AD-A090 465 SRI INTERNATIONAL MENLO PARK CA ARTIFICIAL INTELLIGE--ETC F/G 15/4
A STUDY OF KNOWLEDGE-BASED SYSTEMS FOR PHOTO INTERPRETATION.(U)
JUN 80 R O DUDA, T D GARVEY DAAK70-78-C-0114

UNCLASSIFIED ETL-0235 NL

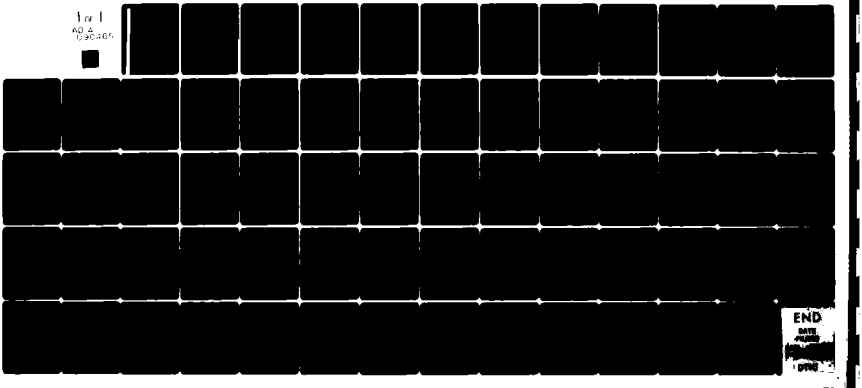

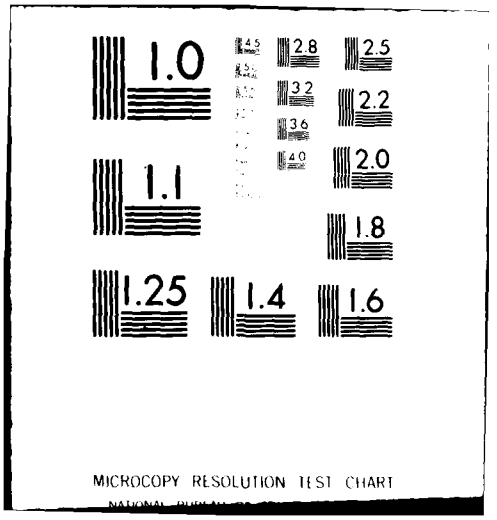
1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0
MICROCOPY RESOLUTION TEST CHART

AD A090465

ETL-0235

LEVEL II 5

A STUDY OF KNOWLEDGE-BASED SYSTEMS FOR PHOTO INTERPRETATION

Richard O. Duda
Thomas D. Garvey
SRI International
333 Ravenswood Avenue
Menlo Park, California 94025

DTIC
ELECTE
OCT 15 1980
D

June 1980

Final Report

Approved for public release; distribution unlimited.

Prepared for:

U.S. Army Engineer Topographic Laboratories
Fort Belvoir, Virginia 22060

DDC FILE COPY

80 10 8 005

NOTICES

Destroy this report when no longer needed.
Do not return it to the originator.

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

The citation in this report of trade names of commercially available products does not constitute official endorsement or approval of the use of such products.

(9) Final rept.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

(19) **REPORT DOCUMENTATION PAGE**

READ INSTRUCTIONS BEFORE COMPLETING FORM

(18) 1. REPORT NUMBER: ETL-0235	2. GOVT ACCESSION NO.: AD-A090 465
3. RECIPIENT'S CATALOG NUMBER:	
4. TITLE (and Subtitle): A STUDY OF KNOWLEDGE-BASED SYSTEMS FOR PHOTO INTERPRETATION	5. TYPE OF REPORT & PERIOD COVERED: Contract Report
	6. PERFORMING ORG. REPORT NUMBER: SRI Project 7864
7. AUTHOR(s): Richard O. Duda, Thomas D. Garvey	8. CONTRACT OR GRANT NUMBER(s): DAAK70-78-C-0114
9. PERFORMING ORGANIZATION NAME AND ADDRESS: SRI International, Artificial Intelligence Center, Computer Science and Technology Division, 333 Ravenswood Ave., Menlo Park, CA 94025	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: 1-1661
11. CONTROLLING OFFICE NAME AND ADDRESS: U. S. Army Engineer Topographic Laboratories, Fort Belvoir, Virginia 22060	12. REPORT DATE: June 1980
	13. NUMBER OF PAGES: v + 59
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)	15. SECURITY CLASS. (of this report): Unclassified
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

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)

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report discusses applications of knowledge-based programming techniques to a selection of photo interpretation tasks. The technology of knowledge-based programming is reviewed, several relevant photo interpretation problems are described, and the most promising applications are discussed in detail. Recommendations are presented for developing knowledge-based systems as expert consultant programs to assist military geographic intelligence analysts and to aid users of advanced photo interpretation tools.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

410667

PREFACE

This report was prepared under Contract DAAK70-78-C-0114 for U. S. Army Engineer Topographic Laboratories, Fort Belvoir, Virginia. The contractor monitor is Dr. Robert D. Leighty.

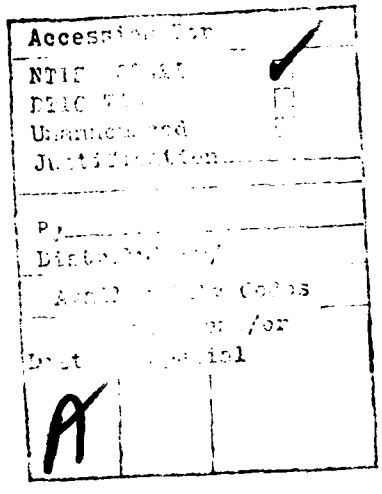

CONTENTS

PREFACE ii

LIST OF TABLES v

I INTRODUCTION 1

A. Knowledge-Based Systems: A Brief Summary 1

B. Present PI Programs 2

C. Report Content and Sequence 6

II BACKGROUND ON KNOWLEDGE-BASED SYSTEMS TECHNOLOGY 7

A. Introduction 7

B. A Classification of Knowledge-Based Systems 9

C. General Remarks 22

III SELECTED PHOTO INTERPRETATION TASKS AND REQUIREMENTS . . . 24

A. DLMS/VO 24

B. MGI (Military Geographic Information) 26

C. CAPIR 27

D. Structure Identification 29

IV APPLICATION OF KNOWLEDGE-BASED TECHNIQUES TO PI TASKS . . . 31

A. Overview 31

B. MGI 33

C. DLMS 35

D. VO 37

E. Automated PI/Cartographic Workstations (CAPIR) . . . 38

F. Monitoring and Identification of Structures and Operations 40

G. Recommendations 41

V Conclusions 43

APPENDICES

A An Overview of Prospector 44

B Perceptual Reasoning 50

REFERENCES 53

TABLES

1 Feature Identification Codes for DLMS 4

2 A Classification of Knowledge-Based Systems 10

3 KBS Applications Matrix 31

I INTRODUCTION

Over the past several years, researchers in artificial intelligence (AI) have developed a number of interesting programs that have come to be known as knowledge-based systems (KBS). Working in specialized problem domains, these programs have achieved levels of performance that approach (and in some cases equal) the skill of expert humans.

This report discusses applications of knowledge-based programming techniques to selected photo interpretation (PI*) tasks. We begin with an overview of knowledge-based programming technology and a description of several potentially relevant PI problems. We then examine the most promising applications and present our recommendations for developing knowledge-based programs to assist military geographic-intelligence (MGI) analysts and to aid users of advanced PI tools.

A. Knowledge-Based Systems: A Brief Summary

A KBS has certain characteristics that differentiate it from more traditional computer programs. First and foremost, there is an externally accessible knowledge base containing information pertaining specifically to the application area in which the KBS is meant to operate. The system is usually structured so that performance can be improved merely by expanding or modifying the knowledge base. A separate computing procedure is used to draw conclusions (inferences) from material in the knowledge base. The inference path along which any particular conclusion is reached is determined by the contents of the knowledge base and by the relevant problem data. A KBS can typically "explain" the reasoning process that culminated in a conclusion by tracing this path.

---

* PI will also be used as an abbreviation for photo interpreter.

KBSs enjoy certain advantages over more traditional computer systems. These include the following:

* Modularity of the knowledge base
* The ability to understand and explain the reasoning process
* Ease of modification
* A possible choice of reasoning mechanisms
* Relatively simple control structures
* Incremental addition of knowledge
* The ability to change the knowledge base independently of the computational structures.

The primary functions of a KBS fall roughly into any of three categories: tutorial (an instruction program); consultant (an interactive system that works with a user to enhance or extend his capabilities); autonomous (a system that uses a knowledge-based approach to solve problems automatically). While these categorizations are not absolute, and a given KBS may cross the boundaries, they form a useful trichotomy for the purposes of this report. We shall discuss each role as it applies to selected PI problems.

KBSs have shown significant performance in a number of diverse problem areas. Among these are mass spectrum analysis, chemical synthesis, medical diagnosis, consultation for mineral exploration, intelligent terminal interfacing, and assessment of air defense situations. Further discussions of these applications will appear in Section II.

B. Present PI Programs

In the course of this investigation visits were made to the Defense Mapping Agency in St. Louis (DMA) and to the U.S. Army Engineer Topographic Laboratories at Ft. Belvoir (ETL). Discussions with photo interpreters and analysts led us to focus our investigation on a few high-priority programs. These included the Digital Landmass Simulation (DLMS) and Vertical Obstruction (VO) programs, Military Geographic Intelligence (MGI), the Computer-Assisted Photo Interpretation Research

(CAPIR) program (along with other advanced computer aids), and monitoring and tasks involving structure identification.

DLMS embodies the development of a geographic data base of surface information used for generating low-resolution, synthetic radar displays. All features that would result in significant radar reflections are mapped and categorized into one of the 255 possible classes shown in Table 1*. The analyst begins work in a new area by first familiarizing himself with collected maps, imagery, and collateral data. He then begins outlining and identifying regions of an aerial photograph, transferring this information to a maplike manuscript. In the period during which he works on an area (anywhere up to two months), he will normally use only a single photograph (or stereo pair) and will be responsible for identifying all significant subareas. As currently performed, the process employs only minimal computer aids.

---

* Table 1 was reproduced from a copy supplied by ETL.

Table 1

Feature Identification Codes for DLMS

Feature Identification	Code No.
Industry	
The industrial category is comprised of the area and facilities, in ng buildings, utilized by those establishments engaged in the extrac or materials, the processing of these materials and the production of *mediate and finished products.*	
Extraction Industry (General)	101
Quarry	102
Gas/Oil Derrick	103
Offshore Platform	104
Offshore Platform with Derrick	105
Mine Shaft Superstructure	106
Processing Industry (General)	110
Chemical Processing Plant	111
Metal Processing Plant	112
Sewage Treatment Plant	113
Evaporator Mining (Settling Basins, Sludge Ponds)	114
Cone Plant	115
Blast Furnace	116
Refinery	120
Catalytic Cracker	121
Flare Pipe	122
Power Plants (General)	130
Hydro-electric Power Plant	135
Thermal Power Plant	136
Transformer Yard	137
Substation	138
Solar Energy Electrical Collection Panels	145
Solar Energy Heat Collection Panels	146
Heavy Fabrication Industry (General)	150
With Flat Roof	151
With Flat Roof and Monitor	152
With Gable Roof (Pitched)	153
With Gable Roof (Pitched) and Monitor	154
With Sawtooth Roof	155
With Curved Roof	156
Light Fabrication Industry (General)	160
With Flat Roof	161
With Flat Roof and Monitor	162
With Gable Roof (Pitched)	163
With Gable Roof (Pitched) and Monitor	164
With Sawtooth Roof	165
With Curved Roof	166
Disposal (General)	170
Scrap Yard	171
Metal Ore Slag Dump	172
Tailings, Waste Piles	173
Strip Mine	174
Associated Structures (General)	180
Building	181
Smokestack	182
Conveyor	183
Pumping Station	184
Bridge Crane	185
Rotating Crane	186
Rotating Crane on Tower	187
Cooling Tower	188
Hopper	189
Dredge, Powershovel, Dragline	190
Transportation	
The transportation category consists of the area and facilities utili ving materials and people from place to place on land.	
Railroads (General)	201
Elevated	202
Railroad Yards/Sidings (Moderate to Heavy Activity)	203
Railroad Yards/Sidings (Light Activity)	204
Railroad Yards/Sidings (Empty)	205
Storage and Repair Building	220
Railroad Terminal Building	221
Railroad Station	222
Railroad Control Towers	223
Roundhouse	224
Roads (General)	240
Elevated	244
Causeways	245
Bridges (General)	260
Suspension	262
Cantilever	263
Arch	264
Truss	265
Moveable Span	266
Bridge Towers	267
Conduits (General)	280
Pipelines (Above Ground)	281
Aqueducts	282
Penstocks, Flumes	283
Associated Structures (General)	290
Commercial/Recreational	
The commercial category consists of the area and buildings where the ir business activities and recreational facilities of the urban area a ducted.	
Commercial Buildings (General)	301
With Flat Roof	302
Circular with Flat Roof	303
With Gable Roof	304
With Curved Roof	305
Recreational Activities (General)	320
Enclosed Stadium	321
Open-ended Stadium	322
Dome Stadium	323
Grandstand	324
Athletic Field	325
Amusement Park	330
Roller Coaster	331
Ferris Wheel	332
Artificial Mountain	334
Display Signs (General)	360
Advertising Billboards	361
Scoreboard	363
Overhead Highway Sign	364
Associated Structures (General)	380
Drive-In Theater Screens	382
Residential/Agricultural	
The residential category is comprised of the area and associated ildings where the civilian or non-military population live.	
Multi-family Dwellings (General)	401
Apartment/Hotels with Flat Roof	402
Apartments/Hotels with Gable Roof	403
Single Family Dwellings (General)	420
Mobile Home	421
Agricultural category is comprised of the structures involved with agr ral activities.	
Agricultural Buildings (General)	430
Stockyard/Holding Pen	433
Windmill (Truss)	434
Windmill (Solid)	435
Associated Structures (General)	450
Cemetery Buildings	453
Communications/Transmission Facilities and Tower Types	
Communication facilities are used for transmitting information from s to place. The communication category includes telephone, telegraph) facilities as well as other forms of electronic features such as pow pylons, and associated structures.	
Communication Towers (General)	501
Radio/Television, Tower Type "A"	511
Radio/Television, Tower Type "I"	512
Microwave Tower, Type "A"	520
Microwave Tower, Type "I"	521
Miscellaneous Towers (General)	530
Observation Tower	531
Tower on Structure	532
Power Transmission Towers (General)	540
Powerline Pylons, Type "A"	541
Powerline Pylons, Type "H"	542
Powerline Pylons, Type "I"	543
Powerline Pylons, Type "Y"	544
Associated Structures	560
Communication Buildings	561
Governmental and Institutional	
The governmental and institutional category is comprised of the area facilities, primarily buildings, which constitute the seat of legal, admin ive or other governmental functions of a country or political subdivisi ell as public service institutions such as universities, churches and ospitals.	
Governmental (General)	601
Capital Building	602
Administration Building	603
Prison	604
Palace	605
Castle	606
Institutional (General)	610
School	620
With Flat Roof	621
With Gable Roof	622
Hospital	630
With Flat Roof	631
With Gable Roof	632
Observatory	640
With Dome Roof	641
Houses of Religious Worship (General)	650
Associated Structures (General)	680
Steeple	681
Monument/Obelisk	682
Arch	683
Pyramid	684
Military/Civil Installations	
The military/civil category consists of the area and facilities utili the air, naval and ground forces for waging war, for training and for ilitering and transporting personnel and material, as well as the area an ilities used for transporting non-military goods and personnel by sea a air.	
Airport/Airbase (General)	701
Control Tower	702
Terminal/Base Operations	703
Hangar with Flat Roof	704
Hangar with Curved Roof	705
Runways and Taxiways	706
Aircraft Parking Areas	707
Airport/Airbase Electronic Navigation Aids (General)	710
Radar Reflectors, Uni-Directional	711
Radar Reflectors, Bi-Directional	712
Radar Reflectors, Omni-Directional	713
VOR/VORTAC TACAN Facility	714
Radar Antenna with Radome	715
Radar Antenna, Tower Mounted with Radome	716
Radar Antenna	717
Radar Antenna, Tower Mounted	718
Radar Antenna, Vehicle Mounted	719
Approach Lights Framework	720
GCA Facility	721
Runway Approach Lighting Systems (General)	725
Maritime Features (General)	750
Breakwater/Jetty	751
Wharf/Pier	752
Drydock	753
Lock	754
Off-Shore Loading Facilities	755
Exposed Wreck	756
Maritime Navigation Aids (General)	760
Navigation Light-Ship (Permanently Moored)	761
Lighthouse	765
Ground (General)	770
Barracks, Curved Roof	771
Barracks, Flat Roof	772
Barracks, Gable Roof	773
Motor Pools	774
Garage, Curved Roof	775
Garage, Flat Roof	776
Garage, Gable Roof	777
Depots	778
Associated Structures (General)	790
Administration Building	791
Engine Test Cells	797
Wind Tunnels	798
Storage	
The storage category is comprised of the area and facilities for han handling liquids or gases, bulk solids, and finished products.	
Tanks (General)	801
Cylindrical, Flat Top	802
Cylindrical, Dome Top	803
Cylindrical, Peaked/Conical Top	804
Cylindrical, Peaked/Conical Top, Tower Mounted	805
Spherical	806
Spherical with Column Support	807
Blimp	810
Ballet	811
Telescoping Gasholder (Gasometer)	812
Closed Storage (General)	820
Silo	821
Grain Elevator	822
Grain Bin	823
Water Tower (Building)	824
Ordnance Storage Mounds	825
Open Storage (General)	830
Mineral	831
Associated Features (General)	860
Warehouses	861
Vehicle Storage Area	862
Vehicle Parking Area	863
Aircraft Storage Area	864
Ship Storage Area	865
Landforms, Vegetation and Miscellaneous Features	
The landforms and vegetation category is comprised of the area which closely describes the surface landscape characteristics or natural ory. This category will also include such features as levees, walls or es.	
Ground Surface (General)	901
Soil	902
Sand/Desert	903
Sand Dunes	90
Marsh/Swamp	90
Rice Paddies	909
Smooth Solid Rock	910
Boulder Field/Lava	911
Rocky Rough Surface	912
Dry Lake	913
Mud/Tidal Flats	914
Surface Features (General)	920
Levees/Embankments	921
Walls	922
Cliffs	923
Dams	925
Water Intake Towers	926
Fences	927
Salt Water (General)	930
Salt Water, Sea State	931
Salt Water, Sea State, Subject to Ice	932
Salt Water, Subject to Ice	933
Salt Pans	934
Fresh Water (General)	940
Fresh Water, Sea State	941
Fresh Water, Sea State, Subject to Ice	942
Fresh Water, Subject to Ice	943
Waterfall	944
Vegetation (General)	950
Orchards/Mangroves	951
Deciduous Trees	952
Evergreen Trees (including Mangrove and Nipa)	953
Mixed Trees (Deciduous and Evergreen)	954
Tundra	955
Vineyards	956
Snow/Ice Areas (General)	960
Permanent Snow	961
Permanent Ice -- Glacier/Ice Cap	962
Fresh Ice (Temporary)	963
Polar Ice Pack	964
Regional Features (General)	980

The VO program attempts to determine and catalog the height of obstructions in the area.

The purpose of the MGI system is to analyze information concerning geographic areas of interest and to produce "factor overlays" for topographic maps that show relevant characteristics, such as soil type, geology, vegetation, slopes, drainage, and roads. Information from these overlays is then combined through mathematical and qualitative models to generate synthetic intelligence products. Examples of the latter are speed and cross-country movement capabilities (CCM) for different vehicles, locations of potential helicopter landing zones, fields of fire, and intervisibility. Synthesized products are usually developed in response to specific requirements.

CAPIR is a research program in support of interactive PI work stations. This will provide monoscopic and stereoscopic computer interface work stations, graphic display interfaces, and a digital image work station for development of automatic and semiautomatic image analysis capabilities. Various research objectives are being pursued under CAPIR. These include point positioning and mensuration, creation and maintenance of geographic-information systems, fusion of multisource data, extraction of elevation data, computer-assisted decision-making, and semiautomatic pattern recognition.

Although of less direct interest to ETL, the task of identifying structures and monitoring associated processes is of sufficient importance that it was also addressed in this report. This activity requires that an interpreter or analyst infer the purpose of structures detected in aerial photography, basing his conclusions on such supporting information as adjacent, identifiable structures collateral information. Once a structure has been identified, it may be necessary to observe its operations so as to deduce pertinant activities. For example, an increase in the number of operations at a military airfield may indicate an impending action involving aircraft based at the installation.

C. Report Content and Sequence

The purpose of this study was to evaluate the applicability of knowledge-based systems techniques to problems in the foregoing areas of photo interpretation. Accordingly, the remainder of this report is organized as follows. To provide the required background we present a review of KBS technology in Section II. In Section III we examine 13 specific applications of this technology to problems in DLMS, VO, MGI and CAPIR, evaluating them with respect to their technical risk and potential value. We conclude in Section IV with a brief summary of our recommendations.

II BACKGROUND ON KNOWLEDGE-BASED SYSTEMS TECHNOLOGY*

A. Introduction

The traditional view of a computer program is that it is a series of instructions for processing data. This places major emphasis upon the processing procedures or algorithms, and upon the data structures upon which they operate [39] **. This perspective also dominated early research in artificial intelligence (AI), which was directed at developing very general procedures for such tasks as search, planning, problem-solving, and theorem-proving [45]. However, in recent years there has been a growing awareness among AI researchers that, if effective application programs are to be developed, general approaches that apply to many different domains must be supplemented by extensive special knowledge about the particular domain of the program. As Goldstein and Papert put it, "Today there has been a shift in paradigm. The fundamental problem of understanding intelligence is not the identification of a few powerful techniques, but rather the question of how to represent large amounts of knowledge in a fashion that permits their effective use and interaction." [31]

Along with this shift in emphasis, general methods have been developed for representing knowledge in both declarative and procedural forms, so that it can be employed in a variety of ways. Three general methods for representing knowledge in explicit, declarative form have found wide application--predicate calculus, frame systems, and production systems. A tutorial exposition of the principles underlying these approaches is given by Barnett and Bernstein [2].

---

* The survey material presented in this section for background information was prepared jointly for this project and for SRI Project 7825.

** References are listed in alphabetical order at the end of this report.

The predicate calculus is a formal, logical system that was particularly favored by the early developers of theorem-proving programs [33, 50, 61]. Here both general knowledge (such as "All men are mortal") and specific knowledge (such as "Socrates is a man") are represented in a uniform fashion as assertions or called well-formed formulas (WFFS). General-purpose theorem-proving methods are used to deduce the consequences of this knowledge.

Frame systems (usually attributed to Minsky [46]) represent knowledge in more highly structured ways to make many of the deductions immediate properties of the representation. In particular, there is automatic acquisition of those properties that an individual inherits from his membership in one or more classes of objects. In addition to being more efficient than general theorem-proving procedures, frame systems simplify the treatment of "inconsistencies" caused by exceptions to general rules. Several methods have been developed to represent structured objects, including semantic networks [8, 38], units [67], and frame languages [7, 60]. The relations among these representations are clearly discussed by Nilsson [51].

Production systems provide a general computational formalism with three major components--a global data base of assertions that represent the facts for a specific problem, a set of production rules that read and write on this data base, and a rule interpreter (also called the executive or the cognitive engine) that selects the rules to be applied [17, 47]. General knowledge about the domain is encoded in the production rules, which are also referred to as the knowledge base. These rules have the form <situation> --> <action>, and can be employed in two different modes. In the antecedent mode (also known as the event-driven, data-driven, or forward-chaining mode) a rule is interpreted to mean that if the specified situation is observed in the global data base, the corresponding action can be taken. In the consequent mode (also known as the goal-driven, hypothesis-driven, or backward-chaining mode) a rule is interpreted to mean that if a certain action is desired, the system should try to establish the corresponding situation.

The production system formalism has many similar variants, including pattern-directed inference systems [37] and rule-based systems [71]. Besides showing the clear separation of the knowledge base from the rule interpreter, they all exploit a modular organization that lets the knowledge base be developed incrementally and be used in a variety of ways. This general approach has been favored in most of the work aimed at developing knowledge-based expert systems, programs that can match the performance of a human expert in some specialized domain [44].

B. A Classification of Knowledge-Based Systems

To better understand both the capabilities and limitations of knowledge-based systems, it is helpful to examine a number of specific systems that have been created to solve particular problems. At first glance, the only common characteristic of the various knowledge-based systems developed to date seems to be their reliance on explicitly encoded knowledge. Groupings based on the method of knowledge representation (e.g., production rules) or on the domain of application (e.g., medicine) reveal little additional commonality.

However, some insight can be obtained by grouping the systems accor in to their general function. Table 2 shows such a classification of 33 well-known knowledge-based systems into 10 general categories. In this section we define each of these categories and describe representative programs in each category.

Table 2

A Classification of Knowledge-Based Systems

Function	Domain	System	Reference
Search	Chemistry	DENDRAL	[24]
	Chemistry	SECHS	[75]
	Chemistry	SYNCHEM	[29]
Problem Solving and Planning	Circuit anal.	EL	[64]
	Genetics	MOLGEN	[66]
	Mechanics	MECHO	[12]
	Programming	PECOS	[4]
Diagnosis	Medicine	PIP	[53]
	Medicine	CASNET	[72]
	Medicine	INTERNIST	[56]
	Medicine	MYCIN	[63]
	Engineering	SACON	[6]
	Geology	PROSPECTOR	[36]
Machine Perception	Acoustics	HASP (SU/X)	[48]
	Imagery	ACRONYM	[9]
	Chemistry	CRYSALIS (SU/P)	[23]
	Elec. Warfare	MSIS	[27]
Simulation	Econometrics	OIL	[15]
CAI	Electronics	SOPHIE	[10]
	Medicine	GUIDON	[14]
Learning	Chemistry	Meta-DENDRAL	[11]
	Agriculture	INDUCE	[19]
	Mathematics	AM	[40]
Intelligent Assistants	Messages	RITA	[71]
	Business	COMEX	[65]
	Operations Res.	TESTBED	[62]
Knowledge Acquisition	Diagnosis	TEIRESIAS	[18]
	Diagnosis	EMYCIN	[69]
	Diagnosis	EXPERT	[73]
	Diagnosis	KAS	[58]
System Building	---	ROSIE	[71]
	---	AGE	[49]
	---	HEARSAY III	---

1. Search: DENDRAL

Many problems in graph theory, game theory, and other areas of discrete mathematics can be posed as search problems. These problems are characterized by the existence (at least in principle) of a systematic method for generating possible solutions, as well as a systematic method for testing acceptability. As regards genuinely interesting problems, the number of possible solutions is so great that exhaustive search is completely infeasible. Any device that significantly reduces the amount of search required (preferably without compromising the quality of the solution) is called a heuristic, and a search strategy guided by heuristics is called heuristic search.

Since the study of heuristic search was one of the earliest activities in artificial intelligence, it is natural that DENDRAL--one of the first knowledge-based systems--was concerned with using knowledge to limit search. Begun in 1965 by Feigenbaum, Lederberg, and their colleagues at Stanford University [24], DENDRAL generates plausible structural representatives of organic molecules from mass spectrogram data, nuclear magnetic resonance data, and additional constraints provided by the user. The program runs in a plan-generate-test sequence, (1) deriving necessary constraints on the molecular structure, (2) systematically generating structures that satisfy those constraints, and (3) testing the proposed structures by predicting the mass spectrogram and rejecting those that disagree with the experimental results. The knowledge needed for Step 2 is encoded as an ingenious special procedure. The knowledge needed for Steps 1 and 3 is encoded as tables of production rules, which turns out to be very compatible with the way chemists think about the rules of mass spectrometry [11].

For the molecular families that are covered by these empirical rules, the program is said to surpass even expert chemists in speed and accuracy. The results obtained with it have been published in some 25 papers in chemistry journals [25].

2. Problem Solving and Planning: EL

An important class of logical problems concerns the decomposition of a given problem into a set of simpler subproblems. Here a solution, to be useful, must be constructive; it usually consists of a sequence of actions that will achieve some goal. Typical examples of such problems are theorem-proving, program synthesis and robot planning.

EL is a knowledge-based system for the steady-state analysis of resistor-diode-transistor circuits [64]. It uses production rules to represent general principles, such as Kirchhoff's laws and Ohm's law, as well as the characteristics of types of devices. Facts about the particular circuit being analyzed are represented as assertions in an associative data base. The rule interpreter is written in a special language called ARS (Antecedent Reasoning System). As its name implies, ARS supports the use of rules in the antecedent mode, in which the factual assertions trigger the rules. The actions of the rules create new assertions, which in turn trigger additional rules.

An important property of ARS is its ability to make conjectures when no additional direct deductions are possible, and to keep track of those conjectures and any conclusions dependent upon them, should subsequently detected contradictions require their revision. This ability allows the analysis of a circuit to proceed when the conducting or nonconducting states of its nonlinear devices are unclear. It also permits the user to modify the circuit and see the effects of changes without having to reanalyze the entire circuit. In doing this backtracking, ARS makes use of facilities that it can also employ to provide explanations of its reasoning.

Unfortunately, much of the performance of EL (and of other planning systems as well) depends more upon the rule interpreter than on the knowledge in the rule base. An obvious suggestion is to represent knowledge about strategies in declarative form, perhaps as strategy rules. Indeed, this approach has been investigated to some extent [17], but its advantages remain to be demonstrated convincingly.

3. Diagnosis: INTERNIST, MYCIN and Prospector

The general diagnosis problem is one of classifying an object or event on the basis of possibly uncertain information about its characteristics. The categories may or may not be mutually exclusive; the data may be acquired sequentially or in parallel. In a formal sense, diagnosis problems can always be posed as problems in statistical decision theory whose solution usually requires estimation of a multivariate probability function from vast amounts of data. The knowledge-based-systems approach to such problems effectively substitutes the knowledge and judgment of expert humans for this unknown function.

Several impressive knowledge-based systems have been developed for various problems in medical diagnosis. The INTERNIST program developed by Pople and Myers at the University of Pittsburgh uses information from some 4000 possible manifestations to diagnose problems of internal medicine that can involve multiple instances of about 500 different disease types [56]. The program contains a large taxonomy of disease types, together with rules that link manifestations to these types, as well as an ingenious control procedure for narrowing down the disease classes that explain the manifestations. In many tests it has demonstrated the capability of correctly diagnosing multiple disease cases that are described in medical journals as being particularly difficult.

Another well-known medical diagnosis program is the MYCIN system developed by Shortliffe at Stanford [63]. MYCIN is a consultation program designed to diagnose bacterial infections and to recommend antibiotic therapy. It is organized around the systematic use of a large collection of rules that link patient data to infection hypotheses. Formulas based on a probability-like theory of certainty are used to accommodate the inexact nature of the relevant medical knowledge. The modularity provided by the use of rules to express all such knowledge is exploited in several ways. From a system development standpoint, it allows long-term incremental development of the system by

continual expansion and refinement of the rule base. Operationally, the program obtains information from a user by the simple strategy of chaining backward through the rules. In addition, this lets the program furnish simple but very useful explanations of its reasoning by stating the rules it is using.

As its developers, we are particularly familiar with the Prospector program for mineral exploration ([21, 22, 36]; a reprint of [22] is included in Appendix A). Developed for the U. S. Geological Survey, Prospector contains models of different kinds of ore deposits. Each model was developed by interviewing a geologist who is an authority for that particular class of deposits, and by translating his knowledge of the associations between field-observable evidence and relevant geological hypotheses into a structured collection of rules. Once obtained, a model can be used for several purposes: (1) evaluation of the favorability of a geologic district for that kind of ore; (2) evaluation of the favorability of a particular exploration site; (3) evaluation of the most favorable drilling sites. In addition, the explanatory facilities of the system provide the rationale for conclusions, suggestions as to those data that would be most valuable for further exploration, and informal education about a class of ore deposits.

Prospector employs a combination of artificial intelligence techniques to perform these tasks. MYCIN-like rules are used to link evidence to hypotheses--with the advantages of modularity, explicability, and uniformity that this approach provides. Uncertainty in both the evidence and the rules is accommodated through the use of Bayesian probability theory. Any piece of evidence or hypothesis corresponds to an assertion that a certain situation exists; these assertions are represented as partitioned semantic networks [38]. This representation enables the system (1) to recognize and exploit general taxonomic relations, (2) to interconnect different models automatically, and (3) to connect user-supplied information to the models. A mixed-initiative control strategy is employed whereby the user can either let

the system use a backward-chaining strategy to gather information, or interrupt the program to select different goals or volunteer relevant information. Tests of its accuracy have repeatedly shown that the evaluations made by Prospector correspond faithfully to the evaluations of the geologists who created the models, the numerical scores typically agreeing to within 7 percent [28].

4. Machine Perception: MSIS

The problem of programming computers to recognize objects and events from raw sensory data has proved to be one of the most difficult problems in artificial intelligence ([34, 59]). It combines the combinatorial problems encountered in search and planning with the problems of uncertainty encountered in diagnosis. Furthermore, where other knowledge-based systems are intended for specialized applications, perception programs are often asked to cope with a great diversity of situations. Thus, few machine perception systems possess anything resembling expert capabilities*.

Furthermore, the sheer volume of raw input data usually makes efficient data processing mandatory, which discourages the use of programs that must interpret declarative knowledge. However, as Table 2 indicates, some interesting attempts have been made to use knowledge-based-system ideas in dealing with perceptual problems. This is a particularly attractive approach when there are strong prior models of what to expect in the data, and when the sensors measure phenomena not readily perceivable by humans.

The recently developed MSIS (Multisensor Integration System) program meets both of these criteria ([27];see also Appendix B). MSIS is intended to interpret information furnished by various sensors to determine the disposition and operation of hostile air defense units.

---

* Important exceptions exist in constrained situations, particularly in the classification of isolated words [74] and isolated two-dimensional objects, such as printed characters [35], blood cells [57], and certain industrial parts [1].

The program runs in an anticipate-plan-interpret cycle. In the anticipation phase it uses general knowledge about the location of threats (which are mobile and interrelated) to identify critical areas of missing information. In the planning phase it uses knowledge about both the potential threats and the sensors to allocate and deploy sensor resources. In the interpretation phase it uses knowledge encoded as production rules to infer probable weapon operation from information about such things as past status, environmental conditions, and distance from the aircraft. Although this effort is still at an early stage of development, it is clear that such knowledge is essential for effective performance.

5. Simulation: OIL

There are various ways in which knowledge-based systems can be interfaced to more conventional computer programs to produce interesting hybrid systems. Among the programs we have mentioned, both DENDRAL and MSIS have major computational components. One particularly promising combination is a knowledge-based program and a simulation program. The former utilizes judgmental knowledge, the latter analytical knowledge.

An interesting early attempt at such a hybrid system was the OIL program developed by Coles at SRI [15].* This program used a qualitative, rule-based political model of the circumstances in which the OPEC countries might choose to raise or lower their prices to drive a quantitative, econometric model of global oil demand, supply, and revenue. The 90 rules in the political model were derived primarily from the news media and included numerical factors to specify the time interval within which the action is expected to take place, the credibility of the source, the importance and the plausibility of the rule. Primarily a feasibility study, the program did demonstrate the validity of the general hybrid-system concept.

---

* SOPHIE [10] is another knowledge-based system employing a simulation model. However, since the goal of the program was computer-aided instruction rather than intelligent use of simulation, we have classified SOPHIE under the CAI category.

6. Computer-Aided Instruction: GUIDON

Three types of traditional computer-aided instruction (CAI) are often distinguished: frame-oriented drill-and-practice programs (which are unrelated to frame-based representations); games and simulations (usually used to teach diagnosis); exploratory systems that allow the student to experiment freely and learn by doing. Among the limitations of these programs are their inability to conduct dialogues with the student in natural language, to respond to unanticipated replies, inability to diagnose the student's errors, or to improve with experience [52]. The potential applicability of artificial intelligence techniques to solving these problems was outlined by Carbonell over ten years ago [13], and a variety of approaches have been subsequently explored. To the extent that knowledge of the subject matter is required for a solution, knowledge-based systems have an obvious contribution to make to CAI.

The GUIDON system developed by Clancey at Stanford exploits the MYCIN knowledge base about meningitis and bacteremia, for the teaching of both facts and problem-solving strategies [14]. MYCIN's 450 diagnostic rules were not modified, but were augmented by an additional 200 rules that included methods for guiding the dialogue with the student, presenting diagnostic strategies, constructing a student model, and responding to the student's initiative. Thus, by replacing MYCIN's rules with a separate set of MYCIN-style rules used in the PUFF program for diagnosing pulmonary disease, Clancey was able to use GUIDON to tutor students about pulmonary-function analysis. Although the natural-language capability of GUIDON is limited, the error diagnosis assumes that the student is using a subset of MYCIN's rules and that the system does not include any learning facilities. GUIDON has demonstrated convincingly that a knowledge-based system can be effectively exploited for teaching about its knowledge base.

7. Learning: Meta-DENDRAL

The idea of including learning mechanisms in a program so that its performance can improve with experience has always been an interesting topic of research in artificial intelligence. Indeed, the ability to learn is universally considered to be a hallmark of intelligence. Since the construction of a knowledge base is one of the most time-consuming tasks in building a knowledge-based system, it may seem surprising that there have been relatively few attempts to design a system that can learn what it needs to know.

A significant exception to this rule, however, is the Meta-DENDRAL program, which can learn new rules of mass spectrometry for DENDRAL [11]. To develop these rules, Meta-DENDRAL uses many examples that show the relative abundance of the masses of molecular fragments obtained when molecules of known structure are bombarded in a mass spectrometer. The learning procedure is not one of purely statistical association, but is guided by a simple theory of mass spectrometry that places acceptability constraints upon rules. The program runs in a plan-generate-test cycle that begins with examination of a large group of examples to identify feasible processes of molecular bond-breaking that could produce the observed masses. A rule generator uses the theory to propose rules that are subsequently screened by heuristic criteria of generality and specificity. The screened rules are modified and tested against the data, and the acceptable new rules are then added to the knowledge base. In addition to rediscovering known mass spectrometry rules, for three closely related families of structures Meta-DENDRAL has discovered new rules that have been published in the chemical literature. While not all knowledge-based systems can employ such techniques, we believe that learning procedures will play an increasingly prominent role in the future.

8. Intelligent Assistants: RITA

Many of the knowledge-based systems we have discussed utilize their abilities to recognize patterns and explain their actions. One

way of putting this to use is in consultation programs. Another way is through programs that act as assistants or agents.

Implemented on a PDP-11/45 minicomputer, RITA (the Rand Intelligent Terminal Agent) is a system for building rule-based "user agents" that can relieve a user from much of the tedium in interfacing with external data systems [71]. These agents can perform such tasks as storing, retrieving, and editing data on local files; handling dialogs with other computer systems accessible by computer networks; initiating and monitoring several external jobs in parallel; and explaining the actions that are being taken. RITA typically employs antecedent reasoning, but the user can specify that certain rules are to be used only for consequent reasoning. RITA has been used in several different applications, perhaps the most ambitious being a 1000-rule system that serves as the Naval Warfare Simulation System (NWSS) man-machine interface at the Naval Ocean Systems Command (NOSC).

9. Knowledge Acquisition: KAS

As an alternative between the extremes of complete handcrafting of the knowledge base and totally automated learning, several researchers have explored the development of various tools that can facilitate the process of knowledge acquisition. The most ambitious of these is Davis's TEIRESIAS system that employs knowledge about the MYCIN system to supervise interaction with an expert in building or augmenting a MYCIN rule set [18].

While giving the expert direct access to the program is an appropriate ultimate goal, it is greatly complicated by the fact that the expert usually does not have sufficient understanding of the representational mechanisms employed by the program to appreciate the consequences of the many choices before him. An alternative is to retain the services of a computer scientist who understands those mechanisms and to provide him with specific tools matched to the knowledge-acquisition process.

KAS, the Knowledge Acquisition System developed by Reboh for Prospector, is an example of such a system [58]. Prospector employs various kinds of networks to represent knowledge--rule networks for expressing judgmental knowledge, semantic networks for expressing the meaning of the propositions employed in the rules, and taxonomic networks for representing static knowledge about the relations among terms in the domain. The core of KAS is a network editor. Its primitive operations allow it to create, modify, or delete various kinds of nodes and arcs. It knows about the various mechanisms employed by Prospector, protects the user against certain kinds of syntactic errors, and includes a bookkeeping system that keeps track of incomplete constructs. Whenever he desires, the user can turn control over to KAS, which will systematically question him to fill in the missing parts of the structures. A semantic network matcher gives the user a limited ability to edit by content rather than form. Finally, being embedded in Prospector, KAS lets the user determine the effects of his changes by permitting controlled execution of the program. Although specialized for Prospector, KAS provides powerful assistance in the time-consuming task of developing the knowledge base.

10. System Building: AGE

Many of the knowledge-based systems that have been built--especially in the area of diagnosis--have a generally similar structure. In particular, rule-based systems all have a collection of rules that comprise the knowledge base, a rule interpreter, and a global data base of assertions about the particular case being diagnosed. Several researchers have illustrated the generality of their systems by showing that they can be applied to another domain merely by removing the rules for a given domain and substituting rules for the new one ([30, 69]).

However, every domain has its own peculiarities and, despite the good intentions of system builders, these features inevitably influence the design of a system. Thus, a serious attempt to build a significant new knowledge-based system almost always requires changes in

all parts of the system. Recognizing these facts, researchers have recently begun developing what amount to programming languages for building expert systems. While these languages are just beginning to be used and will undoubtedly undergo further development, they promise to reduce significantly the programming effort needed to develop a new system.

We have already mentioned ARS and RITA, which were early languages of this type. AGE is a good example of the most recent efforts in this direction. Specifically designed to allow the implementation of a broad spectrum of knowledge-based systems, AGE provides the designer with a set of separate interconnectable, pre-programmed modules for implementing the knowledge base, interpreter, and data base. Furthermore, AGE offers several ways to escape to the host programming language to implement arbitrary procedures. Thus, the knowledge base can be represented either as sets of production rules or as frame systems (called units), or both options may be combined. For the interpreter, AGE supplies standard procedures for forward-chaining and backward-chaining, plus convenient ways to implement other strategies. The standard global data base is a so-called blackboard system [41]. Finally, AGE contains knowledge about its own facilities and procedures. A tutor subsystem allows the user to browse through this on-line manual, while a design subsystem provides on-line advice on the use of AGE itself.

The design of any programming language always involves a compromise between convenience, generality and efficiency (in time or space). Clearly, systems like AGE are attempts to gain convenience, generality and design time efficiency in exchange for relatively modest additional cost in space and run time efficiency. While the community has not had sufficient experience with such systems to be able to assess their real value, their existence bears witness to increasing maturity in the practice of knowledge-based systems. We fully anticipate that they will play a significant role in future developments.

C. General Remarks

While the systems we have described address a rather wide variety of problems, the KBS methodology is better adapted to some problems than to others. In particular, the KBS approach is inappropriate in areas where there are well-defined, effective mathematical procedures for solving the problem. Generalizing from the systems just described, we conclude this section with a few general remarks on the characteristics of problems that are best suited to a KBS approach*.

First of all, a KBS requires a knowledge base, which usually means that there must be at least one human expert who is generally acknowledged to perform a given task very well. Furthermore, the primary source of the expert's exceptional performance must be the special knowledge he possesses. Finally, the expert must be able to explain his knowledge and the methods he employs in applying it to particular problems. These characteristics describe many scientific or technical areas, particularly where there is a tradition of consultation. However, they are typically absent in the case of most perceptual problems, where experts are unusual, the knowledge is extensive but difficult to pinpoint, and understanding of the perceptual process is minimal.

A second important requirement is that the application domain should be well bounded. Even though KBSs have been built that contain impressively large knowledge bases, all of the successful systems have focused on specialized, self-contained problem domains. Thus, while a mathematician possesses specialized knowledge, the amount of general knowledge of the world that is needed to create a consultant mathematician is far beyond the current state of the art.

A third requirement is that there be a suitable user community. This usually implies a sufficient number of professional practitioners who are familiar with a domain's terminology and general methodology, even if they lack the expert's knowledge of some specialty. This

---

* These and related considerations are discussed in greater detail in [2] and [16].

requirement leads to other of economic and sociological factors that must normally be considered for genuinely effective use of any computer system.

When these conditions are met, a properly designed knowledge-based system can offer several significant advantages. The modular structure of the system allows incremental development, with assurance that increased capabilities can be achieved by systematic modifications of or additions to the knowledge base. The provision of a separate interpreter means that the strategies used for deploying that knowledge can be varied for different applications. Since both the knowledge base and the interpreter are accessible, flexible forms of explanation are readily obtainable. This makes knowledge-based systems natural prospects for intelligent computer-aided instruction. Finally, the uniform methods used for representation encourage the development of self-modifying or learning systems that can assume at least part of the burden of building the knowledge base.

III SELECTED PHOTO INTERPRETATION TASKS AND REQUIREMENTS

In this section we will provide a brief description of those PI tasks that were judged to be highly relevant to DMA and ETL programs. These include DLMS/VO, MGI, CAPIR, and the general task of identifying the nature and purpose of structures. The discussions in this section will concentrate on those aspects of each task area that are particularly appropriate to the application of KBS technology.

A. DLMS/VO

A major effort under way at DMA is compilation of the DLMS data base, to be used primarily for radar simulations in the training of pilots. The existence of this data base means that pilots can simulate flights over areas that would otherwise be restricted for political or economic reasons. Since surface materials, the existence of planar surfaces, and surface slopes are critical factors that will determine the strength of a radar reflection, these are consequently key elements in compiling the description of areas for DLMS. In the course of compiling a "feature analysis data table" for an area, the analyst will categorize features into one of 255 different classifications (see Table 1).

1. DLMS Analysis Procedures

An analyst will typically work in a small group concentrating on a particular geographic region. It will usually take them two months to complete work on an area. Prior to starting analysis of aerial imagery, the analyst will receive relevant maps and other collateral information that will enable him to become familiar with the area. Analysis normally involves only a single rectified photograph or stereo pair of photographs and appropriate maps.

The analyst prepares a "manuscript" containing outlines of broad area features, tracings of linear features, and the positions of point features. Keyed to this manuscript is a set of descriptive data, consisting of assigned feature classifications, produced as a separate document that is entered into a computer data base via an optical character recognizer. The feature manuscript is digitized and then compared with the descriptive data for the role of verifying consistency. Any errors detected are corrected later by the analyst. The turnaround time from input to error returns may be one or two days.

While the classification process is less concerned with the identity of a structure than its composition, it is frequently the case that identifying a structure will allow a determination of its composition to be made with greater confidence. In fact, the analysts are typically better at identifying the feature than its composition.

According to analysts engaged in DLMS work, there will usually be no more than five or six instances during the analysis of a given area that the PI will be unable to identify a single significant feature. Such cases will occur more frequently at the beginning stages of work on an area; the PIs become more adept as their familiarity with an area increases. Difficult identifications are frequently resolved by group discussion.

2. DLMS Analyst Training

New analysts attend a cartographic training school for two to six months. Of that period, approximately four weeks are devoted specifically to DLMS.

There are clear potential advantages to be derived from an intelligent, automated training system. It could allow greater depth in training and could adapt to the student's needs and abilities. Moreover, if properly designed, such a program could be readily modified to keep abreast of changes in the DLMS program.

B. MGI (Military Geographic Information)

MGI is processed information on the military significance of natural and manmade features of an area.

1. MGI Production

The interpretation of source material for the purpose of determining the effect of terrain on military personnel, equipment, and material is a very labor intensive, time consuming process. In an effort to reduce production time and provide consistently better terrain information, ETL is developing a series of terrain analysts' procedural guides on various subjects (see, for example, [43, 70, 68, 55]). The products of these guides are factor overlays, i.e., graphic representations of terrain information drawn on mylar film and registered to 1:50,000 scale topographic maps. These factor overlays for terrain subjects such as soil, vegetation, slope, and geology comprise a preformatted data base to support the Army's new Topographic Support System (TSS).

Production of factor overlays requires the analyst to systematically accumulate data from all available sources by progressive interpretation and inference. The information from these overlays can then be combined (with the aid of a mathematical model, when available) into a synthetic product called a "topic graphic." Examples of these topic graphics include cross-country movement maps, lines of communication maps and others.

A good example of a typical synthesis procedure is contained in the guide for cross-country movement (CCM). This procedure instructs the terrain analyst constructing a CCM topic graphic to (1) select the soil, slope, vegetation, drainage, and surface configuration factor overlays from the database; (2) trace them manually onto one complex overlay; (3) run the combined data through a preprogrammed calculator; and (4) classify the resulting speeds into movement categories on the redrawn complex overlay that becomes the CCM topic graphic. This topic graphic or map will then indicate the suitability of terrain for the

movement of infantry and vehicles. In this way, densely forested, steeply sloping terrain with loose soil might appear on the CCM map as a "no go" area for jeeps, but as a "5 kph" area for infantry.

While current procedural guides emphasize adherence to specified procedures, it is apparent that in many cases an analyst deriving information in response to a specific request could alter or circumvent certain procedures in the interests of efficiency. ETL researchers anticipate future use of optical synthesizers and microprocessors to help automate certain aspects of the synthesis task.

C. CAPIR

CAPIR is a program of research aimed at providing a stereoscopic work station that makes available a collection of up-to-date, automated routines and software capabilities for the cartographer/PI [42]. Further development of this program, which incorporates extensions of existing tools and capabilities, can be expected on the future.

CAPIR exemplifies the progress being made toward automated work stations. Meanwhile, other relevant research is under way at ETL in the Digital Image Analysis Laboratory (DIAL) and at various centers around the country under support of the DARPA Image Understanding Program [5]. The objective of this latter project is development of a test-bed facility at DMA.

1. CAPIR Capabilities

The CAPIR facility incorporates three major capabilities:

* A prototype computer-interfaced stereoscopic work station with graphic displays.
* A computer-interfaced monoscopic work station with graphic output.
* A digital image work station with a solid-state camera for sampling images and a video processor for image display.

a. Stereoscopic Work Station

The stereo work station is designed to interface the PI analyzing high-resolution aerial photography with the computer. It will allow stereo viewing for high-precision point positioning, mensuration, and direct stereo digitization with all outputs in a real-world ground coordinate system.

A key aspect of this component will be the ability to use combining optics to superimpose high-resolution graphics over the image. This will enable the PI to view his digitization as he proceeds. Equally important, it will allow overlaying previously created geographic information on an aerial photo, thereby permitting data base verification, editing, revision, and integration with other information sources.

The work station also includes an external graphics display and a CRT terminal for conventional communication with the host computer.

b. Monoscopic Work Station

The monoscopic work station consists of a digitizing tablet and a graphic terminal. Data being digitized can be displayed in real time. This allows direct, on-line data entry and permits various validation checks by the computer.

c. Digital Image Work Station

This is designed primarily to support future research in semiautomatic pattern recognition, with the goal of shifting certain decision-making responsibilities from the human user to the computer. The work station consists of a solid-state camera for sampling images, a video processor for storage and manipulation of image data and display on a color monitor, and a standard CRT terminal.

2. Initial Utilization of CAPIR Capabilities

Specific investigations are being started with the aim of utilizing the capabilities of the CAPIR work stations. These include providing improved point positioning and mensuration; development of procedures for creating, displaying, and maintaining of desired geographic data bases, such as MGI; techniques for fusion of multisource geographic and collateral data; extraction of elevation data; research in computer-assisted decision making; and development of capabilities for semiautomatic pattern recognition.

It appears that the efficient use of complex systems such as CAPIR will require the availability of suitable training and consulting aids. It is preferable that these be available on-line and that they be tightly integrated into the system software. It is also necessary that the training systems be flexible for easy adaptation as work station configurations and capabilities change. These criteria are strong motivations for the use of KBS technology in providing such training systems.

D. Structure Identification

The procedures followed by a PI in identifying various structures, determining their use, and monitoring their operation did not constitute a primary focus of this project. However, ETL personnel considered the topic sufficiently important to be discussed at least briefly in this report.

It is frequently not possible to identify structures in an aerial photograph solely by their appearance. Instead, the identity and function of structures must often be inferred from associated structures. For example, to determine the type of industry present in a heavily built-up area may require that it first be classified as an extraction industry (such as mining), a processing industry (such as a foundry), or a fabrication industry (such as a cannery). Answers to such questions as, "Are there large piles of bulk material?" or "Are there significant power distribution facilities?" could provide the

basis for this initial classification. Once a structure was determined to be associated with, say, a processing industry, further questions would lead to more detailed classifications, such as mechanical, chemical, or heat-processing. As classification becomes narrower and the questions more specific, increasingly confident conclusions as to the nature of a structure become possible.

This methodology lends itself readily to a KBS approach of the type developed for Prospector. The breadth of knowledge required by an expert PI argues persuasively for an automated interactive system.

IV APPLICATION OF KNOWLEDGE-BASED TECHNIQUES TO PI TASKS

A. Overview

In this section, we shall discuss the general applicability of knowledge-based approaches to selected PI problems. We shall indicate both the potential benefits and the apparent technical risk involved in each application. Our findings are summarized in Table 3, which will be used as a guide for the rest of the section.

Table 3

KBS Applications Matrix

MODES:	TRAINING		CONSULTATION		AUTONOMOUS	
TASKS	Risk	Benefit	Risk	Benefit	Risk	Benefit
MGI	Low	Mod.	Low	Mod.	High	High
DLMS	Mod.	Mod.	Low-Mod.	Low	Mod.	Mod.-High
Target ID	Mod.	Mod.-High	Mod.	Mod.	High	High
PI Tools	Mod.	Mod.	Low	Low	High	High
VO	Low	Low	--	--	--	--

Table 3 lists selected PI tasks at the left, cross-indexed with modes of KBS usage along the top. The application of knowledge-based

techniques has been roughly partitioned into three categories: training programs, consultation programs to assist PIs in creating their products, and autonomous programs for performing tasks or subtasks automatically.

A few general comments are in order before we discuss the individual applications. The tasks listed in Table 3 range in order of complexity from highest at the top to lowest at the bottom. More complex tasks typically require larger knowledge bases, greater inference capabilities, and correspondingly higher investments in system development. Accordingly, the payoff for a successful assault upon a complex area will tend to be high, as will the technical risk involved. The utility of the approach tends to diminish as the knowledge base becomes smaller or less well defined.

The use of a knowledge-based approach offers several inherent advantages. The construction of such a program requires collecting, systematizing, and encoding the body of relevant knowledge of the area. This is usually a beneficial exercise, as it compels close scrutiny of knowledge and techniques that may have evolved over a period of time through the experience of a large group of users. Once a KBS has been developed for a particular application, it is usually quite straightforward to modify and extend its capabilities by making changes in its knowledge base, rather than by reprogramming the system itself.

1. KBS for Training

A KBS developed for training new personnel will offer several advantages common to most computer-aided instruction systems, such as standardized curricula and instruction, better utilization of teachers, and the ability to use the system in a self-paced mode. A knowledge-based approach promises several additional benefits, including ease of modifying the curriculum; the potential for the program to "understand" the sources of student errors by examining its own knowledge base; and the ability to access the same instruction system in a computer-based consultant mode when the student graduates to actual performance of the work.

2. KBS for Consultation

The most successful KBSs have been developed for consultation applications; if the applications are appropriate the technical risk is minimal. Among the benefits of developing a KBS for this purpose are the codifying of a body of knowledge; supplying that specialized knowledge to many users when it is needed; alerting users to important factors in making decisions; furnishing explanations; and providing easy ways to extend the system by expanding its knowledge base. Finally, the development of consultation systems may be an effective step toward developing autonomous systems.

3. KBS in Autonomous Operations

In general, fully autonomous systems will offer the highest benefits, but pose the biggest risks. The most realistic use of autonomous operations will be in performing certain subtasks to assist a user who evaluates and coordinates the complete operation. We shall now discuss several applications in detail.

B. MGI

1. Training

MGI is currently under development; its techniques will undergo frequent change and it is likely ultimately to employ large numbers of analysts. Their ability to deliver standard, accurate products efficiently, using the most up-to-date approaches, would be significantly enhanced if they could be trained in a standardized manner, and if new techniques and procedures could be introduced rapidly.

The analysis and synthesis skills required to develop, maintain and use MGI databases vary considerably; these skills require knowledge of specialized subject areas (such as geology and vegetation), as well as of the procedures required to develop specific MGI products. A KBS might be developed for several possible applications, ranging from

a complete system, designed to test and enhance a student's expertise in specific MGI-related skills and to introduce him to the necessary detailed procedures, to specialized systems designed to update a working analyst in an unfamiliar field.

The risk involved in this application is proportional to the amount of nonsystematic specialized knowledge required. For example, a system designed to teach geology to a novice would likely entail a very high technical risk, while one aimed at instructing a geologist in the application of his skills to MGI tasks would entail a much lower risk. Since detailed procedural guides [43, 70, 68, 55] are available for many operations, a training system for introducing an analyst to MGI procedures would involve the least risk.

2. Creation of MGI Products

An MGI consultation program might come in many forms: as an extension of a training system that would lead a user through an analysis procedure; as an intelligent assistant that could keep track of references, collateral materials, and such items as image chips or the representation of terrain features on a map; as a mixed initiative program that could tailor MGI procedures to specific task requirements. Such systems promise moderate to high payoffs both for analysis tasks, in which factor overlays are produced from different input sources, and for synthesis tasks in which factors are used to create complex topic graphics that provide such information as the speed of cross-country movements through a given region.

Many synthetic MGI products are developed in response to specific requirements (e.g., determining fields of fire in a particular area) and may not require all the output generated by standard procedures. Products for such applications might be developed by means of special procedures or subsequences of standard procedures, that could be determined at the time the synthesis was performed. A consultation system with ready access to on-line information sources and with the ability to fashion such special-purpose procedures could greatly reduce the time required to develop MGI products.

The technical risk for such a system could be low for two reasons: a relatively simple system is within the grasp of current KBS technology--an "optimizing" system capable of creating special-purpose procedures appears possible with a relatively small investment; the systematic knowledge needed by such systems appears to be readily available.

3. Automatic Operation

The automatic production of MGI-analysis products, while promising substantial benefits, appears to be well beyond the current state of the AI art. However, the production of synthesis products might be well within the current AI capabilities.

The generation of complex factor overlays for computing synthetic data such as the speed of cross-country movement, could be partially or fully automated with techniques already available. This would enable analysts to obtain accurate information much more rapidly than is now possible, as well as to provide intermediate information for use by a more sophisticated program that for example, might handle the entire planning of a mission. In addition, automatic approaches could allow the enrichment of "standard" MGI products with up-to-date aerial reconnaisance information or other data. This would allow the MGI product to more faithfully represent dynamic aspects of environmental features beyond the scope of current capabilities.

C. DLMS

1. Training

Perhaps the greatest benefit in applying KBS techniques to problems associated with DLMS production will be obtained in the area of training. The requirements include standardized instruction of new analysts and the occasional reorientation of working analysts. The breadth of geographic areas encompassed by DLMS means that analysts will tend to spend a period of six to eight weeks familiarizing themselves and working with a particular region, and then have to switch to a

totally different geographic area. A training system that could familiarize the analyst with key features of the new area could reduce his transition time considerably. A more sophisticated program could compare characteristics of the region just completed with those of the new area and focus attention on the major differences.

In the near future, the DLMS data base is likely to be expanded to allow the simulation of sensors other than radar. A knowledge-based training system that could incorporate general sensor characteristics would help analysts adapt quickly to these new applications.

Once again, it is evident that the technical risks involved in attempting to develop these KBS applications will depend both on the size and breadth of the knowledge base required and on the degree to which it has already been systematized. Building a KBS for the training of analysts would entail moderate risk, primarily because the current course materials appear to lack the essential structure. While most analysts use logical, systematic procedures, these appear to be a result of experience rather than instruction. A KBS for reorienting an analyst to a new geographic area would represent a somewhat lower risk, as the knowledge base, while necessarily broad, would be compiled as a collection of locally coherent, geographically oriented subbases. Once such a system has been constructed, its refinement to provide only difference information would entail relatively minor extensions.

It appears that the highest technical risk would be connected with development of a training system for analysts who want to extend DLMS to handle additional types of sensors.

2. DLMS Consultation

A consultation system for DLMS analysts appears to offer relatively small benefits. This is due primarily to two reasons: working analysts become very familiar with their geographic areas and do not need frequent access to external resources; analysts are seldom unable to identify essential characteristics of cultural or cartographic

features in their imagery (six or seven such cases in two months is a high figure according to one analyst). Apparently the most useful application of KBS methodology in a consultation role would be to reorient a user to new areas, as previously described.

3. Automatic Generation of DLMS Products

The ability automatically to generate accurate DLMS products would free a large number of analysts for other tasks. In principle, a knowledge-based automatic system could specialize in the production of a variety of specific-sensor-oriented products.

The technical risk of developing such a system now is high, as it requires image-processing capabilities that are beyond the current state of the art. The most promising approach to such an automatic system would be to develop a prescreening program that would classify those areas of the image for which it was able to make unambiguous decisions. Regions of the image that it could not identify with confidence would be indicated to the user, who would then assume responsibility for conclusive identification. As techniques become available for recognizing relevant features automatically, they could be introduced incrementally into the system.

A KBS approach would facilitate this incremental expansion of automatic capabilities. Furthermore, a KBS would allow the user to query program in the case of error, and to understand (and possibly repair) the source of difficulty.

D. VO

The development of the VO product appears to require simpler analytical capabilities than does DLMS. In particular, VO seems to represent a problem more amenable to visual search techniques, rather than to logical analysis. For this reason, we feel that the benefits of applying KBS techniques to creation of the VO data base will be fairly small.

E. Automated PI/Cartographic Workstations (CAPIR)

1. Training

While the advantages of a KBS approach for training cartographers are likely to be insubstantial at the present time, real benefits are likely to accrue as more complex equipment and computer software become part of the mapmaker's inventory of tools. CAPIR represents one of the latest in a series of efforts to provide computer automation tools for the cartographer and photo interpreter. Major uses for the CAPIR work stations include support of MGI and DLMS tasks. As such automated work stations as CAPIR continue to be developed, the need for intelligent on-line training systems will become increasingly acute.

The adoption of a KBS approach to these systems offers several benefits: (1) training styles can be standardized for all such systems, making transition from one to another smoother and more efficient; (2) the training system, if constructed so as to provide only instruction and advice relevant to the analyst's current state of knowledge, could also comprise a general user interface, ready to supply information about system operation as needed; (3) a sophisticated KBS could quietly monitor the user's operations and provide advice about improving interaction efficiency, and ultimately even furnish data for use in redesigning the system to enhance its man/machine capabilities.

Using a knowledge-based training system would help familiarize the user with facilities of a type likely to become more prevalent as time passes. Such a system would easily be expanded and modified to handle changes in the capabilities of automated tools.

While the technical risk of a KBS instruction program would be very high if it were required to monitor the user, it would be relatively low for development of an initial instruction interface. Such an interface could be constructed in a local, incremental fashion and expanded interactively along with the increase in capabilities of the overall system. Furthermore, it is appropriate that such a development should commence now. This would help ensure that the

training systems will keep pace with the development of the tools. It would also help guarantee that the "hooks" for inserting such aids into future systems will have been included from the very beginning.

2. Consultation Systems for Advanced PI Tools

A system to assist users in obtaining maximum utility from advanced tools will be a practical necessity. Such systems could take a number of forms, including: (1) an aid to following the procedures for using the tool itself (this would be a variant of instruction systems just described); (2) aids for securing full advantages of the tool in performing specialized tasks (e.g., a program to aid the DLMS analyst using a CAPIR work station); (3) an aid for general task performance (this would really be a consultation/training aid developed for the specific task, such as an MGI consultation system supported by CAPIR.

As the complexity of PI aids increases, so will the need for online systems to support their operation. To obtain maximum result from minimum investment, development of on-line systems should begin as soon as possible. Since the requisite research and development work will produce many intermediate programs, besides which frequent modification of the tools themselves can be expected, the flexibility offered by a knowledge-based approach will return high dividends.

3. Automated PI Aids

As additional computer power is made available to automated work stations, they will be expected to handle more and more routine tasks automatically. A key capability of the analyst is his ability to make decisions on the basis of many sources of knowledge (e.g., imagery, maps, collateral data), and the success of automatic programs will also depend on such capabilities. A KBS capable of determining what information is needed for an operation, and where to obtain it, seems the best approach to this problem. As the technical risk of a completely automatic system is quite high, it must again be emphasized that the proper development strategy is to begin automating simple tasks

in an interactive KBS-mediated environment. Representative work along these lines has already begun (see, for example, [3]), and constitutes the focus of a major DARPA research program in image understanding that is being tracked by DMA. The results of these efforts should be integrated as soon as possible and, in fact, such integration is being carried out for DMA by SRI [26].

F. Monitoring and Identification of Structures and Operations

Although we did not have the opportunity to speak directly to photo interpreters involved in the production of intelligence products, such as determining the nature of activities in remotely sensed sites, our discussions with other photo interpreters indicated that knowledge-based approaches to problems in this area could result in important new capabilities. Accordingly, we shall now touch upon these applications here.

1. Training

Analysts responsible for detecting and monitoring such activities as the operations of industrial installations in remote areas typically build up their expertise over many years. They become adept at recognizing vital or highly relevant pieces of information and at drawing conclusions concerning the nature of operations based on such information. If these specialists could impart efficiently and effectively to novices the logical processes they use in their work, the result would be greater overall efficiency and improved accuracy of intelligence products.

2. Consultation for Structure Identification

A KBS to aid analysts in identifying the nature and use of structures would appear to have high utility, particularly when the type of operations being conducted must be inferred from indirect information. For example, to determine the nature of an industrial area, one might proceed by first generally classifying the industry as

extraction (e.g., mining, oil drilling), processing (e.g., steel production), or fabrication (e.g., shipbuilding). These classifications could be made on the basis of such characteristics as the presence of piles of raw materials, power lines, cooling towers, tanks, and so forth. Further breakdown within the classes would be accomplished by similar procedures. The approach used by Prospector, which would be germane to this problem, could be extended to allow the system to show the user some representative examples of different structures instead of relying solely on verbal descriptions. While the development of such a system would likely entail a low-to-moderate technical risk, the investment required to compile the knowledge base could be fairly large.

3. Automatic Structure Identification

This application, while highly desirable, is well beyond the current state of the art. Even a fairly simple system capable of identifying a few distinctive structures reliably is not yet practicable. Nevertheless, the potential benefits, not only for analysts, but also for such autonomous systems as cruise missiles and fire-and-forget weapons, might be high enough to warrant beginning efforts to study the problem.

G. Recommendations

The two areas that seen likely to derive the greatest benefit from KBS approaches are MGI and advanced work stations, such as CAPIR. In the case of MGI, there appears to be a sizable body of knowledge available from which a consultation system could be developed. Even though such systems as CAPIR are in a relative state of infancy, it appears that a KBS could grow along with the tools themselves.

In both cases the appropriate R & D strategy for maximizing investment return and minimizing risk would be to begin compiling a knowledge base for use by a consultant program, while simultaneously developing the knowledge-based consultant system itself.

Since the results of work on an interactive consultant program have proved the utility of this approach, it would appear reasonable to begin automating some of the simpler tasks, thus increasing the power of the system. Similarly, as the knowledge base becomes more extensive, tutorial programs would be developed. Proceeding in this manner would expedite the compilation of comprehensive, expert-consultant programs capable of substantially enhancing the analyst's own capabilities.

V Conclusions

Our investigations led us to conclude that MGI and CAPIR offer the highest prospective return on the investment required for developing a specialized KBS. We suggest that the appropriate research strategy is to initiate development of a knowledge-based interactive consultant system for each of these two areas; as the utility of the approach is confirmed, some of the simpler tasks can be automated to increase the man/machine capabilities. Tutorial skills may then be added as the knowledge base is developed.

In the area of DLMS, it appears that the best application of KBS techniques would be for training new analysts, as well as for familiarizing analysts with new geographic areas. It is our feeling that the VO program is not likely to benefit much from KBS. Probably the best use of KBS technology in structure identification would be in developing an interactive consultant to guide an analyst's inference processes whenever he is unable to identify a given structure himself.

Appendix A

An Overview of Prospector

SRI International

RULE-BASED MODELING OF ORE DEPOSITS FOR MINERAL EXPLORATION

By: Richard O. Duda, Staff Scientist
Peter E. Hart, Director
Artificial Intelligence Center
Computer Science and Technology Division

November 27, 1979

For presentation at the Annual Meeting of the American Association for the Advancement of Science, San Francisco, California, January 4, 1980

This work was jointly sponsored by the U. S. Geological Survey under Contract No. 14-08-0001-15985, 14-08-0001-17226, 14-08-0001-17227, and 14-08-0001-17296, and by the National Science Foundation under Grant AER 77-04499. Any opinions, findings, conclusions, or recommendations expressed in this document are those of the authors and do not necessarily reflect the views of either the U. S. Geological Survey or the National Science Foundation.

333 Ravenswood Ave • Menlo Park, California 94025
(415) 326-6200 • Cable SRI INTL MPK • TWX 910-373-1246

A major activity in the field of economic geology is the development of models of various classes of ore deposits. Models intended for exploration describe the distribution of ore minerals and their associated petrological and structural features within a broader regional and petrotectonic context. They provide the explorationist with a target that is significantly larger than the ore body, but not so large as to lose all discriminatory power. By organizing observations into a coherent pattern, models aid the interpretation of existing data, suggest what additional data would be most valuable, and guide expensive drilling decisions. In addition, they play a major role in regional resource evaluation by defining the regional characteristics that are favorable (or unfavorable) for specific types of ore deposits.

Because earth scientists do not understand all the complex processes associated with ore deposition, much of what is known depends upon empirical observations. While many deposits have been studied intensively, the number of thoroughly explored deposits of any given type is usually rather small. Furthermore, many types of greatest current interest (such as unconformity-related vein-type uranium deposits and epithermal bulk silver deposits) have been discovered only in the past few years. Thus, the development of a model requires both scientific understanding of the physical and chemical processes of ore deposition and geological judgment based on informed experience.

Since 1976, computer scientists at SRI have been working with consulting economic geologists from the United States and Canada to encode ore deposit models in a computer program called Prospector. The primary goal of this work is to provide a field geologist who is exploring a particular site with computer-based consultation to determine such things as (a) which model best fits the available field data, (b) where the most favorable drilling sites are located, (c) what additional data would be most helpful in reaching firmer conclusions, and (d) what is the basis for these conclusions and recommendations. In performing these tasks, Prospector employs both probabilistic and logical reasoning procedures developed from research in artificial

intelligence. These procedures interpret the field data using the information encoded in formal ore deposit models. The encoded models are hierarchically organized data structures that specify the factors (or hypotheses) that must be considered (such as the petrotectonic setting, regional environment, and local zones of alteration and mineralization), the field evidence relevant to these factors, and how the evidence and the hypotheses are related.

The relations between field evidence and geological hypotheses are often probabilistic; certain observations can suggest or cast doubt on certain hypotheses without being conclusive. This situation characterizes many diagnosis problems, and artificial intelligence researchers working on problems in medical diagnosis have developed very effective "rule-based" methods for their solution. This methodology is also employed in Prospector, where the links between evidence and hypotheses are expressed by separate and distinct rules. The degree of belief in any hypothesis is measured by an associated probability value, which initially is typically quite small. The rules that link evidence to an hypothesis have associated likelihood ratio values which are used to update beliefs as evidence (which may be uncertain) is acquired.

This methodology for representing models has several advantages. The computer system that interprets the rules can be developed independently of the models. Old models can be refined and new models created merely by modifying old rules or adding new ones. The rules can be systematically examined by the program to find the missing or uncertain evidence that would be most useful in reaching firmer conclusions. Since each rule has a specific purpose, rule-based models support an explanation system that can tell the user of the program why any particular piece of evidence is important and why any particular conclusions have been reached. Finally, although the models embody the subjective judgments of their authors, their application to specific problems is done uniformly and objectively, and the basis for any conclusions is open to--indeed invites--public examination.

The current system contains nine ore deposit models and over 900 rules, together with a taxonomy of over 400 rock and mineral terms that are used to link input data to the rules. These models, their authors, and their sizes are tabulated below. In addition, special models have been developed by V. F. Hollister and A. N. Campbell to aid in the selection of drilling targets for porphyry copper and porphyry molybdenum deposits. Using these models, Prospector combines geophysical, geochemical, and geological data derived from digitized exploration maps to produce an output map that indicates the most favorable drilling sites. While many more models must be developed and encoded to span the spectrum of deposits of economic interest, the feasibility of the approach has been amply demonstrated.

Deposit Type	Author	Number of Rules
Mississippi-Valley-type lead/zinc	N. Campbell	20
Kuroko-type massive sulfide	C. F. Park, Jr.	34
Komatiitic-type nickel sulfide	A. J. Naldrett	72
Butte-type porphyry copper	M. T. Einaudi	104
Yerington-type porphyry copper	M. T. Einaudi	143
Island-arc-type porphyry copper	D. Cox	116
Roll-front sandstone uranium	R. I. Rackley	139
Regional roll-front uranium	R. I. Rackley	132
Grants-type sandstone uranium	S. S. Adams	141

In addition to the applications mentioned above, we believe that Prospector can contribute significantly to the science of economic geology. In particular, once ore deposit models are encoded with the precision required by a computer program, it becomes possible to test them quantitatively and to make objective comparisons between competing models. We have completed initial quantitative testing of several of our models to determine the degree to which the encoding reflects their authors' intentions. While it is not possible to summarize those numerical results here in a way that is both precise and meaningful, we can say that in most instances the evaluations produced have been very close to those supplied independently by the authors. Detailed information on the system, the models, and these tests is given in the references.

REFERENCES

1. R. O. Duda, P. E. Hart, P. Barrett, J. G. Gaschnig, K. Konolige, R. Reboh and J. Slocum, "Development of the Prospector Consultation System for Mineral Exploration," Final Report, SRI Projects 5821 and 6415, SRI International, Menlo Park, California (October 1978).

2. R. O. Duda, P. E. Hart, K. Konolige and R. Reboh, "A Computer-Based Consultant for Mineral Exploration," Final Report, SRI Project 6415, SRI International, Menlo Park, California (September 1979).

3. J. Gaschnig, "The Development of Uranium Exploration Models for the Prospector Consultant System," Final Report, SRI Project 7856 (December 1979).

Appendix B

Perceptual Reasoning

PERCEPTUAL REASONING IN A HOSTILE ENVIRONMENT*

Thomas D. Garvey and Martin A. Fischler

Artificial Intelligence Center
SRI International, Menlo Park, CA 94025

ABSTRACT

The thesis of this paper is that perception requires reasoning mechanisms beyond those typically employed in deductive systems. We briefly present some arguments to support this contention, and then offer a framework for a system capable of perceptual reasoning, using sensor-derived information, to survive in a hostile environment. Some of these ideas have been incorporated in a computer program and tested in a simulated environment; a summary of this work and current results are included.

I INTRODUCTION

Living organisms routinely satisfy critical needs such as recognizing threats, potential mates, food sources, and navigable areas, by extracting relevant information from huge quantities of data assimilated by their senses. How are such "relevant" data detected?

We suggest that a reasoning approach that capitalizes on the goal-oriented nature of perception is necessary to define and recognize relevant data. Perception can be characterized as imposing an interpretation on sensory data, within a context defined by a set of loosely specified models. The ability to select appropriate models and match them to physical situations appears to require capabilities beyond those provided by such "standard" paradigms as logical deduction or probabilistic reasoning.

The need for extended reasoning techniques for perception is due to certain critical aspects of the problem, several of which we summarize here:

- The validity of a perceptual inference (interpretation) is determined solely by the adequacy of the interpretation for successfully carrying out some desired interaction with the environment (as opposed to verification within a "closed" formal axiomatic system).
- Since it is impossible to abstractly model the complete physical environment, the degree to which purely abstract reasoning will be satisfactory is limited. Instead, perception requires tight interaction between modeling/hypothesizing, experimenting (accessing information from the environment), and reasoning/verifying.
- Reasoning processes that embody concepts from physics, geometry, topology, causation, and temporal and spatial ordering are critical components of any attempt to "understand" an ongoing physical situation. Explicit representations appropriate to these concepts are necessary for a perceptual system that must provide this understanding. These representations are incommensurate and it is not reasonable to attempt to force them into a single monolithic model.
- There is typically no single, absolutely correct interpretation for sensory data. What is necessary is a "maximally consistent" interpretation, leading to the concept of perception as an optimization problem [1, 2] rather than a deductive problem.

Research in perception and image processing at SRI and elsewhere has addressed many of these issues. An early effort focused upon the goal-directed aspect of perception to develop a program capable of planning and executing special-purpose strategies for locating objects in office scenes [3]. Research addressing interpretation as an optimization problem includes [1, 2, 4]. Current research on an expert system for image interpretation [5] has also considered the strategy-related aspects of determining location in situations involving uncertainty.

The most recent work (at SRI) on perceptual reasoning has addressed the problem of assessing the status of a hostile air-defense environment on the basis of information received from a variety of controllable sensors [6]. This work led us to attempt to formulate a theory of perceptual reasoning that highlighted explicit reasoning processes and that dealt with those aspects of perception just described. In the following section, we will use this work as a vehicle to illustrate a paradigm for perceptual reasoning.

---

* This work was supported by the Defense Advanced Research Projects Agency under Contracts MDA903-79-C-0588, F33615-77-C-1250, and F33615-80-C-1110.

** This paper was accepted for presentation at the First Annual National Conference on Artificial Intelligence, sponsored by The American Association for Artificial Intelligence, to be held August 19-21, 1980 at Stanford, California.

II PERCEPTUAL REASONING IN A SURVIVAL SITUATION

The specific problem addressed was to design a system able to interpret the disposition and operation (i.e., the order of battle or OB) of hostile air-defense units, based on information supplied by sensors carried aboard a penetrating aircraft [6]. The situation may be summarized as follows. A friendly aircraft is faced with the task of penetrating hostile airspace en route to a target behind enemy lines. Along the way, the aircraft will be threatened by a dense network of surface-to-air missiles (SAMs) and antiaircraft artillery (AAAs). The likelihood of safe penetration and return is directly related to the quality of acquired or deduced information about the defense systems.

Partial information is furnished by an initial OB, listing known threats at, say, one hour before the flight. Additional knowledge is available in the form of descriptions of enemy equipment, typical deployments, and standard operating procedures. Since the prior OB will not be completely accurate, the information must be augmented with real-time sensory data. The OB forms the starting point for this augmentation.

The explicit goal of the overall system is to produce and maintain an accurate OB, detecting and identifying each threat prior to entering its lethal envelope. The density of threats means that this goal will result in conflicting subgoals, from which selection must then be made to ensure that critical data will be received. This must be accomplished by integrating data from imperfect sensors with prior knowledge. The paradigm that was developed for this task is summarized below:

(1) Available knowledge is used to create an hypothesized OB that anticipates the developing situation.

(2) A plan that attempts to allocate sensors to detect or verify the presence of threats, in an optimal way, is constructed. Sensors are then allocated and operated.

(3) Information returned from the sensors is interpreted in the context established by the anticipated situation. Interpretations modify the current OB, and the process is iterated.

We will briefly discuss each of these steps.

A. Anticipation (Hypothesis Formation; Model Selection)

In the anticipation phase, the system calls upon a variety of types of internally available knowledge to hypothesize the current OB. The specific goal of this step is to determine where critical knowledge is lacking and to produce a set of information requests; these will become goals for the resource allocation procedure in the experimental planning phase. The primary knowledge source is the initial OB, which lists the geographic locations of all threats known prior to flight time. Since these threat systems are mobile and the situation is changing rapidly, the system determines probable new locations for each threat by comparing data about the mobility of weapon vehicles with a terrain map of the area and computing "uncertainty regions" likely to contain the threats.

Partial information about systems of threats (i.e., collections of weapons deployed in mutually supporting ways) is used by a program that performs "geometric reasoning" to predict the location of missing components. The anticipation phase attempts to reduce the difficult problem of detecting all threats to the simpler problem of verifying hypothesized threats.

B. Experimentation (Accessing Information from the Environment)

The goal of this step is to access information needed to detect or verify the presence of threats inferred in the anticipation step, but not available in the "internal" knowledge base of the system. In general, it might be necessary to define and execute one or more experiments to extract this needed information from the environment. In the more limited context of model instantiation by "passive" sensing, the problem reduces to that of allocating sensor resources to maximize the overall utility of the system; sensing is a specific instance of the more general process of experimentation.

First the list of data-acquisition goals is ordered, based on the current state of information about each threat and its lethality. The allocator attempts to assign (a time sliced segment of) a sensor to satisfy each request based on the expected performance of the sensor for that task.

Sensor detection capabilities are modeled by a matrix of conditional probabilities. These represent the likelihood that the sensor will correctly identify each threat type, given that at least one instance thereof is in the sensor's field of view. This matrix represents performance under optimal environmental conditions (for the sensor) and is modified for suboptimal conditions by means of a specialized procedure. This representation is compact and circumvents the need to store complete, explicit models describing sensor operation in all possible situations. Similar models describe each sensor's identification and location capabilities.

The sensor models are used to compute the utility of allocating each sensor to each of the highest priority threats. These utilities form the basis for the final allocation, which is carried out by a straightforward optimization routine. At the same time, the program determines how the sensor should be directed (for example, by pointing or tuning). Appropriate control commands are then sent to the simulated sensors.

C. Interpretation (Hypothesis Validation; Model Instantiation)

In this phase, the program attempts to interpret sensor data in the context of threats that were anticipated earlier. It first tries to

determine whether sensor data are consistent with specifically anticipated threats, then with general weapon types expected in the area. Since sensor data are inherently ambiguous (particularly if environmental conditions are suboptimal), this step attempts to determine the most likely interpretation.

Inference techniques used for interpretation include production rule procedures, probabilistic computations, and geometric reasoning. Production rules are used to infer probable weapon operation (e.g., target tracking, missile guidance), on the basis of such information as past status, environmental conditions, and distance from the aircraft. Probabilistic updating of identification likelihoods is based on the consistency of actual sensor data with expected data, and on agreement (or disagreement) among sensors with overlapping coverage. Geometric reasoning introduces a concept of global consistency to improve identification by comparing inferred identifications and locations of threat system components with geometric models of typical, known system deployments.

The interpretation phase brings a great deal of a priori knowledge to bear on the problem of determining the most likely threats the sensors are responding to. This results in much better identifications than those produced by the sensors alone. Confident identifications are entered into the OB and the entire process is continued.

D. Performance

An experimental test of the system, using a simulated threat environment, allowed a comparison between two modes of operation--an "undirected" mode and one based on perceptual reasoning. A scoring technique that measured the effectiveness with which the system detected, identified, and located hostile systems in a timely fashion was used to grade performance. The ability of the perceptual reasoning system to use external knowledge sources effectively, and to integrate information from multiple sensors, produced superior capabilities under this measure. These capabilities showed themselves even more prominently in situations where environmental conditions tended to degrade sensor performance, rendering it critical that attention be focused sharply.

III DISCUSSION

Our approach to perceptual reasoning suggests that the problem of perception actually involves the solution of a variety of distinct types of subproblems, rather than repeated instances of the same general problem. The system we described utilizes a nonmonolithic collection of representations and reasoning techniques, tailored to specific subproblems. These techniques include both logical deduction and probabilistic reasoning approaches, as well as procedures capable of geometric reasoning and subjective inference.

We have discussed several key aspects of the general problem of perceptual reasoning, including the assertion that perception is goal oriented, and inductive and interpretative rather than deductive and descriptive; that because complete modeling of the physical world is not practical, "experimentation" is a critical aspect of perception; and finally, that multiple representations and corresponding reasoning techniques, rather than a single monolithic approach, are required.

The specific system discussed above constitutes an attempt to address the reasoning requirements of perception in a systematic way and, to our knowledge, represents one of the few attempts to do so. While systems that truly interact with the physical world in an intelligent manner will certainly assume a variety of forms, we believe they will all ultimately have to resolve those aspects of the problem that have been described here.

REFERENCES

1. M. A. Fischler and R. A. Elschlager, "The Representation and Matching of Pictorial Structures," IEEE Transactions on Computers, Vol. C-22, No. 1, pp. 67-92 (January 1973)

2. H. G. Barrow and J. M. Tenenbaum, "MSYS: A System for Reasoning About Scenes," Artificial Intelligence Center Technical Note No. 121, SRI International, Menlo Park, California (1976)

3. T. D. Garvey, "Perceptual Strategies for Purposive Vision," Artificial Intelligence Center Technical Note No. 117, SRI International, Menlo Park, California (1976)

4. A. Rosenfeld, "Iterative Methods in Image Analysis," Pattern Recognition, Vol. 10, No. 4, pp. 181-187 (1978)

5. R. C. Bolles et al., "The SRI Road Expert: Image-to-Database Correspondence," Proc. DARPA Image Understanding Workshop, Science Applications, Inc., Report No. SAI-79-814-WA, pp. 163-174 (November 1978)

6. T. D. Garvey and M. A. Fischler, "Machine-Intelligence-Based Multisensor ESM System," Technical Report AFAL-TR-79-1162, Air Force Avionics Laboratory, Wright-Patterson AFB, OH 45433 (Final Report DARPA Project 3383) (October 1979)

REFERENCES

1. G. J. Agin, "Computer Vision Systems for Industrial Inspection and Assembly," Computer, Vol. 13, pp. 11-20 (May 1980).

2. J. A. Barnett and M. I. Bernstein, "Knowledge-Based Systems: A Tutorial," Report No. TM-(L)-5903/000/00, Contract No. MDA 904-76-C-0343, System Development Corporation, Santa Monica, California (June 1977).

3. H. G. Barrow et al., "Interactive Aids for Cartography and Photo Interpretation," Progress Report, October 1977, Proc. Image Understanding Workshop, pp. 111-127 (October 1977).

4. D. R. Barstow, "An Experiment in Knowledge-based Automatic Programming," Artificial Intelligence, Vol. 12, pp. 7-119 (August 1979).

5. L. S. Baumann, ed., Proc. Image Understanding Workshop, Report Number SAI-81-170-WA, Science Applications, Inc., Arlington, Virginia, April 1980.

6. J. S. Bennett and R. S. Engelmore, "SACON: A Knowledge-Based Consultant for Structural Analysis," Proc. Sixth International Joint Conference on Artificial Intelligence, pp. 47-49, Tokyo, Japan (August 20-23 1979).

7. D. G. Bobrow and T. Winograd, "An Overview of KRL, a Knowledge Representation Language," Cognitive Science, Vol. 1, pp. 3-46 (January 1977).

8. R. J. Brachman, "What's in a Concept: Structural Foundations for Semantic Networks," BBN Report No. 3433, Bolt Beranek and Newman Inc., Cambridge, Massachusetts (October 1976).

9. R. A. Brooks, R. Greiner and T. O. Binford, "The ACRONYM Model-Based Vision System," Proc. Sixth International Joint Conference on Artificial Intelligence, pp. 105-113, Tokyo, Japan (August 20-23 1979).

10. J. S. Brown, R. R. Burton, and A. G. Bell, "SOPHIE: A Sophisticated Instructional Environment for Teaching Electronic Troubleshooting (An Example of AI in CAI)," Final Report, Contract F41609-73-C-006, Bolt Beranek and Newman Inc., Cambridge, Massachusetts (March 1974).

11. B. G. Buchanan and T. M. Mitchell, "Model-Directed Learning of Production Rules," in Pattern-Directed Inference Systems, D. A. Waterman and F. Hayes-Roth, eds., pp.297-312 (Academic Press, New York, 1978).

12. A. Bundy et al., "Solving Mechanics Problems Using Meta-Level Inference," in Expert Systems in the Micro Electronic Age, D. Michie, ed., pp. 50-64 (Edinburgh University Press, Edinburgh, Scotland, 1979).

13. J. R. Carbonell, "AI in CAI: An Artificial-Intelligence Approach to Computer-Aided Instruction," IEEE Trans. Man-Machine Systems, Vol. MMS-11, pp. 190-202 (December 1970).

14. W. J. Clancey, E. H. Shortliffe, and B. G. Buchanan, "Intelligent Computer-Aided Instruction for Medical Diagnosis," Proc. 3rd Computer Application in Medical Care, pp. 175-183 (1979).

15. L. S. Coles, et al., "Methodologies for Heuristic Modeling," Final Report, SRI Project 3372, Contract 8719/62-8033-74WR, Stanford Research Institute, Menlo Park, California (April 1975).

16. R. Davis and J. King, "An Overview of Production Systems," in Machine Intelligence 8, E. W. Elcock and D. Michie, eds., pp. 300-332 (John Wiley and Sons, Inc., New York, 1976).

17. R. Davis and B. G. Buchanan, "Meta-Level Knowledge: Overview and Applications," Proc. Fifth International Joint Conference on Artificial Intelligence, pp. 920-927, Cambridge, Massachusetts (August 22-25 1977).

18. R. Davis, "Interactive Transfer of Expertise: Acquisition of New Inference Rules," Artificial Intelligence, Vol. 12, pp. 121-157 (August 1979).

19. T. G. Dietterich and R. S. Michalski, "Learning and Generalization of Characteristic Descriptions," Technical Report, Department of Computer Science, University of Illinois, Urbana, Illinois (1979).

20. R. O. Duda, P. E. Hart, N. J. Nilsson, and G. L. Sutherland, "Semantic Network Representations in Rule-Based Inference Systems," in Pattern-Directed Inference Systems, D. A. Waterman and F. Hayes-Roth, eds., pp. 203-221 (Academic Press, New York, 1978).

21. R. Duda, J. Gaschnig, and P. Hart, "Model Design in the Prospector Consultant System for Mineral Exploraion," in Expert Systems in the Micro Electronic Age, D. Michie, ed., pp. 153-167 (Edinburgh University Press, Edinburgh, Scotland, 1979).

22. R. O. Duda, "The Prospector System for Mineral Exploration," Final Report, SRI Project 8172, Artificial Intelligence Center, SRI International, Menlo Park, California (April 1980).

23. R. Engelmore and A. Terry, "Structure and Function of the Crysalis System," Proc. Sixth International Joint Conference on Artificial Intelligence, pp. 250-256, Tokyo, Japan (August 20-23 1979).

24. E. A. Feigenbaum, B. G. Buchanan, and J. Lederberg, "On Generality and Problem Solving: a Case Study Using the DENDRAL Program," in Machine Intelligence 6, B. Meltzer and D. Michie, eds., pp. 165-190 (Edinburgh University Press, Edinburgh, Scotland, 1971).

25. E. A. Feigenbaum, "Themes and Case Studies of Knowledge Engineering," in Expert Systems in the Micro Electronic Age, D. Michie, ed., pp. 3-25 (Edinburgh University Press, Edinburgh, Scotland, 1979).

26. M. A. Fischler et al., "Detection of Roads and Linear Structures in Aerial Imagery," Proc. Image Understanding Workshop, pp. 87-100, (November 1979).

27. T. D. Garvey and M. A. Fischler, "Machine-Intelligence-Based Multisensor ESM System," Technical Report AFAL-TR-79-1162, SRI International, Menlo Park, California (October 1979).

28. J. Gaschnig, "Development of Uranium Exploration Moels for the Prospector Consultant System," Final Report, SRI Project 7856, Artificial Intelligence Center, SRI International, Menlo Park, California (March 1980).

29. H. L. Gelernter et al., "Empirical Explorations of SYNCHEM," Science, pp. 1041-1049 (September 9, 1977).

30. R. Goldberg and S. M. Weiss, "An Experimental Transformation of a Large Expert System Knowledge-Base," Working Paper, Department of Computer Science, Rutgers University, New Brunswick, New Jersey (March 1980).

31. I. Goldstein and S. Papert, "Artificial Intelligence, Language, and the Study of Knowledge," Cognitive Science, Vol. 1, pp. 84-123 (January 1977).

32. G. A. Gorry, "Research on Expert Systems," MAC Technical Memorandum 56, Project MAC, Massachusetts Institute of Technology, Cambridge, Massachusetts (December 1974).

33. C. Green and B. Raphael, "The Use of Theorem Proving Techniques in Question Answering Systems," Proc. ACM 23rd Nat. Conf., pp. 169-181 (Brandon Systems Press, Princeton, New Jersey, 1968).

34. A. Hansen and E. Riseman, eds., Computer Vision Systems (Academic Press, New York, 1978).

35. L. D. Harmon, "Automatic Recognition of Print and Script," Proc. IEEE, Vol. 60, pp. 1165-1176 (October 1972).

36. P. E. Hart, R. O. Duda and M. T. Einaudi, "Computer-Based Consultation System for Mineral Exploration," in Computer Methods for the 80's, Alfred Weiss, ed., pp. 127-140 (American Institute of Mining, Metallurgical, and Petroleum Engineers, New York, 1979).

37. F. Hayes-Roth, D. A. Waterman, and D. B. Lenat, "Principles of Pattern-Directed Inference Systems," in Pattern-Directed Inference Systems, D. A. Waterman and F. Hayes-Roth, eds., pp. 577-601 (Academic Press, New York, 1978).

38. G. G. Hendrix, "Encoding Knowledge in Partitioned Networks," in Associative Networks--The Representation and Use of Knowledge in Computers, N. V. Findler, ed., pp. 51-92 (Academic Press, New York, 1979).

39. D. E. Knuth, The Art of Computer Programming: Fundamental Algorithms (Addison Wesley, New York, 1969).

40. D. Lenat, "Automated Theory Formation in Mathematics," Proc. Fifth International Joint Conference on Artificial Intelligence, pp. 833-842, Cambridge, Massachusetts (August 22-25 1977).

41. V. R. Lesser and L. D. Erman, "A Retrospective View of the HEARSAY-II Architecture," Proc. Fifth International Joint Conference on Artificial Intelligence, pp. 790-800, Cambridge, Massachusetts (August 22-25 1977).

42. G. E. Lukes, "Computer-Assisted Photo Interpretation Research (CAPIR): A Prospectus," Research Institute, U. S. Army Engineer Topographic Laboratories, Fort Belvoir, Virginia 22060 (November 1979).

43. J. A. Messmore, T. C. Vogel and A. R. Pearson, "Terrain Analysis Procedural Guide for Vegetation," ETL-0178, U. S. Army Corps of Engineers, Engineer Topographic Laboratories, Fort Belvoir, Virginia 22060 (March 1979).

44. D. Michie, ed., Expert Systems in the Micro Electronic Age (Edinburgh University Press, Edinburgh, Scotland, 1979).

45. M. Minsky, "Steps Toward Artificial Intelligence," Proc. IRE, Vol. 49, pp. 8-30 (January 1961).

46. M. Minsky, "A Framework for Representing Knowledge," in The Psychology of Computer Vision, P. H. Winston, ed., pp. 211-277 (McGraw-Hill Book Co., New York, 1975).

47. A. Newell and H. A. Simon, Human Problem Solving (Prentice-Hall, Englewood Cliffs, New Jersey, 1972).

48. H. P. Nii and E. A. Feigenbaum, "Rule-Based Understanding of

Signals," in Pattern-Directed Inference Systems, D. A. Waterman and F. Hayes-Roth, eds., pp. 483-501 (Academic Press, New York, 1978).

49. H. P. Nii and N. Aiello, "AGE (Attempt to Generalize): A Knowledge-Based Program for Building Knowledge-Based Programs, Proc. Sixth International Joint Conference on Artificial Intelligence, pp. 645-655, Tokyo, Japan (August 20-23 1979).

50. N. J. Nilsson, Problem Solving Methods in Artificial Intelligence (McGraw-Hill Book Co., New York, 1971).

51. N. J. Nilsson, Principles of Artificial Intelligence (Tioga Publishing Co., Palo Alto, California, 1980).

52. T. O'Shea, "Intelligent Systems in Education," in Expert Systems, pp. 9/1-9/28, 69th Infotech State of the Art Conference, London, England (April 16-18 1980).

53. S. G. Pauker, G. A. Gorry, J. P. Kassirer, and W. B. Schwartz, "Towards the Simulation of Clinical Cognition," American Journal of Medicine, Vol. 60, pp. 981-996 (June 1976).

54. A. R. Pearson, T. C. Vogel, and J. S. Wright, "An Integrated Terrain Analysis System," U. S. Army Engineer Topographic Laboratories, Fort Belvoir, Virginia.

55. A. R. Pearson and J. S. Wright, "Synthesis Guide for Cross-Country Movement," Report No. 4 in ETL series, Guides for Army Terrain Analysts, U. S. Army Corps of Engineers, Engineer Topographic Laboratories, Fort Belvoir, Virginia (June 1979).

56. H. E. Pople, Jr., et al., "DIALOG: A Model of Diagnostic Logic for Internal Medicine,"Proc. Fourth International Joint Conference on Artificial Intelligence, pp. 848-855, Tbilisi, Georgia, USSR (September 1975).

57. K. Preston, Jr., "Biomedical Image Processing," in Advances in Digital Image Processing, P. Stucki, ed., pp. 125-146 (Plenum Press, New York, 1979).

58. R. Reboh, "The Knowledge Acquisition System," in "A Computer-Based Consultant for Mineral Exploration," R. O. Duda et al., Final Report, SRI Project 6415, Artificial Intelligence Center, SRI International, Menlo Park, California (September 1979).

59. D. R. Reddy, "Speech Recognition by Machine: A Review," Proc. IEEE, Vol. 64, pp. 501-531 (April 1976).

60. R. B. Roberts and I. P. Goldstein, "The FRL Primer," Memo 408, MIT Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts (1977).

61. J. A. Robinson, "A Machine-Oriented Logic Based on the Resolution Principle," J. ACM, Vol. 12, pp. 23-41 (January 1965).

62. S. Rosenberg, "An Intelligent Support System for Energy Resources in the United States," Proc. Third Biennial Conference of the Canadian Society for Computational Studies of Intelligence, pp. 34-40, Victoria, British Columbia, Canada (May 14-16 1980).

63. E. H. Shortliffe, Computer-Based Medical Consultations: MYCIN (American Elsevier, New York, 1976).

64. R. M. Stallman and G. J. Sussman, "Forward Reasoning and Dependency-Directed Backtracking in a System for Computer-Aided Circuit Analysis," Artificial Intelligence, Vol. 9, pp. 135-196 (1977).

65. J. L. Stansfield, "COMEX: A Support System for a Commodities Expert," Memo No. 423, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts (August 1977).

66. M. J. Stefik, "Inferring DNA Structures from Segmentation Data," Artificial Intelligence, Vol. 11, pp. 85-114 (August 1978).

67. M. J. Stefik, "An Examination of a Frame-Structured Representation System," Proc. Sixth International Conference on Artificial Intelligence, pp. 845-852, Tokyo, Japan (August 1979).

68. J. Tazelaar, "Terrain Analysis Procedural Guide for Geology," ETL-0207, U. S. Army Corps of Engineers, Engineer Topographic Laboratories, Fort Belvoir, Virginia (November 1979).

69. W. van Melle, "A Domain-Independent Production-Rule System for Consultation Programs," Proc. Sixth International Joint Conference on Artificial Intelligence, pp. 923-925, Tokyo, Japan (August 20-23 1979).

70. T. C. Vogel, "Terrain Analysis Procedural Guide for Roads and Related Structures," ETL-0205 U. S. Army Corps of Engineers, Engineer Topographic Laboratories, Fort Belvoir, Virginia (October 1979).

71. D. A. Waterman, "User-Oriented Systems for Capturing Expertise: A Rule-Based Approach," in Expert Systems in the Micro Electronic Age, D. Michie, ed., pp. 26-34 (Edinburgh University Press, Edinburgh, Scotland, 1979).

72. S. Weiss, C. Kulikowski, and A. Safir, "Glaucoma Consultation by Computer," Computers in Biology and Medicine, Vol. 8, pp. 25-40 (1978).

73. S. M. Weiss and C. A. Kulikowski, "EXPERT: A System for Developing Consultation Models," Proc. Sixth International Joint Conference on Artificial Intelligence, pp. 923-925, Tokyo, Japan (August 20-23 1979).

74. J. R. Welch, "Automatic Speech Recognition--Putting it to Work in Industry," Computer, Vol. 13, pp. 65-73 (May 1980).

75. W. T. Wipke, "Computer Planning of Research in Organic Chemistry," in Computers in Chemical Education and Research, E. V. Ludena, N. H. Sabelli, and A. C. Wahl, eds., pp. 381-391 (Plenum Press, New York 1976).

DAT
FILM