle=32.401825

g1_garage - Object number 3
Center at x=1001 y=960
Area=254 (feet)
Height=18 Width=13 (feet)
Width to Height Ratio=0.750000 Angle=58.962204

g1_garage - Object number 4
Center at x=938 y=985
Area=506 (feet)
Height=0 Width=0 (feet)
Width to Height Ratio=0.750000 Angle=67.772232

g1_garage - Object number 5
Center at x=1102 y=1001
Area=218 (feet)
Height=22 Width=4 (feet)
Width to Height Ratio=0.200000 Angle=53.548679

g1_garage - Object number 6
Center at x=992 y=1016
Area=182 (feet)
Height=27 Width=18 (feet)
Width to Height Ratio=0.666667 Angle=71.563690

g1_garage - Object number 7
Center at x=1093 y=1052
Area=218 (feet)
Height=22 Width=18 (feet)
Width to Height Ratio=0.800000 Angle=87.855263

g1_garage - Object number 8
Center at x=1124 y=1066
Area=146 (feet)
Height=13 Width=9 (feet)
Width to Height Ratio=0.666667 Angle=32.401825

Section 2

REPRES> The descriptor was:
Name--041 house (single family dwelling)
Number of channels--1
Thresholds 185 255
Areas 54 839 (feet)
Widths 6 45 (feet)
Heights 9 67 (feet)
Ratios 0.300000 1.000000
Angles -90.000000 90.000000

041 house (single family dwelling) - Object number 1
Center at x=1010 y=920
Area=315 (feet)
Height=31 Width=18 (feet)
Width to Height Ratio=0.571429 Angle=84.632584

041 house (single family dwelling) - Object number 2
Center at x=1109 y=942

Area=672 (feet)
Height=54 Width=31 (feet)
Width to Height Ratio=0.583333 Angle=-3.658417

041 house (single family dwelling) - Object number 3
Center at x=944 y=1066
Area=612 (feet)
Height=54 Width=22 (feet)
Width to Height Ratio=0.416667 Angle=-1.597931

041 house (single family dwelling) - Object number 4
Center at x=987 y=1064
Area=546 (feet)
Height=27 Width=22 (feet)
Width to Height Ratio=0.833333 Angle=-87.948303

041 house (single family dwelling) - Object number 5
Center at x=1086 y=1109
Area=618 (feet)
Height=40 Width=31 (feet)
Width to Height Ratio=0.777778 Angle=12.047958

Section 3

REPRES> The descriptor was:
Name--041 house (single family dwelling)Number of channels--1
Thresholds 185 255Areas 54 839 (feet)
Widths 6 45 (feet)Heights 9 67 (feet)
Ratios 0.300000 1.000000
Angles -90.000000 90.000000

041 house (single family dwelling) - Object number 1
Center at x=1010 y=920
Area=315 (feet)
Height=31
Width=18 (feet)
Width to Height Ratio=0.571429
Angle=84.632584

041 house (single family dwelling) - Object number 2
Center at x=1109 y=942
Area=672 (feet)
Height=54
Width=31 (feet)
Width to Height Ratio=0.583333

Angle=-3.658417

041 house (single family dwelling) - Object number 3
Center at x=944 y=1066
Area=612 (feet)
Height=54
Width=22 (feet)
Width to Height Ratio=0.416667
Angle=-1.597931

041 house (single family dwelling) - Object number 4
Center at x=987 y=1064
Area=546 (feet)
Height=27
Width=22 (feet)
Width to Height Ratio=0.833333
Angle=-87.948303

041 house (single family dwelling) - Object number 5
Center at x=1086 y=1109
Area=618 (feet)
Height=40
Width=31 (feet)
Width to Height Ratio=0.777778
Angle=12.047958

REPRES> The descriptor was:
Name--069 primary road
Number of channels--1
Thresholds 75 175 Areas 398 39998 (feet)
Widths 0 202 (feet)
Heights 0 202 (feet)
Ratios 0.000000 1.000000
Angles -90.000000 90.000000

069 primary road - Object number 1
Center at x=1028 y=1003
Area=4911 (feet)
Height=58
Width=45 (feet)
Width to Height Ratio=0.769231
Angle=-38.635651

REPRES> The descriptor was:
Name--097 footpath/trail
Number of channels--1
Thresholds 75 175 Areas 4 39998 (feet)

Widths 0 81 (feet)Heights 0 81 (feet)
Ratios 0.000000 1.000000
Angles -90.000000 90.000000

097 footpath/trail - Object number 1
Center at x=295823 y=590735
Area=2 (feet)
Height=4
Width=51 (feet)
Width to Height Ratio=0.000000
Angle=0.000000

097 footpath/trail - Object number 2
Center at x=295823 y=1180559 Area=2 (feet)
Height=9
Width=112 (feet)
Width to Height Ratio=0.500000
Angle=0.000000

097 footpath/trail - Object number 3
Center at x=295823 y=295823
Area=2 (feet)
Height=58
Width=65 (feet)
Width to Height Ratio=0.769231
Angle=0.000000

097 footpath/trail - Object number 4
Center at x=911 y=911
Area=294914 (feet) Height=0
Width=94 (feet)
Width to Height Ratio=0.769231
Angle=0.000000

097 footpath/trail - Object number 5
Center at x=974 y=924
Area=2 (feet)
Height=22
Width=4 (feet)
Width to Height Ratio=0.000000
Angle=0.000000

097 footpath/trail - Object number 6
Center at x=989 y=931
Area=2 (feet)
Height=9
Width=96 (feet)
Width to Height Ratio=0.500000
Angle=0.000000

097 footpath/trail - Object number 7
Center at x=1010 y=938
Area=2 (feet)
Height=18
Width=83 (feet)
Width to Height Ratio=0.250000
Angle=0.000000

097 footpath/trail - Object number 8
Center at x=996 y=947
Area=11 (feet)
Height=4
Width=121 (feet)
Width to Height Ratio=0.000000
Angle=0.000000

097 footpath/trail - Object number 9
Center at x=1005 y=951
Area=6 (feet)
Height=0
Width=166 (feet)
Width to Height Ratio=0.000000
Angle=0.000000

097 footpath/trail - Object number 10
Center at x=1014 y=953
Area=9 (feet)
Height=4
Width=128 (feet)
Width to Height Ratio=0.000000
Angle=58.250000

097 footpath/trail - Object number 11
Center at x=1088 y=967
Area=49 (feet)

Height=9
Width=166 (feet)
Width to Height Ratio=1.000000
Angle=-75.126465

097 footpath/trail - Object number 12
Center at x=913 y=962
Area=9 (feet)
Height=0
Width=153 (feet)
Width to Height Ratio=1.000000
Angle=58.280190

097 footpath/trail - Object number 13
Center at x=1082 y=969
Area=9 (feet)
Height=4
Width=209 (feet)
Width to Height Ratio=0.000000
Angle=0.000000

097 footpath/trail - Object number 14
Center at x=1109 y=978
Area=58 (feet)
Height=13
Width=162 (feet) Width to Height Ratio=0.333333
Angle=7.338737

097 footpath/trail - Object number 15
Center at x=1124 y=980
Area=63 (feet)
Height=13
Width=9 (feet)
Width to Height Ratio=0.333333
Angle=-12.780914

097 footpath/trail - Object number 16
Center at x=1093 y=983
Area=9 (feet)
Height=0
Width=213 (feet)
Width to Height Ratio=0.333333
Angle=58.280190

097 footpath/trail - Object number 17

Center at x=911 y=985
Area=6 (feet)
Height=0
Width=22 (feet)
Width to Height Ratio=0.333333
Angle=0.000000

097 footpath/trail - Object number 18
Center at x=917 y=987
Area=6 (feet)
Height=0
Width=198 (feet)
Width to Height Ratio=0.333333
Angle=0.000000

097 footpath/trail - Object number 19
Center at x=929 y=987
Area=11 (feet)
Height=4
Width=40 (feet)
Width to Height Ratio=0.000000
Angle=-67.498344

097 footpath/trail - Object number 20
Center at x=922 y=1007
Area=69 (feet)
Height=31
Width=204 (feet)
Width to Height Ratio=0.142857
Angle=-11.926117

097 footpath/trail - Object number 21
Center at x=942 y=996
Area=6 (feet)
Height=0
Width=0 (feet)
Width to Height Ratio=0.142857
Angle=0.000000

097 footpath/trail - Object number 22
Center at x=1122 y=998
Area=22 (feet)
Height=4
Width=0 (feet)

Width to Height Ratio=0.000000
Angle=84.847198

097 footpath/trail - Object number 23
Center at x=933 y=1014
Area=72 (feet)
Height=36
Width=0 (feet)
Width to Height Ratio=0.000000
Angle=-2.196289

097 footpath/trail - Object number 24
Center at x=951 y=998
Area=5 (feet)
Height=4
Width=0 (feet)
Width to Height Ratio=0.000000
Angle=0.000000

097 footpath/trail - Object number 25
Center at x=962 y=1023
Area=39 (feet)
Height=-2147483639
Width=2 (feet)
Width to Height Ratio=0.500000
Angle=0.000000

097 footpath/trail - Object number 26
Center at x=976 y=1005
Area=5 (feet)
Height=-2147483644
Width=0 (feet)
Width to Height Ratio=0.000000
Angle=0.000000

097 footpath/trail - Object number 27
Center at x=915 y=1007
Area=3 (feet)
Height=-2147483644
Width=-2147352570 (feet)
Width to Height Ratio=0.
Angle=0.

097 footpath/trail - Object number 28
Center at x=994 y=1032
Area=94 (feet)
Height=-2147483626
Width=131085 (feet)
Width to Height Ratio=0.600000
Angle=0.

097 footpath/trail - Object number 29
Center at x=1077 y=1039
Area=5 (feet)
Height=4
Width=131072 (feet)
Width to Height Ratio=0.
Angle=0.

097 footpath/trail - Object number 30
Center at x=1077 y=1064
Area=29 (feet)
Height=-2147483639
Width=131072 (feet)
Width to Height Ratio=0.
Angle=-1.168732

097 footpath/trail - Object number 31
Center at x=1120 y=1073
Area=51 (feet)
Height=4
Width=131072 (feet)
Width to Height Ratio=0.
Angle=-62.641972

097 footpath/trail - Object number 32
Center at x=920 y=1124
Area=26 (feet)
Height=-2147483630
Width=-2147352576 (feet)
Width to Height Ratio=0.
Angle=23.219749

097 footpath/trail - Object number 33
Center at x=933 y=1109
Area=3 (feet)
Height=-2147483644
Width=131076 (feet)
Width to Height Ratio=1.
Angle=0.

097 footpath/trail - Object number 34
Center at x=951 y=1115
Area=3 (feet)
Height=4
Width=131072 (feet)
Width to Height Ratio=0.
Angle=90.

Vita

Timothy Dwayne Phillips was born on 19 November 1958 in Amite, Louisiana. The second of three sons of a Church of Christ Minister, he moved about the country living in Galatin, Tennessee, Dallas, Texas, Southern California, and Loranger, Louisiana. He graduated from Loranger High School Loranger, Louisiana in June 1976.

He first attended Southeastern Louisiana University in Hammond, Louisiana. He transferred to Louisiana State University in the fall of 1978. He graduated from L.S.U. with a B.S. in Electrical Engineering in May 1980.

He went on to work for the U.S. Department of Defense as a civilian employee. He travelled extensively throughout the U.S. and overseas.

He returned to L.S.U. in the fall of 1983 and graduated with the M.S. degree in Electrical Engineering in December 1984. He finished all class work towards a PhD and the general exam in May 1986. He returned to work for the Defense Department shortly thereafter. He continued to research and work on his doctoral degree until he earned it in August 1990.

DOCTORAL EXAMINATION AND DISSERTATION REPORT

Candidate: Timothy Dwayne Phillips

Major Field: Electrical Engineering

Title of Dissertation: A New Approach to Analyzing High Resolution Aerial Photographs of Urban Areas

Approved:

Charles Harlow

Major Professor and Chairman

F. Allen Hemby

Dean of the Graduate School

EXAMINING COMMITTEE:

Donald H. [illegible]

Subhash Kak

Ahmed Elamawy

[illegible]

[illegible]

Date of Examination:

20 March 1990