

AD-A033 348

HUGHES RESEARCH LABS MALIBU CALIF
UNSTRUCTURED CONTROL AND COMMUNICATION PROCESSES IN REAL WORLD --ETC(U)
OCT 76 B L BULLOCK F44620-74-C-0054

UNCLASSIFIED

AFOSR-TR-76-1229


ML

| OF |
AD
A033348



END

DATE
FILMED
2-77



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

3

ADA033348

UNSTRUCTURED CONTROL AND COMMUNICATION PROCESSES IN REAL WORLD SCENE ANALYSIS

Bruce L. Bullock

Hughes Research Laboratories
3011 Malibu Canyon Road
Malibu, CA 90265

October 1976

Contract F44620-74-C-0054

Interim Scientific Report

For Period 1 April 1975 Through 31 March 1976

Prepared For

Lt. Colonel George McKemmie

Directorate of Mathematical and Information Sciences

Air Force Office of Scientific Research (AFSC)

Bolling Air Force Base

Washington, DC 20332

DDC
RECEIVED
DEC 13 1976
A

Approved for public release;
distribution unlimited.

AIR FORCE OFFICE OF SCIENTIFIC RESEARCH (AFSC)
NOTICE OF TRANSMITTAL TO DDC
This technical report has been reviewed and is
approved for public release IAW AFR 190-12 (7b).
Distribution is unlimited.

A. D. BLOSE
Technical Information Officer

TABLE OF CONTENTS

Section	Page
I. INTRODUCTION AND ABSTRACT	7
A. Introduction	7
B. Research Abstract	7
II. SUMMARY OF RESEARCH PROGRESS TO DATE AND PLANNED RESEARCH	9
A. Implementation of Functionally Complete System . .	9
B. Development of the "TINY EYE" computer Program	18
C. Computer Program Specification by Pictures	20
III. PUBLICATIONS, REPORTS, AND CONFERENCES . . .	23
IV. SCHEDULED ACTIVITIES	27
V. PREVIOUS PUBLICATIONS	31
VI. TECHNOLOGY TRANSFER	33
APPENDIX: WORKSHOP POSITION PAPER	35

[illegible]

approved for public release;
distribution unlimited.

D D C

DEC 18 1976

A

LIST OF ILLUSTRATIONS

Figure		Page
1	Model matching system organization	9
2	Edge feature extraction process	10
3	Texture prominence detection process	11
4	Noisy minimal spanning tree matching	11
5	Edge-line-vertex feature extraction process	12
6	Road line finding	13
7	Complex scene vertices	14
8	Model matching: Linear transformation	15
9	Model matching: Nonlinear transformation	16
10	Texture prominence processor results	17
11		19
12	Box program	22

SECTION I

INTRODUCTION AND ABSTRACT

A. Introduction

The Exploratory Studies Department of Hughes Research Laboratories (HRL) has since 1970 conducted an extensive research program in scene analysis. An effort to develop a theoretical basis for scene analysis has been supported by the Air Force Office of Scientific Research (AFOSR) since 1973. The long-term goal of this program is to develop technology that can derive useful information from complex real-world scenes. The emphasis has been on real-world imagery and the development of complete scene-analysis systems. Past efforts have concentrated on artificial or greatly simplified imagery, and have usually developed piecemeal algorithms that do not necessarily contribute to the solution of real problems.

The HRL program is unique in that it is one of the few attempts to deal squarely with the problems of real-scene systems. The importance and timeliness of the program is underscored by the number of military image-processing applications upon which it can impact. These include reconnaissance and surveillance tasks (such as target cueing, target screening, and change detection) and tactical image-processing tasks (such as missile and RPV target acquisition, guidance, and terminal homing). Recent results on the pictorial specification of scene analysis programming can also impact problems in software production.

B. Research Abstract

The chief result of the AFOSR program was to implement a complete system for complex outdoor scene analysis, with the objective of developing an understanding and methodology of control and communication. The system embraces all levels of image analysis,

from low-level numeric data to high-level symbolic data. The means of implementation has been a pattern-directed production system. The results described in this report include the performance of the present system, the use of a simplified system to aid the development of a theoretical approach, and initial results on the development of a symbolic method for programming specification.

SECTION II

SUMMARY OF RESEARCH PROGRESS TO DATE AND PLANNED RESEARCH

This section will briefly review the results obtained during this latest contract period.

A. Implementation of Functionally Complete System

A prime task defined for this year was to construct, from analysis modules and control techniques previously developed, a functional scene analysis system. One such system has been completed: a model-based system to find structure in outdoor scenes. It is described in detail in the paper, "Finding Structure in Outdoor Scenes," referenced in Section III.

A block diagram of the basic system is shown in Figure 1. Figure 2 shows a block diagram of the "edge-vertex-line" process to extract features from a real-world scene. The block diagram in Figure 3 shows the process for detecting prominent texture in complex scenes. Finally, a diagram of the model-matching system is shown in Figure 4.




Figure 1. Model matching system organization.




Figure 2. Edge feature extraction process.




Figure 3. Texture prominence detection process.




Figure 4. Noisy minimal spanning tree matching.

Some performance results of the present system are shown in Figures 5 through 10. Figure 5 shows the ability of the edge-line-vertex process to correctly find the vertices in a road scene. Figure 6 shows a road correctly extracted from a particularly noisy, cluttered scene. Figure 7 shows some early results of vertex finding in a complex scene. The entire system is shown, in Figures 8 and 9, being used to match two scenes. One scene is used as a reference model and is






Figure 5. Edge-line-vertex feature extraction process.


4633-13




GENERATED LINES



HOUGH SPACE
(THRESHOLD)



EDGE POINTS



ORIGINAL IMAGE

Figure 6. Road line finding.

matched against the sensor scene to derive information on position change. Figure 10 shows an example of the rotational invariance the process for detecting texture prominence.

14






Figure 8. Model matching: Linear transformation.

During the 1976 AFOSR program this basic system will be exercised in many ways to investigate fundamental questions about production system control for scene analysis. This level of sophistication in a production system is unique in current scene analysis research. The current results encourage us to expect that the development of the production control framework will have a tremendous impact on the important emerging area of outdoor scene analysis. (See Section IV, Pattern-Directed Inference Workshop.) We expect that this effort will result, in addition to the paper mentioned previously, in a series of future publications dealing with system organization and the extraction of features in real-world outdoor scenes.

VERTEX MATCHING				
REFERENCE		TEST MATCH		
NODE	(X, Y)	NODE	(U, V)	DISTANCE
1	(19, 115)	1	(24, 135)	1.41
2	(122, 104)	2	(129, 124)	1.41
3	(101, 67)	3	(108, 88)	2.24
4	(12, 75)	4	(17, 96)	2.24
5	(138, 175)	5	(143, 195)	1.41
6	(187, 168)	6	(193, 188)	0.00
7	(170, 143)	7	(177, 162)	1.41
8	(123, 149)	8	(130, 168)	1.00
9	(156, 120)	9	(162, 140)	1.00
10	(195, 115)	10	(203, 136)	2.24
11	(182, 98)	11	(188, 118)	1.00
12	(144, 102)	12	(150, 122)	1.00



REFERENCE PERSPECTIVE



TEST PERSPECTIVE

16

COORDINATE TRANSFORMATION	
	ACTUAL ESTIMATED
ROTATION ANGLE	1.32° 0.28°
X TRANSLATION	5.43 6.86
Y TRANSLATION	-16.18 -18.85
SCALE	1.02 1.00

AIMPOINT ESTIMATION	
ACTUAL AIMPOINT	155, 158 X
ESTIMATED AIMPOINT	156.5, 159.2 O

RESULTS



REFERENCE MST

TEST MST

Figure 9. Model matching: Nonlinear transformation.





Figure 10. Texture prominence processor results.

B. Development of the TINY EYE Computer Program

The development of the model-matching system has promoted investigation of basic concepts for system organization and control, and feature extraction. We have also gained a better perspective on what is necessary to construct a functional system that encompasses low-level analysis and high-level symbolic computation. This has aided us in developing an abstract theoretical basis for specifying scene-analysis programming intent. In this section a tool that has been useful in developing this system will be described. Some initial results are shown in the next section.

The software program for the previous system was quite large, more than 100K of PDP-10 memory. This complexity was necessary, of course, to achieve the generality to deal successfully with real outdoor scenes. However, the feeling was that the theory might be more easily developed by using an "abstract" version of the complete complex system.

Thus, we were motivated to construct a program that has all the basic components of the previous system, is fully functional on simple scenes, but is small and modular enough to allow convenient experimentation. The resulting program is called "TINY EYE" and is programmed in LISP instead of FORTRAN, for flexibility. It can be written in eight pages of code, and requires only about 20K of core above the LISP interpreter.

The current system consists of the following components: edge detection, line finding and fitting, and vertex finding. An example of a terminal session with this system is shown in Figure 11, representing a very simple artificial image of an L-shaped object. The output in the next several lines shows the detected edge points, the detected lines in terms of slope and y intercept, and the x and y location of the vertices. The current version of this program is written in a strictly vertical, hierarchical organization.

```

*(LINE)

THE INPUT SCENE IS:
0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0.
0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0.
0. 0. 0. 0. 0. 0. 0. 0. 10. 10. 10. 10. 10. 10. 0.
0. 0. 0. 0. 0. 0. 0. 0. 10. 10. 10. 10. 10. 10. 0.
0. 0. 0. 0. 0. 0. 0. 0. 10. 10. 10. 10. 10. 10. 0.
0. 0. 0. 0. 0. 0. 0. 0. 10. 10. 10. 10. 10. 10. 0.
0. 0. 0. 0. 0. 0. 0. 0. 10. 10. 10. 10. 10. 10. 0.
0. 0. 0. 0. 0. 0. 0. 0. 10. 10. 10. 10. 10. 10. 0.
0. 0. 0. 0. 0. 0. 0. 0. 10. 10. 10. 10. 10. 10. 0.
0. 0. 10. 10. 10. 10. 10. 10. 10. 10. 10. 10. 10. 0.
0. 0. 10. 10. 10. 10. 10. 10. 10. 10. 10. 10. 10. 0.
0. 0. 10. 10. 10. 10. 10. 10. 10. 10. 10. 10. 10. 0.
0. 0. 10. 10. 10. 10. 10. 10. 10. 10. 10. 10. 10. 0.
0. 0. 10. 10. 10. 10. 10. 10. 10. 10. 10. 10. 10. 0.
0. 0. 10. 10. 10. 10. 10. 10. 10. 10. 10. 10. 10. 0.
0. 0. 10. 10. 10. 10. 10. 10. 10. 10. 10. 10. 10. 0.
0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0. 0.

```

THE EDGE POINTS ARE:

```

(DEFPROP POINTS
(PPOINTS
(1.5707949 3. 9.)
(1.5707949 4. 9.)
(1.5707949 5. 9.)
(1.5707949 6. 9.)
(1.5707949 7. 9.)
(0.78539747 8. 8.)
(0.0 8. 7.)
(0.0 8. 6.)
(0.0 8. 5.)
(0.0 8. 4.)
(0.0 8. 3.)
(0.78539747 9. 2.)
(1.5707949 9. 1.)
(1.5707949 10. 1.)
(1.5707949 11. 1.)
(1.5707949 12. 1.)
(1.5707949 13. 1.)
(2.3561924 14. 2.)
(3.1415929 15. 2.)
(3.1415929 15. 3.)
(3.1415929 15. 4.)
(3.1415929 15. 5.)
(3.1415929 15. 6.)
(3.1415929 15. 7.)
(3.1415929 15. 8.)
(3.1415929 15. 9.)
(3.1415929 15. 10.)
(3.1415929 15. 11.)
(3.1415929 15. 12.)
(3.1415929 15. 13.)
(3.9269873 14. 14.)
(4.7123848 14. 15.)
(4.7123848 13. 15.)
(4.7123848 12. 15.)
(4.7123848 11. 15.)
(4.7123848 10. 15.)
(4.7123848 9. 15.)
(4.7123848 8. 15.)
(4.7123848 7. 15.)
(4.7123848 6. 15.)
(4.7123848 5. 15.)
(4.7123848 4. 15.)
(4.7123848 3. 15.)
(-0.78539747 2. 14.)
(0.0 1. 14.)
(0.0 1. 13.)
(0.0 1. 12.)
(0.0 1. 11.)
(0.0 1. 10.)
(0.78539747 2. 9.))
VALUE)

```

THE EDGE LINES ARE:

```

(DEFPROP LINES
(LINES
(1. -1. 0.)
(0. 15. 1.)
(1. -15. 0.)
(0. 1. 1.)
(1. -8. 0.)
(0. 9. 1.))
VALUE)

```

THE VERTICES ARE:

```

X      Y
1.     15.
15.    15.
15.    1.
8.     9.
0.     9.
NIL

```

Figure 11.
Terminal output showing
edge points, lines, and
vertices.

During the 1976 research effort, this software system will be used to explore methods of control of the basic production system. This will be easy because the present system has functional modularity. As the present production system interpreter is also in LISP, no language-to-language interface will be necessary. The system will also be applied in developing a formalism to describe scene analysis programs.

Initial work on this formalism is described next, and it is hoped that an implementation can be constructed for it during the 1976 program. This will use a production interpreter for graph rules, which will be interfaced with the TINY EYE program. Results of this research should impact not only the development of scene analysis, but also has the potential to lead to new and efficient means of pictorial program specification and debugging that will impact software engineering and program generation for any large software system.

C. Computer Program Specification by Pictures

There are two key obstacles to major progress in scene analysis and computer vision. One is the lack of knowledge about the necessary basic techniques, and the second is the difficulty in implementing a collection of techniques large enough to deal with complex scene problems. The previous work on this AFOSR program has made progress towards better understanding of the basic techniques, while a portion of the research on production systems has been directed overcoming difficulties in implementation. These problems also relate to the more general areas of computer science that deal with software engineering, program language development, and automatic programming. The initial progress made in this area provided the foundation for Tasks 3 and 4 in the follow-on proposal.

Experience shows that it is difficult to implement and debug complex scene-analysis programs. This limits the amount of knowledge and progress that can be achieved by one person. The problem has to do with the difficulty in perceiving the action and difficulty of

control in most programs. The goal of this programming research was to find a method of program specification that was transparently obvious and hopefully far beyond previous attempts.

The approach was to use the highly pictorial domain of scene analysis, and devise programs for simple tasks. The first task was that of recognizing a block. As crude as the initial results are, they are much easier to understand and are shorter than FORTRAN code for the equivalent task. A continuation of this effort should provide some very interesting results that can potentially affect software construction problems beyond scene analysis.

- Symbols

Several symbols need to be defined for use in the program.

An image subwindow is represented by a square  and scanning the window is shown by . A line is shown schematically in a window . A surface intensity assignment is made $\textcircled{A}$. The intensity onto sides of a line is thus $\boxed{\textcircled{A} / \textcircled{B}}$. An angle assignment is shown $\boxed{/\alpha}$.

- Predicates

Simple predicates on the symbols can also be specified.

If $\textcircled{A} = \textcircled{B} = \textcircled{C} \rightarrow \text{LINE} \longrightarrow$

If $\textcircled{A} = \textcircled{B} \rightarrow \text{EDGE} \longleftarrow\longleftarrow\longleftarrow$

- Program

A simple box program is shown in Figure 12.

A further abstraction of this program form is now being investigated. This uses predicate overlays on a typical member of the input scene class to be analyzed and also uses a finite-state transition diagram to specify the order in which the analysis should proceed.

BEGIN SEE BOX




Figure 12. Box program.

SECTION III

PUBLICATIONS, REPORTS AND CONFERENCES

The following publications have been completed during this reporting period:

1. "Unstructured Control Concepts in Scene Analysis"
2. "Finding Structure in Outdoor Scenes"
3. "Footprints: A Representation for Restructured Motion in Outdoor Scenes"

The abstracts of these publications appear on the next three pages, as well as a listing of where they have appeared.

As a result of these publications, AFOSR sponsored research has appeared at the following conferences:

1. Eighth Annual Southeastern Symposium on System Theory.
2. 1976 Joint Workshop on Pattern Recognition and Artificial Intelligence.
3. EIA Sixth Annual Symposium on Automatic Imagery Pattern Recognition.
4. Third International Joint Conference on Pattern Recognition (future).

"Unstructured Control Concepts in Scene Analysis," Bruce L. Bullock.*

Abstract Traditional scene analysis programs have used relatively simple control and interconnection schemes. As progress

*This paper appeared in: (a) Proceedings of the 8th Annual Southeastern Symposium on System Theory, University of Tennessee, 1976; (b) Hughes Research Laboratories Research Report 497, June 1976 (distributed to AFOSR mailing list).

is made on the analysis of complex outdoor scenes, it is becoming clear that more sophisticated strategies are required. This report describes the evaluation of control strategies that have been utilized for scene analysis and proposes a generalized control structure adequate for outdoor scene analysis. This generalized structure is based on a production system framework. The proposed mechanism allows the traditional vertical top-down and bottom-up organizations, heterarchical organization as well as a new horizontal organization which was introduced to enable collection of multiple interpretations necessary for outdoor scenes.

"Finding Structure in Outdoor Scenes," Bruce L. Bullock, et. al.*

Abstract This report describes the current status of a system for finding complex structures in outdoor scenes and determining their orientation. An extensive design tradeoff is described, as well as representative results on road and object scenes. The actual scene analysis uses models of both line and texture features.

"Footprints: A Representation for Restricted Motion in Outdoor Scenes," Bruce L. Bullock and Sahibsingh A. Dudani.**

Abstract This paper introduces some new concepts for scene analysis with motion, for example in situations where one obtains

* This paper appeared in: (a) IEEE Proc. of the 1976 Joint Workshop on Pattern Recognition and Artificial Intelligence, Hyannis, Mass. June, 1976; (b) Academic Press, book to be published during 1977 of selected papers from above conference; (c) EIA Sixth Annual Symposium of Automatic Imagery Pattern Recognition, June, College Park, Maryland 1976; (d) Hughes Research Laboratories, Research Report 498, July 1976 (distributed to AFOSR mailing list).

** This paper will appear in: (a) Proceedings of the Third International Joint Conference on Pattern Recognition, Coronado, Calif, Nov. 8-11, 1976; (b) soon to be released Hughes Research Laboratories Research Report, 1976, (to be distributed to AFOSR mailing list); (c) Revised version will be submitted to IEEE publication.

continuously changing images at different ranges during a motion. It is assumed that the allowed motion is restricted to changes in range only and that the needed analysis is goal directed. A representation, called a footprint, is used to describe an object of interest at a given range. A scene model is then defined as a composition of object footprints present in the scene. An analysis strategy for outdoor scenes is outlined using the object footprint representation.

SECTION IV

SCHEDULED ACTIVITIES

The following activities are either scheduled or anticipated.

1. Third International Joint Conference on Pattern Recognition, paper will be presented in November 1976.
2. NSF Vision Workshop, 29 October 1976. B. Bullock has been invited to participate in an NSF-sponsored workshop on computer vision. This will occur at the University of Maryland. No written papers will be submitted, only oral presentations. The presentation that has been prepared will be an overview of vision activities at HRL. A major portion of this will cover AFOSR-sponsored research.
3. Rand symposium on Pattern Directed Inference (University of Hawaii, January 1977). B. Bullock has been invited to present a paper at a workshop on pattern directed inference techniques in artificial intelligence (AI). Pattern directed inference is a synonym in (AI) research for production systems.

Production systems are currently an "in vogue" area of research in AI. The Bullock, AFOSR-sponsored paper, tentatively accepted, will be the only paper related to scene analysis. It appears as if the HRL program jumped on the production system bandwagon about two years ahead of any other vision group. It will be recalled that the use of production-system control structures have been the core of the last two years of AFOSR research. The submitted abstract for this workshop appears below.

"Procedure Driven Scene Analysis," B. Bullock.

Abstract This paper describes an investigation to determine the utility of the production system framework when applied to the problem of scene analysis. Central to the investigation is the implementation of a production-driven system to locate and describe simple objects in outdoor scenes. The performance of the system is illustrated on several real scenes. The organization and control mechanisms are also described, emphasizing the issues of two-dimensional productions and the use of relational constraints in rule selection. Examples from the implementation are also used to illustrate tradeoffs between production and procedure at several levels and the use of productions to interface between the numeric-spatial-geometric bottlenecks. Finally, the efficiency and programming ease of the production-driven system are compared to a procedure-based system. The abstracts of this workshop will appear in the ACM SIGART publication, the full papers will appear in a published proceedings.

4. Office of Naval Research-Sponsored Workshop on Computer Vision

B. Bullock has been invited to participate in a year long (Delphi-like) study of computer vision. This has been organized by Prof. Ed Reisman of the University of Massachusetts and includes about 20 researchers in computer vision. A position paper has been submitted. These will be reviewed by each participant. The final result will be a workshop at which papers will appear in a printed proceedings (possible hardbound book). A copy of the submitted workshop position paper appears in Appendix I.

5. International Joint Conference on Artificial Intelligence,
MIT, August 1977.

It is planned to submit a paper to this conference on work performed on the AFOSR contract.

SECTION V
PREVIOUS PUBLICATIONS

The following is a current list of previous publications from the HRL group.

1. "Complex Information Processing - A Guided Bibliography," B. Bullock, Hughes Research Laboratory (HRL) Research Report 444, August 1971.
2. "Pattern Recognition Methodology," B. Bullock, HRL Research Report 450, December 1971.
3. "The HRL LISP Scene-Analysis Language Processor," B. Bullock, HRL Research Report 483, January 1974.
4. "Scene Analysis Research Progress," B. Bullock, presented at E.I.A. Symposium on Automatic Imagery Pattern Recognition, January 1974, Washington D.C. *
5. "Automatic Recognition of 3-D Objects from TV Images," S. Dudani, 2nd Annual Computer Science Conference, February 1974, Detroit, (also HRL Research Report 491).
6. "Distance Weighted K-Nearest Neighbor Decision Rule," S. Dudani, 2nd International Joint Conference on Pattern Recognition, Denmark, August 1974.
7. "An Interest Operator For Automatic Target Selection," Research Report 488, B. Bullock, August 1974.
8. "The Performance of Edge Operators on Images with Texture," B. Bullock, Research Report 489, September 1974, (submitted to Computer Graphics and Image Processing). *

* Publications resulting from AFOSR sponsored research.

9. "Real World Scene Analysis in Perspective," presented at National Convention of the Association for Computing Machinery, October 1975, Minneapolis.*

* Publications resulting from AFOSR sponsored research.

SECTION VI

TECHNOLOGY TRANSFER

HRL is in an excellent position to promote rapid transfer of technology to important application areas. The role of HRL is to perform basic and applied research in many technologies. Close ties between HRL and the many operating divisions of the Hughes Aircraft Company encourage the distribution and application of the HRL work. This type of immediate application of research often requires a period of years when the research and application groups are not parts of a common, closely tied organization.

During the period of this contract, the following technology transfers have been accomplished:

1. The Hughes Missile Systems Group has been using the HRL scene analysis software system for the development of improved correlation guidance techniques for the Hughes Maverick Missile.
2. An unsolicited proposal was accepted by the Naval Surface Weapons Center (NSWC). This proposal incorporated the results on edge detection and the simplified edge-directionality operator developed on the 1974 AFOSR contract for target classification and aimpoint selection relevant to a high-energy laser program. This program was in cooperation with the Hughes Electro-Optical Systems Division. Also important in this proposal was the use of moment methods for shape analysis developed by S. Dudani (now at HRL) in his thesis on the Ohio State AFOSR Grant. Work on this contract has been successfully completed. A proposal for a brass-board system is now pending.

4. An unsolicited proposal to DARPA/STO has been funded to look at terminal homing systems using alternative non-correlation techniques developed on this program. The first year of this program has been successfully completed. A follow-on contract is now pending.

APPENDIX
WORKSHOP POSITION PAPER
(See Section IV.)

POSITION PAPER

PERSPECTIVE TOWARDS OUTDOOR SCENES

Bruce L. Bullock
Hughes Research Laboratories,
Malibu, Cal.

I. INTRODUCTION

The view expressed in this paper is based primarily on recent experience on two projects, the first concerned with the development of general scene analysis techniques, concentrating on issues of system organization, control, and user interaction. The second project is really a collection of investigations to develop very goal-directed systems for a range of military applications. In spite of the apparent differences between these two projects, they both share a common interest in the ability to analyze complex objects in outdoor scenes, irrespective of environmental context and goal dynamics. The position described below will reflect this interest while loosely adhering to the topics and order of the suggested position questions.

II. SEGMENTATION

There are at least seven fundamental methodological "tools" associated with segmentation: 1. edge detection by template differencing, 2. edge detection by functional approximation, 3. region growing, 4. detection of macro and microstructure using clustered features, 5. guidance from defocused images, 6. guidance from glancing, 7. utilization of multisensor data. Each of these has independently grown quite sophisticated for dealing with different subproblems of segmentation, yet segmentation capability as a whole remains in trouble in complex outdoor scenes. A significant contribution to progress in low-level vision can be made at this point, not by trading off one technique against the other, but by carefully examining how the best independent elements of each can be integrated into a single

segmentation process that is better than anything currently available. This has not been achieved in any system, including the existing cluster-based systems, to any degree of generality. Although not an easy problem, it is certainly within the reach of any group willing to carefully implement a system of greater complexity than any previous segmentation process.

Among the tools listed above, glancing and multisensory data are probably the least understood. The idea of "glancing" at a scene to find interesting areas certainly has merit in outdoor applications. This technique is used by several groups, including ours, to find possible areas of man-made structure, based on texture. Once such an area is found, it is concentrated on by a structure analysis process. A variety of such operators, optimized to cue on different aspects, could be used to make a preliminary pass at the scene to determine, as the next step, what analysis operators should be applied and how and where.

It is advantageous for the scene analysis community to continue generalizing its view of segmentation, applying it to multisensory data. This is especially true for the development of systems relevant to military applications. An important factor in military sensor systems is all-weather capability. Because of this, very few sensors provide white light intensity data, and even fewer provide red-blue-green color. On the other hand, range is frequently available, and many new systems will be multisensory. It has been argued that work on r-b-g color is relevant to multisensor processing. Unfortunately, the sensor wavelengths are often far apart, presenting new difficulties because of different resolutions and reflectance peculiarities. Realistically, reflectance data from nonwhite light sensors is just as messy as for white light, and usually has less resolution. Thus, range stands out

as one of the most useful sources of information because of its invariance to these problems. Little is known about the availability of texture in range data, but it may also be a good source of information as better analysis techniques become available.

Connected-object segmentation is the most frequently used approach for organizing scene analysis systems. There is growing evidence that additional approaches are necessary, particularly in the outdoor scene environment. From an intuitive view, it seems doubtful if the demand for connectedness associated with segmentation is realistic in most outdoor scenes. Further, it seems doubtful that a complete segmentation, or even every edge or region, is necessary for the understanding of outdoor scenes. The idea of using "distinguished" features, from multiple sensors if possible, has been an initial step towards new approaches. Our group is currently investigating extensions using collections of isolated features from multiple sensors, feature locations from glancing "interest" operators, and available "obvious" region and line-intersection data together in one representation.

III. REPRESENTATIONS

Beyond the problem of low-level analysis, the real barrier to progress in scene analysis is not control or system organization, but the lack of adequate scene representations. The question is not just how a region or line should be characterized, but whether the currently available tools are even adequate for natural objects. Several generalizations of existing representational techniques have been suggested in the literature, but have not been experimentally evaluated. One suggestion is the use of multiple representations in terms of alternative sensors, operators, contexts, and resolutions. A second suggestion is the extension of mathematical approximations (a la splines, etc.). A third is to abandon traditional segmentation, as mentioned above, and base models on nonhomogeneous collections of multifeature

predicates, positioned in a controlled, nonhomogeneous manner. A fourth possibility is the use of iconic models and verification vision. In reality, it is not clear that these will be enough. For example, an adequate model for a complex natural object has never been constructed, even at the conceptual level, using any of these ideas. Unlike segmentation, real progress in this area is thus hampered by the lack of pure invention.

Another important problem associated with representations is their interface to the rest of the system. In many areas of AI there is an interface problem at the numeric/symbolic level. In scene analysis there is the additional barrier at the geometric (spatial)/symbolic level. One interface mechanism that can be used, at least at the geometric-symbolic level, is the use of graph rewriting rules operating in a production system framework. Such rules can specify the spatial relations of numeric operators, can be operated on themselves as symbolic entities, and can specify what numerical or symbolic form is to result.

IV. SYSTEMS

It is clear that parallelism will become an integral part of scene analysis systems, if only to achieve high throughput for the complex processes now evolving. Because of the crude state of current attempts it is less clear whether or not parallelism will affect the methodology itself. The "parallel" algorithms that frequently find their way into publication show amazingly little creative effort to rise above the "micro" level of the problem. The few facilities that have even crude parallel processors (e.g., C.mmp), have understandably, and unfortunately, been bogged down in upgrading the state of the art in operating systems and control. Therefore, a serious attempt from within the vision community, with the right perspective, could have dramatic impact on the whole issue of parallelism. The approach

should be towards the entire system, at all levels of knowledge, rather than, as in the past, attempts at constructing only piecemeal algorithms. The emphasis should not be on hardware, for the hardware problems can all be solved to some level of satisfaction. The real effort should be concerned with programming languages for efficiently and effectively specifying representations, processes, and control. There are certainly many seeds in current multiprocess, AI-language, relaxation-process, production-system, and network research. But these are all only crude starts.

Based on the alternative segmentation ideas discussed above, it is easy to propose systems that make good use of semantic knowledge and multisensor concepts. The primary low-level components can be viewed as glancing operators and analysis (segmentation) operators. The glancing operators are completely knowledge-independent. The system is initialized by selecting glancing operators and sensor modalities based on the user's goal and the available semantic knowledge. The glancing operators provide multiple views of the scene that can be interpreted using semantic knowledge to specify where and which analysis operators are to be employed, and how their parameters (knowledge-dependent) are to be set. Such a scheme could easily incorporate the use of relaxation methods to drive the process to equilibrium.

An interesting system organization, hinted at earlier, that can be used to construct such a system is a variation on a production system framework. There has been a fad associated with the use of production systems in other areas of AI, exclusive of vision. Productions can be used in scene analysis to provide a powerful symbolic programmable control framework for vision modules and semantic knowledge. They are attractive because they can be used to easily embed domain-dependent semantic knowledge (allowing application of one system to several domains) as well as geometric/spatial knowledge, as described previously.

V. RESEARCH DIRECTIONS

The principle limiting factor in outdoor scene analysis is the sad state of current representational capabilities. A good representation should be able to model natural shapes, texture, and three-dimensional information. If sufficiently rich representations for natural scenes were devised, then, rather than the present shotgun approach, there would be some direction to research in analysis techniques to collect the fundamental units.

Another important problem, at a more basic level, is the lack of a good implementation language for scene analysis. There is no existing language with these desirable features: efficient numeric computation, symbolic computation, clean syntax, basic scene analysis primitives, good debugging and editing facilities, reasonable portability, and good documentation. The development of such a language would dramatically affect the rate of progress, standardization of "working" modules, and exchange of capabilities between groups.

Our own group has found that it is vital to have a mixture of both applied and general research. No group, particularly receiving funding from the military community, should be allowed to work exclusively in either domain.

VI. EVALUATION

The problem of evaluating system performance has two levels, formal and informal. Most would probably agree that scene analysis is still in its early stages of research. During the early stage the best form of evaluation is "informal," since the first problem is to get anything to work at all. Rather than divert precious research dollars to evaluate performance, a better strategy would simply be to promote standard image formats so groups could at least exchange raw data.

Complex issues, such as the comparative performance of basic operators, can then be avoided by comparing them in the context of an entire system. The formal level is best ignored until it cannot be avoided. The evaluation of sensor systems at the formal level involves, moreover, the extensive collection of extremely large application-dependent data bases. These, in turn, must be processed to get statistically meaningful results. Both these steps are costly and nonproductive of research goals.

When good vision-system components are available, their application will probably be limited more by the speed with which adequate data bases can be acquired for models, rather than by component performance or system reconfigurability.

Restricted domain systems, capable of understanding man-made structures in outdoor scenes, will be operational in small boxes within five years.

18 AFOSR TR-76-1229	9 Interim repl.
6 UNSTRUCTURED CONTROL AND COMMUNICATION PROCESSES IN REAL WORLD SCENE ANALYSIS.	1 Apr 75-31 Mar 76
10 Bruce L. Bullock	15 F4462/14-C-0054
9. PERFORMING ORGANIZATION NAME AND ADDRESS Hughes Research Laboratories 3011 Malibu Canyon Rd Malibu, CA 90265	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 61102F 2364/A2
11. CONTROLLING OFFICE NAME AND ADDRESS Air Force Office of Scientific Research/NM Bolling AFB, Washington, DC 20332	12. REPORT NUMBER 11 Petri-76
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) 12 38p.	13. NUMBER OF PAGES 43 15. SECURITY CLASS. (of this report) UNCLASSIFIED
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.	
17. DISTRIBUTION STATEMENT (of the abstract entered in block 20, if different from Report)	
18. SUPPLEMENTARY NOTES	
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Scene analysis pattern recognition production-systems program specifications	
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) → This report describes research results on the use of flexible control structures in scene analysis. The primary result during this contract period was the implementation of a complete system for complex outdoor scene analysis. This scene analysis system is implemented with a flexible control structure called a pattern-directed production system. The results described include the performance of the → next page	

20 Abstract

cont
→

present system, the use of a simplified system to aid the development of a theoretical approach, and initial results on the development of a symbolic method for program specification.

↑