

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comment regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE March 1996	3. REPORT TYPE AND DATES COVERED Final, 23 January 1995-22 January 1996		
4. TITLE AND SUBTITLE High Speed Computation for Visualization			5. FUNDING NUMBERS DAAH04-95-1-0068	
6. AUTHOR(S) Allen R. Hanson				
7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(ES) University of Massachusetts Computer Science Department Box 34610 - Lederle GRC Amherst, MA 01003-4610			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSORING / MONITORING AGENCY REPORT NUMBER ARO 34351.1 -RT-RIP	
11. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12 b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This document is the final report on contract number DAAH04-95-1-0068 awarded to the Computer Vision Laboratory at the University of Massachusetts for the purchase of a high performance visualization workstation. It describes the equipment purchased under this contract and contains brief status summaries of two DOD-sponsored research programs making the heaviest use of the equipment. In the first of these, the DEADALUS Battlefield Visualization System is being designed to produce real-time 3D graphical visualizations of the evolving battlefield situation on realistic terrain models derived from aerial imagery. In the second, a system called the Schema Learning System is being developed to automatically assemble task-specific object recognition programs from existing image understanding algorithms through the use of modern machine learning techniques.				
14. SUBJECT TERMS			15. NUMBER OF PAGES	
19960524 008			16. PRICE CODE	
17. SECURITY CLASSIFICATION OR REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

High Speed Computation for Visualization

Final Progress Report

Allen R. Hanson

May 7, 1996

U.S. Army Research Office

DURIP Grant Number DAAH04-95-1-0068

University of Massachusetts, Amherst

Approved for Public Release;
Distribution Unlimited

TABLE OF CONTENTS

1. Introduction	4
2. The Daedalus Battlefield Visualization System.....	4
2.1 DAEDALUS - Battlefield Awareness: Rationale and Concept	4
2.2. Approach	7
3. Learning Object Recognition Control Strategies	9
3.1 Introduction	9
3.2 Application Domain : The RADIUS aerial image interpretation project.....	10
3.3 Evaluation: Information Fusion to Recognize & Model Complex 3D Buildings	11
3.3.1 Data Representations.....	11
3.3.2 The Visual Procedure Library	12
3.3.3 Preliminary Results: Detecting Rooftops.....	13
3.3.4 Future Experimental Plans: 3D Building Reconstruction	14
4. Paper References	14

LIST OF TABLES

Table 1. Features defined for each level of representation
in the visual procedure library for recognizing rooftops.....12

1. Introduction

In (date), the Computer Vision Laboratory at the University of Massachusetts was awarded \$75,000 (out of a requested \$150,000) to support acquisition of a scientific visualization workstation to support ongoing federally funded research programs in the laboratory. Under this grant, and with matching money from the University, the laboratory purchased a Silicon Graphics ONYX Extreme workstation with one R4400 processor, 64 MB of memory, two 2 GB SCSI-2 fast and wide differential internal disks, and a CD drive, and a 21 inch color monitor.

This system became fully functional in May, 1995. Since its installation, this workstation has been heavily used to support two major research thrusts:

1. The Daedalus Battlefield Visualization System, supported under ARPA contract number DAAL02-91-K-0047; this effort currently supports two graduate students, one undergraduate research assistant, and 40% of a post-doctoral researcher.

and

2. Learning Object and Scene Recognition Strategies, supported under ARPA contract number F30602-94-C-0042; this effort currently supports one graduate student, 40% of a post-doctoral researcher and 60% of a professional systems programmer.

The technical objective and research approaches for each of these programs are briefly discussed in the next two sections of this report.

2. The Daedalus Battlefield Visualization System

2.1 DAEDALUS - Battlefield Awareness: Rationale and Concept

The motivating goal of the Daedalus project is battlefield awareness. The problem of providing timely situational awareness for air and ground combat operations is a recurring theme in modern warfare that impacts both force effectiveness and the need to reduce fratricide. Because of communication bandwidth constraints and the dangers to scouts, continuous live views of the battlefield are impractical. As a result, our forces often use stale information for the planning and conduct of their combat missions.

There are currently DOD programs underway that will provide large amounts of information that must be properly digested, fused, and presented in a real-time manner for consumption by the battlefield commander. These include the Unmanned Air Vehicle (UAV) program in which the Tier2-Plus and Tier3-Minus system will produce approximately a terabyte of data per day per air vehicle of ground images. Terrain will be able to be reconstructed at high resolution of any area of interest. Similarly, the Unmanned Ground Vehicle Demo II program is developing the ability to deploy multiple scout vehicles before and during a battle that will asynchronously transmit reconnaissance information, including ATR image chips, back to the command center.

The introduction of these new remote sensors and reconnaissance systems is likely to exacerbate this problem unless novel means are provided for commanders and individual

soldiers to track and instantly comprehend dynamic battlefield situations. This implies that the relevant information must be presented in a coherent, digestible, and friendly form.

The DAEDALUS system is being designed to produce real-time 3D graphical visualizations of the evolving battlefield situation. The system we are developing has two components. The first rapidly builds an extended terrain model from a variety of possible inputs, and stores the model in a real-time visualization data base. The second retrieves the extended terrain model along with any additional data such as real time ATR and friend-or-foe information (e.g. via UGV scout vehicles) and displays the information in a visual environment that enables the user to fly/drive/walk through the battlefield. For example, ground-support pilots can view the battlefield from the perspective of close support aircraft while, at the same time, Army commanders can view the battlefield from the perspective of a foot soldier. Such a system will improve training and tactics, and would certainly decrease the likelihood of friendly fire accidents by enabling command personnel from different units and services to view the dynamic battlefield from their different perspectives.

It is important to note that we use the term "extended" terrain model to mean those terrain databases that contain:

- highly accurate ground elevation maps,
- 3D models of buildings and other cultural features,
- automatically classified terrain and ground cover types,
- changes automatically detected over time, and
- interface links to civilian and military databases (such as DTED Level II and ITD, Landsat, and SPOT).

All of these components are necessary to construct realistic fly/drive/walkthrough visualizations (more about this below). The link to other databases allows the use of information such as road networks, waterways and lakes, ground cover, etc. that these databases provide. This information provides very powerful contextual cues and feature sets for the visual routines that perform functional classification and which actively construct the model of a particular area.

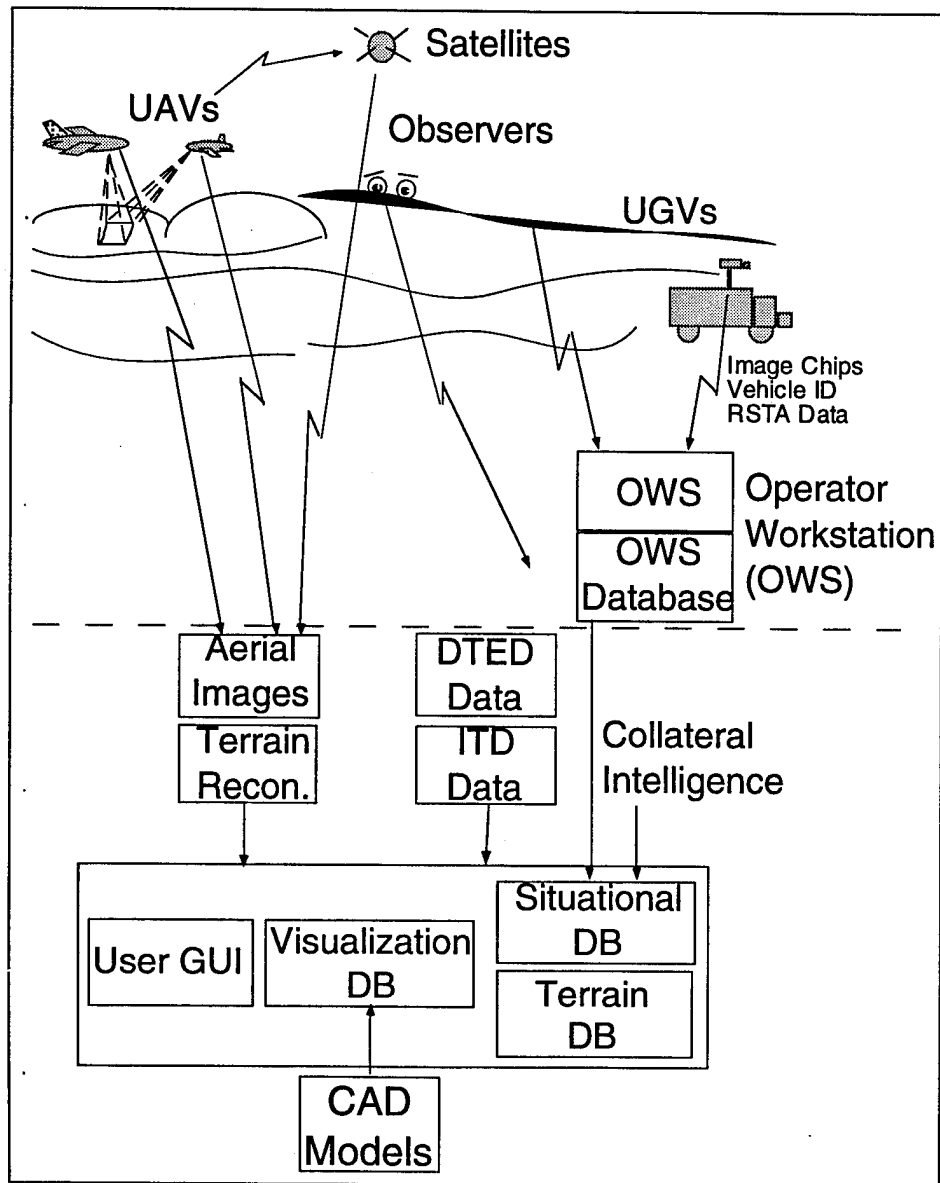


Figure 1. An Overview of DAEDALUS. Information from many sources is fused into a dynamic graphical representation of the battlefield situation.

The visualizations are based on two critical sources of information:

- (1) very high resolution extended terrain maps produced from UAV fly-over images and other sources as discussed earlier, and
- (2) live situational information from sensors located on unmanned ground and air vehicles, human observer data, and other sources of intelligence.

By fusing all the available information, including both remote sensor and human input, the state of the battlefield can be represented in a dramatic, easy to interpret, 3D representation. An overview of the DAEDALUS concept is shown in Figure 1; in this figure, elements above the dotted line represent the battlefield, while items below the line correspond to the components of the DAEDALUS system proper.

Until the most recent generation of machines, it was necessary to specify the fly-through trajectory in advance; rendering the flythrough could take hours or days. With new computational engines such as the SGI Onyx processor, this type of visualization can be produced for any viewpoint or trajectory in near real-time (i.e. in seconds). New visualization tools based on this technology are becoming part of the military planning process.

Real-time rendering of terrain views provides battlefield commanders a powerful new tool for the planning of operations, as well as an extremely effective interface for briefing the plan to those who will implement it. However, while the advantages of these planning tools are just beginning to be exploited, the staleness of the data they rely on inhibit their use for real-time monitoring and assessment of combat operations. The need for a real-time battlefield representation given remote sensors and the previously mentioned communication bandwidth constraints provide the basic design constraints.

2.2. Approach

The approach integrates sensor data from multiple remote sensors into a single, consistent 3D model of the battlefield. As combat vehicles come into the range of deployed sensors, their location, direction and velocity, and IFF status are overlaid onto a high-resolution terrain model of the battlefield. This allows the commander to control the visualization for flying around the battlefield and viewing the situation from any perspective, without being restricted by the original sensor positions. Because the information about vehicle movements is transformed onto the rendered terrain database in an integrated and consistent manner - despite asynchronous information updates from multiple UGV scout and other ground sensors - and thus confusion regarding the perspective of individual views is mitigated.

The visualization program being employed is VGIS (Virtual Geographical Information System) being jointly developed by ARL and Georgia Tech. This system is designed to provide real-time visualizations on state-of-the-art SGI graphics engines. The system is based on a tessellated representation of the terrain map, allow rapid rendering from a specific point of view. Fly/drive-through visualization trajectories are also supported by VGIS.

The terrain map is produced by a terrain reconstruction system being developed at UMass that is designed to produce accurate elevation maps from image pairs taken from oblique viewing angles, from widely separated viewpoints, and from sensors with large baseline-to-height (i.e. distance from sensor baseline midpoint to a nominal point on the terrain) ratios. The UMass system has been shown to produce accurate terrain maps under these condition, even though they are known to present unique problems for traditional stereo systems. The system has been used to produce detailed and highly accurate terrain maps of the ARPA UGV sites at Martin-Marietta and at Fort Hood and these terrain maps have been interfaced to the VGIS system. The CAD models of military vehicles used in the initial simulations were acquired from ARL and placed on the terrain; both the CAD models and the terrain were then rendered in real-time in VGIS and a fly-through scenario was developed. Figure 2 shows a typical Daedalus screen containing vehicles placed on reconstructed terrain (note that the original image is in color).

The initial version of DAEDALUS has a number of limitations which are currently being addressed; these are described very briefly below.

1. *Improving visual realism*

Once the terrain elevation map is produced, one of the aerial images is 'draped' over it (i.e. the elevation map is texture mapped with the image) to produce the final version used in the fly-through. While in many cases such a visualization may be satisfactory, there are some

situations (i.e. viewpoints) where it significantly lacks realism or for which there is a complete absence of required information.

Consider the case where a battlefield is monitored by a downward looking sensor on a UAV and visualized from the perspective of a foot soldier or vehicle. Because the data are taken and visualized from widely different viewpoints, the rendered terrain may not look realistic.




Figure 2. A RSTA query for one of the vehicle models in the simulation brings up the most surveillance imagery chips associated with that vehicle. Note that in this demonstration system, there is a mismatch between the actual vehicle and the RSTA chip displayed in the lower left corner. No CAD models of APC vehicles were available at the time of the simulation.

For example if the battlefield contains trees, the UAV would see a canopy, whereas a soldier would see tree trunks and branches. By classifying the terrain before it is placed in the terrain data base, it is possible to replace complex 3D objects like forests with a 3D graphical model which can be visualized from any viewpoint. The general solution to these visualization problems lies in the proper application of existing image understanding techniques.

The mechanisms for improving visual realism have a number of interesting overlapping subproblems which are being explored:

- (a) Context-based Cultural Object Recognition and Modeling.
- (b) Recognition of Natural Objects.
- (c) Modeling Natural Objects.

2. Use of Collateral Databases.

There are several existing databases which could be effectively used in this effort in a number of ways. For example, information from DTED/ITD could be used to cue recognition processes for both cultural and natural features and objects. Pre-existing low resolution digital elevation maps can be used as a 'seed' map for construction of a higher resolution DEM.

3. Interface to UGV/UAV/Surveillance Sensors

Each scout vehicle is equipped with visible light, FLIR, and possibly spot range sensors. In the future, RSTA (Reconnaissance, Surveillance, and Target Acquisition) systems currently being developed under ARPA programs will be able to detect, identify and track combat vehicles. The output of these systems is a highly compact representation of battlefield activities which can be easily transmitted over the reduced bandwidth communication channels and added to the visualization component of Daedalus.

3. Learning Object Recognition Control Strategies

3.1 Introduction

Although the field of image understanding has made significant advances over the past 20 years, we have not yet developed a theoretical or practical understanding of how these new image understanding algorithms can be combined into functioning systems. As a result, although the library of image understanding procedures, representations and theorems keeps growing, there are very few applications of IU technology in the real world. Fortunately, advances are now taking place that apply machine learning technology to the integration and control of IU algorithms; we believe these techniques will allow the automatic creation of customized, robust systems for special-purpose applications.

For the past three years, the Computer Vision Laboratory at the University of Massachusetts has been developing the Schema Learning System (SLS), a system that automatically assembles task-specific object recognition programs from existing IU algorithms. Developed under an ARPA-sponsored contract on learning in machine vision (administered by Rome Labs), SLS brings together two emerging technologies -- image understanding and machine learning -- to form a computer vision system that automatically learns object and scene recognition strategies.

In many ways, the stage has been set for the Schema Learning System by the research of the past twenty years. Computer vision researchers have been dividing naturally - without any global consensus or mandate - into 10 or 20 subfields with smaller, better defined problems. This has led to clear focus and the development of theories and techniques for the different subdisciplines. There are now several good and improving algorithms for edge and line extraction (straight and curved), stereo analysis, tracking, depth from motion (two-frame and multi-frame), shape recovery, 2D and 3D object recognition, 3D pose determination, and knowledge-based focus of attention, and this list can easily be extended. Indeed, computer vision researchers have made more progress

than most outside the field (and many inside) are aware of. This state of affairs is due primarily to our inability to easily produce highly visible results in the form of integrated task-specific systems. It is as though we have the supplies and tools to build a house, but lack the architectural drawings.

Much of what makes the integration problem difficult is that the most effective combinations of algorithms are object or context dependent. Some objects, for example, have distinct colors that can be used to focus attention on particular parts of an image, while others have easily identifiable substructures, repetitive textures, or other properties that help us to recognize them and place them in space. Unfortunately, the features and techniques needed vary from object to object and context to context, so that many visual tasks require specialized solutions, even within constrained domains such as ATR, aerial image interpretation, or visual inspection. This is a great limitation to the wide and flexible use of computer vision technology, because successful vision systems must be redesigned and/or hand-tuned for each new application. To the extent that it is successful, SLS holds the promise of overcoming this limitation and allowing the automatic customization of control strategies for reliable and autonomous performance.

The need for systems that adapt to the user and context without requiring extensive explicit programming and customization is especially high in some areas of vision such as ATR. Vision systems that learn and adapt are one of the most important directions in IU research right now. This reflects an overall trend -- to make intelligent systems that do not need to be fully and painfully programmed. It is the only way for us to develop vision systems for the military that are robust and easy to use in many different tasks.

3.2 Application Domain : The RADIUS aerial image interpretation project

Abstract theories of learning can only be tested in the context of an application domain; to the extent that the application domain is typical of critical, real-world problems, the evaluation of the learning system is that much more meaningful. The Schema Learning System (SLS) will be tested on problems and data from the ARPA-funded (via the Army Topographic Engineering Center (TEC)) RADIUS aerial image interpretation project.

The goal of the RADIUS project is to supply model supported exploitation (MSE) tools to aid analysts in interpreting aerial and satellite images. Typical tasks that might be automated include (1) recognizing functional areas such as parking lots or storage depots, (2) constructing 3D models of buildings and other objects of interest, and (3) detecting significant changes between new images and previously constructed models. The source data are generally overlapping visible-light images from approximately known viewpoints, although the interpretation of synthetic aperture radar (SAR) images and interferometric synthetic aperture radar (IFSAR) 3D images is of increasing interest.

The RADIUS "tools" are meant to perform common recognition and modeling tasks for the image analyst, freeing him/her to interpret particularly complex parts of a scene. In other words, these tools are recognition programs that are customized to recognize a particular object class (such as buildings or bridges) within a given context. Currently, these customized procedures are manually constructed by university researchers, based on their experience with a limited set of (unclassified) test images. If successful, SLS represents a methodology by which customized recognition procedures could be automatically learned from examples, without the delay or expense of a human researcher. Not only would this be an interesting demonstration of learning in vision, it would prepare the way for image analysts to get more support than a program like

RADIUS can otherwise provide. (It would also allow image analysts to train recognition procedures on classified images that cannot be distributed to university labs.)

3.3 Evaluation: Information Fusion to Recognize & Model Complex 3D Buildings

To evaluate SLS (including various machine learning techniques applied to the control of IU algorithms), we are training it to accomplish tasks from the RADIUS project. The first task is to construct 3D models of buildings from overlapping visible-light images. This task is typical of RADIUS (and other) interpretation tasks in that there are many relevant sources of information, and consequently many applicable IU techniques. Initial 2D data can be extracted in the form of corners or straight lines or edges, and this data can be grouped into closed polygons that might represent buildings. Alternatively, the image can be segmented into regions that may (or may not) support the edge/corner/line data. Finding the right combination of techniques (and compensating for missing data due to noise, shadows or occlusions) is currently the job of the IU researcher.

In addition, 3D information about hypothesized buildings can also be acquired from several sources. By matching corresponding lines in multiple (overlapping) images, their 3D positions and orientations can be determined. Alternatively, correlation-based stereo matching of overlapping images can produce a dense elevation map; by fitting surfaces to the projections of polygons on the elevation data, the position and orientation of roofs and other surfaces can again be determined. Shadows are yet another source of 3D information in visible-images.

The challenge for SLS is to automatically learn a control policy that selects which techniques to use to construct building models and determines how to combine their (sometimes contradictory) results. If successful, these experiments will demonstrate SLS's ability to customize control strategies for complex, real-world tasks.

3.3.1 Data Representations

As has already been discussed, we assigned SLS the task of finding rooftops in aerial images of Ft. Hood, a task that was copied from the RADIUS project task domain. On each trial, the system is given a subimage containing one or sometimes two buildings, and a set of 3D line segments. SLS is also given a visual procedure library that defines eight levels of representation and nine visual procedures. The levels of representation correspond to images, sets of 3D line segments, parallel line pairs, corners (i.e. vertices of orthogonal lines), line groups and polygons. Because much of the power of SLS lies in its ability to distinguish good data from bad based on feature measurements, Table 1 gives the set of features defined for each level of representation.

Representation | Features

Image	Avg. Intensity, SD Edginess, Speckle
LinePairSet	Count, Avg. Contrast
Parallel Line Pair	angle, overlap, shadowness surface fit, distance
Corner (L-Junction)	angle, gap, shadowness surface, fit, scale
Line Group	Count, Parallel Count Corner Count
Polygon	Edge Cover % Worst Edge Cover Avg. Perimeter Contrast Worst Side Contrast Plane fit error (intensity data) scale, shadowness

Table 1. Features defined for each level of representation in the visual procedure library for recognizing rooftops.

3.3.2 The Visual Procedure Library

The visual procedures employed are meant to approximate some of the techniques being used by researchers in the RADIUS project. The 3D line segments were computed off-line and filtered according to the known height of the ground plane. Eight other visual procedures are available. The rectilinear line grouping (RLGS) procedure computes relationships between nearby line segments and uses information about the camera pose (available for all RADIUS images) to remove the effects of perspective distortion from orthogonal and parallel relations. The SelectParallel and SelectCorner procedures select parallel line pairs and corners out of the relations computed by RLGS.

The grouping procedures (GroupParallel and GroupLJnct) take a pair of parallel lines (or a corner) and form a group out of all the lines that share a significant relation to one of the lines in the original pair. This results in a small group of line segments in which the Graph Matching procedure can search for a rectangle of lines. Alternatively, given a pair of parallel or orthogonal line segments, the Par2Polygon and Corner2Polygon algorithms go back to the original image data and apply edge extraction and edge linking operators top down, in order to look for evidence of additional lines that might complete the rectangle. Finally, the Polygon2Roof procedure serves no purpose other than to give SLS a way to designate a particular polygon as a roof.

At first glance, the visual procedure library would appear to have only four paths to the goal, which would make SLS's task fairly easy. The procedure for selecting corners, however, typically finds on the order of fifty to one hundred corners per image, while the procedure for finding parallel line pairs typically finds twice that many. As a result, there are approximately five hundred polygons that SLS might detect in most images, and most of the work that SLS does is in selecting which hypotheses -- in terms of which corners, parallel pairs, and polygons -- to pursue.

3.3.3 Preliminary Results: Detecting Rooftops

SLS was tested on a set of ten image fragments taken from the RADIUS project's images of Ft. Hood. The training and evaluation was carried out using the ground truth data developed for the (self) evaluation of UMass' hand-crafted building detector and 3D reconstruction system.

SLS was tested using a "leave one out" methodology. On each trial, SLS was trained on data from nine images, and then the control policy it learned would be applied to the tenth image. This was repeated ten times, each time holding a different image out of the training set for testing. Figure 4 shows one of the rooftops extracted by SLS. Over ten trials, SLS found a polygon that corresponded to a true roof surface nine times; in the tenth trial, SLS confused part of the roof boundary with shadow line that was near to (and parallel with) the true edge of the roof, creating an incorrect roof hypothesis.

The control policies learned by SLS did not always take a straight path to the solution. Although they always preferred finding corners to parallel lines, they would often select one corner as being the most promising, use it to generate a polygon hypothesis, and then decide that the polygon was not as good as it thought it would be. The system would then backtrack, find the next most promising corner, and try again. In general, the system created ten to twenty polygon hypotheses (out of several hundred possible ones) before finding the polygon it finally selects to be the rooftop.




Figure 4. One of the ten aerial images of buildings at Ft. Hood, and the roof hypothesis (shown in white) SLS learned to find in it.

Significantly, SLS can adapt quickly to the introduction of new procedures or features. The first time we tested the system on the Ft. Hood data, it succeeded in only about half the trials. Looking at its mistakes, we realized it was often confusing shadows with the edges of the buildings that cast them. We then added a shadowness feature to the parallel pair, corner and polygon representations, and immediately SLS's performance improved. In general, we suspect that this is how SLS will be used -- as a system to which people add knowledge until it is able to accomplish the assigned task. Ironically, it could therefore be viewed as a very good knowledge engineering tool.

3.3.4 Future Experimental Plans: 3D Building Reconstruction

The goals of the RADIUS project go beyond simply recognizing objects in aerial images and determining their image position. One of the goals of RADIUS is to provide the image analyst with tools that automatically construct 3D models of buildings from overlapping aerial views. Although more thorough testing of SLS on the 2D recognition task is also planned, the primary emphasis in the future will be on getting SLS to learn control policies for 3D building reconstruction.

Although there are clues to 3D structure in monocular images that SLS is not currently taking advantage of (such as the sun angle and length of shadows), we believe that what will make 3D building reconstruction far more effective is the depth information provided by overlapping aerial views. The UMass terrain reconstruction constructs dense digital elevation maps (DEMs) accurate to within a meter from pairs of images, even when those images were taken with wide baselines at highly oblique angles. This type of dense, 3D data, in combination with the 3D lines computed in the RADIUS system, should make it possible to reconstruct highly accurate building models. These procedures, along with additional procedures for fitting planes and complex surfaces to DEM data, should give SLS many alternative strategies for constructing 3D building models.

SLS's task will be to combine the new 3D procedures with the previous 2D set to form accurate and efficient control policies. The SLS reconstruction experiment are continuing; recent results may be found at URL <http://vis-www.cs.umass.edu/projects/SLS3D.html>.

4. Paper References

Draper, B., Hanson, A., and Riseman, E., "Knowledge-Directed Vision: Control, Learning, and Integration", Special Issue of Proc. IEEE on Signals and Symbols, to appear, 1996.

Draper, B., "Learning Grouping Strategies for 2D and 3D Object Recognition", Proc. ARPA Image Understanding Workshop, Vol. II, pp. 1447-1454, Palm Springs, CA, February 1996.

Hanson, A., and Riseman, E., "Progress in Computer Vision at the University of Massachusetts", Proc. ARPA Image Understanding Workshop, Vol. I, pp. 5-14, Palm Springs, CA, February 1996.

Riseman, E., Schultz, H., W. Jonathan Lim, Draper, B., and Hanson, A., "Daedalus Battlefield Visualization System", Proc. ARPA Image Understanding Workshop, Vol. I, pp. 491- 500, Palm Springs, CA, February 1996.

REPORT OF INVENTIONS AND SUBCONTRACTS

(Pursuant to "Patent Rights" Contract Clause) (See Instructions on Reverse Side.)

Form Approved
OMB No. 0704-0297
Expires Jun 30, 1992*

Public reporting burden for this collection of information is estimated to average 5 minutes per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0297), Washington, DC 20503.

1. NAME OF CONTRACTOR / SUBCONTRACTOR University of Massachusetts Grant and Contract Admin. Amherst, MA 01003		2. NAME OF GOVERNMENT PRIME CONTRACTOR same as 1a		3. TYPE OF REPORT (X one) a. INTERIM <input checked="" type="checkbox"/> b. FINAL	
c. CONTRACT NUMBER DAAH04-95-1-0068		d. AWARD DATE (YYMMDD) 95 01 23		4. REPORTING PERIOD (YYMMDD) a. FROM 95 01 23 b. TO 96 01 22	

SECTION I - SUBJECT INVENTIONS

5. "SUBJECT INVENTIONS" REQUIRED TO BE REPORTED BY CONTRACTOR / SUBCONTRACTOR (If "None," so state) NONE

a. NAME(S) OF INVENTOR(S) (Last, First, MI)	b. TITLE OF INVENTION(S)	c. DISCLOSURE NO., PATENT APPLICATION SERIAL NO. OR PATENT NO.	d. ELECTION TO FILE PATENT APPLICATIONS		e. CONFIRMATORY INSTRUMENT OR ASSIGNMENT FORWARDED TO CONTRACTING OFFICER
			(1) United States (a) Yes (b) No	(2) Foreign (a) Yes (b) No	

6. EMPLOYER OF INVENTOR(S) NOT EMPLOYED BY CONTRACTOR / SUBCONTRACTOR		g. ELECTED FOREIGN COUNTRIES IN WHICH A PATENT APPLICATION WILL BE FILED (2) Foreign Countries of Patent Application	
(1) (a) Name of Inventor (Last, First, MI)	(2) (a) Name of Inventor (Last, First, MI)	(1) Title of Invention	
(b) Name of Employer	(b) Name of Employer		
(c) Address of Employer (Include ZIP Code)	(c) Address of Employer (Include ZIP Code)		

SECTION II - SUBCONTRACTS (Containing a "Patent Rights" clause)

6. SUBCONTRACTS AWARDED BY CONTRACTOR / SUBCONTRACTOR (If "None," so state) NONE

a. NAME OF SUBCONTRACTOR(S)	b. ADDRESS (Include ZIP Code)	c. SUBCONTRACT NO.(S)	d. DEAR "PATENT RIGHTS":		e. DESCRIPTION OF WORK TO BE PERFORMED UNDER SUBCONTRACT(S)	f. SUBCONTRACT DATES (YYMMDD)	
			(1) Clause Number	(2) Date (YYMM)		(1) Award	(2) Estimated Completion

SECTION III - CERTIFICATION

7. CERTIFICATION OF REPORT BY CONTRACTOR / SUBCONTRACTOR (Not required if Small Business or Non-Profit organization.) (X appropriate box)

a. NAME OF AUTHORIZED CONTRACTOR / SUBCONTRACTOR OFFICIAL (Last, First, MI) Allen R. Hanson		c. I certify that the reporting party has procedures for prompt identification and timely disclosure of "Subject Inventions," that such procedures have been followed and that all "Subject Inventions" have been reported.	
b. TITLE Principal Investigator	Carol Sprague Acting Associate Director, OGCA	d. SIGNATURE OF PI <i>Allen R. Hanson</i>	e. DATE SIGNED 4/9/96
		Contracting Official <i>Carol Sprague</i>	