

Applications of Robotics and Artificial Intelligence to Reduce Risk and Improve Effectiveness

By National Research Council

Get any book for free on: www.Abika.com

Contents

[Acknowledgements and Contents](#)

1. [Background](#)
 2. [Summary of the Technology](#)
 3. [Criteria for Selection of Applications](#)
 4. [Recommended Applications and Priorities](#)
 5. [Implementation of Recommended Applications](#)
 6. [Other Considerations](#)
 7. [Recommendations](#)
-
- [Appendix: State of the Art and Predictions for Artificial Intelligence and Robotics](#)
 - [Glossary of Acronyms](#)

APPLICATIONS OF ROBOTICS AND ARTIFICIAL INTELLIGENCE
TO REDUCE RISK AND IMPROVE EFFECTIVENESS

A Study for the United States Army

Committee on Army Robotics and Artificial Intelligence

Manufacturing Studies Board

Commission on Engineering and Technical Systems

National Research Council

NATIONAL ACADEMY PRESS Washington, D.C. 1983

NOTICE: The project that is the subject of this report was approved by the Governing Board of the National Research Council, whose members are drawn from the councils of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine. The members of the committee responsible for the report were chosen for their special competences and with regard for appropriate balance.

This report has been reviewed by a group other than the authors according to procedures approved by a Report Review Committee consisting of members of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine.

The National Research Council was established by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purpose of furthering knowledge and of advising the federal government. The Council operates in accordance with general policies determined by the Academy under the authority of its congressional charter of 1863, which establishes the Academy as a private, nonprofit, self-governing membership corporation. The Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in the conduct of their services to the government, the public, and the scientific and engineering communities. It is administered jointly by both Academies and the Institute of Medicine. The National Academy of Engineering and the Institute of Medicine were established in 1964 and 1970, respectively, under the charter of the National Academy of Sciences.

This report represents work under contract number MDA 903-82-C-0351 between the U.S. Department of the Army and the National Academy of Sciences.

A limited number of copies are available from:
Manufacturing Studies Board

Get any book for free on: www.Abika.com

National Academy of Sciences

2101 Constitution Avenue, N.W.

Washington, D.C. 20418

Printed in the United States of America

ii

COMMITTEE ON ARMY ROBOTICS AND ARTIFICIAL INTELLIGENCE

WALTER ABEL, Senior Fellow for Technology, Emhart Corporation, Chairman

J. MICHAEL BRADY, Artificial Intelligence Laboratory, Massachusetts Institute of Technology

LT. GENERAL HOWARD H. COOKSEY (Retired), Cooksey Corporation

STEVEN DUBOWSKY, Professor of Mechanical Engineering, Massachusetts Institute of Technology

MAURICE J. DUNNE, Vice President, Product Planning, Unimation, Incorporated

MARGARET A. EASTWOOD, Director, Integrated Factory Controls, GCA Industrial Systems Group

COLONEL FREDERICK W. FOX (Retired)

LESTER GERHARDT, Chairman, Electrical, Computer and Systems Engineering Department, Rensselaer Polytechnic Institute

DAVID GROSSMAN, Manager of Automation Research, T. J. Watson Research Center, IBM Corporation

GENERAL JOHN R. GUTHRIE (Retired), Association of the U.S. Army

TENHO R. HUKKALA, System Planning Corporation

LAVEEN KANAL, Department of Computer Science, University of Maryland

WENDY LEHNERT, Department of Computer and Information Sciences, University of Massachusetts

CHARLES ROSEN, Chief Scientist and Director, Machine Intelligence Corporation

PHILIPP F. SCHWEIZER, Manager, Intelligent Systems, Westinghouse R&D Center

JOHN M. SHEA, Project Manager, XMCO, Incorporated

NRC BOARD ON ARMY SCIENCE AND TECHNOLOGY LIAISONS

ARDEN L. BEMENT, Vice President, Technology Resources, TRW, Incorporated

WALTER B. LABERGE, Vice President, Planning and Technology, Lockheed Missile and Space Company

MANUFACTURING STUDIES BOARD LIAISON

ROGER NAGEL, Director, Institute for Robotics, Lehigh University

iii

MANUFACTURING STUDIES BOARD

GEORGE S. ANSELL, Chairman, Dean of Engineering, Rensselaer Polytechnic Institute, Troy, New York

ANDERSON ASHBURN, Editor, AMERICAN MACHINIST, New York, New York

AVAK AVAKIAN, Vice President, GTE Sylvania Systems Group, Waltham, Massachusetts

DANIEL BERG, Provost, Science and Technology, Carnegie-Mellon University, Pittsburgh, Pennsylvania

ERICH BLOCH, Vice President - Technical Personnel Development, IBM Corporation, White Plains, New York

IRVING BLUESTONE, Professor of Labor Studies, Wayne State University, Detroit, Michigan

DONALD C. BURNHAM, Retired Chairman, Westinghouse Electric Corporation

BARBARA A. BURNS, Manufacturing Technology Group Engineer, Lockheed Georgia Company, Marietta, Georgia

JOHN K. CASTLE, President, Donaldson, Lufkin and Jenrette, Inc., New York, New York

ROBERT H. ELMAN, Group Vice President, AMCA International Corporation, Hanover, New Hampshire

JOSEPH ENGELBERGER, President, Unimation Incorporated, Danbury, Connecticut

ELLIOTT M. ESTES, Retired President, General Motors Corporation, Detroit, Michigan

W. PAUL FRECH, Vice President of Operations, Lockheed Corporation, Burbank, California

BELA GOLD, Director, Research Program in Industrial Economics, Case Western Reserve University, Cleveland, Ohio

DALE B. HARTMAN, Director of Manufacturing Technology, Hughes Aircraft Company, Los Angeles, California

MICHAEL HUMENIK, JR., Director, Manufacturing Process Laboratory, Ford Motor Company, Detroit, Michigan

ROBERT B. KURTZ, Retired Vice President, General Electric Corporation, Fairfield, Connecticut

M. EUGENE MERCHANT, Principal Scientist, Manufacturing Research, Cincinnati Milacron, Incorporated, Cincinnati, Ohio

ROY MONTANA, General Manager, Bethpage Operation Center, Grumman Aerospace Corporation, Bethpage, New York

ROGER NAGEL, Director, Institute for Robotics, Lehigh University, Bethlehem, Pennsylvania

REGINALD NEWELL, Director of Research, International Association of Machinists and Aerospace Workers, Washington, D.C.

BERNARD M. SALLOT, Director, Professional and Government Activities, Society of Manufacturing Engineers, Dearborn, Michigan

WICKHAM SKINNER, Harvard Business School, Cambridge, Massachusetts

ALVIN STEIN, Parker Chapin Flattau and Klimpl, New York, New York

ACKNOWLEDGMENTS

While the committee is ultimately responsible for the content of this report, a number of other people gave valuable information and insights during the research and analysis. Without them, this would be a poorer report.

Dr. Roger Nagel, Director of the Institute for Robotics, Lehigh University, wrote most of the appendix. He is to be commended for a thorough job.

Dr. Frank Verderame, Assistant Director for Research Programs, Department of the Army, in the important role of project monitor, offered guidance to the committee and provided background information. Also providing information on Army plans and programs were Lt. Colonel Henry Langendorf, Soldier Support Center; Dr. Robert Leighty, Army Topographic Laboratories; Mr.

Kent Schluskel, Foreign Science and Technology Center; Dr. James Gault, Army Research Office; Dr. Stanley Halpin, Army Research Institute; and Colonel Philip Sobocinski, Office of the Surgeon General.

Dr. William Isler, Defense Advanced Research Projects Agency, was a contributor at all meetings. In addition, E. H. Chaves of ESL Inc., Charles Garvey and Dennis Gulakowaki, both of XMCO, and Carl Ruoff of the Jet Propulsion Laboratory all participated in the committee's second or third meetings. Mr. Chavea is responsible for the discussion of industry's implementation experience in Chapter 6.

Stephen Merrill, Center for Strategic and International Studies, and Harold Davidson, Department of the Army, served as consultants to the committee and assisted in gathering information.

Joel Goldhar, Executive Director of the study through January 1983 and currently Director of Engineering, Illinois Institute of Technology, got the study off to a good start. Janice Greene, Staff Officer, provided support throughout the committee's work and was instrumental in preparing the final draft of the report. This report would not

v

have been possible without the administrative work of Staff Associate Georgene Menk and assistants Patricia Ducy, Donna Reifsnider, and Fran Shaw.

Two boards within the National Research Council reviewed the report: the Manufacturing Studies Board, under Executive Director George Kuper, and the Board on Army Science and Technology, under Executive Director Dennis Miller.

vi

CONTENTS

1.	BACKGROUND	1
	Approach, 1	
	Prior Studies, 2	
	Contribution of This Report, 4	
2.	SUMMARY OF THE TECHNOLOGY	5
	Definitions, 5	
	Research Issues, 6	
3.	CRITERIA FOR SELECTION OF APPLICATIONS	10
	Reasons for Applying Robotics and Artificial Intelligence, 10	
	Combining Short-term and Long-term Objectives, 11	

	Planning for Growth, 11	
	Selecting Applications to Advance Particular Technologies, 12	
4.	RECOMMENDED APPLICATIONS AND PRIORITIES	14
	An Initial List, 14	
	Automatic Loader of Ammunition in Tanks, 16	
	Sentry/Surveillance Robot, 18	
	Intelligent Maintenance, Diagnosis, and Repair System, 20	
	Expert Systems for Army Medical Applications, 22	
	Flexible Material-Handling Modules, 24	
	Automated Battalion Information Management System, 26	
5.	IMPLEMENTATION OF RECOMMENDED APPLICATIONS	28
	Measures of Effectiveness, 31	
6.	OTHER CONSIDERATIONS	35
	Shortage of Experts, 33	
	Operator-Friendly Systems, 34	
	Coordination of Existing Programs, 35	
	Available Technology, 35	
	Getting Started, 35	
	Focus for AI and Robotics, 36	
	Implementation Difficulties, 36	

vii

	CONTENTS (continued)	
7.	RECOMMENDATIONS	39
	Start Using Available Technology Now, 39	
	Criteria: Short-Term, Useful Applications with Planned Upgrades, 40	
	Specific Recommended Applications, 40	
	Visibility and Coordination of Military AI/Robotics, 41	
APPENDIX:	STATE OF THE ART AND PREDICTIONS FOR ARTIFICIAL INTELLIGENCE AND ROBOTICS	42
	Industrial Robots: Fundamental Concepts, 42	
	Research Issues in Industrial Robots, 46	
	Artificial Intelligence, 58	
	State of the Art and Predictions, 69	
	References, 87	

GLOSSARY OF ACRONYMS

90

1 BACKGROUND

Throughout its history, the Army has been manpower-intensive in most of its systems. The combination of demographic changes (fewer young men), changed battlefield scenarios, and advanced technologies in improved robotics, computers, and artificial intelligence (AI) suggests both a need and an opportunity to multiply the effectiveness of Army personnel. Not only can these technologies reduce manpower requirements, they can also replace personnel in hazardous areas, multiply combat power, improve efficiency, and augment capabilities.

The Deputy Chief of Staff for Research, Development and Acquisition authorized the National Research Council to form a committee to review the state of AI and robotics technology, predict developments, and recommend Army applications of AI and robotics. This Committee on Army Robotics and Artificial Intelligence brought together experts with military, industrial, and academic research experience.

APPROACH

The committee began its work with a detailed review of the state of the art in robotics and artificial intelligence as well as with predictions of how the technology will develop during the next 5- and 10-year periods. This review is summarized in Chapter 2 and in its entirety forms the appendix of this report. It is the foundation of the committee's recommendations for selecting and implementing of applications.

The committee used its review of technology and information on Army doctrine, prior reports on Army applications of AI and robotics, and its combined military, university, and industrial experience to develop criteria for selecting applications and to recommend specific applications that it considers of value to the Army and the country. For each application recommended, the committee was asked to report the expected effects on personnel, skills, and equipment, as well as to provide an implementation strategy incorporating priorities, costs, timing, and a measure of effectiveness.

PRIOR STUDIES

As background to its efforts, the committee was briefed on and reviewed three studies completed during 1982 on Army robotics and artificial intelligence:

- D. R. Brown, et al., R&D Plan for Army Applications of AI/Robotics, SRI International, May 1982 (Contract No. DAAK70-81-C-0250, U.S. Army Engineer Topographic Laboratories).
- Army Plan for AI/Robotics Technology Demonstrators, Department of the Army, June 1982.
- Report of the Army Science Board Ad Hoc Subgroup on Artificial Intelligence and Robotics, Army Science Board, September 1982.

Each contributes to the base of knowledge regarding these expanding new technologies and offers insights into potential applications to enhance the Army's combat capabilities. Their conclusions are briefly reviewed here to place the contribution of this particular report in a proper context.

R&D Plan for Army Applications of AI/Robotics

The report by SRI cites as the primary motivation for the application of AI and robotics to Army systems the need to conserve manpower in both combat and noncombat operations. It covers more than 100 possible Army applications of AI and robotics, classified into combat, combat support, and combat service support categories. Many of the applications, though listed as distinct, could easily be drawn together to serve as generic applications. The report focuses on the need to document justification for the value of AI and robotics in Army applications in general, but the committee found that it lacked sufficient detail for ranking the many applications to pursue those of greatest interest and potential payoff.

From the 100 specific concepts that the SRI study considered, 10 broad categories of application were selected. An example from each of these 10 categories was chosen for further study to identify technology gaps and provide the basis for the research plan recommended by the study.

Included in that plan were 5 fundamental research areas, 97 specific research topics, and 8 system considerations. Most potential applications were judged to require advancement of the technology base (basic research and exploratory development) before advanced development could begin. In fact, the study estimated that development on only four could be started in the next 10 years, and two would require deferral of development until the year 2000.

2

A briefing on the Army Proposed Plan was given to the committee at its initial meeting. The report identified five projects for application of AI or robotics technology to demonstrate the Army's ability to exploit AI and robotics:

- Robotic Reconnaissance Vehicle with Terrain Analysis,
- Automated Ammunition Supply Point (ASP),
- Intelligent Integrated Vehicle Electronics,
- AI-Based Maintenance Tutor,
- AI-Based Medical System Development.

Of these five proposed demonstrations, technical availability assessments placed one in the near term, one in the mid-to-far term, and the other three in the far term. Cost estimates and schedules appear optimistic to this committee, considering that much of the effort was neither funded nor programmed at that time.

Report of the Army Science board
Ad Hoc Subgroup on Artificial Intelligence and Robotics

The Army Science Board Ad Hoc Subgroup was established to provide an assessment of the state of the art of AI and robotics as fast-track technologies and of their potential to meet Army needs. It concentrated its efforts on those aspects with which it could deal rapidly and relatively completely; it also considered the five Army demonstrators and supported them.

The report grouped the five demonstrators into two categories: proceed as is or proceed with modification. The subgroup recommended changes to the maintenance tutor and the medical system, and recommended that the other three demonstrators proceed as planned. Other battlefield technology topics recommended were automatic (robotic) weapons, automatic pattern recognition, and expert support systems.

Noting that the introduction of technology into weapon systems could be hampered by management problems, the subgroup recommended establishing a single dedicated proponent of AI and robotics in the Department of the Army, giving preference to existing equipment and technology, and creating an oversight committee from the Army's materiel developer and user communities.

The subgroup tied its recommendations to the five technology thrusts that the Army has designated to receive the majority of research and development funds (lines 6.1, 6.2, and 6.3a of the budget) during the next five-year funding period:

- Very Intelligent Surveillance and Target Acquisition,
- Distributed C31,

3

- Self-Contained Munitions,
- Soldier/Machine Interface,
- Biotechnology.

CONTRIBUTION OF THIS REPORT

This committee is indebted to the foregoing efforts for the base they provide, a base which this report attempts to expand. Our recommendations are founded on a comprehensive assessment of the state of the art and forecasts of technology growth over the next 10 years. The details of that assessment are contained in the Appendix. We hope that our recommendations to the Army will provide a realistic technical assessment that will enable the Army, in turn, to concentrate its efforts in areas offering the most potential return.

No two groups considering possible AI and robotics applications will have identical lists of priorities. This committee used the combination of Army needs and the direction of technology development as a guide in narrowing the list of possible applications. The National Research Council is unique in the diversity of backgrounds of the experts it brings together. The members of this Committee on Army Robotics and Artificial Intelligence have among them 248 years of industry experience, 110 years in academia, and 184 years in government. The recommendations in this report are the consensus of the committee, drawing on those years of experience.

We agree with the authors of studies we have reviewed that AI and robotics technologies offer great potential to save lives, money, and resources and to improve Army effectiveness. This report will

- support the need for ongoing work in these high-risk, high-technology fields that offer such great promise for the country's future security
- help channel Army efforts into the most effective areas,
- build understanding of what AI and robotics can offer within the broad groups in the Army that will need to work with these technologies ,
- provide realistic information on what AI and robotics technology can do now and the directions in which research is heading.

4

2 SUMMARY OF THE TECHNOLOGY

DEFINITIONS

We used the Robot Institute of America's definition of a robot as

a reprogrammable multi-function manipulator designed to move material, parts, tools, or specialized devices through variable programmed motions for the performance of a variety of tasks.

The main components of a robot are

- the mechanical manipulator, which is a set of links that determine the work envelope of the robot and the ability to orient the hand;
- the actuation mechanisms, which are hydraulic, pneumatic, or electric;
- the controller, usually a computer, which controls motion by communicating with the actuation mechanism.

The robot can be augmented by the addition of

- end effectors, or "hands";

- sensors, for performing measurements as required to sense the environment, including electromagnetic (visual, infrared, ultraviolet, radar, radio, etc.), acoustic, tactile, force, torque, spectographic, and many others.
- other "intelligent" functions, such as understanding speech, problem solving, goal seeking, and commonsense reasoning.

None of these, strictly speaking, is part of the robot itself.

This chapter is a summary of the detailed report on the state of the art and predictions for AI and robotics technology contained in the appendix.

5

Artificial intelligence, as defined in SRI International's R&D Plan for Army Applications of AI/Robotics, is

the part of computer science that is concerned with symbol-manipulation processes that produce intelligent action. By "intelligent action" is meant an act or decision that is goal-oriented, arrived at by an understandable chain or symbolic analysis and reasoning steps, and is one in which knowledge of the world informs and guides the reasoning.

The functions or subfields of artificial intelligence are

- natural-language understanding; that is, understanding English or another noncomputer language;
- image understanding; that is, the ability to identify what is in a picture or scene;
- expert systems, which codify human experience and use it to guide actions or answer questions;
- knowledge acquisition and representation;
- heuristic search, a method of looking at a problem and selecting a path to the solution;
- deductive reasoning;
- planning, which entails an initial plan for finding a solution, then monitoring progress.

As this infant field develops, the list of subfields will expand. Artificial intelligence is the application of advanced computer systems and software to these areas, with "intelligent behavior" as the intended result.

RESEARCH ISSUES

The categories of robotics research receiving the most effort are

- improvement of mechanical systems, including manipulation design, actuation systems, end effectors, and locomotion;
- improvement of sensors to enable the robot to react to changes in its environment;

- creation of more sophisticated control systems that can handle dexterity, locomotion, and sensors, while being user friendly.

In artificial intelligence, expert systems is the area of research closest to being ready to move from the laboratory to initial commercial use.

6

Mechanical Systems: Manipulator and Actuation

Research on the kinematics of design, models of dynamic behavior, and alternative design structures, joints, and force programming is leading to highly accurate new robot structures. This research will lead to robots capable of applying force and torque with speed and accuracy and will transform today's heavy, rigid, single robotic arms into more lightweight, ultimately more flexible arms capable of coordinated motion.

Research on end effectors--the hands attached to a robot--seeks to improve dexterity, enabling robots to handle a variety of parts or tools in complex situations. Two goals are the quick-change hand and the dexterous hand. The robot would be able to change a quick-change hand by itself, attaching the means of transmitting power as well as the physical hand to the arm.

Although the dexterous hand is beyond the current state of the art, there are some interesting present approaches. One is a variable finger selection; another is the use of materials that will produce signals proportional to surface pressures. This is coupled with research in microelectronics to analyze and summarize the signals from these multisensored fingers for decision-making outputs.

Early attention to locomotion has led to a large number of robots in current use mounted on tracks or an overhead gantry. Progress has recently been made on a six-legged walking robot that is stable on three legs.

A middle ground between tracked and unconstrained vehicles is a wire-guided vehicle used in plants. These vehicles have onboard microprocessors that communicate with a central control computer at stations placed along the factory floor. The vehicles travel along a wire network that is kept free of permanent obstacles; bumper sensors prevent collisions with temporary obstacles.

Sensors

The purpose of sensors is to give the robot adaptive behavior--that is, the ability to respond to changes in its environment. Vision and tactile sensors have received the lion's share of research effort. While tactile sensors are still fairly primitive, vision systems are already commercially available.

Vision systems enable robots to perform the following types of tasks:

- identification or verification of objects,

- location of objects and their orientation,
- inspection, navigation and scene analysis,
- guidance of the servo mechanism, which controls position through feedback.

7

- The first three tasks can be performed by today's commercial systems. Three-dimensional vision systems are at present rudimentary.

Tactile sensors are just beginning to be commercialized. Within the next few years, force-sensing wrists and techniques for controlling them will be available for such tasks as tightening nuts, inserting shafts, and packing objects. More research will be needed before they can work in other than benign environments.

Control Systems

The underlying research issue in control systems is to broaden the scope of the robot to include dexterous hands, locomotion, sensors, and the ability to perform new complex tasks.

Robots are typically programmed by either the lead-through or the teach-box method. In the former the controller samples the location of each of the robot's axes several times per second, while a person manipulates the robot through the desired motions. The teach-box method enables the operator to use buttons, toggle switches, or a joy stick to move the robot.

Programming languages for robots have long been under research. Early robot languages have combined language statements with use of a teach box. Second-generation robot languages, which resemble the standard structured computer language, have only recently become commercially available. It is these second-generation robot languages that create the potential to build intelligent robots.

Expert Systems

Artificial intelligence has generated several concepts that have led to the development of important practical systems. A subset of these systems has been called expert systems. As the name suggests, an expert system (ES) encodes deep expertise in a narrow domain of human specialty. Several expert systems have been constructed whose behavior surpasses that of humans. Examples include the MIT Macsyma system (symbolic mathematics), the Digital Equipment Corporation R-1 system (configuring VAX computers), the Schlumberger dipmeter analyzer (oil well logs), and various medical expert systems, including PUFF (pulmonary function diagnosis) in regular use at San Francisco Hospital. Expert systems' behavior in research laboratories and the civilian sector is cause for optimism in the military sector.

One can consider expert-systems support not only at the corps and division levels but also for battalions and regiments. As envisioned in the Air Land Battle 2000 scenario, battalion and regimental formations will be operating in forward battle areas in a dispersed manner. Expert-system support at this level will be particularly helpful in increasing combat effectiveness

through flexibility and adaptability to varied, complex situations and improved survivability of men and machines.

8

Although there is cause for optimism, current expert systems have significant limitations and require intensive basic research if the technology is to be successfully transferred from the university laboratory to make rugged operational systems.

- Present expert systems support only narrow domains of expertise. As the domain of application becomes broader, the number of alternative courses of action increases exponentially and effectiveness decreases exponentially. Though research is addressing this issue, practical expert systems are likely to be severely restricted in their domain for the next 5 years.
- Only limited knowledge-representation languages for data and relations are available.
- The input and output of most expert systems are inflexible and not in English (or any other natural language).
- Expert systems still require laborious construction--approximately 10 man-years for a sizable one.
- Because present expert systems need one domain expert in control to maintain consistency in the knowledge data base, they have only a single perspective on a problem.
- Many expert systems are difficult to operate.

9

3 CRITERIA FOR SELECTION OF APPLICATIONS

The committee spent a great deal of time developing criteria for the selection of Army applications of robotics and artificial intelligence. These criteria were essential in guiding the work of the committee; but beyond that, they are more broadly applicable to future decisions by the Army as well as by others. The criteria for selecting applications reflect both the immediate technological benefits and the attitudinal and managerial considerations that will affect the ultimate widespread acceptance of the technology.

REASONS FOR APPLYING ROBOTICS AND ARTIFICIAL INTELLIGENCE

The introduction of robotics and artificial intelligence technology into the Army can result in a number of benefits, among them the following:

- improved combat capabilities,

- minimized exposure of personnel to hazardous environments,
- increased mission flexibility,
- increased system reliability
- reduced unit/life-cycle costs,
- reduced manpower requirements,
- simplified training.

In selecting applications from the much larger list of possibilities, the committee not only looked for opportunities to achieve those benefits but also sought affirmative answers to the following questions:

- Will it perform, in the near term, an essential task for the Army.
- Can its initial version be implemented in 2 to 3 years?
- Can it be readily upgraded as more sophisticated technology becomes available?
- Does it tie in with existing, related programs, including programs of the other services?

10

- Will it use the best technology available in the scientific community?

These considerations should help to ensure initial acceptance and continuing success with these promising developing technologies.

COMBINING SHORT-TERM AND LONG-TERM OBJECTIVES

Initial short-term implementation should provide a basis for future upgrading and growth as the user gains experience and confidence in working with equipment using robotics and AI technology. To this end the Army's program should be carefully integrated and include short-term, achievable objectives with growth projected to meet long-term requirements.

As a result; some of the applications chosen may at first appear to be implementable in the short term by other existing technologies with lower cost and ease. However, such short-term expediency may cause unwarranted and unintended delay in the ultimately more cost-effective application of new developing robot technologies. To prevent this problem, short-term applications should be

- applied to existing, highly visible systems,
- reasonably afforded within the Army's projected budget,
- within the state of the art, requiring development and engineering rather than invention or research,
- able to demonstrate an effective solution to a critical Army need ,
- achievable within 2 to 3 years,
- not redundant with efforts in DARPA or the other services.

On the other hand, the committee considered long-term applications to be important vehicles for advancing research in these technologies and, in some cases, for introducing useful applications of robotics and artificial intelligence. These more advanced applications would ultimately, at reduced cost, assist in meeting the changing requirements of the modern battlefield envisioned in the Army's Air Land Battle 2000 concept.

The principle that guided the committee's selection of applications, therefore, was to combine short-term and long-term benefits; that is, to select applications that can be implemented quickly to meet a current need and, in addition, can be upgraded over the next 10 years in ways that advance the state of the art and perform more complex functions for the Army.

PLANNING FOR GROWTH

For the near term, using state of the art technology and assuming that a demonstration program starts in 1 1/2 to 2 years and continues for 2 years, the committee recommends that projects be selected based not

11

only on what is commercially available now but also on technology that is likely to become available within the next 2 years.

During the next 4 to 5 years, while the Army is developing its demonstration systems, annual expenditures by university, industrial, government, and nonprofit laboratories for R&D and for initial applications will probably exceed several hundred million dollars per year worldwide. To be timely and cost effective, Army demonstration systems should be designed in such a way that these developments can be incorporated without discarding earlier versions.

It is therefore of the utmost importance to specify, at the outset, maximum feasible computer processor (and memory) power for each application. Industry experience has shown that the major deterrent to updating and improving performance and functions has been the choice of the "smallest" processor to meet only the initial functional and performance objectives.

It is at least as important to ensure that this growth potential be protected during development of the initial applications. Both industry and the Army have known programmers with a propensity to expand operating and other systems until they occupy the entire capacity of design processor and memory.

Robots are currently being developed that incorporate external sensors permitting modification of the sequence of motions, the path, and manipulative activities of the robot in an adaptive manner. The status of the "dumb, deaf, and blind" robot is being raised to that approaching an "intelligent" automaton. This upgraded system can automatically cope with changes in its reasonably constrained environment.

The earliest adaptive robot systems are just beginning to be incorporated into production lines. Most of these Systems are presently in an advanced development stage, worked on by application engineers for early introduction into production facilities. Such Systems, called third-generation robot Systems, are expected to supplement the second-generation robot Systems (having programmable control but lacking sensors) in the next 2 to 3 years. Shortly thereafter, as more and more assembly operations are automated, they are likely to become the dominant class of robot Systems. In view of these technological developments, the Army demonstration Systems should, at the very least, be based on the third-generation robot Systems capable of being readily upgraded with minimum change in the internal hardware configuration, relying on future additions of readily interfaceable external sensors and software.

SELECTING APPLICATIONS TO ADVANCE PARTICULAR TECHNOLOGIES

In addition to considering the benefits that result from applying robotics and artificial intelligence, the Army has the opportunity to use its choice of applications to take an active role in advancing

12

particular technologies. Because robotics and AI are developing rapidly, the committee believes that Army should support a range of component technologies.

The two fields are at present separate, and the possible applications can be divided into those that are primarily *robotics* and those that are primarily *artificial intelligence*. The robotics applications can be further divided into those that primarily advance *end-effector* (hand) technology and those that primarily advance *sensor* technology.

The AI applications can be divided into a number of types, of which the furthest developed is expert systems. The committee limited its consideration of AI applications to expert systems, in keeping with its goal of short-term implementation of limited aspects. The primary technology for expert systems is cognition.

Each of these areas--effectors, sensors, and cognition--is an important source of technology for the Army and for this country's industrial base. To encourage R&D in these areas and to enable the Army to have some initial experience in each area, the committee agreed to recommend three applications, one directed at each.

13

4 RECOMMENDED APPLICATIONS AND PRIORITIES

The committee used the criteria described in Chapter 3 to develop an initial list of 10 possible Army applications of robotics and artificial intelligence. These were discussed at length and narrowed to six applications that met the criteria, three of which are strongly recommended.

Many hours of committee discussion are reflected in the following list. The committee found it impossible to match the large numbers of possible applications and criteria in any systematic way. No two groups applying the criteria would arrive at identical lists of Army projects to recommend. The applications recommended below are eminently worthwhile in the judgment of the committee. They clearly address current Army needs, offer short-term benefits, are likely to give Army personnel some positive early experiences with the technology, and are capable of being upgraded.

AN INITIAL LIST

With these considerations in mind, the committee developed the following list of 10 potential applications of robotics and artificial intelligence. Not all of these applications are recommended by the committee; this list is the result of the committee's first effort to narrow down the vast number of possible applications to those most likely to meet the criteria described earlier.

- **Automatic Loader of Ammunition in Tanks.** This system would require development of a robot arm with minimum degrees of freedom for use within the tank. The arm would be capable of acquiring rounds from a magazine or rack and loading them into the gun, with a vision system to provide the means to correct for imprecise positioning of rounds and gun and tactile or force sensors to ensure adequate acquisition.
- **Sentry Robot.** A portable unattended sentry device would detect and report the presence of personnel or vehicles within a designated area or along a specified route. The device would also be capable of sensing the presence of nuclear, biological, and chemical contaminants.

14

- **Flexible Material-Handling Modules.** Adaptive robots mounted on wheeled or tracked vehicles would identify and acquire packages or pallets to load or unload. There are so many potential applications for material-handling systems that material-handling robots are likely to become as ubiquitous as the jeep in the Army supply system, with applications in forward as well as rear areas.
- **Robotic Refueling of Vehicles.** A wheeled robot fitted with an appropriate fuel dispenser (a tool for inserting into a fuel inlet) could automatically refuel a variety of vehicles.
- **Counter-Mine System.** Adaptive robots mounted on wheeled or tracked vehicles could be fitted with specialized sensors and probing or digging tools to find and dispose of buried mines. Vehicles could be remotely controlled in the teleoperator mode.
- **Robot Reconnaissance Vehicle.** The remotely controlled reconnaissance vehicle that the Army is considering as a major demonstration project could be fitted with one or more external robot arms and equipped with vision and other sensors. This would expand

the utility of the system to perform manipulative functions in forward, exposed areas, such as retrieval of disabled equipment; sampling and handling nuclear, biological, and chemically active materials (NBC); and limited decontamination.

- *Airborne Surveillance Robot*. A semiautonomous aerial platform fitted with sensors could observe large areas, provide weather data, detect and identify targets, and measure levels of NBC contamination.
- *Intelligent Maintenance, Diagnosis, and Repair System*. An ES, specialized for a particular piece of equipment, would give advice to the relatively untrained on how to operate, diagnose, maintain, and repair relatively complex electronic, mechanical, or electromechanical equipment. It would also act as a record of repairs, maintenance procedures, and other information for each major item of equipment.
- *Medical Expert System*. This system would give advice on the diagnosis and evacuation of wounded personnel. A trained but not necessarily professional operator would enter relevant information (after prompting by the system) regarding the condition of the wounded individual, including any results of initial medical examination. The system would logically evaluate the relative seriousness of the wound and suggest disposition and priority. This system could be improved by having available a complete past medical record of the individual to be entered into the system prior to asking for its advice.
- *Battalion Information Management System*. This system would provide guidance and assistance in situation assessment, planning, and decisionmaking. Included would be the automatic or semiautomatic production of situation maps, plans, orders, and status reports. It also would include guidance for operator actions in response to specific situations or conditions.

Although this list represents a considerable reduction from the many possible applications that have been conceived, a further narrowing is needed. Knowledgeable researchers and other resources are in such short supply that Army efforts in AI and robotics should

be well thought out and focused. The remainder of this chapter presents in more detail the functions, requisite technology, and expected benefits of the committee's top six priorities.

As noted in Chapter 3, the committee recommends that the Army fund three demonstration projects, one in each of the areas of effectors, sensors, and cognition. This committee's consensus is that, at a minimum, the following projects should be funded:

1. automatic loader of ammunition in tanks (effectors),
2. sentry robot (sensors),
3. intelligent maintenance, diagnosis, and repair system (cognition).

These applications all meet the criteria listed on pages 10-11: they meet a current Army need, demonstrations are feasible within 2 to 3 years, and the systems can be readily upgraded. Together, these applications are strongly recommended for funding.

The committee also found the following applications to meet its criteria. If funding is available, these are also recommended:

4. medical expert system (cognition),
5. flexible material-handling modules (effectors) ,
6. battalion information management system (cognition).

As to the remaining applications, robotic refueling of vehicles is an example of a flexible material-handling module (priority 5) and the airborne surveillance robot is an upgraded version of the sentry robot (priority 2). The reconnaissance vehicle is not in this committee 's recommended list because a demonstration is not likely to be possible within 2 years. The counter-mine vehicle is not recommended because the problem seems better suited to a less expensive, lower-technology solution.

AUTOMATIC LOADER OF AMMUNITION IN TANKS

At present the four-man crew of a U.S. tank consists of a commander, a gunner, a driver, and a loader. The loader receives verbal instructions to load a particular type of ammunition; he then manually selects the designated type of ammunition from a rack, lifts it into position, inserts it into the breech, completes the preparation for firing, and reports the cannon's readiness to fire. The gunner, who has been tracking the intended target, has control of firing the cannon. When fired, the hot, spent casing is automatically ejected and is later disposed of, as convenient, by the loader. The loader occasionally unloads and restores unfired cartridges onto the rack.

With appropriate design of the complete ammunition loading system, these functions can be automated. The committee recommends the use of state-of-the-art robotics to effect this automation, eliminating one

man (the loader) from the crew, and potentially increasing the firing rate of the cannon, now limited by the loader's physical capabilities.

Functional Requirements

The major functional requirements of the system are

- A computer-controlled, fully programmable, servoed robot designed for the special purpose of ammunition selection and loading. Its configuration, size, number of degrees of freedom, type of drive (hydraulic or electric), load capacity, speed precision,

and grippers or hands would be engineered specifically for the purpose as part of the overall system design. Computer power in its controller would be adequate for interfacing with vision, tactile, and other sensors, and for communicating with other computers in the tank. Provisions would be made to introduce additional processing power in the future by leaving some empty "slots" in the processor cage. The principles of design for such a robot are now known, and the major requirement, after setting its specifications, is good engineering. A working prototype should take 1-1/2 to 2 years to produce.

- A simple machine vision system designed to perform the functions of locating the selected type of ammunition in a magazine or rack, guiding the robot to acquire the round, and guiding the robot to insert the round into the breech. Although it is certainly possible to design a more specialized and highly constrained system, the proposed adaptive robot system provides for greater flexibility in operation and reduction of constraints, and will enable more advanced functional capabilities in the future. The principles of designing an appropriate vision system are now available; the design for this purpose should not be difficult. Simplifying constraints such as colored, bar code, or other markings on the tips of shells and breech would eliminate tedious processing to obtain useful imagery for interpretation. Other sensory capabilities (e.g., tactile and force) could readily be added to the system if necessary, for confirming acquisitions and insertions. The robot computer could be programmed to accommodate all these sensors.
- An ammunition storage rack (or, preferably, magazine) designed to facilitate both bulk loading into the tank and acquisition of selected ammunition by the robot gripper. It may even have an auxiliary electromechanical device that would push selected ammunition forward to permit easy acquisition by the robot, such action controlled by the robot computer.
- Robot and vision computers integrated and interfaced with the fire control computer under control of the commander or gunner. This local computer network is intended for use in later developments when further automation of the tank is contemplated. However, it could even be used in the short term to ensure that the type of ammunition loaded is the same type that is indexed in the fire control computer.

17

Benefits

The near term advantages (2 to 5 years) foreseen are

- elimination of one crew member (the loader) and automation of a difficult, physically exhausting task that contributes little to the overall skills of the people who perform it;
- potential increase in fire power by reducing loading time;
- the availability of a test bed for further development and implementation of more advanced systems and increased familiarity of personnel with computer-controlled devices;
- simplification of communications between commander, gunner, and loader, which may lead to direct control by the tank commander and potential reduction of errors during the heat of combat;

- Army experience with computer control, especially of robot systems.

In the long term, if concurrent developments in automated tracking using advanced sensors occur, it may be feasible to eliminate the gunner, reducing the crew to a commander and a driver. This would make possible two-shift operations with two two-man crews operating and maintaining the tank over a 24-hour period, a considerable increase in operating time for very important equipment. Mechanization of the ammunition-loading function and an integrated computer network in place are prerequisites for this development.

A potential tank of the future could be unmanned--a tank controlled by a teleoperator from a remote post or hovering aircraft. The tank would be semiautonomous; that is, it could maneuver, load rounds, track targets, and take evasive action to a limited degree by itself, but its actions would be supervised by a remote commander who would initiate new actions to be carried out by internally stored computer programs. Eliminating people on board the tank could lead to highly improved performance, now limited by human physical endurance and safety. The tank would become an unmanned combat vehicle, smaller, lighter, faster, with far less armor and more maneuverable--essentially a mobile cannon with highly sophisticated control and target acquisition systems.

SENTRY/SURVEILLANCE ROBOT

The modern battlefield, as described in *Air Land Battle 2000*, will be characterized by considerable movement, large areas of operations in a variety of environments, and the potential use of increasingly sophisticated and lethal weapons throughout the area of conflict. Opposing forces will rarely be engaged in the classical sense--that is, along orderly, distinct lines. Clear differentiation between rear and forward areas will not be possible. The implications are that there will be insufficient manpower available to observe and survey the myriad of possible avenues by which hostile forces and weapons may threaten friendly forces.

18

Initially using the concepts and hardware developed in the Remotely Monitored Battlefield Sensor System (REMBASS), a surveillance/ sentry robotic system would provide a capability to detect intrusion in specified areas--either in remote areas along key routes of communication or on the perimeter of friendly force emplacements. Such a system would apply artificial intelligence technology to integrate data collected by a variety of sensors--seismic, infrared, acoustic, magnetic, visual, etc.--to facilitate event identification, recording, and reporting. The device could also monitor NBC sensors, as well as operate within an NBC-contaminated area.

Initially, the system would be stationary but portable, with an antenna on an elevated mast near a sensor field or layout. It can build on sentry robots that are currently available for use in industry. Ultimately, the system would be mobile. Either navigation sensors would provide mobility along predetermined routes or the vehicle would be airborne; the decision should be made as the technology progresses. Also, the mobile system would employ onboard as well as remote sensors.

Functional Requirements

The proposed initial, portable system would require

- A fully programmable, computer-operated controller (with transmit/receive capabilities) that would interface with the remote sensors and process the sensor data to enable automated recognition (object detection, identification, and location). This effort would entail matching the various VHF radio links from existing or developmental remote sensors at a "smart" console to permit integration and interpretation of the data received.
- A secure communications link from the controller to a tactical operations center that would permit remote read-out of sensor data upon command from the tactical operations center. This communications link would also provide the tactical operations center the capability of turning the controller (or parts of it) on or off.

Later versions of the system would have the attributes described above, with the additional features of mobility and onboard sensors. In this case, the sentry/surveillance robot would become part of a teleoperated vehicular platform, either traversing a programmed, repetitive route or proceeding in advance of manned systems to provide early warning of an enemy presence.

Benefits

The principal near-term advantages are

- to provide a test bed for exploiting AI technology in a surveillance/sentry application, using available sensors adapted to

19

special algorithms that would minimize false alarms and speed up the process of detection, identification, and location.

- to permit a savings in the manpower required for monitoring sensor alarms and interpreting readings, while providing 24-hour-a-day, all-weather coverage.
- to provide a capability for operating a surveillance/sentry system under NBC conditions or to warn of the presence of NBC contaminants.

The far-term mobile system would be invaluable in providing surveillance/sentry coverage in the vicinity of critical or sensitive temporary field facilities, such as high-level headquarters or special weapons storage areas.

INTELLIGENT MAINTENANCE, DIAGNOSIS, AND REPAIR SYSTEM

Expert Systems applications in automatic test equipment (ATE) can range from the equipment design stage to work in the field. Expert systems incorporating structural models of pieces of

equipment can be used in equipment design to simplify subsequent trouble shooting and maintenance.

In the field, expert systems can guide the soldier in expedient field repairs. At the depot, expert systems can perform extensive diagnosis, guide repair, and help train new mechanics.

In the diagnostic mode it would instruct the operator not only in the sequence of tests and how to run them, but also in the visual or aural features to look for and their proper sequence.

In the maintenance mode the system would describe the sequence of tests or examinations that should be performed and what to expect at each step.

In the repair mode the system would guide the operator on the correct tools, the precise method of disassembly, the required replacement parts and assemblies by name and identification numbers, and the proper procedure for reassembly. After repair the maintenance mode can be exercised to ensure by appropriate tests that repair has, in fact, been effected without disabling any other necessary function.

In any of the above operations the system would record the repairs, maintenance procedures, or conditions experienced by that piece of equipment. Users would thus have access to essential readiness information without needing bulky, hard-to-maintain maintenance records.

Current Projects and Experience

Some current Army and defense projects concerned with ATE are

- VTRONICS, a set of projects for onboard, embedded sensing of vehicular malfunctions with built-in test equipment (BITE);

20

- VIMAD, Voice Interactive Maintenance Aiding Device, which is external to the vehicle;
- Hawk missile computer-aided instruction for maintenance and repair.

Electronic malfunctions have been the subject of the most research, and electronics is now the most reliable aspect of the systems. Not much work has been done to reduce mechanical or software malfunctions. During wartime, however, such systems will need to be survivable under fire as well as be reliable under normal conditions.

For ground combat vehicles around 1990, a BITE diagnostic capability to tell the status of the vehicle power train is planned. In one development power train system, the critical information is normally portrayed either by cues via a series of gauges or by a digital readout. Malfunctions can be diagnosed through these cues and displays. The individual is prompted to push buttons to go through a sequence of displays.

An existing Army project concerns a helicopter cockpit display diagnostic system. One purpose of the project was to study audible information versus visual display. For example, the response to the FUEL command is to state the amount of fuel or flying time left; the AMMO command tells the operator how much ammunition is left. One reason for using speech output is that monitoring visual displays distracts attention from flying.

A lot of work has been done in the Army on maintenance and repair training, but computer-assisted instruction (CAI) and artificial intelligence could greatly reduce training time. For example, the M1 tank requires 60,000 pages of technical manuals to describe how to repair breakdowns.

The Army has planned for an AI maintenance tutor that would become a maintenance aid, but it is not yet funded. Under the VIMAD project supported by DARPA, a helmet with a small television receiver optically linked to a cathode ray tube (CRT) screen is being investigated as an aid to maintenance. Computer-generated video disk information is relayed.

An individual working inside the turret of an M1 tank, for example, cannot at present easily flip through the pages of the repair manual. With VIMAD, using a transmitter, receiver, floppy disk, and voice recognition capability, the individual can converse with the system to get information from the data base. The system allows a 19-word vocabulary for each of three individuals. The system has a 100-word capability to access more information from the main system and provides a combination of audio cues and visual prompts.

Any Army diagnostic system should be easily understood by any operator, regardless of maintenance background ("user friendly"). Choosing from alternatives presented in a menu approach, for example, is not necessarily easy for a semiliterate person.

Recommended Projects for Expert Systems in ATE

We propose that the following projects be supported as soon as possible:

- Interactive, mixed-media manuals for training and repair. Manuals should employ state-of-the-art video disk and display technology. The MIT Arcmac project, supported by the Office of Naval Research, illustrates this approach.
- Development of expert systems to trouble-shoot the 50 to 100 most common failures of important pieces of equipment. The system should incorporate simple diagnostic cues, be capable of fixed format (stylized, nonnatural) interaction, and emphasize quick fixes to operational machinery. The project should be oriented toward mechanical devices to complement the substantial array of existing electronic ATE. Projects in this category should be ready for operational use by 1987.
- Longer-term development of expert systems for ATE of more complex mechanical and electromechanical equipment. The systems in this category are intended for use at depots near battle lines. They are less oriented to quick fixes and incorporate preventive maintenance with more intelligent trouble shooting. They do not

aim for the sophisticated expertise of a highly qualified technician or mechanic. The emphasis is on (1) determining whether it is feasible to fix this piece of equipment, (2) determining how long it will take to fix, (3) determining if limited resources would be better used to fix other pieces of equipment, and (4) laying out a suitable process for fixing the equipment.

- The trouble-shooting systems recommended above rely on human sensors, exactly like MYCIN and Prospector. MYCIN is an expert system for diagnosing and treating infectious diseases that was developed at Stanford University. Prospector, developed at SRI International, is an expert system to aid in exploration for minerals. Parallel, longer-term efforts should be started to incorporate automatic sensors into the trouble-shooting expert systems recommended above.

EXPERT SYSTEMS FOR ARMY MEDICAL APPLICATIONS

Expert systems for various areas of medicine are being extensively studied at a number of institutions in the United States. These include

- rule-based systems at Stanford (MYCIN) and Rutgers (for glaucoma) ,
- Bayesian statistical systems (for computer-assisted diagnosis of abdominal pain),
- cognitive model systems (for internal medicine, nephrology, and cholestasis) ,
- knowledge management systems for diagnosis of neurological problems at Maryland.

22

Current Army activities to apply robotics and artificial intelligence in the medical area are described in the Army Medical Department's AI/Robotics plan, which was prepared with the help of the Academy of Health Sciences, San Antonio. This plan was presented to this committee by the U.S. Army Medical Research and Development Command (AMRDC).

Current Army Activities

Purdue University's Bioengineering Laboratory has an Army contract to study the concept of a "dog-tag chip" that will assist identification of injured personnel. The goal for this device is to assist in the display of patient symptoms for rapid casualty identification and triage. AMRDC noted that visual identification of casualties in chemical and biological warfare may be very difficult because of the heavy duty garb that will be worn.

Airborne or other remote interrogation of the dog-tag chip, its use in self-aid and buddy-aid modes, and use of logic trees on the chip for chemical warfare casualties are being examined by the Army. Other areas of AI and robotics listed in the U.S. AMRDC plan are training, systems for increased realism, and a "smart aideman" expert system, the latter being a "pure" application of expert systems to assist in early diagnosis.

Medical Environments, Functions, and Payoffs

Medical environments likely to be encountered in the Army are

- routine nonbattle, general illnesses, and disease;
- battle injuries, shock/trauma;
- epidemics;
- chemical;
- radiation;
- bacteriological.

In a battle area, a medical diagnosis paramedic aide machine would

- speed up diagnosis by paramedic and provide productivity increase, noninvasive sensing, and triage;
- suggest the best drugs to give for a condition, subject to patient allergies;
- suggest priority, disposition, and radio sensor signals on a radio link to field hospital, if necessary to consult physician.

At forward aid stations, in addition to routine diagnostic help, the device might infer patterns of illness on the basis of reports from local areas, track patient condition over time, and teach paramedics the nature of conditions occurring in that particular area that may differ from their prior experience.

Payoffs would include increasing soldiers' likelihood of survival and the consequent boost to morale through the knowledge that efforts to save them were being assisted by the latest technology. Note that the automated battalion information management system, described below, will involve building a large planning model, which could include medicine.

Recommended Medical Expert Systems

In view of existing technology, a more aggressive dog-tag chip program than that already under way at Purdue University is advocated. The Army should contract with some commercial company currently making wristwatch monitors to develop a demonstration model Army body monitor and not worry if the development gets out into the public domain. Wristwatch monitors of pulse rate, temperatures, etc., are listed in catalogs such as the one from Edmund Scientific.

Technology for low-level digital communication with cryptography is also available. As a prerequisite to the smart dog-tag, the Army may wish to make use of this technology in various Army systems more mundane than the smart dog-tag chip. Cryptography can ensure that information on a smart dog-tag is not susceptible to interception.

Collection of data on noninvasive new and old sensors and related methods of statistical analysis to determine their efficiency in monitoring casualty/injury conditions should be the subject of a longer term study. The study should create a data base that relates medical diagnosis and sensor capabilities.

The development of AI expert systems aimed at providing computer consulting for nonbattle and battle-area Army medicine and paramedical training are long-term projects that could be undertaken in collaboration with military and university hospitals. For example, the emergency room or shock/trauma unit of a civilian hospital could be used in beginning studies. Correlation of the patient's current condition with past medical history as recorded on a soldier's dog-tag chip would be one result available from an expert system. Paramedic skills may or may not require a slight increase, depending on how well the AI aid is designed. It does seem that the same number of paramedics should be able to accomplish more.

FLEXIBLE MATERIAL-HANDLING MODULES

Most robot applications in industry today are directly related to material handling. These include loading and unloading machines, palletizing, feeding parts for other automation equipment, and presenting parts for inspection.

Material handling in Army operations has many similar applications, which, at the very least, involve a great number of repetitive operations and often require working under hazardous conditions. It is proposed to make use of state-of-the-art robotics to develop a

24

multifunctional, material-handling robotic module that can be readily adapted for many Army functions serving both rear echelon and front line supply needs.

An ammunition resupply robot could select, prepare, acquire, move, load, or unload ammunition at forward weapon sites to reduce exposure of personnel or in rear storage areas to reduce personnel requirements and provide 24-hour capability.

For general use, a robot mounted on a wheeled base is recommended so that the human operator can maneuver the robot into position and then initiate a stored computer program that it will execute without continuous supervision. With present technology constraints on the necessary vision system, it would be necessary to have a bar-code identifying insignia affixed to every package or object in a known position. State-of-the-art pattern recognition devices can then be mounted on the robot arm to identify an object or package for sorting and verification. Future technological advance would reduce the need for identifying insignia.

The proposed robot to refuel vehicles is actually an instance of a material-handling module. It would be mounted on wheels and equipped with vision. The operator would position the robot in the proximate location, where it would then use a fuel dispenser without exposing the crew. Special gas tank caps would be required to facilitate insertion and dispensing of fuel by the robot.

Functional Requirements

The module would be a fully programmable, servo-driven robot with advanced controller capable of interfacing with a vision module, other sensor modules, and teleoperator control. It

would include a teach-box programmer to provide the simplest programming capability by unit-level nonspecialists. The teleoperator would provide the operator with the ability to operate the robot on one-at-a-time tasks that do not require repetitive operations or are too difficult to program for automatic operation.

The robot module base would be designed to be readily mounted on a truck, a trailer, or a weapons carrier, or emplaced on a rigid pad or even firmly embedded in the ground. It would be desirable to engineer several different sizes with different load capacities but operating with identical controllers.

High speed and precision would be desirable but not mandatory. Trade-offs for ruggedness, simplicity, maintainability, and cost should be considered seriously.

Provision would be made for readily interchangeable end effectors, or "hands." Each application would have a specialized end effector, which could be a gripper or tool. The particular requirements of the task or mission would specify which set of effectors accompany the robot.

Some near-term advantages are

- In supply logistics the module could stack such items as packages or ammunition, from either trucks or supply depots, where standard pallet operations are not available or feasible. Many personnel engaged in all forms of moving supplies and munitions would become acquainted with and adept at the use of this strength-enhancing, labor-saving tool. Reduction of staff and elimination of many repetitive and fatiguing operations would result. Key personnel would be time-shared, since a single operator could set up and supervise several robot systems.
- In front line and other hazardous activities, the robot module, after programming, could operate autonomously or under supervisory control from a safe location. Ammunition and fuel resupply for tanks serviced by a robot mounted on a protected vehicle is a typical example. Handling hazardous chemical or nuclear objects or material could be performed remotely. Retrieving and delivering objects under fire may be possible with appropriate remote-controlled vehicles.
- When personnel become familiar and experienced with these systems, they will probably generate and jury-rig a robot to perform new operations creatively. This system is meant to be a general-purpose helper.

The long-range advantages include the following:

- With the future addition of a wide range of sensors, including vision, tactile, force, and torque, the robot module becomes part of an intelligent robot system, enlarging its field of application to parallel many intended uses of systems in industry. With specialized tools, maintenance, repair, reassembly, testing, and other normal functions to maintain

sophisticated weapon systems, all become possible, especially under hazardous conditions.

- The proposed module can be readily duplicated at reasonable cost and serve at many experimental sites for evaluation and development into practical tools. It will undoubtedly uncover needs requiring advanced capabilities that can be added without complete redesign.

AUTOMATED BATTALION INFORMATION MANAGEMENT SYSTEM

Combat operations in a modern army require vast amounts of information of varying completeness, timeliness, and accuracy. Included are operational and logistic reports on the status of friendly and enemy forces and their functional capabilities, tactical analyses, weather, terrain, and intelligence input from sensors and from human sources. The information is often inconsistent and fragmentary but in sufficient quantity to lead to information overload, requiring sorting,

26

classification, and distribution before it can be used. Getting the information to the appropriate people in a timely fashion and in a usable form is a major problem.

A battalion forward command post is usually staffed by officers having responsibility for operations, intelligence, and fire support. These officers are seconded by enlisted personnel with significantly less schooling and experience. Other battalion staff officers assist, but they do not carry the main burden. The battalion executive officer usually positions himself where he can best support the ongoing operation. Together, these men simultaneously fight the current battle and plan the next operation. Thus, efforts must be made to alleviate fatigue and stress. There is a consequent need for automated decision aids.

Expert systems for combat support could assist greatly. It appears that information sources consist currently of hand-written, repeatedly copied reports and that intelligence operations integration is degraded because of information overload and because information is inconsistent. Thus, while capable of intuitive judgments that machines do poorly, officers find it difficult to integrate unsorted and unrelated information, are limited in their ability to examine alternatives, and are slow to recognize erroneous information. Decisionmaking in tense situations is spontaneous and potentially erroneous.

Capturing the knowledge of an officer, even in a highly domain-restricted situation such as a forward command post, is difficult. Even though they strain the state of the art, expert systems for combat support have such potential payoff in increasing combat effectiveness that they should receive high priority and be begun immediately. The following sequence of projects can be identified:

- how to capture and deploy knowledge and duties of the operations, intelligence, logistics, and fire-support officers into operations, intelligence, logistics, and fire-support expert systems to aid these officers;
- how to automate screening messages and establishing priorities to reduce information overload;
- how to integrate the operations of the expert systems to support the command;
- how to integrate general information with detailed information about the particular situation at hand; for example, how supplemental experts for multisensor reconnaissance and intelligence, topographic mapping, situation mapping, and other functions such as night attack and air assault can be used to adapt the general battalion expert system to the particular battle situation.

5 IMPLEMENTATION OF RECOMMENDED APPLICATIONS

For the applications recommended in Chapter 4, the committee made gross estimates of the time, cost, and technical complexity/risk associated with each. The results of those deliberations are summarized in this chapter.

The matrix on the following pages was developed to present the committee ' s proposed implementation plan. For each candidate, the matrix shows the estimated time and man-years of effort from initiation of contractual effort until demonstration of the concept by a bread- or brass-board model, gross estimates of costs for a single contractor, projected payoff, relative technical complexity, remarks, and, finally, recommended priority in which projects should be undertaken. In light of constrained funding and even more strictly limited technical capacity, we recommend that one candidate in each of the three areas--effectors, sensors, and cognition--be undertaken now. The recommended top-priority applications are the automatic loader of ammunition in tanks (effectors), the sentry/surveillance robot (sensors), and the intelligent maintenance, diagnosis, and repair system (cognition).

While the committee agreed that it would be preferable in all cases for at least two firms to undertake R&D simultaneously, it recognized that constrained funding would probably preclude such action. Cost estimates in the matrix, therefore, represent the committee ' s estimate of the costs of a single contractor based on the number of man years of a fully supported senior engineer. Believing that the Army was in far better position to estimate its administrative, in-house, and testing costs, the committee limited its cost estimates to those of the contractor.

After extensive discussion, the committee chose \$200,000 as a reasonable and representative estimate of the cost of a fully burdened industrial man-year for a senior engineer. The estimated costs for contractor effort for different supported man-year costs can be calculated. The estimates given are for demonstrators, not for production models.

OBJECT	TECHNICAL COMPLEXITY	TIME FRAME	COSTS (M=Million) (MY=Man-year)	PAYOFF
Automatic Loader of Ammunition in Tanks	5 (low)	2 to 2-1/2 yr	\$4-\$5 M/system/contractor 5 MY/yr (senior people)	Reduce crew from 4 to 3 to 2
Surveillance/Sentry Robot	2 (high)	2 to 3 yr	\$5 M (2 or more sensor modalities) 15 MY \$5 M/yr	Improve rear area surveillance Augment ability of existing personnel
Intelligent Maintenance, Diagnosis, and Repair System	4	2 yr	\$1 M; 2 MY/yr for 2 yr	Don't abandon when bandaid fixes are possible
		3 yr	\$2 to 2 1/2 M; 3 MY/yr for 3 yr (prototype)	Amplify mechanic's skills
		3 yr	\$2 to 2-1/2 M; 3 MY/yr for 3 yr <hr/> Total 22 MY	

Medical Expert System	1 (most)	3 yr	15 MY	Morale, efficiency, survival Reduce training load for medics
Flexible Material-Handling Modules	6 (least)	2-3 yr	\$4-5 M/ system/ contractor 15 MY (5 MY/yr)	Performance (productivity) Safety Training
Battlefield Information Management System	3	5 or more yr	\$2-1/2 to 3 M/ yr for 2 MY/yr for each of 4 ES and 1 overseer Total \$12-1/2 to 15 M for 5 yr (9 MY/yr for 5 yr = 45 MY)	Reduce staff by duplicating staff functions Facilitate 24 operations In time, provide improved data for decisionmaking

MEASURES OF EFFECTIVENESS

The committee had considerable difficulty in attempting to develop useful measures of effectiveness because such measures appear to be meaningful only as applied to a specific application. Even then, the benefits of applying robotics and artificial intelligence are often difficult to quantify at this early stage. How, for example, does one measure the value of a human life or of increments in the probability of success in battle?

Therefore, instead of attempting to develop quantitative measures that strain credibility, the committee offers general guidelines against which to measure the worthiness of proposed applications of robotics and artificial intelligence. These guidelines are grouped according to their intended effect.

People

- Reduced danger or improved environment
- Reduced skill level or training requirements

- Improved survivability

Mission

- Improved productivity or reduced manpower requirements
- Military advantage
- New opportunities
- Enhanced capability to conduct 24-hour per day operations
- Improved RAMS (reliability, availability, maintainability, and supportability)

Material

- Reduced cost

The final item, reduced cost, is not the only one that can be assigned a quantitative value. A reduced need for training, for example, should result in reduced training costs. Similarly, improvements in RAMS should reduce life-cycle costs because of diminished need for repair parts, reduced maintenance costs stemming from greater mean time between failure, and reduced maintenance man-hours per maintenance action. However, meaningful estimates with acceptable levels of confidence would require large volumes of experience data that simply are not available at this early stage in the development of a new and revolutionary technology.

Military advantage is probably the ultimate measure of effectiveness. For example, if it could be shown through modeling or gaming that investment in a system meant the difference between winning or losing, that system could be described as infinitely cost effective.

The committee simply does not have access to sufficient pertinent information to make other than a subjective judgment of the effectiveness of its proposed applications at this time. Further, because each application is to be implemented progressively, such measures will change over time. Finally, because the final versions of the applications require substantial research and development, the committee, despite its collective experience, can provide only the gross estimates of probable costs and payoffs contained in the matrix.

What, then, can the committee say about measuring the effectiveness of the proposed applications? First, that in its collective judgment, the recommended applications provide sound benefits for the Army and second, that these benefits will stem from more than one of the nine areas listed above.

A possible precedent to consider is the manner in which DOD funded the Very High Speed Integrated Circuits (VHSIC) program. It was considered an area of great promise that warranted funding as a matter of highest priority; applications were sought and found later on, after the research was well under way. Similarly, there is little question that we have barely begun to scratch the surface in identifying high-payoff applications of robotics and artificial intelligence technology.

6 OTHER CONSIDERATIONS

In the course of its studies, the committee identified a number of important considerations that can be expected to bear heavily on the Army's decisions on future applications of robotics and AI technology. These considerations, discussed in the paragraphs that follow, apply more generally than to the specific topics covered in the previous chapters.

SHORTAGE OF EXPERTS

Probably the most important single consideration at this time is that there are far too few research experts in the areas of robotics and artificial intelligence. Most of those available to the Army for their applications are clustered in a few universities where some 70 professors with an average of 4 to 5 (apprentice) students apiece represent the bulk of existing technical expertise. There are appreciably fewer qualified practitioners in military service. As a result, despite the fact that additional funding in these areas is required, it must be allocated with great care to ensure that recipients have the capability to spend the money wisely and effectively. For example, SRI is unable to accept more money for some branches of AI because its technical capacity is already fully committed.

Similarly, there is a critical shortage of military experts in the domains to be captured by expert systems. In particular, it is difficult to find the military officers required to participate in the design and development of complex expert systems, such as those required for division and corps tactical operations centers.

Both factors underline the need for an Army-university partnership in educating qualified individuals in order to expand the research and development base as soon as possible. They also appear to indicate a need for some sort of centralized coordination, to ensure that optimum use is made of the limited human and fiscal resources available.

OPERATOR-FREINDLY SYSTEMS

The creation of operator-friendly systems is essential to the successful spread of this technology. A truly operator-friendly system will appeal to all levels of people, especially under adverse conditions. In addition, these systems will facilitate the important task of getting novices acquainted with and accustomed to using robots and robotic systems. Not only will this lead to the critically needed confidence that comes from hands-on experience, but it will also demonstrate the reality of what can be done now and point the way toward more advanced applications of the future.

The importance of operator-friendly hardware has been recognized by the military since World War II, when the studies of aircraft accidents identified a number of pilot errors caused by the design of the plane. Since then, military R&D has included the analysis of human factors in the

design of new technologies. Expected benefits include fewer accidents, improved performance, reduced production costs, lower training costs, and improved implementation.

Operator-friendly systems are of particular importance to the military because the objective is to ensure proper use of the systems under less than favorable conditions. In most cases the environmental conditions in which the robot will be expected to operate are more severe than those currently experienced in industrial applications. Furthermore, in times of crisis the robot may need to be operated by or work with personnel that are not fully trained. Careful design of the hardware and software can reduce training, maintenance, and repair costs. It can also ensure that the expected benefits are more likely to be achieved.

In some environments, such as tanks, humans and robots will be working in close quarters. If there is hostility or difficulty with the robotic system, or if the maneuvers require too much space or movement, the system will not work effectively. In a crisis, there may not be a second chance or an available backup for a system failure, so the man-machine combination must work effectively and quickly.

Essential to any operator-friendly system are high levels of reliability, availability, and maintainability, and redundant fail-safe provisions. With the many hostile environments, it will be of basic importance to assure adequate redundancy in components and systems. What are the backups? What happens when power fails? Can muscle power operate the system?

As military equipment becomes increasingly complex, its operation and maintenance will compete with industry for scarce mechanical and computer skills. This shortage of experts and trained skilled workers can be ameliorated by robotic applications, such as maintenance and repair aids.

COORDINATION OF EXISTING PROGRAMS

The committee is concerned that specific efforts be made to guard against reinventing the wheel. With so many programs in the armed services, it appears to outsiders that many activities are repeated because each particular area wants its own activity. The Army should have some means of knowing the programs in the other services that could have application to Army needs. The committee has learned that the Joint Laboratory Directors, operating under the aegis of the Joint Logistics Commanders, have begun to address this important need. Any steps that foster communication in this area are to be welcomed.

AVAILABLE TECHNOLOGY

There are already a number of successful applications of robotics in use in industry. Such applications as spot welding, arc welding, palletizing, and spray painting are not exotic and are proven successes. The Army can improve its operations immediately by taking advantage of commercially proven systems for production and maintenance in its depots.

GETTING STARTED

The Army will experience the same growing problems that industry has experienced. Outside of a few areas like robotic spot welding of automobiles and robotic unloading of die casting machines, there has been much talk about robotic applications but only slow growth. There is evidence that implementation of robotics projects will now move at a much faster pace. The Army should bear in mind, however, that getting a dynamic technological program going almost invariably requires more time and money than its developers originally plan.

These technologies will cause a savings in manpower, though not necessarily for the initial thrust. Experience and training will be needed in all areas--operators, maintenance personnel, supervisors, and managers. Once the new systems are understood by all levels, then the savings will be realized. In many cases this savings will take the form of more output per unit. In addition, the savings will compound as the systems grow with technology additions as well as familiarity.

An important by-product following the initial learning period will be the motivation of individuals. Being master of a phase of new technology gives one an accomplishment and ability that can be the base for growth within the existing employment area or for selling personal ability and knowledge outside the area--in short, a ladder for growth and personal development.

35

FOCUS FOR AI AND ROBOTICS

The committee has noted that the Army has identified the five technology thrusts of Very Intelligent Surveillance and Target Acquisition (VISTA),

Distributed Command, Control, Communications and Intelligence, Self-Contained Munitions, Soldier-Machine Interface, Biotechnology.

These are areas to which it intends to devote its research and exploratory development efforts.

Robotics and artificial intelligence technology is not designated as a separate high-priority thrust. It is possible to relate specific robotics/AI applications to one or more of the technology thrusts, as the Army Science Board Ad Hoc Group on Artificial Intelligence and Robotics did in its report. However, the danger remains that robotics and AI efforts--particularly where they do not fall clearly under the mantle of one of the chosen five--will be considered lower priority, with the attendant implications of reduced funding and support. Failure to identify robotics and AI as a special thrust may also contribute to the lack of focus in management and diffusion of effort and funding noted elsewhere in this report.

IMPLEMENTATION DIFFICULTIES

In addition to technical barriers that might normally be expected, several misconceptions have continually clouded industry's technology development and ongoing research in artificial

intelligence. Unrealistic expectations combined with problems inherent in any new technology have created barriers to easy implementation. Based on recent industrial experiences, the Army can expect these to include

- Unrealistic expectations of the technology's capabilities. In an extremely narrow context, some expert systems outperform humans (e.g., MACSYMA), but certainly no machine exhibits the commonsense facility of humans at this time. Machines cannot outperform humans in a general sense, and that may never be possible. Further, the belief that such systems will bail out current or impending disasters in more conventional system developments that are presently under way is almost always erroneous.
- The technology is not readily learned. The notion that "this is nothing more than smart software" continually demonstrates the naiveté of first impressions. Current experience in industry refutes this contention. A seemingly simple concept of knowledge acquisition,

36

simply having an expert state his rules of thumb, is currently an intricate art and so complex as to defy automatic techniques. It is, and will remain for some time, a research area.

- Expectations often dramatically exceed what is possible. This is particularly true of the times estimated for development. Performance of the systems has often lagged because of such problems as classification restrictions or a lack of available expertise.
- Desire for quick success. Very often the political goals are not consonant with the technical goals, thereby increasing the risk associated with developing an expert system by placing unrealistic time constraints on the staff.
- University goals versus the goals of industry. Top research universities are motivated to gain new knowledge, develop researchers, publish papers and dissertations, and establish a vehicle for the perpetuation of these. The goals of a responsive industrial unit are to build a system or provide a service that results in a usable, functioning system in an acceptable time to meet the needs of the customer for use by practitioners. Because of this diversity of purpose, much of the software and hardware developed is not easily transferable, and costly transformations have been required.
- Fear of not succeeding. This is as detrimental to technological progress as in any other art or science. Industry and government have often committed funds to unambitious projects that met inadequate risks in order to prove nothing.
- Calling it AI when it is not or is only loosely related. The expectation that development in this area will be readily funded encourages jumping on bandwagons.
- Lack of credentials. Several people and groups are claiming expertise in AI, though they may not have the rich base upon which research capability is normally developed. Careful credential checking is imperative.
- Technology transfer. The preponderance of practitioners are in the universities and have only recently been moving to industry, primarily to venture activities. Most have never delivered products in the industrial context (e.g., documented with life-cycle

considerations). The transfer of knowledge to industry at large is thus rarely done by those with knowledge of both industry and the technology, which makes the industrialization process more risky.

- *Premature determination of results.* The risk exists of unwittingly predetermining the outcome of decisions that should be made after further research and development. The needed skills simply are not in industry or in the government in the quantities needed to prevent this from happening on occasion.
- *Nontransferable software tools.* Virtually all software knowledge engineering systems and languages are scantily documented and often only supported to the extent possible by the single researcher who originally wrote it. The universities are not in the business to assure proper support of systems for the life-cycle needs of the military and industry, although some of the new AI companies are beginning to support their respective programming environments.

37

- *Lack of standards.* There are no documentation standards or restrictions on useful programming languages or performance indices to assess system performance.
- *Mismatch between needed computer resources and existing machinery.* The symbolic languages and the programs written are more demanding on conventional machines than appears on the surface or is being advertised by some promoters.
- *Knowledge acquisition is an art.* The successful expert systems developed to date are all examples of handcrafted knowledge. As a result, system performance cannot be specified and the concepts of test, integration, reliability, maintainability, testability, and quality assurance in general are very fuzzy notions at this point in the evaluation of the art. A great deal of work is required to quantify or systematically eliminate such notions.
- *Formal programs for education and training do not exist.* The academic centers that have developed the richest base of research activities award the computer science degree to encompass all sub-disciplines. The lengthy apprenticeship required to train knowledge engineers, who form the bridge between the expert and development of an expert system, has not been formalized.

38

7 RECOMMENDATIONS

START USING AVAILABLE TECHNOLOGY NOW

Robotics and artificial intelligence technology can be applied in many areas to perform useful, valuable functions for the Army. As noted in Chapter 3, these technologies can enable the Army to

- improve combat capabilities,
- minimize exposure of personnel to hazardous environments,
- increase mission flexibility,

- increase system reliability,
- reduce unit/life cycle costs,
- reduce manpower requirements,
- simplify training.

Despite the fact that robotics technology is being extensively used by industry (almost \$1 billion introduced worldwide in 1982, with increases expected to compound at an annual rate of at least 30 percent for the next 5 to 10 years), the Army does not have any significant robot hardware or software in the field. The Army's needs for the increased efficiency and cost effectiveness of this new technology surely exceed those of industry when one considers the potential reduction in risk and casualties on the battlefield.

The shrinking manpower base resulting from the decline in the 19-to 21-year-old male population, and the substantial costs of maintaining present Army manpower (approximately 29 percent of the total Army budget in FY 1983), emphasize that a major effort should be made to conserve manpower and reduce battlefield casualties by replacing humans with robotic devices.

The potential benefits of robotics and artificial intelligence are clearly great. It is important that the Army begin as soon as possible so as not to fall further behind. Research knowledge and practical industrial experience are accumulating. The Army can and should begin to take advantage of what is available today.

39

CRITERIA: SHORT-TERM, USEFUL APPLICATIONS WITH PLANNED UPGRADES

The best way for the Army to take advantage of the potential offered by robotics and AI is to undertake some short-term demonstrators that can be progressively upgraded. The initial demonstrators should

- meet clear Army needs,
- be demonstrable within 2 to 3 years,
- use the best state of the art technology available,
- have sufficient computer capacity for upgrades,
- form a base for familiarizing Army personnel--from operators to senior leadership--with these new and revolutionary technologies.

As upgraded, the applications will need to be capable of operating in a hostile environment.

The dual approach of short-term applications with planned upgrades is, in the committee's opinion, the key to the Army's successful adoption of this promising new technology in ways that will improve safety, efficiency, and effectiveness. It is through experience with relatively simple applications that Army personnel will become comfortable with and appreciate the benefits of these new technologies. There are indeed current Army needs that can be met by available robotics and AI technology.

In the Army, as in industry, there is a danger of much talk and little concrete action. We recommend that the Army move quickly to concentrate in a few identified areas and establish those as a base for growth.

SPECIFIC RECOMMENDED APPLICATIONS

The committee recommends that, at a minimum, the Army should fund the three demonstrator programs described in Chapter 4 at the levels described in Chapter 5:

- The Automatic Loader of Ammunition in Tanks, using a robotic arm to replace the human loader of ammunition in a tank. We recommend that two contractors work simultaneously for 2 to 2 1/2 years at a total cost of \$4 to \$5 million per contractor.
- The Surveillance/Sentry Robot, a portable, possibly mobile platform to detect and identify movement of troops. Funded at \$5 million for 2 to 3 years, the robot should be able to include two or more sensor modalities.
- The Intelligent Maintenance, Diagnosis, and Repair System, in its initial form (\$1 million over 2 years), will be an interactive trainer. Within 3 years, for an additional \$5 million, the system should be expanded to diagnose and suggest repairs for common breakdowns, recommend whether or not to repair, and record the repair history of a piece of equipment.

40

If additional funds are available, the other projects described in Chapter 4, the medical expert system, the flexible material-handling modules, and the battalion information management system, are also well worth doing.

VISIBILITY AND COORDINATION OF MILITARY AI/ROBOTICS

Much additional creative work in this area is needed. The committee recommends that the Army provide increased funding for coherent research and exploratory development efforts (lines 6.1 and 6.2 of the budget) and include artificial intelligence and robotics as a special technology thrust.

The Army should aggressively take the lead in pursuing early application of robotics and AI technologies to solve compelling battlefield needs. To assist in coordinating efforts and preventing duplication, it may wish to establish a high-level review board or advisory board for the AI/Robotics program. This body would include representatives from the universities and industry, as well as from the Army, Navy, Air Force, and DARPA. We recommend that the Army consider this idea further.

41

APPENDIX STATE OF THE ART AND PREDICTIONS FOR

ARTIFICIAL INTELLIGENCE AND ROBOTICS

INDUSTRIAL ROBOTS: FUNDAMENTAL CONCEPTS

The term robot conjures up a vision of a mechanical man--that is, some android as viewed in *Star Wars* or other science fiction movies. Industrial robots have no resemblance to these *Star Wars* figures. In reality, robots are largely constrained and defined by what we have so far managed to do with them.

In the last decade the industrial robot (IR) has developed from concept to reality, and robots are now used in factories throughout the world. In lay terms, the industrial robot would be called a mechanical arm. This definition, however, includes almost all factory automation devices that have a moving lever. The Robot Institute of America (RIA) has adopted the following working definition:

A robot is a programmable multifunction device designed to move material, parts, tools, or specialized devices through variable programmed motions for the performance of a variety of tasks.

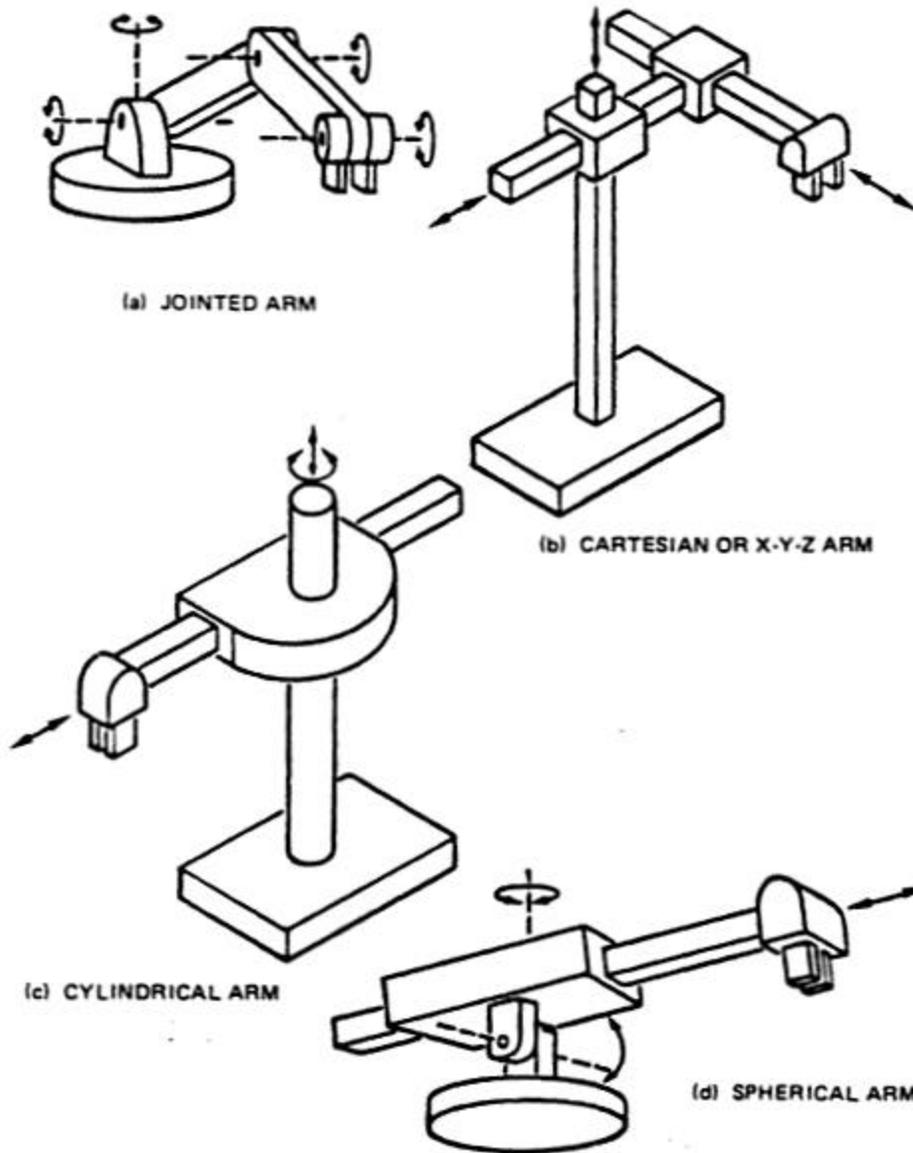
It is generally agreed that the three main components of an industrial robot are the mechanical manipulator, the actuation mechanism, and the controller.

The mechanical manipulator of an IR is made up of a set of axes (either rotary or slide) , typically three to six axes per IR. The first three axes determine the work envelope of the IR, while the last

three deal with the wrist of the IR and the ability to orient the hand. Figure 1 shows the four basic IR configurations. Although these are typical of robot configurations in use today, there are no hard and fast rules that impose these constraints. Many robots are more

The appendix is largely the work of Roger Nagel, Director, Institute for Robotics, Lehigh University. James Albus of the National Bureau of Standards and committee members J. Michael Brady, Stephen Dubowsky, Margaret Eastwood, David Grossman, Laveen Kanal, and Wendy Lehnert also contributed.

FIGURE 1
Four Common Arrangements of Robot Manipulator Joints



restricted in their motions than the six-axis robot. Conversely, robots are sometimes mounted on extra axes such as an x-y table or track to provide an additional one or two axes.

It is important to note at this point that the "hand" of the robot, which is typically a gripper or tool specifically designed for one or more applications, is not a part of a general purpose IR. Hands, or end effectors, are special purpose devices attached to the "wrist" of an IR.

The `actuation mechanism` of an IR is typically either hydraulic, pneumatic, or electric. More important distinctions in capability are based on the ability to employ servo mechanisms, which use feedback control to correct mechanical position, as opposed to nonservo open-loop actuation systems. Surprisingly, nonservo open-loop industrial robots perform many seemingly complex tasks in today's factories.

The `controller` is the device that stores the IR program and, by communications with the actuation mechanism, controls the IR motions. Controllers have undergone extensive evolution as robots have been introduced to the factory floor. The changes have been in the method of programming (human interface) and in the complexity of the programs allowed. In the last three years the trend to computer control (as opposed to plug board and special-purpose devices) has resulted in computer controls on virtually all industrial robots.

The `method of programming` industrial robots has, in the most popular and prevailing usage, not included the use of a language. Languages for robots have, however, long been a research issue and are now appearing in the commercial offerings for industrial robots. We review first the two prevailing programming methods.

Programming by the `lead-through` method is accomplished by a person manipulating a well-counterbalanced robot (or surrogate) through the desired path in space. The program is recorded by the controller, which samples the location of each of the robot's axes several times per second. This method of programming records a continuous path through the work envelope and is most often used for spray painting operations. One major difficulty is the awkwardness of editing these programs to make any necessary changes or corrections.

An additional--and perhaps the most serious--difficulty with the lead-through method is the inability to teach conditional commands, especially those that compute a sensory value. Generally, the control structure is very rudimentary and does not offer the programmer much flexibility. Thus, mistakes or changes usually require completely reprogramming the task, rather than making small changes to an existing program.

Programming by the `teach-box` method employs a special device that allows the programmer/operator to use buttons, toggle switches, or a joy stick to move the robot in its work envelope. Primitive teach boxes allow for the control only in terms of the basic axis motions of the robot, while more advanced teach boxes provide for the use of Cartesian and other coordinate systems.

The program generated by a teach box is an ordered set of points in the workspace of the robot. Each recorded point specifies the location of every axis of the robot, thus providing both position and orienta-

tion. The controller allows the programmer to specify the need to signal or wait for a signal at each point. The signal, typically a binary value, is used to sequence the action of the IR with another device in its environment. Most controllers also now allow the specification of velocity/acceleration between points of the program and indication of whether the point is to be passed through or is a destination for stopping the robot.

Although computer language facilities are not provided with most industrial robots, there is now the limited use of a `subroutine library` in which the routines are written by the vendor and sold as options to the user. For example, we now see `palletizing`, where the robot can follow a set of indices to load or unload pallets.

Limited use of simple sensors (binary valued) is provided by preprogrammed `search routines` that allow the robot to stop a move based on a sensor trip.

Typical advanced industrial robots have a computer control with a keyboard and screen as well as the teach box, although most do not support programming languages. They do permit subdivision of the robot program (sequence of points) into branches. This provides for limited creation of subroutines and is used for error conditions and to store programs for more than one task.

The ability to specify a `relocatable branch` has provided the limited ability to use sensors and to create primitive programs.

Many industrial robots now permit `down-loading` of their programs (and up-loading) over RS232 communication links to other computers. This facility is essential to the creation of flexible manufacturing system (FMS) cells composed of robots and other programmable devices. More difficult than communication of whole programs is communication of parts of a program or locations in the workspace. Current IR controller support of this is at best rudimentary. Yet the ability to communicate such information to a robot during the execution of its program is essential to the creation of `adaptive behavior` in industrial robots.

Some pioneering work in the area was done at McDonnell Douglas, supported by the Air Force Integrated Computer-Aided Manufacturing (ICAM) program. In that effort a Cincinnati Milacron robot was made part of an adaptive cell. One of the major difficulties was the awkwardness of communicating goal points to the robot. The solution lies not in achieving a technical breakthrough, but rather in understanding and standardizing the interface requirements. These issues and others were covered at a National Bureau of Standards (NBS) workshop in January 1980 and again in September 1982 [1].

`Programming languages` for industrial robots have long been a research issue. During the last two years, several robots with an off-line programming language have appeared in the market. Two factors have greatly influenced the development of these languages.

The first is the perceived need to hold a Ph.D., or at least be a trained computer scientist, to use a programming language. This is by no means true, and the advent of the personal computer, as well as the invasion of computers into many unrelated fields, is encouraging. Nonetheless, the fear of computers and of programming them continues.

45

Because robots operate on factory floors, some feel programming languages must be avoided. Again, this is not necessary, as experience with user-friendly systems has shown.

The second factor is the desire to have industrial robots perform complex tasks and exhibit adaptive behavior. When the motions to be performed by the robot must follow complex geometrical paths, as in welding or assembly, it is generally agreed that a language is necessary. Similarly, a cursory look at the person who performs such tasks reveals the high reliance on sensory information. Thus a language is needed both for complex motions and for sensory interaction. This dual need further complicates the language requirements because the community does not yet have enough experience in the use of complex (more than binary) sensors.

These two factors influenced the early robot languages to use a combination of language statements and teach box for developing robot programs. That is, one defines important points in the workspace via the teach-box method and then instructs the robot with language statements controlling interpolation between points and speed. This capability coupled with access to on-line storage and simple sensor (binary) control characterizes the VAL language. VAL, developed by Unimation for the Puma robot, was the first commercially available language. Several similar languages are now available, but each has deficiencies. They are not languages in the classical computer science sense, but they do begin to bridge the gap. In particular they do not have the capability to do arithmetic on location in the workplace, and they do not support computer communication.

A second-generation language capability has appeared in the offering of RAIL and AML by Automatix and IBM, respectively. These resemble the standard structured computer language. RAIL is PASCAL-based, and AML is a new structured language. They contain statements for control of the manipulator and provide the ability to extend the language in a hierarchical fashion. See, for example, the description of a research version of AML in [2].

In a very real sense these languages present the first opportunity to build intelligent robots. That is, they (and others with similar form) offer the necessary building blocks in terms of controller language. The potential for language specification has not yet been realized in the present commercial offerings, which suffer from some temporary implementation-dependent limitations.

Before going on to the topic of intelligent robot systems, we discuss in the next section the current research areas in robotics.

RESEARCH ISSUES IN INDUSTRIAL ROBOTS

As described previously, robots found in industry have mechanical manipulators, actuation mechanisms, and control systems. Research interest raises such potential topics as locomotion, dexterous hands, sensor systems, languages, data bases, and artificial intelligence. Although there are clearly relationships amongst these and other

46

research topics, we will subdivide the research issues into three categories: mechanical systems, sensor systems, and control systems.

In the sections that follow we cover manipulation design, actuation systems, end effectors, and locomotion under the general heading of mechanical systems. We will then review sensor systems as applied to robots--vision, touch, ranging, etc. Finally, we will discuss robot control systems from the simple to the complex, covering languages, communication, data bases, and operating systems. Although the issue of intelligent behavior will be discussed in this section, we reserve for the final section the discussion of the future of truly intelligent robot systems. For a review of research issues with in-depth articles on these subjects see Birk and Kelley [3].

Mechanical Systems

The design of the IR has tended to evolve in an ad hoc fashion. Thus, commercially available industrial robots have a repeatability that ranges up to 0.050 in., but little, if any, information is available about their performance under load or about variations within the work envelope.

Mechanical designers have begun to work on industrial robots. Major research institutes are now working on the kinematics of design, models of dynamic behavior, and alternative design structures. Beyond the study of models and design structure are efforts on direct drive motors, pneumatic servo mechanisms, and the use of tendon arms and hands. These efforts are leading to highly accurate new robot arms. Much of this work in the United States is being done at university laboratories, including those at the Massachusetts Institute of Technology (MIT), Carnegie-Mellon University (CMU), Stanford University, and the University of Utah.

Furthermore, increased accuracy may not always be needed. Thus, compliance in robot joints, programming to apply force (rather than go to a position), and the dynamics of links and joints are also now actively under investigation at Draper Laboratories, the University of Florida, the Jet Propulsion Laboratory (JPL), MIT, and others.

The implications of this research for future industrial robots are that we will have access to models that predict behavior under load (therefore allowing for correction), and we will see new and more stable designs using recursive dynamics to allow speed. The use of robots to apply force and torque or to deal with tools that do so will be possible. Finally, greater accuracy and compliance where desired will be available [4-8].

The method of actuation, design of actuation, and servo systems are of course related to the design and performance dynamics discussed above. However some significant work on new

actuation systems at Carnegie-Mellon University, MIT, and elsewhere promises to provide direct drive motors, servo-control pneumatic systems, and other advantages in power systems.

The end effector of the robot has also been a subject of intensive research. Two fundamental objectives--developing quick-change hands

and developing general-purpose hands--seek to alleviate the constraints on dexterity at the end of a robot arm.

As described earlier, common practice is to design a new end effector for each application. As robots are used in more complex tasks (assembly, for example), the need to handle a variety of parts and tools is unavoidable. For a good discussion of current end-effector technology, see Toepperwein et al. [9].

The quick-change hand is one that the robot can rapidly change itself, thus permitting it to handle a variety of objects. A major impediment to progress in this area is a lack of a standard method of attaching the hand to the arm. This method must provide not only the physical attachment but also the means of transmitting power and control to the hand. If standards were defined, quick-change mechanisms and a family of hand grippers and robot tools would rapidly become available.

The development of a dexterous hand is still a research issue. Many laboratories in this country and abroad are working on three-fingered hands and other configurations. In many cases the individual fingers are themselves jointed manipulators. In the design of a dexterous hand, development of sensors to provide a sense of touch is a prerequisite. Thus, with sensory perception, a dexterous hand becomes the problem of designing three robots (one for each of three fingers) that require coordinated control.

The control technology to use the sensory data, provide coordinated motion, and avoid collision is beyond the state of the art. We will review the sensor and control issues in later sections. The design of dexterous hands is being actively worked on at Stanford, MIT, Rhode Island University, the University of Florida, and other places in the United States. Clearly, not all are attacking the most general problem [10, 11], but by innovation and cooperation with other related fields (such as prosthetics), substantial progress will be made in the near future.

The concept of robot locomotion received much early attention. Current robots are frequently mounted on linear tracks and sometimes have the ability to move in a plane, such as on an overhead gantry. However, these extra degrees of freedom are treated as one or two additional axes, and none of the navigation or obstacle avoidance problems are addressed.

Early researchers built prototype wheeled and legged (walking) robots. The work originated at General Electric, Stanford, and JPL has now expanded, and projects are under way at Tokyo Institute of Technology, Tokyo University. Researchers at Ohio State, Rensselaer

Polytechnic Institute (RPI), and CMU are also now working on wheeled, legged, and in one case single leg locomotion. Perhaps because of the need to deal with the navigational issues in control and the stability problems of a walking robot, progress in this area is expected to be slow [12].

In a recent development, Odetics, a small California-based firm, announced a six-legged robot at a press conference in March 1983. According to the press release, this robot, called a "functionoid," can lift several times its own weight and is stable when standing on

48

only three of its legs. Its legs can be used as arms, and the device can walk over obstacles. Odetics scientists claim to have solved the mathematics of walking, and the functionoid does not use sensors. It is not clear from the press release to what extent the Odetics work is a scientific breakthrough, but further investigation is clearly warranted.

The advent of the wire-guided vehicle (and the painted stripe variety) offers an interesting middle ground between the completely constrained and unconstrained locomotion problems. Wire-guided vehicles or robot carts are now appearing in factories across the world and are especially popular in Europe. These carts, first introduced for transportation of pallets, are now being configured to manipulate and transport material and tools. They are also found delivering mail in an increasing number of offices. The carts have onboard microprocessors and can communicate with a central control computer at predetermined communication centers located along the factory or office floor.

The major navigational problems are avoided by the use of the wire network, which forms a "freeway" on the factory floor. The freeway is a priori free of permanent obstacles. The carts use a bumper sensor (limit switch) to avoid collisions with temporary obstacles, and the central computer provides routing to avoid traffic jams with other carts.

While carts currently perform simple manipulation (compared to that performed by industrial robots), many vendors are investigating the possibility of robots mounted on carts. Although this appears at first glance to present additional accuracy problems (precise self-positioning of carts is still not available), the use of cart location fixturing devices at stations may be possible.

Sensor Systems

The robot without sensors goes through a path in its workspace without regard for any feedback other than that of its joint resolvers. This imposes severe limitations on the tasks it can undertake and makes the cost of fixturing (precisely locating things it is to manipulate) very high. Thus there is great interest in the use of sensors for robots. The phrase most often used is "adaptive behavior," meaning that the robot using sensors will be able to deal properly with changes in its environment.

Of the five human senses--vision, touch, hearing, smell, and taste--vision and touch have received the most attention. Although the Defense Advanced Research Projects Agency

(DARPA) has sponsored work in speech understanding, this work has not been applied extensively to robotics. The senses of smell and taste have been virtually ignored in robot research.

Despite great interest in using sensors, most robotics research lies in the domain of the sensor physics and data reduction to meaningful information, leaving the intelligent use of sensory data to

the artificial intelligence (AI) investigators. We will therefore cover sensors in this chapter and discuss the AI implications later.

Vision Sensors

The use of vision sensors has sparked the most interest by far and is the most active research area. Several robot vision systems, in fact, are on the market today. Tasks for such systems are listed below in order of increasing complexity:

- identification (or verification) of objects or of which of stable states they are in,
- location of objects and their orientation,
- simple inspection tasks (is part complete? cracked?),
- visual servoing (guidance),
- navigation and scene analysis,
- complex inspection.

The commercial systems currently available can handle subsets of the first three tasks. They function by digitizing an image from a video camera and then thresholding the digitized image. Based on techniques invented at SRI and variations thereof, the systems measure a set of features on known objects during a training session. When shown an unknown object, they then measure the same feature set and calculate feature distance to identify the object.

Objects with more than one stable state are trained and labeled separately. Individual feature values or pairs of values are used for orientation and inspection decisions.

While these systems have been successful, there are many limitations because of the use of binary images and feature sets--for example, the inability to deal with overlapped objects. Nevertheless, in the constrained environment of a factory, these systems are valuable tools. For a description of the SRI vision system see Gleason and Aguin [13]; for a variant see Lavin and Lieberman [14].

Not all commercial vision Systems use the SRI approach, but most are limited to binary images because the data in a binary image can be reduced to run length code. This reduction is important because of the need for the robot to use visual data in real time (fractions of a second). Although

one can postulate situations in which more time is available, the usefulness of vision increases as its speed of availability increases.

Gray-scale image operations are being developed that will overcome the speed problems associated with nonbinary vision. Many vision algorithms lend themselves to parallel computation because the same calculation is made in many different areas of the image. Such parallel computations have been introduced on chips by MIT, Hughes, Westinghouse, and others.

Visual servoing is the process of guiding the robot by the use of visual data. The National Bureau of Standards (NBS) has developed a special vision and control system for this purpose. If robots are ever

50

to be truly intelligent, they must be capable of visual guidance. Clearly the speed requirements are very significant.

Vision systems that locate objects in three-dimensional space can do so in several ways. Either structured light and triangulation or stereo vision can be used to simulate the human system. Structured light systems use a shaped (structured) light source and a camera at a fixed angle [15]. Some researchers have also used laser range-finding devices to make an image whose picture elements (pixels) are distances along a known direction. All these methods--stereo vision, structured light, laser range-finding, and others--are used in laboratories for robot guidance.

Some three-dimensional systems are now commercially available. Robot Vision Inc. (formerly Solid Photography), for example, has a commercial product for robot guidance on the market. Limited versions of these approaches and others are being developed for use in robot arc welding and other applications [16].

Special-purpose vision systems have been developed to solve particular problems. Many of the special-purpose systems are designed to simplify the problem and gain speed by attacking a restricted domain of applicability. For example, General Motors has used a version of structured light for accumulating an image with a line scan camera in its Consight system. Rhode Island University has concentrated on the bin picking problem. SRI, Automatix, and others are working on vision for arc welding.

Others such as MIT, University of Maryland, Bell Laboratories, JPL, RPI, and Stanford are concentrating on the special requirements of robot vision systems. They are developing algorithms and chips to achieve faster and cheaper vision computation. There is evidence that they are succeeding. Special-purpose hardware using very large-scale integration (VLSI) techniques is now in the laboratories. One can, we believe, expect vision chips that will release robot vision from the binary and special-purpose world in the near future.

Research in vision, independent of robots, is a well-established field. That literature is too vast to cover here beyond a few general remarks and issues. The reader is referred to the literature on image processing, image understanding, pattern recognition, and image analysis.

Vision research is not limited to binary images but also deals with gray-scale, color, and other multispectral images. In fact, the word "image" is used to avoid the limitation to visual spectra. If we avoid the compression, transmission, and other representation issues, then we can classify vision research as follows:

- Low-level vision involves extracting feature measurements from images. It is called low-level because the operations are not knowledge based. Typical operations are edge detection, threshold selection, and the measurement of various shapes and other features. These are the operations now being reduced to hardware.
- High-level vision is concerned with combining knowledge about objects (shape, size, relationships), expectations about the image (what might be in it), and the purpose of the processing (identifying

objects, detecting changes) to aid in interpreting the image. This high-level information interacts with and helps guide processing. For example, it can suggest where to look for an object and what features to look for.

While research in vision is maturing, much remains to be investigated. Current topics include the speed of algorithms, parallel processing, coarse/fine techniques, incomplete data, and a variety of other extensions to the field. In addition, work is also now addressing such AI questions as

- representing knowledge about objects, particularly shape and spatial relationships;
- developing methods for reasoning about spatial relationships among objects;
- understanding the interaction between low-level information and high-level knowledge and expectations;
- interpreting stereo images, e.g., for range and motion;
- understanding the interaction between an image and other information about the scene, e.g., written descriptions.

Vision research is related to results in VLSI and Ar. While there is much activity, it is difficult to predict specific results that can be expected.

Tactile Sensing

Despite great interest in the use of tactile sensing, the state of the art is relatively primitive. Systems on industrial robots today are limited to detecting contact of the robot and an object by varying versions of the limit-switch concept, or they measure some combination of force and torque vectors that the hand or fingers exert on an object.

While varying versions of the limit-switch concept have been used, the most advanced force/torque sensors for robots have been developed at Draper Laboratories. The remote center of compliance (RCC) developed at Draper Laboratories, which allows passive compliance in the robots' behavior during assembly, has been commercialized by Astek and Lord Kinematics. Draper has in the last few years instrumented the RCC to provide active feedback to the robot. The instrumented remote center compliance (IRCC) represents the state of the art in wrist sensors. It allows robot programs to follow contours, perform: insertions, and incorporate rudimentary touch programming into the control system [17].

IBM and others have begun to put force sensors in the fingers of a robot. With x,y,z strain gauges in each of the fingers, the robot with servoed fingers can now perform simple touch-sensitive tasks. Hitachi has developed a hand using metal contact detectors and pressure-sensitive conductive rubber that can feel for objects and

recognize form. Thus, primitive technology can be applied for useful tasks. However, most of the sophisticated and complex tactile sensors are in laboratory development.

The subject of touch-sensor technology, including a review of research, relevance for robots, work in the laboratory, and predictions of future results, is covered in a survey article by Leon Harmon [18] of Case Western Reserve University Much of that excellent article is summarized below, and we refer the reader to it for a detailed review.

The general needs for sensing in manipulator control are proximity) touch/slip, and force/torque. The following remarks are taken from a discussion on "smart sensors" by Bejcsy [19]:

specific manipulation-related key events are not contained in visual data at all, or can only be obtained from visual data sources indirectly and incompletely and at high cost. These key events are the contact or near-contact events including the dynamics of interaction between the mechanical hand and objects.

The non-visual information is related to controlling the physical interaction, contact or near-contact of the mechanical hand with the environment. This information provides a combination of geometric and dynamic reference data for the control of terminal positioning/orientation and dynamic accommodation/compliance of the mechanical hand.

Although existing industrial robots manage to sense position, proximity, contact, force, and slip with rather primitive techniques, all of these variables plus shape recognition have received extensive attention in research and development laboratories. In some of these areas a new generation of sophistication is beginning to emerge.

Tactile-sensing requirements are not well known, either theoretically or empirically. Most prior wrist, hand, and finger sensors have been simple position and force-feedback indicators. Finger sensors have barely emerged from the level of microswitch limit switches and push-rod axial

travel measurement. Moreover, the relevant technologies are themselves relatively new. For example, force and torque sensing dates back only to 1972, touch/slip are dated to 1966, and proximity sensing is only about 9 years old. We do know that force and pressure sensing are vital elements in touch, though to date, as we have seen, industrial robots employ only simple force feedback. Nevertheless, unless considerable gripper overpressure can be tolerated, slip sensing is essential to proper performance in many manipulation tasks. Information about contact areas, pressure distributions, and their changes over time are needed in order to achieve the most complete and useful tactile sensing.

In contacting, grasping, and manipulating objects, adjustments to gripping forces are required in order to avoid slip and to avoid possibly dangerous forces to both the hand and the workpiece. Besides the need for slip-sensing transducers, there is the requirement that the robot be able to determine at each instant the necessary minimum new force adjustments to prevent slip.

Transducers As of about 1971 the only devices available for tactile sensing were microswitches, pneumatic jets, and (binary) pressure-sensitive pads. These devices served principally as limit switches and provided few means or none for detecting shape, texture, or compliance. Still, such crude devices are used currently.

In the early 1970s the search was already under way for shape detection and for "artificial skin" that could yield tactile information of complexity comparable to the human sense of touch. An obvious methodology for obtaining a continuous measurement of force is potentiometer response to a linear (e.g., spring-loaded rod) displacement. Early sensors in many laboratories used such sensors, and they are still in use today.

Current research lies in the following areas:

- conductive materials and arrays produced with conductive rubbers and polymers;
- semiconductor sensors, such as piezo-electrics;
- electromagnetic, hydraulic, optical, and capacitive sensors.

Outstanding Problems and New Opportunities The two main areas most in need of development are (1) improved tactile sensors and (2) improved integration of touch feedback signals with the effector control system in response to the task-command structure. Sensory feedback problems underlie both areas. More effective comprehensive sensors (device R&D) and the sophisticated interpretation of the sense signals by control structures (system R&D) are needed.

Sensitive, dexterous hands are the greatest challenge for manipulators, just as sensitive, adaptable feet are the greatest challenge for legged locomotion vehicles. Each application area has its own detailed special problems to solve; for example, the design approach for muddy-water object recovery and for delicate handling of unspecified objects in an unstructured environment differ vastly.

Emergent Technology One of the newest developments in touch-sensing technology is that of reticular (Cartesian) arrays using solid-state transduction and attached microcomputer elements that compute three-dimensional shapes. The approach is typified by the research of Marc Raibert, now at CMU, done while he was at JPL (20). Raibert's device is compact and has high resolution; hence, the fingertip is a self-contained "smart finger." See also the work of Hillis at MIT in this area [21]. This is a quantum jump ahead of prior methods, for example, where small arrays of touch sensors use passive substrates and materials such as conductive elastomers. Resolution in such devices has been quite low, and hysteresis a problem.

Sound Sensors

Many researchers are interested in the use of voice recognition sensors for command and control of robot systems. However, we leave out voice systems and review here the use of sound as a sensing mechanism.

In this context, sound systems are used as a method for measuring distance. The Polaroid sonic sensor has been used at NBS and elsewhere as a safety sensor. Sensors mounted on the robot detect intrusions into either the workspace or, more particularly, the path of the robot.

Researchers at Pennsylvania State University have developed a spark gap system that uses multiple microphones to determine the position of the manipulator for calibration purposes.

Several researchers at Carnegie-Mellon University and other locations are working on ultrasonic sensors to be used in the arc welding process.

Control Systems

The underlying research issue in control systems for robots is to broaden the scope of the robot. As the sophistication of the manipulator and its actuation mechanism increases, new demands are made on the control system. The advent of dexterous or smart hands, locomotion, sensors, and new complex tasks all extend the controller capability.

The desires for user-friendly systems, for less user training, and for adaptive behavior further push the robot controller into the world of artificial intelligence. Before discussing intelligent robot systems, we describe some of the issues of computer-controlled robots.

Hierarchical Control/Distributed Computing

Almost all controller research is directed at hierarchies in robot control systems. At the National Bureau of Standards, pioneering research has developed two hierarchies--one for control information and one for sensory data. Integrated at each level, the two hierarchies use the task decomposition approach. That is, commands at each level are broken down into subcommands at the lower level until they represent joint control at the lowest level. In a similar fashion, raw

vision data are at the lowest level, with higher levels representing image primitives, then features, and finally objects [22].

The levels-of-control issue rapidly leads to an interest in distributed computing in order to balance the computing needs and meet the requirements for real-time performance. The use of smart hand or complex sensor systems, such as vision, also mandates distributed computing--again, in order not to overload the control computer and degrade the real-time nature of the robot's behavior.

Distributed computing for robot control systems has taken two paths so far. Automatix, NBS, and others use multiple CPUs from the

same vendor (Intel or Motorola) and perform processor communication in the architecture of the base system.

Others have used nonhomogeneous computer systems. They have had to pay a price in the need to define and build protocols and work within awkward constraints. Examples of this are found in the development of MCL by McDonnell Douglas and in a variety of other firms that have linked vision systems with robots. For a case study of one attempt see Nagel et al. [23].

Major impediments to progress in these areas are the lack of standards for the interfaces needed, the need for advances in distributed computing, and the need for a better definition of the information that must flow. Related research that is not covered here is the work on local area networks.

Data Bases

There is a great interest in robot access to the data bases of CAD/CAM systems. As robot programming moves from the domain of the teach box to that of a language, several new demands for data arise. For example, the programmer needs access to the geometry and physical properties of the parts to be manipulated. In addition, he needs similar data with respect to the machine tools, fixtures, and the robot itself. One possible source for this is the data already captured in CAD/CAM data bases. One can assume that complete geometrical and functional information for the robot itself, the things the robot must manipulate, and the things in its environment are contained in these data bases.

As robot programming evolves, an interest has developed in computer-aided robot programming (CARP) done at interactive graphics terminals. In such a modality the robot motions in manipulating parts would be done in a fashion similar to that used for graphic numerical control programming. Such experiments are under way, and early demonstrations have been shown by Automatix and GCA Corporation.

Furthermore, it is now reasonable to assume the desire to have robots report to shop floor control systems, take orders from cell controllers, and update process planning inventory control systems and the variety of factory control, management, and planning systems now in place or under development. Thus, robot controllers must access other data bases and communicate with other factory systems.

Research on the link to CAD/CAM systems and the other issues above is under way at NBS and other research facilities, but major efforts are needed to achieve results.

Robot Programming Environment

As mentioned earlier, second-generation languages are now available. While the community as a whole does not yet have sufficient experience with them to choose standards, more are clearly needed.

56

Programming advanced robot systems with current languages is reminiscent of programming main-frame computers in assembly language before the advent of operating systems. It is particularly a problem in the use of even the simplest sensor (binary) mechanisms. What are needed are robot operating systems, which would do for robot users what operating systems do for computer users in such areas as input/output and graphics.

To clarify, we define an explicit language as one in which the commands correspond with the underlying machine (in this case a robot/ computer pair). We further define an implicit language as one in which the commands correspond with the task; that is, for an assembly task an insert command would be implied. Use of an implicit language is complicated by the fact that robots perform families of tasks. A robot operating system would be a major step toward implicit languages.

It is far easier to suggest the work above than to write a definition of requirements. Thus, fundamental research is needed in this area. The Autopass system developed at IBM is probably the most relevant accomplishment to date.

The concepts of graphic robot programming and simulation are exciting research issues. The desire for computer-assisted robot programming (CARP) stems from the data base arguments of before and the belief that graphics is a good mechanism for describing motion. These expectations are widely held, and Computervision, Automatix, and other organizations are conducting some research. However, no major efforts appear in the current literature.

Graphic simulation, on the other hand, is now a major topic. Work in this area is motivated by the advent of offline programming languages and the need for fail-safe debugging languages, but other benefits arise in robot cell layout, training mechanisms, and the ability to let the robot stay in production while new programs are developed.

Work on robot simulation is hampered by the lack of standards for the language but is in process at IBM for AML, at McDonnell Douglas for MCL, and at many universities for VAL and is expected to be a commercial product shortly. It is worth noting that simulation of sensor-based robots requires simulation of sensor physics. With the exception of some work at IBM, we are unaware of any efforts in sophisticated simulation.

The use of multiple arms in coordinated (as opposed to sequenced) motion raises the issue of multitasking, collision avoidance, and a variety of programming methodology questions. General Electric, Olivetti, Westinghouse, IBM, and others are pursuing multiarm assembly. However these issues require more attention, even in research that is well under way.

Intelligent Robots

It should be clear by now that robot control has become a complex issue. Controllers dealing with manipulator motion, feedback, complex sensors, data bases, hierarchical control, operating systems, and multitasking must turn to the AI area for further development. In the following section we review briefly the AI field, and in the final section we discuss both robotics and AI issues and the need for expansion of the unified research issues.

ARTIFICIAL INTELLIGENCE¹

The term artificial intelligence is defined in two ways: the first defines the field, and the second describes some of its functions.

1. "Artificial intelligence research is the part of computer science that is concerned with the symbol-manipulation processes that produce intelligent action. By 'intelligent action' is meant an act of decision that is goal-oriented, arrived at by an understandable chain of symbolic analysis and reasoning steps, and is one in which knowledge of the world informs and guides the reasoning" [24].

2. Artificial intelligence is a set of advanced computer software applicable to classes of nondeterministic problems such as natural language understanding, image understanding, expert systems, knowledge acquisition and representation, heuristic search, deductive reasoning, and planning.

If one were to give a name suggestive of the processes involved in all of the above, *knowledge engineering* would be the most appropriate; that is, one carries out knowledge engineering to exhibit intelligent behavior by the computer. For general information on artificial intelligence see references 25-34.

Background

The number of researchers in artificial intelligence is rapidly expanding with the increasing number of applications and potential applications of the technology. This growth is occurring not only in the United States, but worldwide, particularly in Europe and Japan.

Basic research is going on primarily at universities and some research institutes. Originally, the primary research sites were MIT, CMU, Stanford, SRI, and the University of Edinburgh. Now, most major

universities include artificial intelligence in the computer science curriculum.

¹Much of the material in this section summarizes the material in Brown et al. [24].

An increasing number of other organizations either have or are establishing research laboratories for artificial intelligence. Some of them are conducting basic research; others are primarily interested in applications. These organizations include Xerox, Hewlett-Packard, Schlumberger-Fairchild, Hughes, Rand, Perceptronics, Unilever, Philips, Toshiba, and Hamamatsu.

Also emerging are companies that are developing artificial intelligence products. U.S. companies include Teknowledge, Cognitive Systems, Intelligenetics, Artificial Intelligence Corp., Symantec, and Kestrel Institute.

Fundamental issues in artificial intelligence that must be resolved include

- representing the knowledge needed to act intelligently,
- acquiring knowledge and explaining it effectively,
- reasoning: drawing conclusions, making inferences, making decisions ,
- evaluating and choosing among alternatives.

Natural Language Interpretation

Research on interpreting natural language is concerned with developing computer systems that can interact with a person in English (or another nonartificial language). One primary goal is to enable computers to use human languages rather than require humans to use computer languages.

Research is concerned with both written and spoken language. Although many of the problems are independent of the communication medium, the medium itself can present problems. We will first consider written language, then the added problems of speech.

There are many reasons for developing computer systems that can interpret natural-language inputs. They can be grouped into two basic categories: improved human/machine interface and automatic interpretation of written text.

Improving the human/machine interface will make it simple for humans to

- give commands to the computer or robot,
- query data bases,
- conduct a dialogue with an intelligent computer system.

The ability to interpret text automatically will enable the computer to

- produce summaries of texts,
- provide better indexing methods for large bodies of text,
- translate texts automatically or semiautomatically,
- integrate text information with other information.

Current Status

Natural-language understanding systems that interpret individual (independent) sentences about a restricted subject (e.g., data in a data base) are becoming available. These systems are usually constrained to operate on some subset of English grammar, using a limited vocabulary to cover a restricted subject area. Most of these systems have difficulty interpreting sentences within the larger context of an interactive dialogue, but a few of the available systems confront the problem of contextual understanding with promising capability. There are also some systems that can function despite grammatically incorrect sentences and run-on constructions. But even when grammatical constraints are lifted, all commercial systems assume a specific knowledge domain and are designed to operate only within that domain.

Commercial systems providing natural-language access to data bases are becoming available. Given the appropriate data in the area base they can answer questions such as

- Which utility helicopters are mission-ready?
- Which are operational?
- Are any transport helicopters mission-ready?

However, these systems have limitations:

- They must be tailored to the data base and subject area.
- They only accept queries about facts in the data base, not about the contents of the data base--e.g., "What questions can you answer about helicopters?"
- Few Computations can be performed on the data.

In evaluating any given system, it is crucial to consider its ability to handle queries in context. If no contextual processing is to be performed, sentences will often be interpreted to mean something other than what a naive user intends. For example, suppose there is a natural-language query system designed to field questions about air force equipment maintenance, and a user asks "What is the status of squadron A?" If the query is followed by "What utility helicopters are ready?" the utterance will be interpreted as meaning "Which among all the helicopters are

ready?" rather than "Which of the squadron A helicopters are ready?" The system will readily answer the question; it just will not be the question the user thought he was asking.

Data base access systems with more advanced capabilities are still in the research stages. These capabilities include

- easy adaptation to a new data base or new subject area,
- replies to questions about the contents of the data base (e.g., what do you know about tank locations?),
- answers to questions requiring computations (e.g., the time for a ship to get someplace).

60

It is nevertheless impressive to see what can be accomplished within the current state of the art for specific information processing tasks. For example, a natural-language front end to a data base on oil wells has been connected to a graphics system to generate customized maps to aid in oil field exploration. The following sample of input illustrates what the system can do.

Show me a map of all tight wells drilled by Texaco before May 1, 1970, that show oil deeper than 2,000 ft, are themselves deeper than 5,000 ft, are now operated by Shell, are wildcat wells where the operator reported a drilling problem, and have mechanical logs, drill stem tests, and a commercial oil analysis, that were drilled within the area defined by latitude 30 deg 20 min 30 sec to 31:20:30 and 80-81. Scale 2,000 ft.

This system corrects spelling errors, queries the user if the map specifications are incomplete, and allows the user to refer to previous requests in order to generate maps that are similar to previous maps.

This sort of capability cannot be duplicated for many data bases or information processing tasks, but it does show what current technology can accomplish when appropriate problems are tackled.

Research Issues

In addition to extending capabilities of natural-language access to data bases, much of the current research in natural language is directed toward determining the ways in which the context of an utterance contributes to its meaning and toward developing methods for using contextual information when interpreting utterances. For example, consider the following pairs of utterances:

Sam: The lock nut should be tight.

Joe: I've done it.

and

Sam: Has the air filter been removed?

Joe: I've done it.

Although Joe's words are the same in both cases, and both state that some action has been completed, they each refer to different actions--in one case, tightening the lock nut; in the other, removing the air filter. The meanings can only be determined by knowing what has been said and what is happening.

Some of the basic research issues being addressed are

- interpreting extended dialogues and texts (e.g., narratives, written reports) in which the meaning depends on the context;

61

- interpreting indirect or subtle utterances, such as recognizing that "Can you reach the *salt?*" is a request for the salt;
- developing ways of expressing the more subtle meanings of sentences and texts.

Spoken Language

Commercial devices are available for recognizing a limited number of spoken words, generally fewer than 100. These systems are remarkably reliable and very useful for certain applications.

The principal limitations of these systems are that

- they must be trained for each speaker,
- they only recognize words spoken in isolation,
- they recognize only a limited number of words.

Efforts to link isolated word recognition with the natural-language understanding systems are now under way. The result would be a system that, for a limited subject area and a user with some training, would respond to spoken English inputs.

Understanding connected speech (i.e., speech without pauses) with a reasonably large vocabulary will require further basic research in acoustics and linguistics as well as the natural-language issues discussed above.

Generating Information

Computers can be used to present information in various modes, including written language, spoken language, graphics, and pictures. One of the principal concerns in artificial intelligence is to develop methods for tailoring the presentation of information to individuals. The presentation

should take into account the needs, language abilities, and knowledge of the subject area of the person or persons.

In many cases, generation means deciding both what to present and how to present it. For example, consider a repair adviser that leads a person through a repair task. For each step, the adviser must decide which information to give to the person. A very naive person may need considerable detail; a more sophisticated person would be bored by it. There may, for example, be several ways of referring to a tool. If the person knows the tool's name then the name could be used; if not, it might be referred to as "the small red thing next to the toolchest." The decision may extend to other modes of output. For example, if a graphic display is available, a picture of the tool could be drawn rather than a verbal description given.

Current Status

At present, most of the generation work in artificial intelligence is concerned with generating language. Quite a few systems have been developed to produce grammatical English (or other natural language) sentences. However, although a wide range of constructions can be produced, in most cases the choice of which construction (e.g., active or passive voice) is made arbitrarily. A few systems can produce stilted paragraphs about a restricted subject area.

A few researchers have addressed the problems of generating graphical images to express information instead of language. However, many research issues remain in this area.

Research Issues

Some of the basic research issues associated with generating information include

- deciding which grammatical construction to use in a given situation ;
- deciding which words to use to convey a certain idea;
- producing coherent bodies of text, paragraphs, or more;
- tailoring information to fit an individual's needs.

Assimilating Information

Being in any kind of changing environment and interacting with the environment means getting new information. That information must be incorporated into what is already known, tested against it, used to modify it, etc. Since one aspect of intelligence is the ability to cope with a new or changing situation, any intelligent system must be able to assimilate new information about its environment.

Because it is impossible to have complete and consistent information about everything, the ability to assimilate new information also requires the ability to detect and deal with inconsistent and incomplete information.

Expert Systems

The material presented here is designed to provide a simple overview of expert systems technology, its current status, and research issues. The importance of this single topic, however, suggests that it merits a more in-depth review; an excellent one recently published by the NBS is recommended [25].

Expert systems are computer programs that capture human expertise about a specialized subject area. Some applications of expert systems are medical diagnosis (INTERNIST, MYCIN, PUFF), mineral exploration (PROSPECTOR), and diagnosis of equipment failure (DART).

63

The basic technique behind expert Systems is to encode an expert 's knowledge as rules stating the likelihood of a hypothesis based on available evidence. The expert system uses these rules and the avail-able evidence to form hypotheses. If evidence is lacking, the expert system will ask for it.

An example rule might be

IF THE JEEP WILL NOT START

and

THE HORN WILL NOT WORK

and

THE LIGHTS ARE VERY DIM,

then

THE BATTERY IS DEAD,

WITH 90 PERCENT PROBABILITY.

If an expert system has this rule and is told, "the jeep will not start," the system will ask about the horn and lights and decide the likelihood that the battery is dead.

Current Status

Expert systems are being tested in the areas of medicine, molecular genetics, and mineral exploration, to name a few. Within certain limitations these systems appear to perform as well as human experts. There is already at least one commercial product based on expert-system technology.

Each expert system is tailored to the subject area. It requires extensive interviewing of an expert, entering the expert's information into the computer, verifying it, and sometimes writing new computer programs. Extensive research will be required to improve the process of getting the human expert 's knowledge into the computer and to design systems that do not require programming changes for each new subject area.

In general, the following are prerequisites for the success of a knowledge-based expert system:

- There must be at least one human expert acknowledged to perform the task well.
- The primary source of the expert 's exceptional performance must be special knowledge, judgment, and experience.
- The expert must be able to explain the special knowledge and experience and the methods used to apply them to particular problems.
- The task must have a well-bounded domain of applications [25].

Research Issues

Basic research issues in expert systems include

64

- the use of, causal models, i.e., models of how something works to help determine why it has failed;
- techniques for reasoning with incomplete, uncertain, and possibly conflicting information;
- techniques for getting the proper information into rules;
- general-purpose expert systems that can handle a range of similar problems, e.g., work with many different kinds of mechanical equipment.

Planning

Planning is concerned with developing computer Systems that can combine sequences of actions for specific problems. Samples of planning problems include

- placing sensors in a hostile area,
- repairing a jeep,
- launching planes off a carrier,
- conducting combat operations,
- navigating,
- gathering information.

Some planning research is directed towards developing methods for fully automatic planning; other research is on interactive planning, in which the decision making is shared by a combination of the person and the computer. The actions that are planned can be carried out by people, robots, or both.

An artificial intelligence planning system starts with

- knowledge about the initial situation, e.g., partially known terrain in hostile territory;
- facts about the world, e.g., that moving changes location;
- possible actions, e.g., walk, fly, look around, hide;
- available objects, e.g., a platform on wheels, arms, sensors;
- a goal, e.g., installing sensors to detect hostile movements and activity.

The system will produce (either by itself or with guidance from a person) a plan containing these actions and objects that will achieve the goal in this situation.

Current Status

The planning aspects of AI are still in the research stages. The research is both theoretical in developing better methods for expressing knowledge about the world and reasoning about it and more experimental in building systems to demonstrate some of the techniques that have been developed. Most of the experimental systems have been

65

tested on small problems. Recent work at SRI on interactive planning is one attempt to address larger problems by sharing the decisionmaking between the human and machine.

Research Issues

Research issues related to planning include

- reasoning about alternative actions that can be used to accomplish a goal or goals,
- reasoning about action in different situations,
- representing spatial relationships and movements through space and reasoning about them,
- evaluating alternative plans under varying circumstances,
- planning and reasoning with uncertain, incomplete, and inconsistent information,
- reasoning about actions with strict time requirements; for example, some actions may have to be performed sequentially or in parallel or at specific times (e.g., night time),
- replanning quickly and efficiently when the situation changes.

Monitoring Actions and Situations

Another aspect of reasoning is detecting that something significant has occurred (e.g., that an action has been performed or that a situation has changed). The key here is *significant*. Many things take place and are reported to a computer system; not all of them are significant all the time. In fact, the same events may be important to some people and not to others. The problem for an intelligent system is to decide when something is important.

We will consider three types of monitoring: monitoring the execution of planned actions, monitoring situations for change, and recognizing plans.

Execution Monitoring

Associated with planning is *execution monitoring*, that is, following the execution of a plan and replanning (if possible) when problems arise or possibly gathering more information when needed. A monitoring system will look for specific situations to be sure that they have been achieved; for example, it would determine if a piece of equipment has arrived at a location to which it was to have been moved.

We characterize the basic problem as follows: given some new information about the execution of an action or the current situation, determine how that information relates to the plan and expected situation, and then decide if that information signals a problem; if so, identify options available for fixing it. The basic steps are: (1) find the problem (if there is one), (2) decide what is affected,

66

(3) determine alternative ways to fix the problem, and (4) select the best alternative. Methods for fixing a problem include choosing another action to achieve the same goal, trying to achieve some larger goal another way, or deciding to skip the step entirely.

Research in this area is still in the basic stages. At present, most approaches assume a person supplies unsolicited new information about the situation. However, for many problems the system must be able to acquire directly the information needed to be sure a plan is proceeding as expected, instead of relying on volunteered information. Planning to acquire information is a more difficult problem because it requires that the computer system have information about what situations are crucial to a plan's success and be able to detect that those situations hold. Planning too many monitoring tasks could be burdensome; planning too few might result in the failure to detect an unsuccessful execution of the plan.

Situation Monitoring

Situation monitoring entails monitoring reported information in order to detect changes, for example, to detect movements of headquarters or changes in supply routes.

Some research has been devoted to this area, and techniques have been developed for detecting certain types of changes. Procedures can be set to be triggered whenever a certain type of information is inserted into a data base. However, there are still problems associated with specifying the conditions that should trigger them. In general, it is quite difficult to specify what constitutes a change. For example, a change in supply route may not be signaled by a change of one truck's route, but in some cases three trucks could signal a change. A system should not alert a person every time a truck detours, but it should not wait until the entire supply line has changed. Specifying when the change is significant and developing methods for detecting it are

still research issues.

Plan Recognition

Plan recognition is the process of recognizing another's plan from knowledge of the situation and observations of actions. The ability to recognize another's plan is particularly important in adversary situations where actions are planned based on assumptions about the other side's intentions. Plan recognition is also important in natural language generation because a question or statement is often part of some larger task. For example, if a person is told to use a ratchet wrench for some task, the question "What ' s a ratchet wrench?" may be asking "How can I identify a ratchet wrench?" Responding appropriately to the question entails recognizing that having the wrench is part of the person ' s plan to do the task.

67

Research in plan recognition is in early stages and requires further basic research, particularly on the problem of inferring goals and intentions.

Applications-Oriented Research

The general areas of natural-language processing, speech recognition, expert systems, planning, and monitoring suggest the sorts of problems that are studied in artificial intelligence, but they may not, by themselves, suggest the variety of information processing applications that will be possible with AI technology. Some research projects are now consolidating advances in more than one area of AI in order to create sophisticated Systems that better address the information processing needs of industry and the military.

For example, an expert system that understands principles of programming and software design can be used as a programming tutor for students at the introductory level. This illustrates how an expert system can be incorporated in a computer-aided instruction (CAI) system to provide a more sophisticated level of interactive instruction than is currently available.

Programs for CAI can also be enhanced by natural-language processing for instruction in domains that require the ability to answer and ask questions. For example, Socratic teaching methods could be built into a political science tutor when natural-language processing progresses to a robust stage of sophistication and reliability. Even with the current technology, a reading tutor for students with poor literacy skills could be designed for individualized instruction and evaluation-. In fact, the long-neglected area of machine translation could be profitably revisited at this time with an eye toward automated language tutors. Today's language analysis technology could be put to work evaluating student translations of single sentences in restricted knowldomains, and our generation systems could suggest appropriate alternatives to incorrect translations as needed. This task orientation is slightly different from that of an automated translator, yet it would be a valuable application that our current state of the art could tackle effectively.

Systems that incorporate knowledge of plans and monitoring can be applied to the office environment to provide intelligent clerical assistants. Such an automated assistant could keep track of ongoing projects, reminding the user where he is with respect to a particular job and what steps remain to be taken. Some scheduling advice might be given if limited resources (time, secretarial help, necessary supplies) have to be used efficiently. A truly intelligent assistant with natural-language processing abilities could screen electronic mail and generate suggested responses to the more routine items of business at hand ("yes, I can make that meeting"; "I'm sorry I won't be able to make that deadline" ; "no, I don't have access to the technology"). Automated assistants with knowledge of specific procedures could be useful both to novices who are learning the ropes and to more experienced users who simply need to use their time as effectively as possible.

68

While most expert systems today assimilate new knowledge in highly restricted ways, the importance of learning systems should not be overlooked. In the long run, general principles of learning will become critical in designing sophisticated information processing systems that access large quantities of data and work within multiple knowledge domains. As AI moves away from problems within restricted knowledge domains, it will become increasingly important for more powerful systems to integrate and organize new information automatically--i.e., to learn by themselves. We will have to move away from simplistic pattern-matching strategies to the more abstract notions of analogy and precedents. Research on learning is still in its infancy, but we can expect it to become an application-oriented research issue very quickly--within 5 to 10 years, if the field progresses at a healthy pace. Without sufficient research support in this area, our efforts may stagnate in the face of apparent impasses.

With a field that moves as rapidly as AI, it is important to realize that a long-term perspective must be assumed for even the most pragmatic research effort. Even a 2-year project designed to use existing technology may adapt new techniques that become possible during the life of the project. The state of the art is a very lively moving target, and advances can render research publications obsolete in the space of a few months. New Ph.D.s must keep close tabs on their areas of interest to maintain the expertise they worked so hard to establish in graduate school. We must therefore emphasize how dangerous a short view of AI is and how critical it is for the field to maintain a sensitive perspective on long-term progress in all of our research efforts.

STATE OF THE ART AND PREDICTIONS

In the previous sections we have reviewed the state of the art in robotics and artificial intelligence. Clearly, both robotics and artificial intelligence are relatively new fields with diverse and complex research questions. Furthermore, the intersection field--robotics/ artificial intelligence or the intelligent robot--is an embryonic research area. This area is made more complex by the obvious dependence on heretofore unrelated fields, including mechanical design, control, vision sensing, force and touch sensing, and knowledge engineering. Thus, predicting the state of the art 5 and 10 years from now is difficult. Moreover, because predictions for the

near future are likely to be more accurate than those for the more distant future, our 10-year predictions should be treated with particular precaution.

One approach to the problem of prediction is to decouple the fundamental research areas and predict possible developments in each technology area. Such a task is easy only in comparison to the former question; nevertheless, in the following sections we undertake a field-by-field assessment and predictions of 5- and 10-year developments.

In the sections that follow, we develop tables describing the current state of the art and predictions for the next 5- and 10-year periods. Each section contains a short narrative and some general

comments with respect to research funding and researchers working in the problem area. The table at the end of the chapter summarizes the findings.

Mechanical Design of the Manipulator and Actuation Mechanism

The industrial robot is a single mechanical arm with rigid, heavy members and linkages. Actuation of the slide or rotary joints is based on transmission gears, which results in backlash. Joint bearings of conventional design have high friction and stiction, which cause poor robot performance. Thus, with the rare exception of some semiconductor applications that are more accurate, robot repeatability is in the range of 0.1 to 0.005 inches. Robots today operate from fixed locations with little or no mobility (except track mountings or simple wire-guided vehicles) and have a limited work envelope. The operating environment is constrained to the factory floor, and the typical robot is not self-contained but requires an extensive support system with big power supplies.

The factors listed above are reflected in the first column of the table under entry numbers 1 to 11. As shown in the table, on a point by point basis we expect significant improvements within 5 years (column 2) and even more within 10 years (column 3).

Table entries 12 and 13 address the kinematics and dynamics of robots as they are today (column 1) and predict how they will evolve. These issues, while based fundamentally on the mechanical structure of the robot and how it behaves in motion and under load, are clearly intertwined with the issues of manipulator control and computation speed. For example, we do not today have enough computer power in the robot control system to take advantage of kinematic model data.

Thus, while we make some predictions under these headings, they are closely related to the control issues to be addressed later.

The research on mechanical design and actuation mechanisms has been supported by NSF, ONR, and others but is not the main focus of a major funding program at this time. University laboratories such as those at MIT, CMU, Stanford, and the University of Florida at Gainesville

are investigating the manipulator and its kinematics. Locomotion research is continuing at Ohio State, CMU, and RPI. The Jet Propulsion Laboratory, Stanford Research Institute, and Draper Laboratories are also active in some of these areas [3-7].

End-Effector Design

Current industrial robots use many hands, each specifically designed for a different application. As described in the Research section, this has led to research in two directions--one to produce the dexterous hand and the second to produce the quick-change hand. The lack of progress in these areas makes most applications expensive because of the need to design a special hand, and it prohibits others because of a lack of dexterity or the ability to change hands rapidly.

70

Many are also working on hand-based sensor systems; these issues are covered in depth under the topic of sensor systems. Entries 14 and 15 in the table describe current technology hands as simple (open or closed) hands that are rarely servoed--though the IBM RSI is a notable exception, which others are following.

End effectors today are also sometimes tools that are operated by an on/off signal. Today's hands do employ limited sensors and permit rudimentary force programming. As described in the table, we expect progress in the development of quick-change hands to precede the wide use of instrumented dexterous hands.

Research in end effectors is taking place at the University of Utah (based on prior work in prosthetics), the University of Rhode Island, and at most of the locations cited for mechanical design research. References 9-11 are suggested for further details.

Funding of these hand efforts is typically a part of some larger project and is not a major project of any funding agency.

Vision Sensors

As described earlier, vision has been a high-interest area for robotics in both the visual servoing (guidance) and inspection or measurement modality.

Commercial vision systems use binary images and simple features and are restricted to high contrast images. As shown in table entry 16, we expect that VLSI technology, now in research labs at MIT, Hughes, Westinghouse, and others, will be commercialized. In 5 years this will provide real-time edge images, a richer shape-capturing feature set, and will ease the restriction on high-contrast binary images, allowing gray-scale and texture-based objects to be handled. These predictions are conservative. In 10 years we further expect rapid-recognition systems that can handle a limited class of objects in arbitrary orientation. Thus, the visual servoing problem will be routinely achievable.

The use of so-called three-dimensional vision, using stereo, structured light systems, and other vision-based methods to acquire "depth" information, is rudimentary today, as shown in table entry 17. The stereo mapper system at DMA is an exception. This system, which works well on textured terrain such as forests, is ineffective on urban landscapes. A big step forward is expected in the next 5 years. Currently in research labs are systems that extract depth using

- stereo, employing either vision or laser light (MIT, Stanford);
- shape from shading, special light (GE, MIT, SRI);
- gross shape from motion (CMU, MIT, Stanford, University of Minnesota) ;
- shape from structured light systems (GE, GM, NBS).

Commercial systems will market three-dimensional vision systems that will generate a depth map in relatively benign situations. They will be slow, too slow for military rapid response situations in the next 5 years. The algorithms for all these methods for computing

depth are inherently parallel. They can be computed using highly parallel computers specifically designed. A hardware stereo (vision or laser) and shape from motion system is possible in 5 years. One practical problem is lithographic density. Putting a lot of processing on chips of 1 micron density restricts spatial resolution of an image. However, 0.1 micron densities seem feasible in 5 years.

Merely generating a depth map is not the same as seeing. It is also necessary to extract objects and to recognize them from arbitrary orientation. The depth map is likely to be noisy and relatively coarse. It will be possible, for example, to identify a shape as a person, but not to recognize which person. It will recognize a tank, but only determine type if it is significantly different from another.

Tasks that will become feasible with depth data include

- three-dimensional inspection of object surfaces for dents, cracks, etc. that do not affect outline;
- better edge maps and shape, leading to recognition of objects by outline shape, e.g., an automobile.

In 10 years, one can confidently predict

- reliable hardware stereo systems,
- systems capable of determining the movement of an object and maneuvering to avoid it,
- rapid recognition of limited classes of objects from an arbitrary viewpoint.

Vision research is a very active field in the United States (see reference 34). For a survey of vision research, see reference 35. For a review of image understanding, see reference 14. Most

three-dimensional vision research in the United States is funded by the DARPA Image Understanding (IU) program. See, for example, the IU workshop proceedings from DARPA.

Commercial vision systems are marketed by GE, Octek, Automatix, Cognex, Machine Intelligence Corporation, ORS, and others. Government and foundation support of major programs is provided by the Office of Naval Research (ONR), DARPA, Systems Development Foundations (SDF), and NSF.

Many corporations in Japan, including Hitachi, Sony, and Fujitsu, are doing work in this area; there are also several large university efforts (see references 13, 36, 39).

Nonvisual sensors (radar, SAR, FLIR, etc.) have mostly been developed by defense contractors for DARPA, AFOSR, and ONR. The following systems are among those available from Lockheed, TRW, Honeywell, and others:

- synthetic aperture radar (SAR),
- forward looking infrared (FLIR),
- millimeter radar,
- Xray.

For example, the cruise missile uses one-dimensional correlations on radar images. This is rather crude. Capabilities are mostly classified.

Advantages of nonvisual sensing are that they simplify certain problems. For example, it is easy to find hot spots in infrared. Often they correspond to camouflaged targets.

Limitations are that the physics of nonvisual imagery are poorly understood, and algorithms are limited in scope. Two main applications are for seeing large static objects and for automatically navigating certain kinds of terrain.

Research is intense, funding levels are high, and progress will be good. This is entirely an industry effort with DOD sponsorship. However, vision does appear to be the best way forward because it is passive and operators know what visual images mean. This is a serious issue, since trained observers are needed to check results of processing nonvisual images.

Contact/Tactile Sensors

As described earlier, contact/tactile sensors are an important area of robotics development. Although progress has so far been slow, this is an important area for determining

- surface shape, including surface inspection;
- slip computation--how sure the grasp is;
- proximity--how close the hand is to the object;

- force/torque, to control and measure its application.

Robots today are programmed for position only; in rare instances, they can do some rudimentary force programming using a commercial version of the Draper Laboratory IRCC. For the state of the art, see references 18-21 and 37

Current systems suffer from both rudimentary control capability (i.e., touch/no-touch and some vector valued sensors) and limited sensors, with high hysteresis and poor wear and tear. As shown in table entry 18, the next 5 years will see better control techniques (possibly hybrid, as Raibert and Craig [37] suggest) and the development of array sensors with more applications. But the real progress of broad commercialization, a true sense of feel, and the development and understanding of the control/programming issues will take us into the 10-year time frame.

Research in tactile sensing is being done at Ohio State University, MIT, JPL, CMU, Stanford University, the University of Delaware, General Electric in Schenectady, and in France. Force sensing is being done at MIT, Draper, Astek, IBM, and other commercial firms.

Research support is not on a large scale: too few people, not enough money. Nevertheless, this is a critical area for assembly and other complex tasks. A concentrated research program by a major funding agency or agencies would speed progress.

Artificial Intelligence Research

As can be seen from the review of research areas, there are many avenues for combining AI and robotics. The future will see a natural combination and extension of each area into the domain of the other, but to date there are no true joint developments. MIT, Stanford, and CMU are beginning to lead the way in joint efforts, and many others are sure to join in.

The general area of reasoning and AI can be partitioned in many ways, and every taxonomy will result in fuzzy edges and work that resists a comfortable pigeonhole. A large portion of AI research can nevertheless be characterized in terms of advisory Systems that strive to assist users in some information processing task. This research can be categorized as work on expert systems, natural-language data base access, computer-aided instruction (CAL), intelligent tutors, and automated assistants.

A great deal of basic research is conducted without recourse to specific task orientations, and progress at this level penetrates a variety of areas in a myriad of guises. Basic research is conducted on knowledge representation, learning, planning, general problem solving, and memory organization. It is difficult to describe the milestones and research plateaus in these areas without some technical introduction to the issues, which is well beyond the scope of this paper. Problems and issues in these areas tend to be tightly interrelated, so we will highlight some of the more obvious accomplishments in a grossly inadequate overview of basic research topics. For further detail, see reference 38.

Expert systems are specialized systems that work effectively in providing competent analyses within a narrow area of expertise (e.g., oil exploration, diagnosis of infectious diseases, VLSI design, military intelligence, target selection for artillery). A few commercial systems are being customized for specific areas. Typically, current expert systems are restricted in a number of ways. First, the expertise is restricted in a very narrow corpus of knowledge. Examples include pulmonary function disorders, criteria for assessing copper deposits, and configuring certain types of computers. Second, interactions with the outside world and the consequent types of information that can be fed into such expert systems are capable of only a very small number of responses--for example, 1 of 92 drug therapies. Finally, they adopt a single perspective on a problem. Consider, by way of contrast, that trouble-shooting an automobile failure to turn over the starter motor (electrical) suggests a flat battery. The battery is charged by the turning of the fan (part of the hydraulic cooling system). This turns out to be deficient because of a broken fan belt (mechanical).

Table entry 19 summarizes the current state of expert systems and reflects the expectation of their integration with other systems within 5 years and significant improvement within 10 years. Significant work centers are at Stanford, Carnegie-Mellon, Teknowledge, Schlumberger, and a variety of other locations.

74

Natural-language data base access is now limited to queries that address the contents of a specific data base. Some require restricted subsets of English grammar; others can unravel ungrammatical input, run-on sentences, and spelling errors. Some applications handle a limited amount of context-sensitive processing, in which queries are interpreted within the larger context of an interactive dialogue. We are just now seeing the first commercial systems in this area. As table entry 20 shows, we expect sophisticated dialogue capabilities for interactive sessions and better recognition capability for requests the data base cannot handle. More domains will have been tackled, and some work may relate natural-language access capabilities to data base design issues. We should see some efforts to connect expert-system capabilities with natural-language data base access to provide advisory systems that engage in natural-language dialogues in the next 5 years.

In 10 years the line between natural-language data base access and expert systems will be hard to draw. Systems will answer questions and give advice with equal ease but still within well-specified domains and limited task orientations. Key research efforts are at Yale, Cognitive Systems, Teknowledge, Machine Intelligence Corporation, and other locations.

Basic research on automated assistants is now being conducted for a variety of tasks. As shown in table entry 21, this work, which takes place at MIC, SRI, the University of Massachusetts, IBM, and DEC, can be integrated with the other AI technologies. The field is not yet funded to any extent, but commercial interest is growing and should attract funding.

With respect to knowledge representation and memory organization, there are techniques that operate adequately or competently for specific tasks over restricted domains. Most of the

work in learning, planning, and problem solving has been domain-independent, with prototype programs operating in specific domains (e.g., learning by analogy). The domain-dependent work in these areas tends to start from a domain-independent base, augmenting this foundation with semantics and memory structures. As shown in table entry 22, progress is dependent on better understanding of knowledge; its representation is hard to predict.

Control Structure/Programming Methodology

Perhaps the most difficult area of all to cover is the future of control structures and programming methodology. In some sense, all the developments described impinge on this area; new mechanical designs, locomotion, dexterous hands, vision, contact/tactile sensors, and the various AI methodologies all affect the architecture of robot control and will affect the complexity of programming methodology.

In order to treat the subject in an orderly way, we deal first with a logical progression of control structure. Then, possibly with overlap, we deal with the other topics.

75

The most advanced current work in control structures uses multiple microprocessors on a common bus structure. Typically, such robot controllers partition the control problem into levels as follows:

1. Servo control to provide closed-loop feedback control.
2. Coordinate transformation to joint coordinates, and coordinated joint motion.
3. Path planning for simple interpolated (straight line) motion through specified points.
4. Simple language constructs to provide subroutines, lock-step interaction, and binary sensor-based program branches.
5. Structured languages, limited data base control) complex sensor communication, and hierarchical language definitions.

Levels 1 to 3 are common in most servo robots; level 4 is represented by the first-generation languages such as VAL on Unimation robots, while level 5 represents second-generation languages as found in the IBM AML Language, the Automatix RAIL, and at the National Bureau of Standards.

Beyond the first five levels of control are a diversity of directions being pursued to different extents by various groups. Thus, we can expect a number of developments in the next 5 years but clearly will not see them integrated in that time. As shown in table entry 23, we see the following extensions:

- Graphic systems will be used to lay out, program, and simulate robot operations. Such systems are starting to enter the market today from McAuto, Computervision, GCA, and others.
- Hierarchical task-oriented interface languages will be developed on the current structural languages (AML, RAIL, etc.) to allow process planners to program applications.
- Robot operating systems and controllers will be more powerful. They will remove the burden of low-level control over sensors, I/O, and communication; that is, they will do more of what computer operating systems do for their users today.
- Interfaces to other nonhomogeneous computers via developments in local area networks and distributed computing will broaden coordination beyond the lock-step synchronization available today.
- The use of multiple arms, dexterous hands, locomotion mechanisms, and other mechanical advances will foster the definition of a sixth level of control. This will emerge from research labs and be available in some rudimentary form.
- The incorporation of AI technology in the use of expert systems is in the laboratory plans of some now. This, coupled with the use of natural-language front ends and knowledge engineering, will begin the definition of a seventh level of control.
- The linkage of robot control/programming systems with CAD, CAM, and other factory data bases will be made.

Beyond these advances in new areas will be significant improvements in the first five levels as computers get more powerful and cheaper.

76

For example, the use of kinematic and dynamic models discussed in table entries 12 and 13 will affect the first five levels, as will the development and instrumentation of new sensors for resolving robot position.

The research in these areas is growing rapidly. Robotics institutes at major universities--CMU, MIT, Stanford, Florida, Lehigh, Michigan, RPI, and others--are now accelerating their programs under funding from DOD agencies, DARPA, and NSF. As the programs grow, the need for research dollars escalates, but so do the results. Robotics research is expected to expand significantly in the next decade. Commercial firms, both vendors and users, are linking themselves with universities. The list of firms involved includes IBM, Westinghouse, DEC, GE, and many others.

The 10-year time frame is very difficult to predict. This is because of the variety of technologies that must interact and the dependence on the output of a myriad of research opportunities being pursued. However, we feel the following to be conservative estimates.

- Robotics will branch out beyond industrial arms to include a wide scope of automatic equipment. The directions will depend on funding emphasis and other such factors.
- Sensor-based, advanced mechanical, partially locomotive (in restricted domains), somewhat intelligent robots will have been developed.

- Many integration issues and further technological advances will still remain open research questions.

Conclusion

In conclusion, one is forced to observe that the following table describes a technology that is very active--a technology that, while diversifying into many research areas, must be integrated for true success.

For those whose interest is in transferring the technology outside the manufacturing arena, immediate focus on targeted projects appears to be required. Although robotics and AI will be integrated, and the focus on manufacturing will broaden by an evolutionary process, the process will be painfully slow, even when pushed by well-funded initiatives.

Summary State of the Art for Robots and Artificial Intelligence

Now	In 5 Years	In 10 Years
<u>Mechanical Design and Activation of the Manipulator</u>		
1. Single arms with fixed bases	2 or 3 rigidly mounted arms designed to work together	Multiple arms with coordinated motion
2. Heavy; designed to be rigid	Designed to be rigid but lightweight, using composite materials	Designed to be very light-weight and flexible
3. Humanlike mechanical arrangements; linkage systems	No change	Nonlinkage design (e.g., snakes, butterflies)
4. Discrete degrees of freedom (DOF)	No change	Continuous degrees of freedom without discrete joints; flexible elements
5. Simple joints, revolute or sliding; Cincinnati Milacron has one version of the 3-roll wrist now	Flexible joints possible; better discrete joints (e.g., 3-roll wrist)	Flexible joints as above
6. Actuators are electrical, hydraulic, and pneumatic; heavy, low power, often require transmission gears that result in backlash problems	Some improvement: lighter weight, rare-earth motors, direct drive	New actuator concept: distributed actuator (muscle type)

7. Joint bearing, conventional high friction and stiction; poor motion performance	New discrete bearing designs (air bearings); some flexible joints possible	No discrete joints, possibly no bearings: flexible elements, for mobility
8. No absolute accuracy; repeatability 0.1 in. to 0.005 in. except in highly specialized semiconductor applications	Some absolute accuracy is required (for offline programming); good repeatability of 0.005 in. to 0.001 in.	Controlled to micron level as required; also coupled to force and position sensors to give broad functional range
9. Fixed location--some on tracks or wire-guided vehicles; walking, wheeled, and hopping robot mechanisms are now in research labs	Mobility based on wheeled-track vehicles in controlled environment (flat factory floor); rudimentary walking in specific environments	Mobility in semicontrolled environment, better vehicular control, some walking ability
10. Limited work envelopes	More flexible, but constrained envelopes as defined by factors above	Greatly improved work domains by new designs linkages, mobility, as defined above
11. Operate in controlled environment (factories) or with support systems (e.g., underwater applications); not self-contained, umbilical cords, big power unit	Packaging for uncontrolled environments; not self-contained	Possibly self-contained; wider range of environments tolerated (e.g., nuclear hardened)

Now	In 5 Years	In 10 Years
12. The kinematics are a significant computational burden that limits practical performance--real limitation is on <u>real time control and action</u>	New dedicated chips will be available to greatly reduce computational burdens--some slow motion real time possible	Computation not an issue; real time kinematic possible at high speed
13. Dynamics are not considered in robot design and performance. They are basically slow devices operating in "quasistatic" modes. Control systems are on joints only and position only and are relatively primitive. Typically, velocity-dependent and inertial terms ignored. Arms made to run slowly to compensate	Robots will be designed for higher-speed performance with some absolute accuracy. There will be combined force and position control with respect to the workspace rather than joints. Robotic trajectories will be planned for optimal dynamic performance, including the effects of actuator and robot dynamics, and limitations. Adaptive control methods will be available, so the robot will be insensitive and tolerant (dynamically) to its environment and its task	Robots will be high speed and lightweight, with <u>tuned dynamic behavior</u> . Systems will control and exploit their flexibility to achieve high performance. Issues of dynamics and performance in most cases will move to a higher level. Questions of control of individual elements will be transparent, such as the motion of control surfaces in supersonic aircraft is not considered by the pilot

End Effectors

<p>14. Currently grippers and special tools. They are typically binary (open or closed, on or off) and have few or rudimentary sensors; very simple mechanical actions, mostly one DOF such as parallel jaw pneumatically; and rudimentary force control</p>	<p>End effectors with proportional mobility--a hand that can be centered and servoed to fit a wide variety of objects; position and force sensors and limited tactile sensing; several discrete DOF; major emphasis still on grasping or sucking, with limited assembly or quick-change hand availability. Research labs will have developed multifingered hands and demonstrated their use to grasp a variety of three dimensional shapes</p>	<p>Continuous motion, intelligent control and sensing at the wrist, fingers, and fingertips. Beginning to be controlled by vision and other noncontact sensors to perform assembly</p>
<p>15. Quick-change hands are available today on a limited special basis due to a lack of standards for their interconnection to a variety of robots</p>	<p>Development of a standard robot-arm-to-end-effector interface. Commercial availability of a family of hands for tasks such as assembly, using adaptations of current tools and grippers</p>	<p>Specially designed sensor-based robot hands with tools for a family of tasks. All able to fit the standard interface</p>

Now	In 5 Years	In 10 Years
<u>Vision Sensors</u>		
<p>16. Current commercial systems are restricted to binary image and simple features; gray-scale and color are available today only in very restrictive form</p>	<p>VLSI implementation now in labs will be commercialized. This will facilitate edge images from gray-scale data, and richer feature sets will be developed</p>	<p>Systems that permit rapid recognition and provide orientation of limited classes of objects from arbitrary points of view</p>
<p>17. 3-D vision systems, structured light, and stereo approaches to acquiring depth image are rudimentary and only beginning to emerge from laboratories into commercial systems</p>	<p>Laboratory systems of several varieties will be commercially available. They will produce depth maps in controlled situations, but they will be slow, will produce noisy images, and have limited resolution. They will permit 3-D surface inspection and will discriminate objects with large shape differences</p>	<p>Reliable hardware for depth images and systems for tracking and recognizing moving objects</p>

Contact and Tactile Sensing

18. Few robots have force or tactile sensors. The IBM RSI is an exception. Limited use of commercialized RCC and IRCC versions of Draper Research products provide limited control capacity at present

Force-sensing wrists and techniques for programming and controlling force will be available. They are likely to work only in benign situations, but should be able to tighten nuts, insert shafts, pack objects--simple assembly operations. Will not yet be good enough to examine objects by feeling them

Well-established techniques for creating and using these sensors will be developed. Determining shape of object detecting slippage in grip, inspecting for cracks, and programming in the force domain will be possible. Touch sensors will be implemented in hardware, probably using VLSI technology. This will permit all of the above and offer a wider range of force monitoring and compliance operations

Now	In 5 Years	In 10 Years
-----	------------	-------------

Artificial Intelligence

19. Expert systems that work effectively in providing competent analysis within a narrow area of expertise, e.g. oil exploration, medical diagnosis, VLSI design, are being customized and commercialized. They are limited by a narrow body of

Automated design assistance for building and updating expert systems. Formalization of knowledge gathering and integration of graphic displays for use in some applications. Integration with robot control systems and sensors to provide con-

Integrated systems that draw on multiple domains of expertise to formulate problem solutions. Possibly total automation in generating new expert systems for certain domains. Self-diagnosis



<p>21. Automated assistants research is now going on in a variety of tasks, such as word processing, text editing, and office automation</p>	<p>Systems that assist and familiarize users with the capabilities of the system being used</p>	<p>Integrated systems that draw on multiple domains and provide the user with greater task flexibility</p>
<p>22. Knowledge representation in restricted domains is now workable (see entries 19-21). But learning, problem-solving, and planning systems need broader domains.</p>	<p>Increased understanding of trade-offs between domain-independent and domain-dependent techniques</p>	<p>Possibly a notation system that allows formulation of models that are sensitive to domain constraints without having specific commitments to particular domains</p>

Control Structure/Programming Methodology

<p>23. The control hierarchy of robots sometimes implemented on multiple microprocessors has at most 5 levels now.</p> <ol style="list-style-type: none"> 1. Servo control of joints 2. Coordinate transformation and coordinated joint motion. 3. Interpolated path planning for smooth motion paths. 	<p>Individual elements of progress (not all in any one offering) will be developed.</p> <ul style="list-style-type: none"> • Graphical layout of robotic cells and programming will be commercialized • Hierarchical task-oriented interface languages designed for process planners will be developed. 	<p>Levels six and seven as defined in the previous column will permit domain-dependent, sensor-based intelligent robots. Many integration issues and advances to technology will still be open questions. Robotics will broaden in scope beyond manufacturing to limited-domain automatic devices in new areas.</p>
---	---	---

Now	In 5 Years	In 10 Years
<p>4. Simple subroutines, use of sensors, and lock-step coordination</p> <p>5. Rudimentary operating system, structural language, complex sensor interface, hierarchical constructs</p>	<ul style="list-style-type: none"> • Robot operating systems will do more for the user who uses sensors to permit task orientation • Interfaces to other nonhomogeneous computers will broaden coordination beyond lock-step available now • Multiple arm, dexterous hand, locomotive control, and other new mechanical advances will define a sixth level of control and be available • The incorporation of AI technology in the form of expert systems, natural-language front ends, and knowledge representation will define a seventh level of control. • Data bases from CAD, CAM, and other sources will be incorporated to the language and control structure 	

REFERENCES

1. National Bureau of Standards. 1980. Proceedings of NBS/Air Force ICAM Workshop on Robot Interfaces, June 4-6. NBSIR 80-2152.
2. Taylor, R. H., P. D. Summers, and J. M. Meyer. 1982. AML: A Manufacturing Language. International Journal of Robotics Research 1(3):19-41.
3. Birk, J. and R. Kelley, eds. 1980. Research Needed to Advance the State of Knowledge in Robotics. Kingston: Rhode Island University.
4. Roth, B. Kinematic Design for Manipulation, in [3], pp. 110-118.
5. Dubowsky, S. Dynamics for Manipulation, in [3], pp. 119-128.
6. Houston, R. Compliance in Manipulation Links and Joints, in [3], pp. 129-145.
7. Paul, R. P. 1981. Robot Manipulators Mathematics Programming and Control. Cambridge, Mass.: MIT Press.
8. Brady, M. and J. Hollerbach. 1982. Robot Motion: Planning and Control. Cambridge, Mass.: MIT Press.
9. Toepperwein, L. L., M. T. Blackmon, R. Fukui, W. T. Park, and B. Pollard. 1980. ICAM Robotics Applications Guide. Vol. II. Technical Report AFWAL-TR-80-4042.
10. Salisbury, J. K. and J. Craig. 1982. Articulated Hands: Force Control and Kinematic Issues. International Journal of Robotics Research 1(1):4-17.
11. Hollerbach, J. M. 1982. Workshop on Dexterous Hands. MIT AI Memo.
12. Orin, D. E. 1982. Supervisory Control of a Multilegged Robot. International Journal of Robotics Research 1(1):79-91.

-
13. Gleason, G. J. and G. Again. 1979. A Modular Vision System For Sensor Control Manipulation and Inspection. SRI Report, Project 4391. SRI International.
 14. Lavin, M. A. and L. I. Lieberman. 1982. AML/V: An Industrial Machine Vision System. International Journal of Robotics Research 1(3):42-56.

15. Nagel, R. N., et al. 1979. Experiments in Part Acquisition Using Robot Vision. SME Technical Paper MS 79-784.
16. Brady, M. 1982. Computational Approaches to Image Understanding. Computing Surveys 14:4-71.
17. Nevins, J. L., et al. Exploratory Research in Industrial Assembly and Part Mating. Report No. R-1276. Cambridge, Mass.: Charles Stark Draper Laboratory. 193 pp.
18. Harmon, L. D. 1982. Automated Tactile Sensing. International Journal of Robotics Research 1(2):3-32.
19. Bejczy, A. K. 1979. Manipulator Control Automation Using Smart Sensors. Paper delivered at Electro/79 Conference, New York, April 24-26.
20. Raibert, M. H. and J. E. Tanner. 1982. Design and Analysis of a VLSI Tactile Sensor. International Journal of Robotics Research. 1(3):3-18.
21. Hillis, W. D. 1982. A High Resolution Image Touch Sensor. International Journal of Robotics Research. 1(2):33-44.
22. Albus, J. S., A. J. Barbera, M. L. Fitzgerald, R. N. Nagel, G. J. VanderBrug, and T. E. Wheatley. 1980. Measurement and Control Model for Adaptive Robots. Pp. 447-466 in Proceedings, 10th International Symposium on Industrial Robots, Milan, Italy, March 5-7.
23. Nagel, R. N., et al. 1982. Connecting the Puma Robot With the MIC Vision System and Other Sensors. Pp.447-466 in Robot VI Conference Proceedings, Detroit, March 2-4.
24. D. R. Brown, et al. 1982. R&D Plan for Army Applications of AI/Robotics. SRI Project 3736. SRI International. 324 pp.
25. Nau, D. S. 1982. Expert Computer Systems and Their Applicability to Automated Manufacturing. NBSIR 81-2466.
26. Charniak, E., and Y. Wilks, eds. 1976. Computational Semantics: An Introduction to Artificial Intelligence and Natural Language Comprehension. Amsterdam: North Holland Publishing Co.

-
27. Lehnert, W., and M. Ringle, eds. 1982. Strategies for Natural Language Processing. Hillsdale, N.J.: Lawrence Erlbaum Associates.

28. Nilsson, N. J. 1971. *Problem Solving Methods in Artificial Intelligence*. New York: McGraw-Hill.
29. Schank, R., and R. Abelson. 1977. *Scripts, Plans, Goals and Understanding*. Hillsdale, N.J.: Lawrence Erlbaum Associates.
30. Waltz, D. L. 1982. *Artificial Intelligence*. *Scientific American*. 247(4):118-133.
31. Winston, P. H. 1977. *Artificial Intelligence*. Reading, Pa.: Addison Wesley.
32. *Proceedings for the Conference on Applied Natural Language Processing, Santa Monica, Calif., February 1983*.
33. *Proceedings for the Association of Artificial Intelligence Conference on Artificial Intelligence (IJCAI 1969, 1973, 1975, 1977, 1979, 1981)*.
34. Ballard, D. H. and C. M. Brown. 1982. *Computer Vision*. Englewood Cliffs, N.J.: Prentice-Hall.
35. Rosenfeld, A. 1983. *Picture Processing: 1982*. Computer Science Technical Report. College Park: University of Maryland.
36. Dennicoff, M. 1982. *Robotics in Japan*. Washington, D.C.. Office of Naval Research.
37. Raibert, M., and J. Craig. 1981. *Hybrid Controller*. *IEEE Systems Management Cybernetics*.
38. Barr, A., and E. A. Feigenbaum, eds. 1981, 1982. *Handbook of Artificial Intelligence*, vols. I-III. Stanford, Calif.: HeurisTech Press.
39. *State of the Art of Vision in Japan*, *IEEE Computer Magazine* (13) 1980.

GLOSSARY OF ACRONYMS

AFOSR	Air Force Office of Scientific Research
AI	artificial intelligence
AML	manufacturing language developed at IBM
AMRDC	U.S. Army Medical Research and Development Command
ASB	Army Science Board
ASP	Automated Ammunition Supply Point
ATE	automatic test equipment
BITE	built-in test equipment
C ³ I	command, control, communication, and intelligence

CAD/CAM	computer-aided design and manufacturing
CAI	computer-aided instruction
CARP	computer-aided robot programming
CMU	Carnegie-Mellon University
CPU	central processing unit
CRT	cathode ray tube
DARPA	Defense Advanced Research Projects Agency
DART	expert system for the diagnosis of equipment failure
DEC	Digital Equipment Corporation
DMA	Defense Mapping Agency
ES	expert system
FLIR	forward-looking infrared
FMS	flexible manufacturing system
GE	General Electric Company
GM	General Motors Corporation
Hawk-Missile	CAI trainer at Fort Bliss Air Defense School
ICAM	Integrated Computer-Aided Manufacturing program of the U.S. Air Force
IR	industrial robot
IRCC	instrumented remote center of compliance developed at Draper Laboratories
JPL	Jet Propulsion Laboratory
MACSYMA	symbolic mathematics expert system

MCL	computer language developed at McDonnell Douglas
MIC	Machine Intelligence Corporation
MIT	Massachusetts Institute of Technology
MYCIN	production system for diagnosis and treatment of infectious diseases
NBC	nuclear, biological, and chemical
NBS	National Bureau of Standards
NSF	National Science Foundation
ONR	Office of Naval Research
Prospector	expert system to aid in exploration for minerals

PUFF	pulmonary function diagnosis expert system
P ³ I	preplanned product improvement
RAIL	Pascal-based second generation language by IBM
RAMS	reliability, availability, maintainability, and supportability
R&D	research and development
REMBASS	remotely monitored battlefield sensor system
RIA	Robot Institute of America
RPI	Rensselaer Polytechnic Institute
SAR	synthetic aperture radar
SRI	Stanford Research Institute
VAL	language developed by Unimation for Puma robot
VHF	very high frequency
VHSIC	Very High Speed Integrated Circuits
VIMAD	Voice Interactive Maintenance Assistance Development system (supported by DARPA)
VLSI	very large-scale integration
VTRONICS	set of projects for onboard, embedded sensing of vehicular malfunctions with built-in test equipment (BITE)

1 BACKGROUND

Throughout its history, the Army has been manpower-intensive in most of its systems. The combination of demographic changes (fewer young men), changed battlefield scenarios, and advanced technologies in improved robotics, computers, and artificial intelligence (AI) suggests both a need and an opportunity to multiply the effectiveness of Army personnel. Not only can these technologies reduce manpower requirements, they can also replace personnel in hazardous areas, multiply combat power, improve efficiency, and augment capabilities.

The Deputy Chief of Staff for Research, Development and Acquisition authorized the National Research Council to form a committee to review the state of AI and robotics technology, predict developments, and recommend Army applications of AI and robotics. This Committee on Army Robotics and Artificial Intelligence brought together experts with military, industrial, and academic research experience.

APPROACH

The committee began its work with a detailed review of the state of the art in robotics and artificial intelligence as well as with predictions of how the technology will develop during the next 5- and 10-year periods. This review is summarized in Chapter 2 and in its entirety forms the appendix of this report. It is the foundation of the committee's recommendations for selecting and implementing of applications.

The committee used its review of technology and information on Army doctrine, prior reports on Army applications of AI and robotics, and its combined military, university, and industrial experience to develop criteria for selecting applications and to recommend specific applications that it considers of value to the Army and the country. For each application recommended, the committee was asked to report the expected effects on personnel, skills, and equipment, as well as to provide an implementation strategy incorporating priorities, costs, timing, and a measure of effectiveness.

PRIOR STUDIES

As background to its efforts, the committee was briefed on and reviewed three studies completed during 1982 on Army robotics and artificial intelligence:

D. R. Brown, et al., R&D Plan for Army Applications of AI/Robotics, SRI International, May 1982 (Contract No. DAAK70-81-C-0250, U.S. Army Engineer Topographic Laboratories).

Army Plan for AI/Robotics Technology Demonstrators, Department of the Army, June 1982.

Report of the Army Science Board Ad Hoc Subgroup on Artificial Intelligence and Robotics, Army Science Board, September 1982.

Each contributes to the base of knowledge regarding these expanding new technologies and offers insights into potential applications to enhance the Army's combat capabilities. Their conclusions are briefly reviewed here to place the contribution of this particular report in a proper context.

R&D Plan for Army Applications of AI/Robotics

The report by SRI cites as the primary motivation for the application of AI and robotics to Army systems the need to conserve manpower in both combat and noncombat operations. It covers more than 100 possible Army applications of AI and

robotics, classified into combat, combat support, and combat service support categories. Many of the applications, though listed as distinct, could easily be drawn together to serve as generic applications. The report focuses on the need to document justification for the value of AI and robotics in Army applications in general, but the committee found that it lacked sufficient detail for ranking the many applications to pursue those of greatest interest and potential payoff.

From the 100 specific concepts that the SRI study considered, 10 broad categories of application were selected. An example from each of these 10 categories was chosen for further study to identify technology gaps and provide the basis for the research plan recommended by the study.

Included in that plan were 5 fundamental research areas, 97 specific research topics, and 8 system considerations. Most potential applications were judged to require advancement of the technology base (basic research and exploratory development) before advanced development could begin. In fact, the study estimated that development on only four could be started in the next 10 years, and two would require deferral of development until the year 2000.

A briefing on the Army Proposed Plan was given to the committee at its initial meeting. The report identified five projects for application of AI or robotics technology to demonstrate the Army's ability to exploit AI and robotics:

Robotic Reconnaissance Vehicle with Terrain Analysis,

Automated Ammunition Supply Point (ASP),

Intelligent Integrated Vehicle Electronics,

AI-Based Maintenance Tutor,

AI-Based Medical System Development.

Of these five proposed demonstrations, technical availability assessments placed one in the near term, one in the mid-to-far term, and the other three in the far term. Cost estimates and schedules appear optimistic to this committee, considering that much of the effort was neither funded nor programmed at that time.

Report of the Army Science board

Ad Hoc Subgroup on Artificial Intelligence and Robotics

The Army Science Board Ad Hoc Subgroup was established to provide an assessment of the state of the art of AI and robotics as fast-track technologies and of their potential to meet Army needs. It concentrated its efforts on those aspects with which it could deal rapidly and relatively completely; it also considered the five Army demonstrators and supported them.

The report grouped the five demonstrators into two categories:

proceed as is or proceed with modification. The subgroup recommended changes to the maintenance tutor and the medical system, and recommended that the other three demonstrators proceed as planned. Other battlefield technology topics recommended were automatic (robotic) weapons, automatic pattern recognition, and expert support systems.

Noting that the introduction of technology into weapon systems could be hampered by management problems, the subgroup recommended establishing a single dedicated proponent of AI and robotics in the Department of the Army, giving preference to existing equipment and technology, and creating an oversight committee from the Army's materiel developer and user communities.

The subgroup tied its recommendations to the five technology thrusts that the Army has designated to receive the majority of research and development funds (lines 6.1, 6.2, and 6.3a of the budget) during the next five-year funding period:

Very Intelligent Surveillance and Target Acquisition,

Distributed C31,

Self-Contained Munitions,

Soldier/Machine Interface,

Biotechnology.

CONTRIBUTION OF THIS REPORT

This committee is indebted to the foregoing efforts for the base they provide, a base which this report attempts to expand. Our recommendations are founded on a comprehensive assessment of the state of the art and forecasts of technology growth over the

next 10 years. The details of that assessment are contained in the Appendix. We hope that our recommendations to the Army will provide a realistic technical assessment that will enable the Army, in turn, to concentrate its efforts in areas offering the most potential return.

No two groups considering possible AI and robotics applications will have identical lists of priorities. This committee used the combination of Army needs and the direction of technology development as a guide in narrowing the list of possible applications. The National Research Council is unique in the diversity of backgrounds of the experts it brings together. The members of this Committee on Army Robotics and Artificial Intelligence have among them 248 years of industry experience, 110 years in academia, and 184 years in government. The recommendations in this report are the consensus of the committee, drawing on those years of experience.

We agree with the authors of studies we have reviewed that AI and robotics technologies offer great potential to save lives, money, and resources and to improve Army effectiveness. This report will support the need for ongoing work in these high-risk, high-technology fields that offer such great promise for the country's future securityhelp channel Army efforts into the most effective areas,build understanding of what AI and robotics can offer within the broad groups in the Army that will need to work with these technologies ,

provide realistic information on what AI and robotics technology can do now and the directions in which research is heading.

2 SUMMARY OF THE TECHNOLOGY

DEFINITIONS

We used the Robot Institute of America's definition of a robot as

a reprogrammable multi-function manipulator designed to move

material, parts, tools, or specialized devices through variable

programmed motions for the performance of a variety of tasks.

The main components of a robot are the mechanical manipulator, which is a set of links that determine the work envelope of the robot and the ability to orient the hand; the actuation mechanisms, which are hydraulic, pneumatic, or electric; the controller, usually a computer, which controls motion by communicating with the actuation mechanism.

The robot can be augmented by the addition of end effectors, or "hands";

sensors, for performing measurements as required to sense the environment, including electromagnetic (visual, infrared, ultraviolet, radar, radio, etc.), acoustic, tactile, force, torque, spectographic, and many others.

other "intelligent" functions, such as understanding speech, problem solving, goal seeking, and commonsense reasoning.

None of these, strictly speaking, is part of the robot itself.

This chapter is a summary of the detailed report on the state of the art and predictions for AI and robotics technology contained in the appendix.

Artificial intelligence, as defined in SRI International's R&D Plan for Army Applications of AI/Robotics, is the part of computer science that is concerned with symbol-manipulation processes that produce intelligent action. By "intelligent action" is meant an act or decision that is goal-oriented, arrived at by an understandable chain or symbolic analysis and reasoning steps, and is one in which knowledge of the world informs and guides the reasoning.

The functions or subfields of artificial intelligence are natural-language understanding; that is, understanding English or another noncomputer language; image understanding; that is, the ability to identify what is in a picture or scene; expert systems, which codify human experience and use it to guide actions or answer questions; knowledge acquisition and representation;

heuristic search, a method of looking at a problem and selecting a path to the solution; deductive reasoning; planning, which entails an initial plan for finding a solution, then monitoring progress.

As this infant field develops, the list of subfields will expand. Artificial intelligence is the application of advanced

computer systems and software to these areas, with "intelligent behavior" as the intended result.

RESEARCH ISSUES

The categories of robotics research receiving the most effort are

improvement of mechanical systems, including manipulation design, actuation systems, end effectors, and locomotion;

improvement of sensors to enable the robot to react to changes in its environment; creation of more sophisticated control systems that can handle dexterity, locomotion, and sensors, while being user friendly.

In artificial intelligence, expert systems is the area of research closest to being ready to move from the laboratory to initial commercial use.

Research on the kinematics of design, models of dynamic behavior, and alternative design structures, joints, and force programming is leading to highly accurate new robot structures. This research will lead to robots capable of applying force and torque with speed and accuracy and will transform today's heavy, rigid, single

robotic arms into more lightweight, ultimately more flexible arms capable of coordinated motion.

Research on end effectors--the hands attached to a robot--seeks to improve dexterity, enabling robots to handle a variety of parts or tools in complex situations. Two goals are the quick-change hand and the dexterous hand. The robot would be able to change a quick-change hand by itself, attaching the means of transmitting power as well as the physical hand to the arm.

Although the dexterous hand is beyond the current state of the art, there are some interesting present approaches. One is a variable finger selection; another is the use of materials that will produce signals proportional to surface pressures. This is coupled with research in microelectronics to analyze and summarize the signals from these multisensored fingers for decision-making outputs.

Early attention to locomotion has led to a large number of robots in current use mounted on tracks or an overhead gantry. Progress has recently been made on a six-legged walking robot that is stable on three legs.

A middle ground between tracked and unconstrained vehicles is a wire-guided vehicle used in plants. These vehicles have onboard microprocessors that communicate with a central control computer at stations placed along the factory floor. The vehicles travel along a wire network that is kept free of permanent obstacles; bumper sensors prevent collisions with temporary obstacles.

Sensors

The purpose of sensors is to give the robot adaptive behavior--that is, the ability to respond to changes in its environment. Vision and tactile sensors have received the lion's share of research effort. While tactile sensors are still fairly primitive, vision systems are already commercially available.

Vision systems enable robots to perform the following types of tasks:

identification or verification of objects,
location of objects and their orientation,
inspection,
navigation and scene analysis,

guidance of the servo mechanism, which controls position through feedback.

The first three tasks can be performed by today's commercial systems. Three-dimensional vision systems are at present rudimentary.

Tactile sensors are just beginning to be commercialized. Within the next few years, force-sensing wrists and techniques for controlling them will be available for such tasks as tightening nuts, inserting shafts, and packing objects. More research will be needed before they can work in other than benign environments.

Control Systems

The underlying research issue in control systems is to broaden the scope of the robot to include dexterous hands, locomotion, sensors, and the ability to perform new complex tasks.

Robots are typically programmed by either the lead-through or the teach-box method. In the former the controller samples the location of each of the robot's axes several times per second, while a person manipulates the robot through the desired motions. The teach-box method enables the

operator to use buttons, toggle switches, or a joy stick to move the robot.

Programming languages for robots have long been under research. Early robot languages have combined language statements with use of a teach box. Second-generation robot languages, which resemble the standard structured computer language, have only recently become commercially available. It is these second-generation robot languages that create the potential to build intelligent robots.

Expert Systems

Artificial intelligence has generated several concepts that have led to the development of important practical systems. A subset of these systems has been called expert systems. As the name suggests, an expert system (ES) encodes deep expertise in a narrow domain of human specialty. Several expert systems have been constructed whose behavior surpasses that of humans. Examples include the MIT Macsyma system (symbolic mathematics), the Digital Equipment Corporation R-1 system (configuring VAX computers), the Schlumberger dipmeter analyzer (oil well logs), and various medical expert systems, including PUFF (pulmonary function

diagnosis) in regular use at San Francisco Hospital. Expert systems' behavior in research laboratories and the civilian sector is cause for optimism in the military sector.

One can consider expert-systems support not only at the corps and division levels but also for battalions and regiments. As envisioned in the Air Land Battle 2000 scenario, battalion and regimental formations will be operating in forward battle areas in a dispersed manner. Expert-system support at this level will be particularly helpful in increasing combat effectiveness through flexibility and adaptability to varied, complex situations and improved survivability of men and machines.

Although there is cause for optimism, current expert systems have significant limitations and require intensive basic research if the technology is to be successfully transferred from the university laboratory to make rugged operational systems.

Present expert systems support only narrow domains of expertise. As the domain of application becomes broader, the number of alternative courses of action increases

exponentially and effectiveness decreases exponentially. Though research is addressing this issue, practical expert systems are likely to be severely restricted in their domain for the next 5 years.

Only limited knowledge-representation languages for data and relations are available.

The input and output of most expert systems are inflexible and not in English (or any other natural language).

Expert systems still require laborious construction--approximately 10 man-years for a sizable one.

Because present expert systems need one domain expert in control to maintain consistency in the knowledge data base, they have only a single perspective on a problem.

Many expert systems are difficult to operate.

3 CRITERIA FOR SELECTION OF APPLICATIONS

The committee spent a great deal of time developing criteria for the selection of Army applications of robotics and

artificial intelligence. These criteria were essential in guiding the work of the committee; but beyond that, they are more broadly applicable to future decisions by the Army as well as by others. The criteria for selecting applications reflect both the immediate technological benefits and the attitudinal and managerial considerations that will affect the ultimate widespread acceptance of the technology.

REASONS FOR APPLYING ROBOTICS AND ARTIFICIAL INTELLIGENCE

The introduction of robotics and artificial intelligence technology into the Army can result in a number of benefits, among them the following:

improved combat capabilities,
minimized exposure of personnel to hazardous environments,
increased mission flexibility,
increased system reliability
reduced unit/life-cycle costs,
reduced manpower requirements,
simplified training.

In selecting applications from the much larger list of possibilities, the committee not only looked for opportunities to achieve those benefits but also sought affirmative answers to the following questions: Army.

Will it perform, in the near term, an essential task for the

Can its initial version be implemented in 2 to 3 years?

Can it be readily upgraded as more sophisticated technology becomes available?

Does it tie in with existing, related programs, including programs of the other services?

Will it use the best technology available in the scientific community?

These considerations should help to ensure initial acceptance and continuing success with these promising developing technologies.

COMBINING SHORT-TERM AND LONG-TERM OBJECTIVES

Initial short-term implementation should provide a basis for future upgrading and growth as the user gains experience and

confidence in working with equipment using robotics and AI technology. To this end the Army's program should be carefully integrated and include short-term, achievable objectives with growth projected to meet long-term requirements.

As a result; some of the applications chosen may at first appear to be implementable in the short term by other existing technologies with lower cost and ease. However, such short-term expediency may cause unwarranted and unintended delay in the ultimately more cost-effective application of new developing robot technologies. To prevent this problem, short-term applications should be

applied to existing, highly visible systems,

reasonably afforded within the Army's projected budget,

within the state of the art, requiring development and engineering rather than invention or research,

able to demonstrate an effective solution to a critical Army need ,

achievable within 2 to 3 years,

not redundant with efforts in DARPA or the other services.

On the other hand, the committee considered long-term applications to be important vehicles for advancing research in these technologies and, in some cases, for introducing useful applications of robotics and artificial intelligence. These more advanced applications would ultimately, at reduced cost, assist in meeting the changing requirements of the modern battlefield envisioned in the Army's Air Land Battle 2000 concept.

The principle that guided the committee's selection of applications, therefore, was to combine short-term and long-term benefits; that is, to select applications that can be implemented quickly to meet a current need and, in addition, can be upgraded over the next 10 years in ways that advance the state of the art and perform more complex functions for the Army.

PLANNING FOR GROWTH

For the near term, using state of the art technology and assuming that a demonstration program starts in 1 1/2 to 2 years and continues for 2 years, the

committee recommends that projects be selected based not

only on what is commercially available now but also on technology that is likely to become available within the next 2 years.

During the next 4 to 5 years, while the Army is developing its demonstration systems, annual expenditures by university, industrial, government, and nonprofit laboratories for R&D and for initial applications will probably exceed several hundred million dollars per year worldwide. To be timely and cost effective, Army demonstration systems should be designed in such a way that these developments can be incorporated without discarding earlier versions.

It is therefore of the utmost importance to specify, at the outset, maximum feasible computer processor (and memory) power for each application. Industry experience has shown that the major deterrent to updating and improving performance and functions has been the choice of the "smallest" processor to meet only the initial functional and performance objectives.

It is at least as important to ensure that this growth potential be protected during development of the initial applications

Both industry and the Army have known programmers with a propensity to expand operating and other systems until they occupy the entire capacity of design processor and memory.

Robots are currently being developed that incorporate external sensors permitting modification of the sequence of motions, the path, and manipulative activities of the robot in an adaptive manner. The status of the "dumb, deaf, and blind" robot is being raised to that approaching an "intelligent" automaton. This upgraded system can automatically cope with changes in its reasonably constrained environment.

The earliest adaptive robot systems are just beginning to be incorporated into production lines. Most of these systems are presently in an advanced development stage, worked on by application engineers for early introduction into production facilities. Such systems, called third-generation robot systems, are expected to supplement the second-generation robot systems (having programmable control but lacking sensors) in the next 2 to 3 years. Shortly thereafter, as more and more assembly operations are automated, they are likely to become the dominant class of robot systems. In view of these

technological developments, the Army demonstration Systems should, at the very least, be based on the third-generation robot Systems capable of being readily upgraded with minimum change in the internal hardware configuration, relying on future additions of readily interfaceable external sensors and software.

SELECTING APPLICATIONS TO ADVANCE

PARTICULAR TECHNOLOGIES

In addition to considering the benefits that result from applying robotics and artificial intelligence, the Army has the opportunity to use its choice of applications to take an active role in advancing

particular technologies. Because robotics and AI are developing rapidly, the committee believes that Army should support a range of component technologies.

The two fields are at present separate, and the possible applications can be divided into those that are primarily robotics and those that are primarily artificial intelligence. The robotics applications can be further divided into those that primarily advance end-effector (hand)

technology and those that primarily advance sensor technology.

The AI applications can be divided into a number of types, of which the furthest developed is expert systems. The committee limited its consideration of AI applications to expert systems, in keeping with its goal of short-term implementation of limited aspects. The primary technology for expert systems is cognition.

Each of these areas--effectors, sensors, and cognition--is an important source of technology for the Army and for this country's industrial base. To encourage R&D in these areas and to enable the Army to have some initial experience in each area, the committee agreed to recommend three applications, one directed at each.

4 RECOMMENDED APPLICATIONS AND PRIORITIES

The committee used the criteria described in Chapter 3 to develop an initial list of 10 possible Army applications of robotics and artificial intelligence. These were discussed at length and narrowed to six applications that met the criteria, three of which are strongly recommended.

Many hours of committee discussion are reflected in the following list. The

committee found it impossible to match the large numbers of possible applications and criteria in any systematic way. No two groups applying the criteria would arrive at identical lists of Army projects to recommend. The applications recommended below are eminently worthwhile in the judgment of the committee. They clearly address current Army needs, offer short-term benefits, are likely to give Army personnel some positive early experiences with the technology, and are capable of being upgraded.

AN INITIAL LIST

With these considerations in mind, the committee developed the following list of 10 potential applications of robotics and artificial intelligence. Not all of these applications are recommended by the committee; this list is the result of the committee's first effort to narrow down the vast number of possible applications to those most likely to meet the criteria described earlier.

Automatic Loader of Ammunition in Tanks. This system would require development of a robot arm with minimum degrees of freedom for use within the tank. The arm would be capable of acquiring rounds from a magazine

or rack and loading them into the gun, with a vision system to provide the means to correct for imprecise positioning of rounds and gun and tactile or force sensors to ensure adequate acquisition.

Sentry Robot. A portable unattended sentry device would detect and report the presence of personnel or vehicles within a designated area or along a specified route. The device would also be capable of sensing the presence of nuclear, biological, and chemical contaminants.

Flexible Material-Handling Modules. Adaptive robots mounted on wheeled or tracked vehicles would identify and acquire packages or pallets to load or unload. There are so many potential applications for material-handling systems that material-handling robots are likely to become as ubiquitous as the jeep in the Army supply system, with applications in forward as well as rear areas.

Robotic Refueling of Vehicles. A wheeled robot fitted with an appropriate fuel dispenser (a tool for inserting into a fuel inlet) could automatically refuel a variety of vehicles.

Counter-Mine System. Adaptive robots mounted on wheeled or tracked vehicles

could be fitted with specialized sensors and probing or digging tools to find and dispose of buried mines. Vehicles could be remotely controlled in the teleoperator mode.

Robot Reconnaissance Vehicle. The remotely controlled reconnaissance vehicle that the Army is considering as a major demonstration project could be fitted with one or more external robot arms and equipped with vision and other sensors. This would expand the utility of the system to perform manipulative functions in forward, exposed areas, such as retrieval of disabled equipment; sampling and handling nuclear, biological, and chemically active materials (NBC); and limited decontamination.

Airborne Surveillance Robot. A semiautonomous aerial platform fitted with sensors could observe large areas, provide weather data, detect and identify targets, and measure levels of NBC contamination.

Intelligent Maintenance, Diagnosis, and Repair System. An ES, specialized for a particular piece of equipment, would give advice to the relatively untrained on how to operate, diagnose, maintain, and repair relatively complex electronic, mechanical,

or electromechanical equipment. It would also act as a record of repairs, maintenance procedures, and other information for each major item of equipment.

Medical Expert System. This system would give advice on the diagnosis and evacuation of wounded personnel. A trained but not necessarily professional operator would enter relevant information (after prompting by the system) regarding the condition of the wounded individual, including any results of initial medical examination. The system would logically evaluate the relative seriousness of the wound and suggest disposition and priority. This system could be improved by having available a complete past medical record of the individual to be entered into the system prior to asking for its advice.

Battalion Information Management System. This system would provide guidance and assistance in situation assessment, planning, and decisionmaking. Included would be the automatic or semiautomatic production of situation maps, plans, orders, and status reports. It also would include guidance for operator actions in response to specific situations or conditions.

Although this list represents a considerable reduction from the many possible applications that have been conceived, a further narrowing is needed. Knowledgeable researchers and other resources are in such short supply that Army efforts in AI and robotics should be well thought out and focused. The remainder of this chapter presents in more detail the functions, requisite technology, and expected benefits of the committee's top six priorities.

As noted in Chapter 3, the committee recommends that the Army fund three demonstration projects, one in each of the areas of effectors,

sensors, and cognition. This committee's consensus is that, at a minimum, the following projects should be funded:

1. automatic loader of ammunition in tanks (effectors),
2. sentry robot (sensors),
3. intelligent maintenance, diagnosis, and repair system (cognition).

These applications all meet the criteria listed on pages 10-11: they meet a current

Army need, demonstrations are feasible within 2 to 3 years, and the systems can be readily upgraded. Together, these applications are strongly recommended for funding.

The committee also found the following applications to meet its criteria. If funding is available, these are also recommended:

4. medical expert system (cognition),
5. flexible material-handling modules (effectors) ,
6. battalion information management system (cognition).

As to the remaining applications, robotic refueling of vehicles is an example of a flexible material-handling module (priority 5) and the airborne surveillance robot is an upgraded version of the sentry robot (priority 2). The reconnaissance vehicle is not in this committee ' s recommended list because a demonstration is not likely to be possible within 2 years. The counter-mine vehicle is not recommended because the problem seems better suited to a less expensive, lower-technology solution.

AUTOMATIC LOADER OF AMMUNITION IN TANKS

At present the four-man crew of a U.S. tank consists of a commander, a gunner, a driver, and a loader. The loader receives verbal instructions to load a particular type of ammunition; he then manually selects the designated type of ammunition from a rack, lifts it into position, inserts it into the breech, completes the preparation for firing, and reports the cannon's readiness to fire. The gunner, who has been tracking the intended target, has control of firing the cannon. When fired, the hot, spent casing is automatically ejected and is later disposed of, as convenient, by the loader. The loader occasionally unloads and restores unfired cartridges onto the rack.

With appropriate design of the complete ammunition loading system, these functions can be automated. The committee recommends the use of state-of-the-art robotics to effect this automation, eliminating one

man (the loader) from the crew, and potentially increasing the firing rate of the cannon, now limited by the loader's physical capabilities.

Functional Requirements

The major functional requirements of the system are

A computer-controlled, fully programmable, servoed robot designed for the special purpose of ammunition selection and loading. Its configuration, size, number of degrees of freedom, type of drive (hydraulic or electric), load capacity, speed precision, and grippers or hands would be engineered specifically for the purpose as part of the overall system design. Computer power in its controller would be adequate for interfacing with vision, tactile, and other sensors, and for communicating with other computers in the tank. Provisions would be made to introduce additional processing power in the future by leaving some empty "slots" in the processor cage. The principles of design for such a robot are now known, and the major requirement, after setting its specifications, is good engineering. A working prototype should take 1-1/2 to 2 years to produce.

A simple machine vision system designed to perform the functions of locating the selected type of ammunition in a magazine or rack, guiding the robot to acquire the round, and guiding the robot to insert the round into the breech. Although it is certainly possible to design a more specialized and highly constrained system, the proposed adaptive robot system provides

for greater flexibility in operation and reduction of constraints, and will enable more advanced functional capabilities in the future. The principles of designing an appropriate vision system are now available; the design for this purpose should not be difficult. Simplifying constraints such as colored, bar code, or other markings on the tips of shells and breech would eliminate tedious processing to obtain useful imagery for interpretation. Other sensory capabilities (e.g., tactile and force) could readily be added to the system if necessary, for confirming acquisitions and insertions. The robot computer could be programmed to accommodate all these sensors.

An ammunition storage rack (or, preferably, magazine) designed to facilitate both bulk loading into the tank and acquisition of selected ammunition by the robot gripper. It may even have an auxiliary electromechanical device that would push selected ammunition forward to permit easy acquisition by the robot, such action controlled by the robot computer.

Robot and vision computers integrated and interfaced with the fire control computer under control of the commander or gunner. This local computer network is intended for

use in later developments when further automation of the tank is contemplated. However, it could even be used in the short term to ensure that the type of ammunition loaded is the same type that is indexed in the fire control computer.

Benefits

The near term advantages (2 to 5 years) foreseen are

elimination of one crew member (the loader) and automation of a difficult, physically exhausting task that contributes little to the overall skills of the people who perform it;

potential increase in fire power by reducing loading time;

the availability of a test bed for further development and implementation of more advanced systems and increased familiarity of personnel with computer-controlled devices;

simplification of communications between commander, gunner, and loader, which may lead to direct control by the tank commander and potential reduction of errors during the heat of combat;

Army experience with computer control, especially of robot systems.

In the long term, if concurrent developments in automated tracking using advanced sensors occur, it may be feasible to eliminate the gunner, reducing the crew to a commander and a driver. This would make possible two-shift operations with two two-man crews operating and maintaining the tank over a 24-hour period, a considerable increase in operating time for very important equipment. Mechanization of the ammunition-loading function and an integrated computer network in place are prerequisites for this development.

A potential tank of the future could be unmanned--a tank controlled by a teleoperator from a remote post or hovering aircraft. The tank would be semiautonomous; that is, it could maneuver, load rounds, track targets, and take evasive action to a limited degree by itself, but its actions would be supervised by a remote commander who

would initiate new actions to be carried out by internally stored computer programs. Eliminating people on board the tank could lead to highly improved performance, now limited by human physical endurance and

safety. The tank would become an unmanned combat vehicle, smaller, lighter, faster, with far less armor and more maneuverable--essentially a mobile cannon with highly sophisticated control and target acquisition systems.

SENTRY/SURVEILLANCE ROBOT

The modern battlefield, as described in Air Land Battle 2000, will be characterized by considerable movement, large areas of operations in a variety of environments, and the potential use of increasingly sophisticated and lethal weapons throughout the area of conflict. Opposing forces will rarely be engaged in the classical sense--that is, along orderly, distinct lines. Clear differentiation between rear and forward areas will not be possible. The implications are that there will be insufficient manpower available to observe and survey the myriad of possible avenues by which hostile forces and weapons may threaten friendly forces.

Initially using the concepts and hardware developed in the Remotely Monitored Battlefield Sensor System (REMBASS), a surveillance/ sentry robotic system would provide a capability to detect intrusion in specified areas--either in remote areas

along key routes of communication or on the perimeter of friendly force emplacements. Such a system would apply artificial intelligence technology to integrate data collected by a variety of sensors--seismic, infrared, acoustic, magnetic, visual, etc.--to facilitate event identification, recording, and reporting. The device could also monitor NBC sensors, as well as operate within an NBC-contaminated area.

Initially, the system would be stationary but portable, with an antenna on an elevated mast near a sensor field or layout. It can build on sentry robots that are currently available for use in industry. Ultimately, the system would be mobile. Either navigation sensors would provide mobility along predetermined routes or the vehicle would be airborne; the decision should be made as the technology progresses. Also, the mobile system would employ onboard as well as remote sensors.

Functional Requirements

The proposed initial, portable system would require

A fully programmable, computer-operated controller (with transmit/receive capabilities) that would interface with the remote sensors and process the sensor data

to enable automated recognition (object detection, identification, and location). This effort would entail matching the various VHF radio links from existing or developmental remote sensors at a "smart" console to permit integration and interpretation of the data received.

A secure communications link from the controller to a tactical operations center that would permit remote read-out of sensor data upon command from the tactical operations center. This communications link would also provide the tactical operations center the capability of turning the controller (or parts of it) on or off.

Later versions of the system would have the attributes described above, with the additional features of mobility and onboard sensors. In this case, the sentry/surveillance robot would become part of a teleoperated vehicular platform, either traversing a programmed, repetitive route or proceeding in advance of manned systems to provide early warning of an enemy presence.

Benefits

The principal near-term advantages are

to provide a test bed for exploiting AI technology in a surveillance/sentry application, using available sensors adapted to

special algorithms that would minimize false alarms and speed up the process of detection, identification, and location.

to permit a savings in the manpower required for monitoring sensor alarms and interpreting readings, while providing 24-hour-a-day, all-weather coverage.

to provide a capability for operating a surveillance/sentry system under NBC conditions or to warn of the presence of NBC contaminants.

The far-term mobile system would be invaluable in providing surveillance/sentry coverage in the vicinity of critical or sensitive temporary field facilities, such as high-level headquarters or special weapons storage areas.

INTELLIGENT MAINTENANCE, DIAGNOSIS, AND REPAIR SYSTEM

Expert Systems applications in automatic test equipment (ATE) can range from the equipment design stage to work in the field. Expert systems incorporating

structural models of pieces of equipment can be used in equipment design to simplify subsequent trouble shooting and maintenance.

In the field, expert systems can guide the soldier in expedient field repairs. At the depot, expert systems can perform extensive diagnosis, guide repair, and help train new mechanics.

In the diagnostic mode it would instruct the operator not only in the sequence of tests and how to run them, but also in the visual or aural features to look for and their proper sequence.

In the maintenance mode the system would describe the sequence of tests or examinations that should be performed and what to expect at each step.

In the repair mode the system would guide the operator on the correct tools, the precise method of disassembly, the required replacement parts and assemblies by name and identification numbers, and the proper procedure for reassembly. After repair the maintenance mode can be exercised to ensure by appropriate tests that repair has, in fact, been effected without disabling any other necessary function.

In any of the above operations the system would record the repairs, maintenance procedures, or conditions experienced by that piece of equipment. Users would thus have access to essential readiness information without needing bulky, hard-to-maintain maintenance records.

Current Projects and Experience

Some current Army and defense projects concerned with ATE are

VTRONICS, a set of projects for onboard, embedded sensing of vehicular malfunctions with built-in test equipment (BITE);

VIMAD, Voice Interactive Maintenance Aiding Device, which is external to the vehicle;

Hawk missile computer-aided instruction for maintenance and repair.

Electronic malfunctions have been the subject of the most research, and electronics is now the most reliable aspect of the systems. Not much work has been done to reduce mechanical or software malfunctions. During wartime, however, such systems will need to be survivable under fire as well as be reliable under normal conditions.

For ground combat vehicles around 1990, a BITE diagnostic capability to tell the status of the vehicle power train is planned. In one development power train system, the critical information is normally portrayed either by cues via a series of gauges or by a digital readout. Malfunctions can be diagnosed through these cues and displays. The individual is prompted to push buttons to go through a sequence of displays.

An existing Army project concerns a helicopter cockpit display diagnostic system. One purpose of the project was to study audible information versus visual display. For example, the response to the FUEL command is to state the amount of fuel or flying time left; the AMMO command tells the operator how much ammunition is left. One reason for using speech output is that monitoring visual displays distracts attention from flying.

A lot of work has been done in the Army on maintenance and repair training, but computer-assisted instruction (CAI) and artificial intelligence could greatly reduce training time. For example, the M1 tank requires 60,000 pages of technical manuals to describe how to repair breakdowns.

The Army has planned for an AI maintenance tutor that would become a maintenance aid, but it is not yet funded. Under the VIMAD project supported by DARPA, a helmet with a small television receiver optically linked to a cathode ray tube (CRT) screen is being investigated as an aid to maintenance. Computer-generated video disk information is relayed.

An individual working inside the turret of an M1 tank, for example, cannot at present easily flip through the pages of the repair manual. With VIMAD, using a transmitter, receiver, floppy disk, and voice recognition capability, the individual can converse with the system to get information from the data base. The system allows a 19-word vocabulary for each of three individuals. The system has a

100-word capability to access more information from the main system and provides a combination of audio cues and visual prompts.

Any Army diagnostic system should be easily understood by any operator, regardless of maintenance background ("user friendly"). Choosing from alternatives presented in a menu approach, for example, is not necessarily easy for a semiliterate person.

We propose that the following projects be supported as soon as possible:

Interactive, mixed-media manuals for training and repair. Manuals should employ state-of-the-art video disk and display technology. The MIT Arcmac project, supported by the Office of Naval Research, illustrates this approach.

Development of expert systems to trouble-shoot the 50 to 100 most common failures of important pieces of equipment. The system should incorporate simple diagnostic cues, be capable of fixed format (stylized, nonnatural) interaction, and emphasize quick fixes to operational machinery. The project should be oriented toward mechanical devices to complement the substantial array of existing electronic ATE. Projects in this category should be ready for operational use by

1987.

Longer-term development of expert systems for ATE of more complex mechanical and electromechanical equipment. The systems in this category are intended for use at depots near battle lines. They are less oriented to quick fixes and incorporate preventive maintenance with more intelligent trouble shooting. They do not

aim for the sophisticated expertise of a highly qualified technician or mechanic. The emphasis is on (1) determining whether it is feasible to fix this piece of equipment, (2) determining how long it will take to fix, (3) determining if limited resources would be better used to fix other pieces of equipment, and (4) laying out a suitable process for fixing the equipment.

The trouble-shooting systems recommended above rely on human sensors, exactly like MYCIN and Prospector. MYCIN is an expert system for diagnosing and treating infectious diseases that was developed at Stanford University. Prospector, developed at SRI International, is an expert system to aid in exploration for minerals. Parallel, longer-term efforts should be started to incorporate automatic sensors into the trouble-shooting expert systems recommended above.

EXPERT SYSTEMS FOR ARMY MEDICAL APPLICATIONS

Expert systems for various areas of medicine are being extensively studied at a number of institutions in the United States. These include

rule-based systems at Stanford (MYCIN) and Rutgers (for glaucoma) ,

Bayesian statistical systems (for computer-assisted diagnosis of abdominal pain),
cognitive model systems (for internal medicine, nephrology, and cholestasis) ,
knowledge management systems for diagnosis of neurological problems at Maryland.

Current Army activities to apply robotics and artificial intelligence in the medical area are described in the Army Medical Department's AI/Robotics plan, which was prepared with the help of the Academy of Health Sciences, San Antonio. This plan was presented to this committee by the U.S. Army Medical Research and Development Command (AMRDC).

Current Army Activities

Purdue University's Bioengineering Laboratory has an Army contract to study the concept of a "dog-tag chip" that will assist identification of injured personnel. The goal for this device is to assist in the display of patient symptoms for rapid casualty identification and triage. AMRDC noted that visual identification of casualties in chemical and biological warfare may be very difficult because of the heavy duty garb that will be worn.

Airborne or other remote interrogation of the dog-tag chip, its use in self-aid and buddy-aid modes, and use of logic trees on the chip for chemical warfare casualties are being examined by the Army. Other areas of AI and robotics listed in the U.S. AMRDC plan are training, systems for increased realism, and a "smart aideman" expert system, the latter being a "pure" application of expert systems to assist in early diagnosis.

Medical Environments, Functions, and Payoffs Medical environments likely to be encountered in the Army are

routine nonbattle, general illnesses, and disease;

battle injuries, shock/trauma;

epidemics;

chemical;

radiation;

bacteriological.

In a battle area, a medical diagnosis paramedic aide machine would

speed up diagnosis by paramedic and provide productivity increase, noninvasive sensing, and triage;

suggest the best drugs to give for a condition, subject to patient allergies;

suggest priority, disposition, and radio sensor signals on a radio link to field hospital, if necessary to consult physician.

At forward aid stations, in addition to routine diagnostic help, the device might infer patterns of illness on the basis of reports from local areas, track patient condition over time, and teach paramedics the nature of conditions occurring in that particular area that may differ from their prior experience.

Payoffs would include increasing soldiers' likelihood of survival and the consequent boost to morale through the knowledge that efforts

to save them were being assisted by the latest technology. Note that the automated battalion information management system, described below, will involve building a large planning model, which could include medicine.

Recommended Medical Expert Systems

In view of existing technology, a more aggressive dog-tag chip program than that already under way at Purdue University is advocated. The Army should contract with some commercial company currently making wristwatch monitors to develop a demonstration model Army body monitor and not worry if the development gets out into the public domain. Wristwatch monitors of pulse rate, temperatures, etc., are listed in catalogs such as the one from Edmund Scientific.

Technology for low-level digital communication with cryptography is also available. As a prerequisite to the smart dog-tag, the Army may wish to make use of this technology in various Army systems more mundane than the smart dog-tag chip. Cryptography can ensure that information on a smart dog-tag is not susceptible to interception.

Collection of data on noninvasive new and old sensors and related methods of statistical analysis to determine their efficiency in monitoring casualty/injury conditions should be the subject of a longer term study. The study should create

a data base that relates medical diagnosis and sensor capabilities.

The development of AI expert systems aimed at providing computer consulting for nonbattle and battle-area Army medicine and paramedical training are long-term projects that could be undertaken in collaboration with military and university hospitals. For example, the emergency room or shock/trauma unit of a civilian hospital could be used in beginning studies. Correlation of the patient 's current condition with past medical history as recorded on a soldier's dog-tag chip would be one result available from an expert system. Paramedic skills may or may not require a slight increase, depending on how well the AI

aid is designed. It does seem that the same number of paramedics should be able to accomplish more.

FLEXIBLE MATERIAL-HANDLING MODULES

Most robot applications in industry today are directly related to material handling. These include loading and unloading machines, palletizing, feeding parts for other automation equipment, and presenting parts for inspection.

Material handling in Army operations has many similar applications, which, at the very least, involve a great number of repetitive operations and often require working under hazardous conditions. It is proposed to make use of state-of-the-art robotics to develop a

multifunctional, material-handling robotic module that can be readily adapted for many Army functions serving both rear echelon and front line supply needs.

An ammunition resupply robot could select, prepare, acquire, move, load, or unload ammunition at forward weapon sites to reduce exposure of personnel or in rear storage areas to reduce personnel requirements and provide 24-hour capability.

For general use, a robot mounted on a wheeled base is recommended so that the human operator can maneuver the robot into position and then initiate a stored computer program that it will execute without continuous supervision. With present technology constraints on the necessary vision system, it would be necessary to have a bar-code identifying insignia affixed to every package or object in a known position. State-of-the-art

pattern recognition devices can then be mounted on the robot arm to identify an object or package for sorting and verification. Future technological advance would reduce the need for identifying insignia.

The proposed robot to refuel vehicles is actually an instance of a material-handling module. It would be mounted on wheels and equipped with vision. The operator would position the robot in the proximate location, where it would then use a fuel dispenser without exposing the crew. Special gas tank caps would be required to facilitate insertion and dispensing of fuel by the robot.

Functional Requirements

The module would be a fully programmable, servo-driven robot with advanced controller capable of interfacing with a vision module, other sensor modules, and teleoperator control. It would include a teach-box programmer to provide the simplest programming capability by unit-level nonspecialists. The teleoperator would provide the operator with the ability to operate the robot on one-at-a-time tasks that do not require repetitive operations

or are too difficult to program for automatic operation.

The robot module base would be designed to be readily mounted on a truck, a trailer, or a weapons carrier, or emplaced on a rigid pad or even firmly embedded in the ground. It would be desirable to engineer several different sizes with different load capacities but operating with identical controllers.

High speed and precision would be desirable but not mandatory. Trade-offs for ruggedness, simplicity, maintainability, and cost should be considered seriously.

Provision would be made for readily interchangeable end effectors, or "hands." Each application would have a specialized end effector, which could be a gripper or tool. The particular requirements of the task or mission would specify which set of effectors accompany the robot.

Some near-term advantages are

In supply logistics the module could stack such items as packages or ammunition, from either trucks or supply depots, where standard pallet operations are not available or feasible. Many personnel engaged in all forms of moving supplies and

munitions would become acquainted with and adept at the use of this strength-enhancing, labor-saving tool. Reduction of staff and elimination of many repetitive and fatiguing operations would result. Key personnel would be time-shared, since a single operator could set up and supervise several robot systems.

In front line and other hazardous activities, the robot module, after programming, could operate autonomously or under supervisory control from a safe location. Ammunition and fuel resupply for tanks serviced by a robot mounted on a protected vehicle is a typical example. Handling hazardous chemical or nuclear objects or material could be performed remotely. Retrieving and delivering objects under fire may be possible with appropriate remote-controlled vehicles.

When personnel become familiar and experienced with these systems, they will probably generate and jury-rig a robot to perform new operations creatively. This system is meant to be a general-purpose helper.

The long-range advantages include the following:

With the future addition of a wide range of sensors, including vision, tactile, force, and torque, the robot module becomes part of an intelligent robot system, enlarging its field of application to parallel many intended uses of systems in industry. With specialized tools, maintenance, repair, reassembly, testing, and other normal functions to maintain sophisticated weapon systems, all become possible, especially under hazardous conditions.

The proposed module can be readily duplicated at reasonable cost and serve at many experimental sites for evaluation and development into practical tools. It will undoubtedly uncover needs requiring advanced capabilities that can be added without complete redesign.

AUTOMATED BATTALION INFORMATION MANAGEMENT SYSTEM

Combat operations in a modern army require vast amounts of information of varying completeness, timeliness, and accuracy. Included are operational and logistic reports on the status of friendly and enemy forces and their functional capabilities, tactical analyses, weather, terrain, and intelligence input from sensors and from human sources. The information is often

inconsistent and fragmentary but in sufficient quantity to lead to information overload, requiring sorting,

classification, and distribution before it can be used. Getting the information to the appropriate people in a timely fashion and in a usable form is a major problem.

A battalion forward command post is usually staffed by officers having responsibility for operations, intelligence, and fire support. These officers are seconded by enlisted personnel with significantly less schooling and experience. Other battalion staff officers assist, but they do not carry the main burden. The battalion executive officer usually positions himself where he can best support the ongoing operation. Together, these men simultaneously fight the current battle and plan the next operation. Thus, efforts must be made to alleviate fatigue and stress. There is a consequent need for automated decision aids.

Expert systems for combat support could assist greatly. It appears that information sources consist currently of hand-written, repeatedly copied reports and that intelligence operations integration is degraded because of information overload

and because information is inconsistent. Thus, while capable of intuitive judgments that machines do poorly, officers find it difficult to integrate unsorted and unrelated information, are limited in their ability to examine alternatives, and are slow to recognize erroneous information. Decisionmaking in tense situations is spontaneous and potentially erroneous.

Capturing the knowledge of an officer, even in a highly domain-restricted situation such as a forward command post, is difficult. Even though they strain the state of the art, expert systems for combat support have such potential payoff in increasing combat effectiveness that they should receive high priority and be begun immediately. The following sequence of projects can be identified:

how to capture and deploy knowledge and duties of the operations, intelligence, logistics, and fire-support officers into operations, intelligence, logistics, and fire-support expert systems to aid these officers;

how to automate screening messages and establishing priorities to reduce information overload;

how to integrate the operations of the expert systems to support the command;

how to integrate general information with detailed information about the particular situation at hand; for example, how supplemental experts for multisensor reconnaissance and intelligence, topographic mapping, situation mapping, and other functions such as night attack and air assault can be used to adapt the general battalion expert system to the particular battle situation.

5 IMPLEMENTATION OF RECOMMENDED APPLICATIONS

For the applications recommended in Chapter 4, the committee made gross estimates of the time, cost, and technical complexity/risk associated with each. The results of those deliberations are summarized in this chapter.

The matrix on the following pages was developed to present the committee's proposed implementation plan. For each candidate, the matrix shows the estimated time and man-years of effort from initiation of contractual effort until demonstration of the concept by a bread- or brass-board model, gross estimates of costs for a single contractor, projected payoff,

relative technical complexity, remarks, and, finally, recommended priority in which projects should be undertaken. In light of constrained funding and even more strictly limited technical capacity, we recommend that one candidate in each of the three areas--effectors, sensors, and cognition--be undertaken now. The recommended top-priority applications are the automatic loader of ammunition in tanks (effectors), the sentry/surveillance robot (sensors), and the intelligent maintenance, diagnosis, and repair system (cognition).

While the committee agreed that it would be preferable in all cases for at least two firms to undertake R&D simultaneously, it recognized that constrained funding would probably preclude such action. Cost estimates in the matrix, therefore, represent the committee's estimate of the costs of a single contractor based on the number of man years of a fully supported senior engineer. Believing that the Army was in far better position to estimate its administrative, in-house, and testing costs, the committee limited its cost estimates to those of the contractor.

After extensive discussion, the committee chose \$200,000 as a reasonable and representative estimate of the cost of a

fully burdened industrial man-year for a senior engineer. The estimated costs for contractor effort for different supported man-year costs can be calculated. The estimates given are for demonstrators, not for production models.

MEASURES OF EFFECTIVENESS

The committee had considerable difficulty in attempting to develop useful measures of effectiveness because such measures appear to be meaningful only as applied to a specific application. Even then, the benefits of applying robotics and artificial intelligence are often difficult to quantify at this early stage. How, for example, does one measure the value of a human life or of increments in the probability of success in battle?

Therefore, instead of attempting to develop quantitative measures that strain credibility, the committee offers general guidelines against which to measure the worthiness of proposed applications of robotics and artificial intelligence. These guidelines are grouped according to their intended effect.

People

Reduced danger or improved environment

Reduced skill level or training requirements

Improved survivability

Mission

Improved productivity or reduced manpower requirements

Military advantage

New opportunities

Enhanced capability to conduct 24-hour per day operations

Improved RAMS (reliability, availability, maintainability, and supportability)

Material

Reduced cost

The final item, reduced cost, is not the only one that can be assigned a quantitative value. A reduced need for training, for example, should result in reduced training costs. Similarly, improvements in RAMS should reduce life-cycle costs because of diminished need for repair parts, reduced maintenance costs stemming from greater mean time between failure, and reduced maintenance man-hours

per maintenance action. However, meaningful estimates with acceptable levels of confidence would require large volumes of experience data that simply are not available at this early stage in the development of a new and revolutionary technology.

Military advantage is probably the ultimate measure of effectiveness. For example, if it could be shown through modeling or gaming that investment in a system meant the difference between winning or losing, that system could be described as infinitely cost effective.

The committee simply does not have access to sufficient pertinent information to make other than a subjective judgment of the effectiveness of its proposed applications at this time. Further, because each application is to be implemented progressively, such measures will change over time. Finally, because the final versions of the applications require substantial research and development, the committee, despite its collective experience, can provide only the gross estimates of probable costs and payoffs contained in the matrix.

What, then, can the committee say about measuring the effectiveness of the proposed applications? First, that in its collective judgment, the recommended applications provide sound benefits for the Army and second, that these benefits will stem from more than one of the nine areas listed above.

A possible precedent to consider is the manner in which DOD funded the Very High Speed Integrated Circuits (VHSIC) program. It was considered an area of great promise that warranted funding as a matter of highest priority; applications were sought and found later on, after the research was well under way. Similarly, there is little question that we have barely begun to scratch the surface in identifying high-payoff applications of robotics and artificial intelligence technology.

6 OTHER CONSIDERATIONS

In the course of its studies, the committee identified a number of important considerations that can be expected to bear heavily on the Army's decisions on future applications of robotics and AI technology. These considerations, discussed in the paragraphs that follow, apply more

generally than to the specific topics covered in the previous chapters.

SHORTAGE OF EXPERTS

Probably the most important single consideration at this time is that there are far too few research experts in the areas of robotics and artificial intelligence. Most of those available to the Army for their applications are clustered in a few universities where some 70 professors with an average of 4 to 5 (apprentice) students apiece represent the bulk of existing technical expertise. There are appreciably fewer qualified practitioners in military service. As a result, despite the fact that additional funding in these areas is required, it must be allocated with great care to ensure that recipients have the capability to spend the money wisely and effectively. For example, SRI is unable to accept more money for some branches of AI because its technical capacity is already fully committed.

Similarly, there is a critical shortage of military experts in the domains to be captured by expert systems. In particular, it is difficult to find the military officers required to participate in the design and development of complex expert

systems, such as those required for division and corps tactical operations centers.

Both factors underline the need for an Army-university partnership in educating qualified individuals in order to expand the research and development base as soon as possible. They also appear to indicate a need for some sort of centralized coordination, to ensure that optimum use is made of the limited human and fiscal resources available.

The creation of operator-friendly systems is essential to the successful spread of this technology. A truly operator-friendly system will appeal to all levels of people, especially under adverse conditions. In addition, these systems will facilitate the important task of getting novices acquainted with and accustomed to using robots and robotic systems. Not only will this lead to the critically needed confidence that comes from hands-on experience, but it will also demonstrate the reality of what can be done now and point the way toward more advanced applications of the future.

The importance of operator-friendly hardware has been recognized by the

military since World War II, when the studies of aircraft accidents identified a number of pilot errors caused by the design of the plane. Since then, military R&D has included the analysis of human factors in the design of new technologies. Expected benefits include fewer accidents, improved performance, reduced production costs, lower training costs, and improved implementation.

Operator-friendly systems are of particular importance to the military because the objective is to ensure proper use of the systems under less than favorable conditions. In most cases the environmental conditions in which the robot will be expected to operate are more severe than those currently experienced in industrial applications. Furthermore, in times of crisis the robot may need to be operated by or work with personnel that are not fully trained. Careful design of the hardware and software can reduce training, maintenance, and repair costs. It can also ensure that the expected benefits are more likely to be achieved.

In some environments, such as tanks, humans and robots will be working in close quarters. If there is hostility or difficulty with the robotic system, or if

the maneuvers require too much space or movement, the system will not work effectively. In a crisis, there may not be a second chance or an available backup for a system failure, so the man-machine combination must work effectively and quickly.

Essential to any operator-friendly system are high levels of reliability, availability, and maintainability, and redundant fail-safe provisions. With the many hostile environments, it will be of basic importance to assure adequate redundancy in components and systems. What are the backups? What happens when power fails? Can muscle power operate the system?

As military equipment becomes increasingly complex, its operation and maintenance will compete with industry for scarce mechanical and computer skills. This shortage of experts and trained skilled workers can be ameliorated by robotic applications, such as maintenance and repair aids.

The committee is concerned that specific efforts be made to guard against reinventing the wheel. With so many programs in the armed services, it appears to outsiders that many activities are repeated because each particular area wants

its own activity. The Army should have some means of knowing the programs in the other services that could have application to Army needs. The committee has learned that the Joint Laboratory Directors, operating under the aegis of the Joint Logistics Commanders, have begun to address this important need. Any steps that foster communication in this area are to be welcomed.

AVAILABLE TECHNOLOGY

There are already a number of successful applications of robotics in use in industry. Such applications as spot welding, arc welding, palletizing, and spray painting are not exotic and are proven successes. The Army can improve its operations immediately by taking advantage of commercially proven systems for production and maintenance in its depots.

GETTING STARTED

The Army will experience the same growing problems that industry has experienced. Outside of a few areas like robotic spot welding of automobiles and robotic unloading of die casting machines, there has been much talk about robotic applications but only slow growth. There is evidence that implementation of robotics

projects will now move at a much faster pace. The Army should bear in mind, however, that getting a dynamic technological program going almost invariably requires more time and money than its developers originally plan.

These technologies will cause a savings in manpower, though not necessarily for the initial thrust. Experience and training will be needed in all areas--operators, maintenance personnel, supervisors, and managers. Once the new systems are understood by all levels, then the savings will be realized. In many cases this savings will take the form of more output per unit. In addition, the savings will compound as the systems grow with technology additions as well as familiarity.

An important by-product following the initial learning period will be the motivation of individuals. Being master of a phase of new technology gives one an accomplishment and ability that can be the base for growth within the existing employment area or for selling personal ability and knowledge outside the area--in short, a ladder for growth and personal development.

The committee has noted that the Army has identified the five technology thrusts of Very Intelligent Surveillance and Target Acquisition (VISTA),

Distributed Command, Control, Communications and Intelligence, Self-Contained Munitions, Soldier-Machine Interface, Biotechnology.

These are areas to which it intends to devote its research and exploratory development efforts.

Robotics and artificial intelligence technology is not designated as a separate high-priority thrust. It is possible to relate specific robotics/AI applications to one or more of the technology thrusts, as the Army Science Board Ad Hoc Group on Artificial Intelligence and Robotics did in its report. However, the danger remains that robotics and AI efforts--particularly where they do not fall clearly under the mantle of one of the chosen five--will be considered lower priority, with the attendant implications of reduced funding and support. Failure to identify robotics and AI as a special thrust may also contribute to the lack of focus in management and diffusion of effort and funding noted elsewhere in this report.

IMPLEMENTATION DIFFICULTIES

In addition to technical barriers that might normally be expected, several misconceptions have continually clouded industry's technology development and ongoing research in artificial intelligence. Unrealistic expectations combined with problems inherent in any new technology have created barriers to easy implementation. Based on recent industrial experiences, the Army can expect these to include

Unrealistic expectations of the technology's capabilities. In an extremely narrow context, some expert systems outperform humans (e.g., MACSYMA), but certainly no machine exhibits the commonsense facility of humans at this time. Machines cannot outperform humans in a general sense, and that may never be possible. Further, the belief that such systems will bail out current or impending disasters in more conventional system developments that are presently under way is almost always erroneous.

The technology is not readily learned. The notion that "this is nothing more than smart software" continually demonstrates the naiveté of first impressions. Current

experience in industry refutes this contention. A seemingly simple concept of knowledge acquisition,

simply having an expert state his rules of thumb, is currently an intricate art and so complex as to defy automatic techniques. It is, and will remain for some time, a research area.

Expectations often dramatically exceed what is possible. This is particularly true of the times estimated for development. Performance of the systems has often lagged because of such problems as classification restrictions or a lack of available expertise.

Desire for quick success. Very often the political goals are not consonant with the technical goals, thereby increasing the risk associated with developing an expert system by placing unrealistic time constraints on the staff.

University goals versus the goals of industry. Top research universities are motivated to gain new knowledge, develop researchers, publish papers and dissertations, and establish a vehicle for the perpetuation of these. The goals of a responsive industrial unit are to build a system or provide a service that results in

a usable, functioning system in an acceptable time to meet the needs of the customer for use by practitioners. Because of this diversity of purpose, much of the software and hardware developed is not easily transferable, and costly transformations have been required.

Fear of not succeeding. This is as detrimental to technological progress as in any other art or science. Industry and government have often committed funds to unambitious projects that met inadequate risks in order to prove nothing.

Calling it AI when it is not or is only loosely related. The expectation that development in this area will be readily funded encourages jumping on bandwagons.

Lack of credentials. Several people and groups are claiming expertise in AI, though they may not have the rich base upon which research capability is normally developed. Careful credential checking is imperative.

Technology transfer. The preponderance of practitioners are in the universities and have only recently been moving to industry, primarily to venture activities. Most have never delivered products in the industrial context (e.g., documented with life-cycle considerations). The transfer of knowledge

to industry at large is thus rarely done by those with knowledge of both industry and the technology, which makes the industrialization process more risky.

Premature determination of results. The risk exists of unwittingly predetermining the outcome of decisions that should be made

after further research and development. The needed skills simply are not in industry or in the government in the quantities needed to prevent this from happening on occasion.

Nontransferable software tools. Virtually all software knowledge engineering systems and languages are scantily documented and often only supported to the extent possible by the single researcher who originally wrote it. The universities are not in the business to assure proper support of systems for the life-cycle needs of the military and industry, although some of the new AI companies are beginning to support their respective programming environments.

Lack of standards. There are no documentation standards or restrictions on useful programming languages or performance indices to assess system performance.

Mismatch between needed computer resources and existing machinery. The symbolic languages and the programs written are more demanding on conventional machines than appears on the surface or is being advertised by some promoters.

Knowledge acquisition is an art. The successful expert systems developed to date are all examples of handcrafted knowledge. As a result, system performance cannot be specified and the concepts of test, integration, reliability, maintainability, testability, and quality assurance in general are very fuzzy notions at this point in the evaluation of the art. A great deal of work is required to quantify or systematically eliminate such notions.

Formal programs for education and training do not exist. The academic centers that have developed the richest base of research activities award the computer science degree to encompass all sub-disciplines. The lengthy apprenticeship required to train knowledge engineers, who form the bridge between the expert and development of an expert system, has not been formalized.

7 RECOMMENDATIONS

START USING AVAILABLE TECHNOLOGY NOW

Robotics and artificial intelligence technology can be applied in many areas to perform useful, valuable functions for the Army. As noted in Chapter 3, these technologies can enable the Army to

- improve combat capabilities,
- minimize exposure of personnel to hazardous environments,
- increase mission flexibility,
- increase system reliability,
- reduce unit/life cycle costs,
- reduce manpower requirements,
- simplify training.

Despite the fact that robotics technology is being extensively used by industry (almost \$1 billion introduced worldwide in 1982, with increases expected to compound at an annual rate of at least 30 percent for the next 5 to 10 years), the Army does not have any significant robot hardware or software in the field. The Army's needs for the increased efficiency and cost effectiveness of this new technology surely exceed those of industry when one considers the potential reduction in risk and casualties on the battlefield.

The shrinking manpower base resulting from the decline in the 19-to 21-year-old male population, and the substantial costs of maintaining present Army manpower (approximately 29 percent of the total Army budget in FY 1983), emphasize that a major effort should be made to conserve manpower and reduce battlefield casualties by replacing humans with robotic devices.

The potential benefits of robotics and artificial intelligence are clearly great. It is important that the Army begin as soon as possible so as not to fall further behind. Research knowledge and practical industrial experience are accumulating. The Army can and should begin to take advantage of what is available today.

The best way for the Army to take advantage of the potential offered by robotics and AI is to undertake some short-term demonstrators that can be progressively upgraded. The initial demonstrators should meet clear Army needs, be demonstrable within 2 to 3 years,

use the best state of the art technology available,

have sufficient computer capacity for upgrades) form a base for familiarizing Army personnel--from operators to senior

leadership--with these new and revolutionary technologies.

As upgraded, the applications will need to be capable of operating in a hostile environment.

The dual approach of short-term applications with planned upgrades is, in the committee ' s opinion, the key to the Army's successful adoption of this promising new technology in ways that will improve safety, efficiency, and effectiveness. It is through experience with relatively simple applications that Army personnel will become comfortable with and appreciate the benefits of these new technologies. There are indeed current Army needs that can be met by available robotics and AI technology.

In the Army, as in industry, there is a danger of much talk and little concrete action. We recommend that the Army move quickly to concentrate in a few identified areas and establish those as a base for growth.

SPECIFIC RECOMMENDED APPLICATIONS

The committee recommends that, at a minimum, the Army should fund the three

demonstrator programs described in Chapter 4 at the levels described in Chapter 5:

The Automatic Loader of Ammunition in Tanks, using a robotic arm to replace the human loader of ammunition in a tank. We recommend that two contractors work simultaneously for 2 to 2 1/2 years at a total cost of \$4 to \$5 million per contractor.

The Surveillance/Sentry Robot, a portable, possibly mobile platform to detect and identify movement of troops. Funded at \$5 million for 2 to 3 years, the robot should be able to include two or more sensor modalities.

The Intelligent Maintenance, Diagnosis, and Repair System, in its initial form (\$1 million over 2 years), will be an interactive trainer. Within 3 years, for an additional \$5 million, the system should be expanded to diagnose and suggest repairs for common break-downs, recommend whether or not to repair, and record the repair history of a piece of equipment.

If additional funds are available, the other projects described in Chapter 4, the medical expert system, the flexible material-handling modules, and the

battalion information management system, are also well worth doing.

VISIBILITY AND COORDINATION OF MILITARY AI/ROBOTICS

Much additional creative work in this area is needed. The committee recommends that the Army provide increased funding for coherent research and exploratory development efforts (lines 6.1 and 6.2 of the budget) and include artificial intelligence and robotics as a special technology thrust.

The Army should aggressively take the lead in pursuing early application of robotics and AI technologies to solve compelling battlefield needs. To assist in coordinating efforts and preventing duplication, it may wish to establish a high-level review board or advisory board for the AI/Robotics program. This body would include representatives from the universities and industry, as well as from the Army, Navy, Air Force, and DARPA. We recommend that the Army consider this idea further.

APPENDIX

STATE OF THE ART AND PREDICTIONS FOR

ARTIFICIAL INTELLIGENCE AND ROBOTICS

INDUSTRIAL ROBOTS: FUNDAMENTAL CONCEPTS

The term robot conjures up a vision of a mechanical man--that is, some android as viewed in Star Wars or other science fiction movies. Industrial robots have no resemblance to these Star Wars figures. In reality, robots are largely constrained and defined by what we have so far managed to do with them.

In the last decade the industrial robot (IR) has developed from concept to reality, and robots are now used in factories throughout the world. In lay terms, the industrial robot would be called a mechanical arm. This definition, however, includes almost all factory automation devices that have a moving lever. The Robot Institute of America (RIA) has adopted the following working definition:

A robot is a programmable multifunction device designed to move material, parts, tools, or specialized devices through variable programmed motions for the performance of a variety of tasks.

It is generally agreed that the three main components of an industrial robot are the

mechanical manipulator, the actuation mechanism, and the controller.

The mechanical manipulator of an IR is made up of a set of axes (either rotary or slide) , typically three to six axes per IR. The first three axes determine the work envelope of the IR, while the last

three deal with the wrist of the IR and the ability to orient the hand. Figure 1 shows the four basic IR configurations. Although these are typical of robot configurations in use today, there are no hard and fast rules that impose these constraints. Many robots are more

The appendix is largely the work of Roger Nagel, Director, Institute for Robotics, Lehigh University. James Albus of the National Bureau of Standards and committee members J. Michael Brady, Stephen Dubowsky, Margaret Eastwood, David Grossman, Laveen Kanal, and Wendy Lehnert also contributed.

restricted in their motions than the six-axis robot. Conversely, robots are sometimes mounted on extra axes such as an x-y table or track to provide an additional one or two axes.

It is important to note at this point that the "hand" of the robot, which is typically

a gripper or tool specifically designed for one or more applications, is not a part of a general purpose IR. Hands, or end effectors, are special purpose devices attached to the "wrist" of an IR.

The actuation mechanism of an IR is typically either hydraulic, pneumatic, or electric. More important distinctions in capability are based on the ability to employ servo mechanisms, which use feedback control to correct mechanical position, as opposed to nonservo open-loop actuation systems. Surprisingly, nonservo open-loop industrial robots perform many seemingly complex tasks in today's factories.

The controller is the device that stores the IR program and, by communications with the actuation mechanism, controls the IR motions. Controllers have undergone extensive evolution as robots have been introduced to the factory floor. The changes have been in the method of programming (human interface) and in the complexity of the programs allowed. In the last three years the trend to computer control (as opposed to plug board and special-purpose devices) has resulted in computer controls on virtually all industrial robots.

The method of programming industrial robots has, in the most popular and prevailing usage, not included the use of a language. Languages for robots have, however, long been a research issue and are now appearing in the commercial offerings for industrial robots. We review first the two prevailing programming methods.

Programming by the lead-through method is accomplished by a person manipulating a well-counterbalanced robot (or surrogate) through the desired path in space. The program is recorded by the controller, which samples the location of each of the robot's axes several times per second. This method of programming records a continuous path through the work envelope and is most often used for spray painting operations. One major difficulty is the awkwardness of editing these programs to make any necessary changes or corrections.

An additional--and perhaps the most serious--difficulty with the lead-through method is the inability to teach conditional commands, especially those that compute a sensory value. Generally, the control structure is very rudimentary and does not offer the programmer much flexibility. Thus, mistakes or changes usually require completely reprogramming

the task, rather than making small changes to an existing program.

Programming by the teach-box method employs a special device that allows the programmer/operator to use buttons, toggle switches, or a joy stick to move the robot in its work envelope. Primitive teach boxes allow for the control only in terms of the basic axis motions of the robot, while more advanced teach boxes provide for the use of Cartesian and other coordinate systems.

The program generated by a teach box is an ordered set of points in the workspace of the robot. Each recorded point specifies the location of every axis of the robot, thus providing both position and orientation.-

. The controller allows the programmer to specify the need to signal or wait for a signal at each point. The signal, typically a binary value, is used to sequence the action of the IR with another device in its environment. Most controllers also now allow the specification of velocity/acceleration between points of the program and indication of whether the point is to be passed through or is a destination for stopping the robot.

Although computer language facilities are not provided with most industrial robots, there is now the limited use of a subroutine library in which the routines are written by the vendor and sold as options to the user. For example, we now see palletizing, where the robot can follow a set of indices to load or unload pallets.

Limited use of simple sensors (binary valued) is provided by preprogrammed search routines that allow the robot to stop a move based on a sensor trip.

Typical advanced industrial robots have a computer control with a keyboard and screen as well as the teach box, although most do not support programming languages. They do permit subdivision of the robot program (sequence of points) into branches. This provides for limited creation of subroutines and is used for error conditions and to store programs for more than one task.

The ability to specify a relocatable branch has provided the limited ability to use sensors and to create primitive programs.

Many industrial robots now permit down-loading of their programs (and up-loading) over RS232 communication links to other computers. This facility is essential to

the creation of flexible manufacturing system (FMS) cells composed of robots and other programmable devices. More difficult than communication of whole programs is communication of parts of a program or locations in the workspace. Current IR controller support of this is at best rudimentary. Yet the ability to communicate such information to a robot during the execution of its program is essential to the creation of adaptive behavior in industrial robots.

Some pioneering work in the area was done at McDonnell Douglas, supported by the Air Force Integrated Computer-Aided Manufacturing (ICAM) program. In that effort a Cincinnati Milacron robot was made part of an adaptive cell. One of the major difficulties was the awkwardness of communicating goal points to the robot. The solution lies not in achieving a technical breakthrough, but rather in understanding and standardizing the interface requirements. These issues and others were covered at a National Bureau of Standards (NBS) workshop in January 1980 and again in September 1982 [1].

Programming languages for industrial robots have long been a research issue. During the last two years, several robots with an off-

line programming language have appeared in the market. Two factors have greatly influenced the development of these languages.

The first is the perceived need to hold a Ph.D., or at least be a trained computer scientist, to use a programming language. This is by no means true, and the advent of the personal computer, as well as the invasion of computers into many unrelated fields, is encouraging. Nonetheless, the fear of computers and of programming them continues.

Because robots operate on factory floors, some feel programming languages must be avoided. Again, this is not necessary, as experience with user-friendly systems has shown.

The second factor is the desire to have industrial robots perform complex tasks and exhibit adaptive behavior. When the motions to be performed by the robot must follow complex geometrical paths, as in welding or assembly, it is generally agreed that a language is necessary. Similarly, a cursory look at the person who performs such tasks reveals the high reliance on sensory information. Thus a language is needed both for complex motions and for sensory

interaction. This dual need further complicates the language requirements because the community does not yet have enough experience in the use of complex (more than binary) sensors.

These two factors influenced the early robot languages to use a combination of language statements and teach box for developing robot programs. That is, one defines important points in the workspace via the teach-box method and then instructs the robot with language statements controlling interpolation between points and speed. This capability coupled with access to on-line storage and simple sensor (binary) control characterizes the VAL language. VAL, developed by Unimation for the Puma robot, was the first commercially available language. Several similar languages are now available, but each has deficiencies. They are not languages in the classical computer science sense, but they do begin to bridge the gap. In particular they do not have the the capability to do arithmetic on location in the workplace, and they do not support computer communication.

A second-generation language capability has appeared in the offering of RAIL and AML by Automatix and IBM, respectively. These

resemble the standard structured computer language. RAIL is PASCAL-based, and AML is a new structured language. They contain statements for control of the manipulator and provide the ability to extend the language in a hierarchical fashion. See, for example, the description of a research version of AML in [2].

In a very real sense these languages present the first opportunity to build intelligent robots. That is, they (and others with similar form) offer the necessary building blocks in terms of controller language. The potential for language specification has not yet been realized in the present commercial offerings, which suffer from some temporary implementation-dependent limitations.

Before going on to the topic of intelligent robot systems, we discuss in the next section the current research areas in robotics.

RESEARCH ISSUES IN INDUSTRIAL ROBOTS

As described previously, robots found in industry have mechanical manipulators, actuation mechanisms, and control systems. Research interest raises such potential topics as locomotion, dexterous hands, sensor systems, languages, data bases, and

artificial intelligence. Although there are clearly relationships amongst these and other

research topics, we will subdivide the research issues into three categories: mechanical systems, sensor systems, and control systems.

In the sections that follow we cover manipulation design, actuation systems, end effectors, and locomotion under the general heading of mechanical systems. We will then review sensor systems as applied to robots-vision, touch, ranging, etc. Finally, we will discuss robot control systems from the simple to the complex, covering languages, communication, data bases, and operating systems. Although the issue of intelligent behavior will be discussed in this section, we reserve for the final section the discussion of the future of truly intelligent robot systems. For a review of research issues with in-depth articles on these subjects see Birk and Kelley [3].

Mechanical Systems

The design of the IR has tended to evolve in an ad hoc fashion. Thus, commercially available industrial robots have a repeatability that ranges up to 0.050 in., but little, if any, information is

available about their performance under load or about variations within the work envelope.

Mechanical designers have begun to work on industrial robots. Major research institutes are now working on the kinematics of design, models of dynamic behavior, and alternative design structures. Beyond the study of models and design structure are efforts on direct drive motors, pneumatic servo mechanisms, and the use of tendon arms and hands. These efforts are leading to highly accurate new robot arms. Much of this work in the United States is being done at university laboratories, including those at the Massachusetts Institute of Technology (MIT), Carnegie-Mellon University (CMU), Stanford University, and the University of Utah.

Furthermore, increased accuracy may not always be needed. Thus, compliance in robot joints, programming to apply force (rather than go to a position), and the dynamics of links and joints are also now actively under investigation at Draper Laboratories, the University of Florida, the Jet Propulsion Laboratory (JPL), MIT, and others.

The implications of this research for future industrial robots are that we will have access to models that predict behavior under load (therefore allowing for correction), and we will see new and more stable designs using recursive dynamics to allow speed. The use of robots to apply force and torque or to deal with tools that do so will be possible. Finally, greater accuracy and compliance where desired will be available [4-8].

The method of actuation, design of actuation, and servo systems are of course related to the design and performance dynamics discussed above. However some significant work on new actuation systems at Carnegie-Mellon University, MIT, and elsewhere promises to provide direct drive motors, servo-control pneumatic systems, and other advantages in power systems.

The end effector of the robot has also been a subject of intensive research. Two fundamental objectives--developing quick-change hands

and developing general-purpose hands--seek to alleviate the constraints on dexterity at the end of a robot arm.

As described earlier, common practice is to design a new end effector for each

application. As robots are used in more complex tasks (assembly, for example), the need to handle a variety of parts and tools is unavoidable. For a good discussion of current end-effector technology, see Toepperwein et al. [9].

The quick-change hand is one that the robot can rapidly change itself, thus permitting it to handle a variety of objects. A major impediment to progress in this area is a lack of a standard method of attaching the hand to the arm. This method must provide not only the physical attachment but also the means of transmitting power and control to the hand. If standards were defined, quick-change mechanisms and a family of hand grippers and robot tools would rapidly become available.

The development of a dexterous hand is still a research issue. Many laboratories in this country and abroad are working on three-fingered hands and other configurations. In many cases the individual fingers are themselves jointed manipulators. In the design of a dexterous hand, development of sensors to provide a sense of touch is a prerequisite. Thus, with sensory perception, a dexterous hand becomes the problem of designing three

robots (one for each of three fingers) that require coordinated control.

The control technology to use the sensory data, provide coordinated motion, and avoid collision is beyond the state of the art. We will review the sensor and control issues in later sections. The design of dexterous hands is being actively worked on at Stanford, MIT, Rhode Island University, the University of Florida, and other places in the United States. Clearly, not all are attacking the most general problem [10, 11], but by innovation and cooperation with other related fields (such as prosthetics), substantial progress will be made in the near future.

The concept of robot locomotion received much early attention. Current robots are frequently mounted on linear tracks and sometimes have the ability to move in a plane, such as on an overhead gantry. However, these extra degrees of freedom are treated as one or two additional axes, and none of the navigation or obstacle avoidance problems are addressed.

Early researchers built prototype wheeled and legged (walking) robots. The work originated at General Electric, Stanford, and JPL has now expanded, and projects are

under way at Tokyo Institute of Technology, Tokyo University. Researchers at Ohio State, Rensselaer Polytechnic Institute (RPI), and CMU are also now working on wheeled, legged, and in one case single leg locomotion. Perhaps because of the need to deal with the navigational issues in control and the stability problems of a walking robot, progress in this area is expected to be slow [12].

In a recent development, Odetics, a small California-based firm, announced a six-legged robot at a press conference in March 1983. According to the press release, this robot, called a "functionoid," can lift several times its own weight and is stable when standing on

only three of its legs. Its legs can be used as arms, and the device can walk over obstacles. Odetics scientists claim to have solved the mathematics of walking, and the functionoid does not use sensors. It is not clear from the press release to what extent the Odetics work is a scientific breakthrough, but further investigation is clearly warranted.

The advent of the wire-guided vehicle (and the painted stripe variety) offers an interesting middle ground between the

completely constrained and unconstrained locomotion problems. Wire-guided vehicles or robot carts are now appearing in factories across the world and are especially popular in Europe. These carts, first introduced for transportation of pallets, are now being configured to manipulate and transport material and tools. They are also found delivering mail in an increasing number of offices. The carts have onboard microprocessors and can communicate with a central control computer at predetermined communication centers located along the factory or office floor.

The major navigational problems are avoided by the use of the wire network, which forms a "freeway" on the factory floor. The freeway is a priori free of permanent obstacles. The carts use a bumper sensor (limit switch) to avoid collisions with temporary obstacles, and the central computer provides routing to avoid traffic jams with other carts.

While carts currently perform simple manipulation (compared to that performed by industrial robots), many vendors are investigating the possibility of robots mounted on carts. Although this appears at first glance to present additional accuracy problems (precise self-positioning of carts

is still not available), the use of cart location fixturing devices at stations may be possible.

Sensor Systems

The robot without sensors goes through a path in its workspace without regard for any feedback other than that of its joint resolvers. This imposes severe limitations on the tasks it can undertake and makes the cost of fixturing (precisely locating things it is to manipulate) very high. Thus there is great interest in the use of sensors for robots. The phrase most often used is "adaptive behavior," meaning that the robot using sensors will be able to deal properly with changes in its environment.

Of the five human senses--vision, touch, hearing, smell, and taste--vision and touch have received the most attention. Although the Defense Advanced Research Projects Agency (DARPA) has sponsored work in speech understanding, this work has not been applied extensively to robotics. The senses of smell and taste have been virtually ignored in robot research.

Despite great interest in using sensors, most robotics research lies in the domain of the sensor physics and data reduction to

meaningful information, leaving the intelligent use of sensory data to

the artificial intelligence (AI) investigators. We will therefore cover sensors in this chapter and discuss the AI implications later.

Vision Sensors

The use of vision sensors has sparked the most interest by far and is the most active research area. Several robot vision systems, in fact, are on the market today. Tasks for such systems are listed below in order of increasing complexity:

their

identification (or verification) of objects
stable states they are in,

location of objects and their orientation,
simple inspection tasks (is part complete?
visual servoing (guidance), navigation and
scene analysis, complex inspection.

or of which of cracked?) ,

The commercial systems currently available can handle subsets of the first three tasks. They function by digitizing an image from a video camera and then thresholding the digitized image. Based on techniques

invented at SRI and variations thereof, the systems measure a set of features on known objects during a training session. When shown an unknown object, they then measure the same feature set and calculate feature distance to identify the object.

Objects with more than one stable state are trained and labeled separately. Individual feature values or pairs of values are used for orientation and inspection decisions.

While these systems have been successful, there are many limitations because of the use of binary images and feature sets--for example, the inability to deal with overlapped objects. Nevertheless, in the constrained environment of a factory, these systems are valuable tools. For a description of the SRI vision system see Gleason and Aguin [13]; for a variant see Lavin and Lieberman [14].

Not all commercial vision systems use the SRI approach, but most are limited to binary images because the data in a binary image can be reduced to run length code. This reduction is important because of the need for the robot to use visual data in real time (fractions of a second). Although one can postulate situations in which more time is available, the usefulness of vision

increases as its speed of availability increases.

Gray-scale image operations are being developed that will overcome the speed problems associated with nonbinary vision. Many vision algorithms lend themselves to parallel computation because the same calculation is made in many different areas of the image. Such parallel computations have been introduced on chips by MIT, Hughes, Westinghouse, and others.

Visual servoing is the process of guiding the robot by the use of visual data. The National Bureau of Standards (NBS) has developed a special vision and control system for this purpose. If robots are ever to be truly intelligent, they must be capable of visual guidance. Clearly the speed requirements are very significant.

Vision systems that locate objects in three-dimensional space can do so in several ways. Either structured light and triangulation or stereo vision can be used to simulate the human system. Structured light systems use a shaped (structured) light source and a camera at a fixed angle [15]. Some researchers have also used laser range-finding devices to make an image whose picture elements (pixels) are

distances along a known direction. All these methods--stereo vision, structured light, laser range-finding, and others--are used in laboratories for robot guidance.

Some three-dimensional systems are now commercially available. Robot Vision Inc. (formerly Solid Photography), for example, has a commercial product for robot guidance on the market. Limited versions of these approaches and others are being developed for use in robot arc welding and other applications [16].

Special-purpose vision systems have been developed to solve particular problems. Many of the special-purpose systems are designed to simplify the problem and gain speed by attacking a restricted domain of applicability. For example, General Motors has used a version of structured light for accumulating an image with a line scan camera in its Consight system. Rhode Island University has concentrated on the bin picking problem. SRI, Automatix, and others are working on vision for arc welding.

Others such as MIT, University of Maryland, Bell Laboratories, JPL, RPI, and Stanford are concentrating on the special requirements of robot vision systems. They are developing algorithms and chips to

achieve faster and cheaper vision computation. There is evidence that they are succeeding. Special-purpose hardware using very large-scale integration (VLSI) techniques is now in the laboratories. One can, we believe, expect vision chips that will release robot vision from the binary and special-purpose world in the near future.

Research in vision, independent of robots, is a well-established field. That literature is too vast to cover here beyond a few general remarks and issues. The reader is referred to the literature on image processing, image understanding, pattern recognition, and image analysis.

Vision research is not limited to binary images but also deals with gray-scale, color, and other multispectral images. In fact, the word "image" is used to avoid the limitation to visual spectra. If we

avoid the compression, transmission, and other representation issues, then we can classify vision research as follows:

Low-level vision involves extracting feature measurements from images. It is called low-level because the operations are not knowledge based. Typical operations are

edge detection, threshold selection, and the measurement of various shapes and other features. These are the operations now being reduced to hardware.

High-level vision is concerned with combining knowledge about objects (shape, size, relationships), expectations about the image (what might be in it), and the purpose of the processing (identifying objects, detecting changes) to aid in interpreting the image. This high-level information interacts with and helps guide processing. For example, it can suggest where to look for an object and what features to look for.

While research in vision is maturing, much remains to be investigated. Current topics include the speed of algorithms, parallel processing, coarse/fine techniques, incomplete data, and a variety of other extensions to the field. In addition, work is also now addressing such AI questions as

representing knowledge about objects, particularly shape and spatial relationships;

developing methods for reasoning about spatial relationships among objects;

understanding the interaction between low-level information and high-level knowledge and expectations;

interpreting stereo images, e.g., for range and motion;

understanding the interaction between an image and other information about the scene, e.g., written descriptions.

Vision research is related to results in VLSI and Ar. While there is much activity, it is difficult to predict specific results that can be expected.

Tactile Sensing

Despite great interest in the use of tactile sensing, the state of the art is relatively primitive. Systems on industrial robots today are limited to detecting contact of the robot and an object by varying versions of the limit-switch concept, or they measure some combination of force and torque vectors that the hand or fingers exert on an object.

While varying versions of the limit-switch concept have been used, the most advanced force/torque sensors for robots have been developed at Draper Laboratories. The remote center of compliance (RCC) developed

at Draper Laboratories, which allows passive compliance in the robots' behavior during assembly, has been commercialized by Astek and Lord Kinematics. Draper has in the last few years instrumented the RCC to provide active feedback to the robot. The instrumented remote center compliance (IRCC) represents the state of the art in wrist sensors. It allows robot programs to follow contours, perform:

insertions, and incorporate rudimentary touch programming into the control system [17].

IBM and others have begun to put force sensors in the fingers of a robot. With x,y,z strain gauges in each of the fingers, the robot with servoed fingers can now perform simple touch-sensitive tasks. Hitachi has developed a hand using metal contact detectors and pressure-sensitive conductive rubber that can feel for objects and

recognize form. Thus, primitive technology can be applied for useful tasks. However, most of the sophisticated and complex tactile sensors are in laboratory development.

The subject of touch-sensor technology, including a review of research, relevance

for robots, work in the laboratory, and predictions of future results, is covered in a survey article by Leon Harmon [18] of Case Western Reserve University Much of that excellent article is summarized below, and we refer the reader to it for a detailed review.

The general needs for sensing in manipulator control are proximity) touch/slip, and force/torque. The following remarks are taken from a discussion on "smart sensors" by Bejcsy [19]:

specific manipulation-related key events are not contained in visual data at all, or can only be obtained from visual data sources indirectly and incompletely and at high cost. These key events are the contact or near-contact events including the dynamics of interaction between the mechanical hand and objects.

The non-visual information is related to controlling the physical interaction, contact or near-contact of the mechanical hand with the environment. This information provides a combination of geometric and dynamic reference data for the control of terminal positioning/orientation and dynamic accommodation/compliance of the mechanical hand.

Although existing industrial robots manage to sense position, proximity, contact, force, and slip with rather primitive techniques, all of these variables plus shape recognition have received extensive attention in research and development laboratories. In some of these areas a new generation of sophistication is beginning to emerge.

Tactile-sensing requirements are not well known, either theoretically or empirically. Most prior wrist, hand, and finger sensors have been simple position and force-feedback indicators. Finger sensors have barely emerged from the level of microswitch limit switches and push-rod axial travel measurement. Moreover, the relevant technologies are themselves relatively new. For example, force and torque sensing dates back only to 1972, touch/slip are dated to 1966, and proximity sensing is only about 9 years old. We do know that force and pressure sensing are vital elements in touch, though to date, as we have seen, industrial robots employ only simple force feedback. Nevertheless, unless considerable gripper overpressure can be tolerated, slip sensing is essential to proper performance in many manipulation tasks. Information about contact areas, pressure distributions, and their changes

over time are needed in order to achieve the most complete and useful tactile sensing.

In contacting, grasping, and manipulating objects, adjustments to gripping forces are required in order to avoid slip and to avoid possibly dangerous forces to both the hand and the workpiece. Besides the need for slip-sensing transducers, there is the requirement that the robot be able to determine at each instant the necessary minimum new force adjustments to prevent slip.

Transducers As of about 1971 the only devices available for tactile sensing were microswitches, pneumatic jets, and (binary) pressure-sensitive pads. These devices served principally as limit switches and provided few means or none for detecting shape, texture, or compliance. Still, such crude devices are used currently.

In the early 1970s the search was already under way for shape detection and for "artificial skin" that could yield tactile information of complexity comparable to the human sense of touch. An obvious methodology for obtaining a continuous measurement of force is potentiometer

response to a linear (e.g., spring-loaded rod) displacement. Early sensors in many laboratories used such sensors, and they are still in use today.

Current research lies in the following areas:

conductive materials and arrays produced with conductive rubbers and polymers;

semiconductor sensors, such as piezo-electrics;

electromagnetic, hydraulic, optical, and capacitive sensors.

Outstanding Problems and New Opportunities

The two main areas most in need of development are (1) improved tactile sensors and (2) improved integration of touch feedback signals with the effector control system in response to the task-command structure. Sensory feedback problems underlie both areas. More effective comprehensive sensors (device R&D) and the sophisticated interpretation of the sense signals by control structures (system R&D) are needed.

Sensitive, dexterous hands are the greatest challenge for manipulators, just as sensitive, adaptable feet are the greatest

challenge for legged locomotion vehicles. Each application area has its own detailed special problems to solve; for example, the design approach for muddy-water object recovery and for delicate handling of unspecified objects in an unstructured environment differ vastly.

Emergent Technology One of the newest developments in touch-sensing technology is that of reticular (Cartesian) arrays using solid-state transduction and attached microcomputer elements that compute three-dimensional shapes. The approach is typified by the research of Marc Raibert, now at CMU, done while he was at JPL [20]. Raibert's device is compact and has high resolution; hence, the fingertip is a self-contained "smart finger." See also the work of Hillis at MIT in this area [21]. This is a quantum jump ahead of prior methods, for example, where small arrays of touch sensors use passive substrates and materials such as conductive elastomers. Resolution in such devices has been quite low, and hysteresis a problem.

Sound Sensors

Many researchers are interested in the use of voice recognition sensors for command and control of robot systems. However, we

leave out voice systems and review here the use of sound as a sensing mechanism.

In this context, sound systems are used as a method for measuring distance. The Polaroid sonic sensor has been used at NBS and elsewhere as a safety sensor. Sensors mounted on the robot detect intrusions into either the workspace or, more particularly, the path of the robot.

Researchers at Pennsylvania State University have developed a spark gap system that uses multiple microphones to determine the position of the manipulator for calibration purposes.

Several researchers at Carnegie-Mellon University and other locations are working on ultrasonic sensors to be used in the arc welding process.

Control Systems

The underlying research issue in control systems for robots is to broaden the scope of the robot. As the sophistication of the manipulator and its actuation mechanism increases, new demands are made on the control system. The advent of dexterous or smart hands, locomotion, sensors, and new complex tasks all extend the controller capability.

The desires for user-friendly systems, for less user training, and for adaptive behavior further push the robot controller into the world of artificial intelligence. Before discussing intelligent robot systems, we describe some of the issues of computer-controlled robots.

Hierarchical Control/Distributed Computing

Almost all controller research is directed at hierarchies in robot control systems. At the National Bureau of Standards, pioneering research has developed two hierarchies--one for control information and one for sensory data. Integrated at each level, the two hierarchies use the task decomposition approach. That is, commands at each level are broken down into subcommands at the lower level until they represent joint control at the lowest level. In a similar fashion, raw vision data are at the lowest level, with higher levels representing image primitives, then features, and finally objects [22].

The levels-of-control issue rapidly leads to an interest in distributed computing in order to balance the computing needs and meet the requirements for real-time performance. The use of smart hand or complex sensor systems, such as vision,

also mandates distributed computing--again, in order not to overload the control computer and degrade the real-time nature of the robot's behavior.

Distributed computing for robot control systems has taken two paths so far. Automatix, NBS, and others use multiple CPUs from the

same vendor (Intel or Motorola) and perform processor communication in the architecture of the base system.

Others have used nonhomogeneous computer systems. They have had to pay a price in the need to define and build protocols and work within awkward constraints. Examples of this are found in the development of MCL by McDonnell Douglas and in a variety of other firms that have linked vision systems with robots. For a case study of one attempt see Nagel et al. [23].

Major impediments to progress in these areas are the lack of standards for the interfaces needed, the need for advances in distributed computing, and the need for a better definition of the information that must flow. Related research that is not covered here is the work on local area networks.

Data Bases

There is a great interest in robot access to the data bases of CAD/CAM systems. As robot programming moves from the domain of the teach box to that of a language, several new demands for data arise. For example, the programmer needs access to the geometry and physical properties of the parts to be manipulated. In addition, he needs similar data with respect to the machine tools, fixtures, and the robot itself. One possible source for this is the data already captured in CAD/CAM data bases. One can assume that complete geometrical and functional information for the robot itself, the things the robot must manipulate, and the things in its environment are contained in these data bases.

As robot programming evolves, an interest has developed in computer-aided robot programming (CARP) done at interactive graphics terminals. In such a modality the robot motions in manipulating parts would be done in a fashion similar to that used for graphic numerical control programming. Such experiments are under way, and early demonstrations have been shown by Automatix and GCA Corporation.

Furthermore, it is now reasonable to assume the desire to have robots report to shop floor control systems, take orders from cell controllers, and update process planning inventory control systems and the variety of factory control, management, and planning systems now in place or under development. Thus, robot controllers must access other data bases and communicate with other factory systems.

Research on the link to CAD/CAM systems and the other issues above is under way at NBS and other research facilities, but major efforts are needed to achieve results.

Robot Programming Environment

As mentioned earlier, second-generation languages are now available. While the community as a whole does not yet have sufficient experience with them to choose standards, more are clearly needed.

Programming advanced robot systems with current languages is reminiscent of programming main-frame computers in assembly language before the advent of operating systems. It is particularly a problem in the use of even the simplest sensor (binary) mechanisms. What are needed are robot operating systems, which would do for robot users what operating systems do

for computer users in such areas as input/output and graphics.

To clarify, we define an explicit language as one in which the commands correspond with the underlying machine (in this case a robot/ computer pair). We further define an implicit language as one in which the commands correspond with the task; that is, for an assembly task an insert command would be implied. Use of an implicit language is complicated by the fact that robots perform families of tasks. A robot operating system would be a major step toward implicit languages.

It is far easier to suggest the work above than to write a definition of requirements. Thus, fundamental research is needed in this area. The Autopass system developed at IBM is probably the most relevant accomplishment to date.

The concepts of graphic robot programming and simulation are exciting research issues. The desire for computer-assisted robot programming (CARP) stems from the data base arguments of before and the belief that graphics is a good mechanism for describing motion. These expectations are widely held, and Computervision, Automatix, and other organizations are

conducting some research. However, no major efforts appear in the current literature.

Graphic simulation, on the other hand, is now a major topic. Work in this area is motivated by the advent of offline programming languages and the need for fail-safe debugging languages, but other benefits arise in robot cell layout, training mechanisms, and the ability to let the robot stay in production while new programs are developed.

Work on robot simulation is hampered by the lack of standards for the language but is in process at IBM for AML, at McDonnell Douglas for MCL, and at many universities for VAL and is expected to be a commercial product shortly. It is worth noting that simulation of sensor-based robots requires simulation of sensor physics. With the exception of some work at IBM, we are unaware of any efforts in sophisticated simulation.

The use of multiple arms in coordinated (as opposed to sequenced) motion raises the issue of multitasking, collision avoidance, and a variety of programming methodology questions. General Electric, Olivetti, Westinghouse, IBM, and others are pursuing multiarm assembly. However these issues

require more attention, even in research that is well under way.

It should be clear by now that robot control has become a complex issue. Controllers dealing with manipulator motion, feedback, complex sensors, data bases, hierarchical control, operating systems, and multitasking must turn to the AI area for further development. In the following section we review briefly the AI field, and in the final section we discuss both robotics and AI issues and the need for expansion of the unified research issues.

ARTIFICIAL INTELLIGENCE

The term artificial intelligence is defined in two ways: the first defines the field, and the second describes some of its functions.

1. "Artificial intelligence research is the part of computer science that is concerned with the symbol-manipulation processes that produce intelligent action. By 'intelligent action' is meant an act of decision that is goal-oriented, arrived at by an understandable chain of symbolic analysis and reasoning steps, and is one in which knowledge of the world informs and guides the reasoning" [24].

2. Artificial intelligence is a set of advanced computer software applicable to classes of nondeterministic problems such as natural language understanding, image understanding, expert systems, knowledge acquisition and representation, heuristic search, deductive reasoning, and planning.

If one were to give a name suggestive of the processes involved in all of the above, knowledge engineering would be the most appropriate; that is, one carries out knowledge engineering to exhibit intelligent behavior by the computer. For general information on artificial intelligence see references 25-34.

Background

The number of researchers in artificial intelligence is rapidly expanding with the increasing number of applications and potential applications of the technology. This growth is occurring not only in the United States, but worldwide, particularly in Europe and Japan.

Basic research is going on primarily at universities and some research institutes. Originally, the primary research sites were MIT, CMU, Stanford, SRI, and the University of Edinburgh. Now, most major

universities include artificial intelligence in the computer science curriculum.

Much of the material in this section summarizes the material in Brown et al. [24].

An increasing number of other organizations either have or are establishing research laboratories for artificial intelligence. Some of them are conducting basic research; others are primarily interested in applications. These organizations include Xerox, Hewlett-Packard, Schlumberger-Fairchild, Hughes, Rand, Perceptronics, Unilever, Philips, Toshiba, and Hamamatsu.

Also emerging are companies that are developing artificial intelligence products. U.S. companies include Teknowledge, Cognitive Systems, Intelligenetics, Artificial Intelligence Corp., Symantec, and Kestrel Institute.

Fundamental issues in artificial intelligence that must be resolved include representing the knowledge needed to act intelligently,
acquiring knowledge and explaining it effectively,

reasoning: drawing conclusions, making inferences, making decisions ,
evaluating and choosing among alternatives.

Natural Language Interpretation

Research on interpreting natural language is concerned with developing computer systems that can interact with a person in English (or another nonartificial language). One primary goal is to enable computers to use human languages rather than require humans to use computer languages.

Research is concerned with both written and spoken language. Although many of the problems are independent of the communication medium, the medium itself can present problems. We will first consider written language, then the added problems of speech.

There are many reasons for developing computer systems that can interpret natural-language inputs. They can be grouped into two basic categories: improved human/machine interface and automatic interpretation of written text.

Improving the human/machine interface will make it simple for humans to

give commands to the computer or robot,

query data bases,

conduct a dialogue with an intelligent computer system.

The ability to interpret text automatically will enable the computer to

produce summaries of texts,

provide better indexing methods for large bodies of text,

translate texts automatically or semiautomatically,

integrate text information with other information.

Natural-language understanding systems that interpret individual (independent) sentences about a restricted subject (e.g., data in a data base) are becoming available. These systems are usually constrained to operate on some subset of English grammar, using a limited vocabulary to cover a restricted subject area. Most of these systems have difficulty interpreting sentences within the larger context of an interactive dialogue, but a few of the available systems confront the problem of contextual understanding with promising

capability. There are also some systems that can function despite grammatically incorrect sentences and run-on constructions. But even when grammatical constraints are lifted, all commercial systems assume a specific knowledge domain and are designed to operate only within that domain.

Commercial systems providing natural-language access to data bases are becoming available. Given the appropriate data in the area base they can answer questions such as

Which utility helicopters are mission-ready?

Which are operational?

Are any transport helicopters mission-ready?

However, these systems have limitations:

They must be tailored to the data base and subject area.

They only accept queries about facts in the data base, not about the contents of the data base--e.g., "What questions can you answer about helicopters?"

Few Computations can be performed on the data.

In evaluating any given system, it is crucial to consider its ability to handle queries in context. If no contextual processing is to be performed, sentences will often be interpreted to mean something other than what a naive user intends. For example, suppose there is a natural-language query system designed to field questions about air force equipment maintenance, and a user asks "What is the status of squadron A?" If the query is followed by "What utility helicopters are ready?" the utterance will be interpreted as meaning "Which among all the helicopters are ready?" rather than "Which of the squadron A helicopters are ready?" The system will readily answer the question; it just will not be the question the user thought he was asking.

Data base access systems with more advanced capabilities are still in the research stages. These capabilities include

easy adaptation to a new data base or new subject area,

replies to questions about the contents of the data base (e.g., what do you know about tank locations?),

answers to questions requiring computations (e.g., the time for a ship to get someplace).

It is nevertheless impressive to see what can be accomplished within the current state of the art for specific information processing tasks. For example, a natural-language front end to a data base on oil wells has been connected to a graphics system to generate customized maps to aid in oil field exploration. The following sample of input illustrates what the system can do.

Show me a map of all tight wells drilled by Texaco before May 1, 1970, that show oil deeper than 2,000 ft, are themselves deeper than

5,000 ft, are now operated by Shell, are wildcat wells where the operator reported a drilling problem, and have mechanical logs, drill stem tests, and a commercial oil analysis, that were drilled within the area defined by latitude 30 deg 20 min 30 sec to 31:20:30 and 80-81. Scale 2,000 ft.

This system corrects spelling errors, queries the user if the map specifications are incomplete, and allows the user to refer to previous requests in order to

generate maps that are similar to previous maps.

This sort of capability cannot be duplicated for many data bases or information processing tasks, but it does show what current technology can accomplish when appropriate problems are tackled.

Research Issues

In addition to extending capabilities of natural-language access to data bases, much of the current research in natural language is directed toward determining the ways in which the context of an utterance contributes to its meaning and toward developing methods for using contextual information when interpreting utterances. For example, consider the following pairs of utterances:

Sam: The lock nut should be tight.

Joe: I've done it.

and

Sam: Has the air filter been removed?

Joe: I've done it.

Although Joe's words are the same in both cases, and both state that some action has

been completed, they each refer to different actions--in one case, tightening the lock nut; in the other, removing the air filter. The meanings can only be determined by knowing what has been said and what is happening.

Some of the basic research issues being addressed are

interpreting extended dialogues and texts (e.g., narratives, written reports) in which the meaning depends on the context;

interpreting indirect or subtle utterances, such as recognizing

that "Can you reach the *salt*?" is a request for the salt; developing ways of expressing the more subtle meanings of

sentences and texts.

Spoken Language

Commercial devices are available for recognizing a limited number of spoken words, generally fewer than 100. These systems are remarkably reliable and very useful for certain applications.

The principal limitations of these systems are that

they must be trained for each speaker,

they only recognize words spoken in isolation,

they recognize only a limited number of words.

Efforts to link isolated word recognition with the natural-language understanding systems are now under way. The result would be a system that, for a limited subject area and a user with some training, would respond to spoken English inputs.

Understanding connected speech (i.e., speech without pauses) with a reasonably large vocabulary will require further basic research in acoustics and linguistics as well as the natural-language issues discussed above.

Generating Information

Computers can be used to present information in various modes, including written language, spoken language, graphics, and pictures. One of the principal concerns in artificial intelligence is to develop methods for tailoring the presentation of information to individuals. The presentation should take into account the needs, language

abilities, and knowledge of the subject area of the person or persons.

In many cases, generation means deciding both what to present and how to present it. For example, consider a repair adviser that leads a person through a repair task. For each step, the adviser must decide which information to give to the person. A very naive person may need considerable detail; a more sophisticated person would be bored by it. There may, for example, be several ways of referring to a tool. If the person knows the tool's name then the name could be used; if not, it might be referred to as "the small red thing next to the toolchest." The decision may extend to other modes of output. For example, if a graphic display is available, a picture of the tool could be drawn rather than a verbal description given.

Current Status

At present, most of the generation work in artificial intelligence is concerned with generating language. Quite a few systems have been developed to produce grammatical English (or other natural language) sentences. However, although a wide range of constructions can be produced, in most cases the choice of which construction

(e.g., active or passive voice) is made arbitrarily. A few systems can produce stilted paragraphs about a restricted subject area.

A few researchers have addressed the problems of generating graphical images to express information instead of language. However, many research issues remain in this area.

Research Issues

Some of the basic research issues associated with generating information include

deciding which grammatical construction to use in a given situation ;

deciding which words to use to convey a certain idea;

producing coherent bodies of text, paragraphs, or more;

tailoring information to fit an individual's needs.

Assimilating Information

Being in any kind of changing environment and interacting with the environment means getting new information. That information

must be incorporated into what is already known, tested against it, used to modify it, etc. Since one aspect of intelligence is the ability to cope with a new or changing situation, any intelligent system must be able to assimilate new information about its environment.

Because it is impossible to have complete and consistent information about everything, the ability to assimilate new information also requires the ability to detect and deal with inconsistent and incomplete information.

Expert Systems

The material presented here is designed to provide a simple overview of expert systems technology, its current status, and research issues. The importance of this single topic, however, suggests that it merits a more in-depth review; an excellent one recently published by the NBS is recommended [25].

Expert systems are computer programs that capture human expertise about a specialized subject area. Some applications of expert systems are medical diagnosis (INTERNIST, MYCIN, PUFF), mineral exploration (PROSPECTOR), and diagnosis of equipment failure (DART).

The basic technique behind expert Systems is to encode an expert 's knowledge as rules stating the likelihood of a hypothesis based on available evidence. The expert system uses these rules and the available evidence to form hypotheses. If evidence is lacking, the expert system will ask for it.

An example rule might be

IF THE JEEP WILL NOT START

and

THE HORN WILL NOT WORK

and

THE LIGHTS ARE VERY DIM,

then

THE BATTERY IS DEAD,

WITH 90 PERCENT PROBABILITY.

If an expert system has this rule and is told, "the jeep will not start," the system will ask about the horn and lights and decide the likelihood that the battery is dead.

Current Status

Expert systems are being tested in the areas of medicine, molecular genetics, and mineral exploration, to name a few. Within certain limitations these systems appear to perform as well as human experts. There is already at least one commercial product based on expert-system technology.

Each expert system is tailored to the subject area. It requires extensive interviewing of an expert, entering the expert's information into the computer, verifying it, and sometimes writing new computer programs. Extensive research will be required to improve the process of getting the human expert ' s knowledge into the computer and to design systems that do not require programming changes for each new subject area.

In general, the following are prerequisites for the success of a knowledge-based expert system:

There must be at least one human expert acknowledged to perform the task well.

The primary source of the expert ' s exceptional performance must be special knowledge, judgment, and experience.

The expert must be able to explain the special knowledge and experience and the

methods used to apply them to particular problems.

The task must have a well-bounded domain of applications [25].

Research Issues

Basic research issues in expert systems include

the use of, causal models, i.e., models of how something works to help determine why it has failed;

techniques for reasoning with incomplete, uncertain, and possibly conflicting information;

techniques for getting the proper information into rules;

general-purpose expert systems that can handle a range of similar problems, e.g., work with many different kinds of mechanical equipment.

Planning

Planning is concerned with developing computer systems that can combine sequences of actions for specific problems. Samples of planning problems include

placing sensors in a hostile area,
repairing a jeep,
launching planes off a carrier,
conducting combat operations,
navigating,
gathering information.

Some planning research is directed towards developing methods for fully automatic planning; other research is on interactive planning, in which the decision making is shared by a combination of the person and the computer. The actions that are planned can be carried out by people, robots, or both.

An artificial intelligence planning system starts with

knowledge about the initial situation,
e.g., partially known terrain in hostile territory;

facts about the world, e.g., that moving changes location;

possible actions, e.g., walk, fly, look around, hide;

available objects, e.g., a platform on wheels, arms, sensors;

a goal, e.g., installing sensors to detect hostile movements and activity.

The system will produce (either by itself or with guidance from a person) a plan containing these actions and objects that will achieve the goal in this situation.

Current Status

The planning aspects of AI are still in the research stages. The research is both theoretical in developing better methods for expressing knowledge about the world and reasoning about it and more experimental in building systems to demonstrate some of the techniques that have been developed. Most of the experimental systems have been

tested on small problems. Recent work at SRI on interactive planning is one attempt to address larger problems by sharing the decisionmaking between the human and machine.

Research Issues

Research issues related to planning include

reasoning about alternative actions that can be used to accomplish a goal or goals,

reasoning about action in different situations,

representing spatial relationships and movements through space and reasoning about them,

evaluating alternative plans under varying circumstances, planning and reasoning with uncertain, incomplete, and inconsistent information,

reasoning about actions with strict time requirements; for example, some actions may have to be performed sequentially or in parallel or at specific times (e.g., night time),

replanning quickly and efficiently when the situation changes.

Monitoring Actions and Situations

Another aspect of reasoning is detecting that something significant has occurred (e.g., that an action has been performed or that a situation has changed). The key here is significant. Many things take place and are reported to a computer system; not all of them are significant all the time. In

fact, the same events may be important to some people and not to others. The problem for an intelligent system is to decide when something is important.

We will consider three types of monitoring: monitoring the execution of planned actions, monitoring situations for change, and recognizing plans.

Execution Monitoring

Associated with planning is execution monitoring, that is, following the execution of a plan and replanning (if possible) when problems arise or possibly gathering more information when needed. A monitoring system will look for specific situations to be sure that they have been achieved; for example, it would determine if a piece of equipment has arrived at a location to which it was to have been moved.

We characterize the basic problem as follows: given some new information about the execution of an action or the current situation, determine how that information relates to the plan and expected situation, and then decide if that information signals a problem; if so, identify options available for fixing it. The basic steps are:

(1) find the problem (if there is one), (2) decide what is affected,

(3) determine alternative ways to fix the problem, and (4) select the best alternative. Methods for fixing a problem include choosing another action to achieve the same goal, trying to achieve some larger goal another way, or deciding to skip the step entirely.

Research in this area is still in the basic stages. At present, most approaches assume a person supplies unsolicited new information about the situation. However, for many problems the system must be able to acquire directly the information needed to be sure a plan is proceeding as expected, instead of relying on volunteered information. Planning to acquire information is a more difficult problem

because it requires that the computer system have information about what situations are crucial to a plan's success and be able to detect that those situations hold. Planning too many monitoring tasks could be burdensome; planning too few might result in the failure to detect an unsuccessful execution of the plan.

Situation Monitoring

Situation monitoring entails monitoring reported information in order to detect changes, for example, to detect movements of headquarters or changes in supply routes.

Some research has been devoted to this area, and techniques have been developed for detecting certain types of changes. Procedures can be set to be triggered whenever a certain type of information is inserted into a data base. However, there are still problems associated with specifying the conditions that should trigger them. In general, it is quite difficult to specify what constitutes a change. For example, a change in supply route may not be signaled by a change of one truck's route, but in some cases three trucks could signal a change. A system should not alert a person every time a truck detours, but it should not wait until the entire supply line has changed.

Specifying when the change is significant and developing methods for detecting it are still research issues.

Plan Recognition

Plan recognition is the process of recognizing another's plan from knowledge of the situation and observations of

actions. The ability to recognize another's plan is particularly important in adversary situations where actions are planned based on assumptions about the other side's intentions. Plan recognition is also important in natural language generation because a question or statement is often part of some larger task. For example, if a person is told to use a ratchet wrench for some task, the question "What ' s a ratchet wrench?" may be asking "How can I identify a ratchet wrench?" Responding appropriately to the question entails recognizing that having the wrench is part of the person ' s plan to do the task.

Research in plan recognition is in early stages and requires further basic research, particularly on the problem of inferring goals and intentions.

Applications-Oriented Research

The general areas of natural-language processing, speech recognition, expert systems, planning, and monitoring suggest the sorts of problems that are studied in artificial intelligence, but they may not, by themselves, suggest the variety of information processing applications that will be possible with AI technology. Some research projects are now consolidating

advances in more than one area of AI in order to create sophisticated systems that better address the information processing needs of industry and the military.

For example, an expert system that understands principles of programming and software design can be used as a programming tutor for students at the introductory level. This illustrates how an expert system can be incorporated in a computer-aided instruction (CAI) system to provide a more sophisticated level of interactive instruction than is currently available.

Programs for CAI can also be enhanced by natural-language processing for instruction in domains that require the ability to answer and ask questions. For example, Socratic teaching methods could be built into a political science tutor when natural-language processing progresses to a robust stage of sophistication and reliability. Even with the current technology, a reading tutor for students with poor literacy skills could be designed for individualized instruction and evaluation. In fact, the long-neglected area of machine translation could be profitably revisited at this time with an eye toward automated language tutors.

Today's language analysis technology could be put to work evaluating student translations of single sentences in restricted knowldomains, and our generation systems could suggest appropriate alternatives to incorrect translations as needed. This task orientation is slightly different from that of an automated translator, yet it would be a valuable application that our current state of the art could tackle effectively.

Systems that incorporate knowledge of plans and monitoring can be applied to the office environment to provide intelligent clerical assistants. Such an automated assistant could keep track of ongoing projects, reminding the user where he is with respect to a particular job and what steps remain to be taken. Some scheduling advice might be given if limited resources (time, secretarial help, necessary supplies) have to be used efficiently. A truly intelligent assistant with natural-language processing abilities could screen electronic mail and generate suggested responses to the more routine items of business at hand ("yes, I can make that meeting"; "I'm sorry I won't be able to make that deadline" ; "no, I don't have access to the technology"). Automated assistants with knowledge of specific procedures could be useful both to

novices who are learning the ropes and to more experienced users who simply need to use their time as effectively as possible.

While most expert systems today assimilate new knowledge in highly restricted ways, the importance of learning systems should not be overlooked. In the long run, general principles of learning will become critical in designing sophisticated information processing systems that access large quantities of data and work within multiple knowledge domains. As AI moves away from problems within restricted knowledge domains, it will become increasingly important for more powerful systems to integrate and organize new information automatically--i.e., to learn by themselves. We will have to move away from simplistic pattern-matching strategies to the more abstract notions of analogy and precedents. Research on learning is still in its infancy, but we can expect it to become an application-oriented research issue very quickly--within 5 to 10 years, if the field progresses at a healthy pace. Without sufficient research support in this area, our efforts may stagnate in the face of apparent impasses.

With a field that moves as rapidly as AI, it is important to realize that a long-term

perspective must be assumed for even the most pragmatic research effort. Even a 2-year project designed to use existing technology may adapt new techniques that become possible during the life of the project. The state of the art is a very lively moving target, and advances can render research publications obsolete in the space of a few months. New Ph.D.s must keep close tabs on their areas of interest to maintain the expertise they worked so hard to establish in graduate school. We must therefore emphasize how dangerous a short view of AI is and how critical it is for the field to maintain a sensitive perspective on long-term progress in all of our research efforts.

STATE OF THE ART AND PREDICTIONS

In the previous sections we have reviewed the state of the art in robotics and artificial intelligence. Clearly, both robotics and artificial intelligence are relatively new fields with diverse and complex research questions. Furthermore, the intersection field--robotics/artificial intelligence or the intelligent robot--is an embryonic research area. This area is made more complex by the obvious dependence on heretofore unrelated fields, including mechanical design, control,

vision sensing, force and touch sensing, and knowledge engineering. Thus, predicting the state of the art 5 and 10 years from now is difficult. Moreover, because predictions for the near future are likely to be more accurate than those for the more distant future, our 10-year predictions should be treated with particular precaution.

One approach to the problem of prediction is to decouple the fundamental research areas and predict possible developments in each technology area. Such a task is easy only in comparison to the former question; nevertheless, in the following sections we undertake a field-by-field assessment and predictions of 5- and 10-year developments.

In the sections that follow, we develop tables describing the current state of the art and predictions for the next 5- and 10-year periods. Each section contains a short narrative and some general

comments with respect to research funding and researchers working in the problem area. The table at the end of the chapter summarizes the findings.

Mechanical Design of the Manipulator and Actuation Mechanism

The industrial robot is a single mechanical arm with rigid, heavy members and linkages. Actuation of the slide or rotary joints is based on transmission gears, which results in backlash. Joint bearings of conventional design have high friction and stiction, which cause poor robot performance. Thus, with the rare exception of some semiconductor applications that are more accurate, robot repeatability is in the range of 0.1 to 0.005 inches. Robots today operate from fixed locations with little or no mobility (except track mountings or simple wire-guided vehicles) and have a limited work envelope. The operating environment is constrained to the factory floor, and the typical robot is not self-contained but requires an extensive support system with big power supplies.

The factors listed above are reflected in the first column of the table under entry numbers 1 to 11. As shown in the table, on a point by point basis we expect significant improvements within 5 years (column 2) and even more within 10 years (column 3).

Table entries 12 and 13 address the kinematics and dynamics of robots as they are today (column 1) and predict how they will evolve. These issues, while based

fundamentally on the mechanical structure of the robot and how it behaves in motion and under load, are clearly intertwined with the issues of manipulator control and computation speed. For example, we do not today have enough computer power in the robot control system to take advantage of kinematic model data.

Thus, while we make some predictions under these headings, they are closely related to the control issues to be addressed later.

The research on mechanical design and actuation mechanisms has been supported by NSF, ONR, and others but is not the main focus of a major funding program at this time. University laboratories such as those at MIT, CMU, Stanford, and the University of Florida at Gainesville are investigating the manipulator and its kinematics.

Locomotion research is continuing at Ohio State, CMU, and RPI. The Jet Propulsion Laboratory, Stanford Research Institute, and Draper Laboratories are also active in some of these areas [3-7].

End-Effector Design

Current industrial robots use many hands, each specifically designed for a different application. As described in the Research section, this has led to research in two

directions--one to produce the dexterous hand and the second to produce the quick-change hand. The lack of progress in these areas makes most applications expensive because of the need to design a special hand, and it prohibits others because of a lack of dexterity or the ability to change hands rapidly.

Many are also working on hand-based sensor systems; these issues are covered in depth under the topic of sensor systems. Entries 14 and 15 in the table describe current technology hands as simple (open or closed) hands that are rarely servoed--though the IBM RSI is a notable exception, which others are following.

End effectors today are also sometimes tools that are operated by an on/off signal. Today's hands do employ limited sensors and permit rudimentary force programming. As described in the table, we expect progress in the development of quick-change hands to precede the wide use of instrumented dexterous hands.

Research in end effectors is taking place at the University of Utah (based on prior work in prosthetics), the University of Rhode Island, and at most of the locations cited for mechanical design research.

References 9-11 are suggested for further details.

Funding of these hand efforts is typically a part of some larger project and is not a major project of any funding agency.

Vision Sensors

As described earlier, vision has been a high-interest area for robotics in both the visual servoing (guidance) and inspection or measurement modality.

Commercial vision systems use binary images and simple features and are restricted to high contrast images. As shown in table entry 16, we expect that VLSI technology, now in research labs at MIT, Hughes, Westinghouse, and others, will be commercialized. In 5 years this will provide real-time edge images, a richer shape-capturing feature set, and will ease the restriction on high-contrast binary images, allowing gray-scale and texture-based objects to be handled. These predictions are conservative. In 10 years we further expect rapid-recognition systems that can handle a limited class of objects in arbitrary orientation. Thus, the visual servoing problem will be routinely achievable.

The use of so-called three-dimensional vision, using stereo, structured light systems, and other vision-based methods to acquire "depth" information, is rudimentary today, as shown in table entry 17. The stereo mapper system at DMA is an exception. This system, which works well on textured terrain such as forests, is ineffective on urban landscapes. A big step forward is expected in the next 5 years. Currently in research labs are systems that extract depth using

stereo, employing either vision or laser light (MIT, Stanford);

shape from shading, special light (GE, MIT, SRI);

gross shape from motion (CMU, MIT, Stanford, University of Minnesota) ;

shape from structured light systems (GE, GM, NBS).

Commercial systems will market three-dimensional vision systems that will generate a depth map in relatively benign situations. They will be slow, too slow for military rapid response situations in the next 5 years. The algorithms for all these methods for computing

depth are inherently parallel. They can be computed using highly parallel computers specifically designed. A hardware stereo (vision or laser) and shape from motion system is possible in 5 years. One practical problem is lithographic density. Putting a lot of processing on chips of 1 micron density restricts spatial resolution of an image. However, 0.1 micron densities seem feasible in 5 years.

Merely generating a depth map is not the same as seeing. It is also necessary to extract objects and to recognize them from arbitrary orientation. The depth map is likely to be noisy and relatively coarse. It will be possible, for example, to identify a shape as a person, but not to recognize which person. It will recognize a tank, but only determine type if it is significantly different from another.

Tasks that will become feasible with depth data include

three-dimensional inspection of object surfaces for dents, cracks, etc. that do not affect outline;

better edge maps and shape, leading to recognition of objects by outline shape, e.g., an automobile.

In 10 years, one can confidently predict reliable hardware stereo systems, systems capable of determining the movement of an object and maneuvering to avoid it, rapid recognition of limited classes of objects from an arbitrary viewpoint.

Vision research is a very active field in the United States (see reference 34). For a survey of vision research, see reference 35. For a review of image understanding, see reference 14. Most three-dimensional vision research in the United States is funded by the DARPA Image Understanding (IU) program. See, for example, the IU workshop proceedings from DARPA.

Commercial vision systems are marketed by GE, Octek, Automatix,

Cognex, Machine Intelligence Corporation, ORS, and others. Government

and foundation support of major programs is provided by the Office of

Naval Research (ONR), DARPA, Systems Development Foundations (SDF), and

NSF.

Many corporations in Japan, including Hitachi, Sony, and Fujitsu, are doing work in this area; there are also several large university efforts (see references 13, 36, 39).

Nonvisual sensors (radar, SAR, FLIR, etc.) have mostly been developed by defense contractors for DARPA, AFOSR, and ONR. The following systems are among those available from Lockheed, TRW, Honeywell, and others:

synthetic aperture radar (SAR),

forward looking infrared (FLIR),

millimeter radar,

Xray.

For example, the cruise missile uses one-dimensional correlations on radar images. This is rather crude. Capabilities are mostly classified.

Advantages of nonvisual sensing are that they simplify certain problems. For example, it is easy to find hot spots in infrared. Often they correspond to camouflaged targets.

Limitations are that the physics of nonvisual imagery are poorly understood, and algorithms are limited in scope. Two

main applications are for seeing large static objects and for automatically navigating certain kinds of terrain.

Research is intense, funding levels are high, and progress will be good. This is entirely an industry effort with DOD sponsorship. However, vision does appear to be the best way forward because it is passive and operators know what visual images mean. This is a serious issue, since trained observers are needed to check results of processing nonvisual images.

Contact/Tactile Sensors

As described earlier, contact/tactile sensors are an important area of robotics development. Although progress has so far been slow, this is an important area for determining

surface shape, including surface inspection;

slip computation--how sure the grasp is;

proximity--how close the hand is to the object;

force/torque, to control and measure its application.

Robots today are programmed for position only; in rare instances, they can do some rudimentary force programming using a commercial version of the Draper Laboratory IRCC. For the state of the art, see references 18-21 and 37

Current systems suffer from both rudimentary control capability (i.e., touch/no-touch and some vector valued sensors) and limited sensors, with high hysteresis and poor wear and tear. As shown in table entry 18, the next 5 years will see better control techniques (possibly hybrid, as Raibert and Craig [37] suggest) and the development of array sensors with more applications. But the real progress of broad commercialization, a true sense of feel, and the development and understanding of the control/programming issues will take us into the 10-year time frame.

Research in tactile sensing is being done at Ohio State University,

MIT, JPL, CMU, Stanford University, the University of Delaware, General

Electric in Schenectady, and in France. Force sensing is being done at

MIT, Draper, Astek, IBM, and other commercial firms.

Research support is not on a large scale: too few people, not enough money. Nevertheless, this is a critical area for assembly and other complex tasks. A concentrated research program by a major funding agency or agencies would speed progress.

As can be seen from the review of research areas, there are many avenues for combining AI and robotics. The future will see a natural combination and extension of each area into the domain of the other, but to date there are no true joint developments. MIT, Stanford, and CMU are beginning to lead the way in joint efforts, and many others are sure to join in.

The general area of reasoning and AI can be partitioned in many ways, and every taxonomy will result in fuzzy edges and work that resists a comfortable pigeonhole. A large portion of AI research can nevertheless be characterized in terms of advisory Systems that strive to assist users in some information processing task. This research can be categorized as work on expert systems, natural-language data base access, computer-aided instruction (CAL), intelligent tutors, and automated assistants.

A great deal of basic research is conducted without recourse to specific task orientations, and progress at this level penetrates a variety of areas in a myriad of guises. Basic research is conducted on knowledge representation, learning, planning, general problem solving, and memory organization. It is difficult to describe the milestones and research plateaus in these areas without some technical introduction to the issues, which is well beyond the scope of this paper. Problems and issues in these areas tend to be tightly interrelated, so we will highlight some of the more obvious accomplishments in a grossly inadequate overview of basic research topics. For further detail, see reference 38.

Expert systems are specialized systems that work effectively in providing competent analyses within a narrow area of expertise (e.g., oil exploration, diagnosis of infectious diseases, VLSI design, military intelligence, target selection for artillery). A few commercial systems are being customized for specific areas. Typically, current expert systems are restricted in a number of ways. First, the expertise is restricted in a very narrow corpus of knowledge. Examples include pulmonary function disorders, criteria for

assessing copper deposits, and configuring certain types of computers. Second, interactions with the outside world and the consequent types of information that can be fed into such expert systems are capable of only a very small number of responses--for example, 1 of 92 drug therapies. Finally, they adopt a single perspective on a problem. Consider, by way of contrast, that trouble-shooting an automobile failure to turn over the starter motor (electrical) suggests a flat battery. The battery is charged by the turning of the fan (part of the hydraulic cooling system). This turns out to be deficient because of a broken fan belt (mechanical).

Table entry 19 summarizes the current state of expert systems and reflects the expectation of their integration with other systems within 5 years and significant improvement within 10 years. Significant work centers are at Stanford, Carnegie-Mellon, Teknowledge, Schlumberger, and a variety of other locations.

Natural-language data base access is now limited to queries that

address the contents of a specific data base. Some require restricted subsets of English grammar; others can unravel

ungrammatical input, run-on sentences, and spelling errors. Some applications handle a limited amount of context-sensitive processing, in which queries are interpreted within the larger context of an interactive dialogue. We are just now seeing the first commercial systems in this area. As table entry 20 shows, we expect sophisticated dialogue capabilities for interactive sessions and better recognition capability for requests the data base cannot handle. More domains will have been tackled, and some work may relate natural-language access capabilities to data base design issues. We should see some efforts to connect expert-system capabilities with natural-language data base access to provide advisory systems that engage in natural-language dialogues in the next 5 years.

In 10 years the line between natural-language data base access and expert systems will be hard to draw. Systems will answer questions and give advice with equal ease but still within well-specified domains and limited task orientations. Key research efforts are at Yale, Cognitive Systems, Teknowledge, Machine Intelligence Corporation, and other locations.

Basic research on automated assistants is now being conducted for a variety of tasks. As shown in table entry 21, this work, which takes place at MIC, SRI, the University of Massachusetts, IBM, and DEC, can be integrated with the other AI technologies. The field is not yet funded to any extent, but commercial interest is growing and should attract funding.

With respect to knowledge representation and memory organization, there are techniques that operate adequately or competently for specific tasks over restricted domains. Most of the work in learning, planning, and problem solving has been domain-independent, with prototype programs operating in specific domains (e.g., learning by analogy). The domain-dependent work in these areas tends to start from a domain-independent base, augmenting this foundation with semantics and memory structures. As shown in table entry 22, progress is dependent on better understanding of knowledge; its representation is hard to predict.

Control Structure/Programming Methodology

Perhaps the most difficult area of all to cover is the future of control structures and programming methodology. In some sense,

all the developments described impinge on this area; new mechanical designs, locomotion, dexterous hands, vision, contact/tactile sensors, and the various AI methodologies all affect the architecture of robot control and will affect the complexity of programming methodology.

In order to treat the subject in an orderly way, we deal first with a logical progression of control structure. Then, possibly with overlap, we deal with the other topics.

The most advanced current work in control structures uses multiple microprocessors on a common bus structure. Typically, such robot controllers partition the control problem into levels as follows:

1. Servo control to provide closed-loop feedback control.
2. Coordinate transformation to joint coordinates, and coordinated joint motion.
3. Path planning for simple interpolated (straight line) motion through specified points.
4. Simple language constructs to provide subroutines, lock-step interaction, and binary sensor-based program branches.

5. Structured languages, limited data base control) complex sensor communication, and hierarchical language definitions.

Levels 1 to 3 are common in most servo robots; level 4 is represented by the first-generation languages such as VAL on Unimation robots, while level 5 represents second-generation languages as found in the IBM AML Language, the Automatix RAIL, and at the National Bureau of Standards.

Beyond the first five levels of control are a diversity of directions being pursued to different extents by various groups. Thus, we can expect a number of developments in the next 5 years but clearly will not see them integrated in that time. As shown in table entry 23, we see the following extensions:

Graphic systems will be used to lay out, program, and simulate robot operations. Such systems are starting to enter the market today from McAuto, Computervision, GCA, and others.

Hierarchical task-oriented interface languages will be developed on the current structural languages (AML, RAIL, etc.) to allow process planners to program applications.

Robot operating systems and controllers will be more powerful. They will remove the burden of low-level control over sensors, I/O, and communication; that is, they will do more of what computer operating systems do for their users today.

Interfaces to other nonhomogeneous computers via developments in local area networks and distributed computing will broaden coordination beyond the lock-step synchronization available today.

The use of multiple arms, dexterous hands, locomotion mechanisms, and other mechanical advances will foster the definition of a sixth level of control. This will emerge from research labs and be available in some rudimentary form.

The incorporation of AI technology in the use of expert systems is in the laboratory plans of some now. This, coupled with the use of natural-language front ends and knowledge engineering, will begin the definition of a seventh level of control.

The linkage of robot control/programming systems with CAD, CAM, and other factory data bases will be made.

Beyond these advances in new areas will be significant improvements in the first five

levels as computers get more powerful and cheaper.

For example, the use of kinematic and dynamic models discussed in table entries 12 and 13 will affect the first five levels, as will the development and instrumentation of new sensors for resolving robot position.

The research in these areas is growing rapidly. Robotics institutes at major universities--CMU, MIT, Stanford, Florida, Lehigh, Michigan, RPI, and others--are now accelerating their programs under funding from DOD agencies, DARPA, and NSF. As the programs grow, the need for research dollars escalates, but so do the results. Robotics research is expected to expand significantly in the next decade. Commercial firms, both vendors and users, are linking themselves with universities. The list of firms involved includes IBM, Westinghouse, DEC, GE, and many others.

The 10-year time frame is very difficult to predict. This is because of the variety of technologies that must interact and the dependence on the output of a myriad of research opportunities being pursued. However, we feel the following to be conservative estimates.

Robotics will branch out beyond industrial arms to include a wide scope of automatic equipment. The directions will depend on funding emphasis and other such factors.

Sensor-based, advanced mechanical, partially locomotive (in restricted domains), somewhat intelligent robots will have been developed.

Many integration issues and further technological advances will still remain open research questions.

Conclusion

In conclusion, one is forced to observe that the following table describes a technology that is very active--a technology that, while diversifying into many research areas, must be integrated for true success.

For those whose interest is in transferring the technology outside the manufacturing arena, immediate focus on targeted projects appears to be required. Although robotics and AI will be integrated, and the focus on manufacturing will broaden by an evolutionary process, the process will be painfully slow, even when pushed by well-funded initiatives.

Summary State of the Art for Robots and Artificial Intelligence

Now In 5 Years In 10 Years

Mechanical Design and Activation of the Manipulator

1. Single arms with fixed bases
 2. Heavy; designed to be rigid
 3. Humanlike mechanical arrangements; linkage systems
 4. Discrete degrees of freedom
(DOF)
 5. Simple joints, revolute or sliding; Cincinnati Milacron has one version of the 3-roll wrist now
 6. Actuators are electrical, hydraulic, and pneumatic; heavy, low power, often require transmission gears that result in backlash problems
- 2 or 3 rigidly mounted arms designed to work together
- Designed to be rigid but lightweight, using composite materials

No change

No change

Flexible joints possible; better discrete joints (e.g., 3-roll wrist)

Some improvement: lighter weight, rare-earth motors, direct drive

Multiple arms with coordinated motion

Designed to be very lightweight and flexible

Nonlinkage design (e.g., snakes, butterflies)

Continuous degrees of freedom without discrete joints; flexible elements

Flexible joints as above

New actuator concept: distributed actuator (muscle type)

7. Joint bearing, conventional high friction and stiction; poor motion performance

8. No absolute accuracy; repeatability 0.1 in. to 0.005 in. except in highly specialized semiconductor applications

9. Fixed location--some on tracks or wire-guided vehicles; walking, wheeled, and hopping robot mechanisms are now in research labs

10. Limited work envelopes

11. Operate in controlled environment (factories) or with support systems (e.g., underwater applications); not self-contained, umbilical cords, big power unit

New discrete bearing designs (air bearings); some flexible joints possible

Some absolute accuracy is required (for offline pro-gramming); good repeatability of 0.005 in. to

0.001 in.

Mobility based on wheeled-track vehicles in controlled environment (flat factory floor); rudimentary walking in specific environments

More flexible, but constrained envelopes as defined by factors above

Packaging for uncontrolled environments;
not self-contained

No discrete joints, possibly no bearings:
flexible elements, for mobility

Controlled to micron level as required;
also closely coupled to force and position
sensors to give broad functional range

Mobility in semicontrolled environment,
better vehicular control, some walking
ability

Greatly improved work domains by new
designs, linkages, mobility, as defined
above

Possibly self-contained; wider range of
environments tolerated (e.g., nuclear
hardened)

Now In 5 Years In 10 Years

12. The kinematics are a significant
computational burden that limits practical
performance--real limitation is on real
time control and action

13. Dynamics are not considered in robot
design and performance. They are basically
slow devices operating in "quasistatic"

modes. Control systems are on joints only and position only and are relatively primitive. Typically, velocity-dependent and inertial terms ignored. Arms made to run slowly to compensate

New dedicated chips will be available to greatly reduce computational burdens--some slow motion real time possible

Robots will be designed for higher-speed performance with some absolute accuracy. There will be combined force and position control with respect to the workspace rather than joints. Robotic trajectories will be planned for optimal dynamic performance, including the effects of actuator and robot dynamics, and limitations. Adaptive control methods will be available, so the robot will be insensitive and tolerant (dynamically) to its environment and its task

Computation not an issue; real time kinematic possible at high speed

Robots will be high speed and lightweight, with tuned dynamic behavior. Systems will control and exploit their flexibility to achieve high performance. Issues of

dynamics and performance in most cases will move to a higher level. Questions of control of individual elements will be transparent, such as the motion of control surfaces in supersonic aircraft is not considered by the pilot

End Effectors

14 . Currently grippers and special tools. They are, typically

binary (open or closed, on or off) and have few or rudimentary sensors; very simple mechanical actions, mostly one DOF such as parallel jaw pneumatically; and rudimentary force control

15. Quick-change hands are available today on a limited special basis due to a lack of standards for their interconnection to a variety of robots

End effectors with proportional mobility--a hand that can be centered and servoed to fit a wide variety of objects; position and force sensors and limited tactile sensing;

several discrete DOF; major emphasis still on grasping or sucking, with limited assembly or quick-change hand availability. Research labs will have developed multifingered hands and demonstrated their use to grasp a variety of three dimensional shapes

Development of a standard robotarm-to-end-effector interface. Commercial availability of a family of hands for tasks such as assembly, using adaptations of current tools and grippers

Continuous motion, intelligent control and sensing at the wrist, fingers, and fingertips. Beginning to be controlled by vision and other noncontact sensing to perform assembly

Specially designed sensor-based robot hands with tools for a family of tasks. All able to fit the standard interface

Now In 5 Years In 10 Years

Vision Sensors

16. Current commercial systems are restricted to binary image and simple features; gray-scale and color are available today only in very restrictive form

17. 3-D vision systems, structured light, and stereo approaches to acquiring depth image are rudimentary and only beginning to emerge from laboratories into commercial systems

VLSI implementation now in labs will be commercialized. This will facilitate edge images from gray-scale data, and richer feature sets will be developed

Laboratory systems of several varieties will be commercially available. They will produce depth maps in controlled situations, but they will be slow, will produce noisy images, and have limited resolution. They will permit 3-D surface inspection and will discriminate objects with large shape differences

Systems that permit rapid recognition and provide orientation of limited classes of objects from arbitrary points of view

Reliable hardware for depth images and systems for tracking and recognizing moving objects

Contact and Tactile Sensing

18. Few robots have force or tactile sensors. The IBM RSI is an exception. Limited use of commercialized RCC and IRCC versions of Draper Research products provide limited control capacity at present

Force-sensing wrists and techniques for programming and controlling force will be available. They are likely to work only in benign situations, but should be able to tighten nuts, insert shafts, pack objects--simple assembly operations. Will not yet be good enough to examine objects by feeling them

Well-established techniques for creating and using these sensors will be developed. Determining shape of objects, detecting slippage in grip, inspecting for cracks, and programming in the force domain will be possible. Touch sensors will be implemented in hardware, probably using VLSI technology. This will permit all of the above and offer a wider range of force monitoring and compliant operations

Now In 5 Years In 10 Years

Artificial Intelligence

19. Expert systems that work effectively in providing competent analysis within a narrow area of expertise,

e.g. oil exploration, medical diagnosis, VLSI design, are being customized and commercialized. They are limited by a narrow body of simple interactions, and they take a single perspective on the problem. There are no generalized ways to build the expert systems

20. Natural-language data base access methodology is limited to single-shot query systems for specific data bases. Some require restricted subsets of English grammar, but others are more general about input. Commercial systems are just starting to appear

Automated design assistance for building and updating expert systems. Formalization of knowledge gathering and integration of graphic displays for use in some applications. Integration with robot control systems and sensors to provide controlled expertise for limited domains, e.g., arc welding

New sophisticated dialog capabilities for interactive sessions will appear. Some developments will permit the start of natural-language data bases. The connection of expert systems to natural language will begin

Integrated systems that draw on multiple domains of expertise to formulate problem solutions. Possibly total automation in generating new expert systems for certain domains . Self-diagnosing and limited repair of electronic equipment limited repair of electronic equipment

The hard line between natural-language query and expert systems will disappear. Systems will be integrated, but the domain of knowledge will still be restrictive

21. Automated assistants research is now going on in a variety of tasks, such as word processing, text editing, and office automation ion

22. Knowledge representation in restricted domains is now workable (see entries 19-21). But learning, problem-solving, and planning systems need broader domains .

Systems that assist and familiarize users with the capabilities of the system being used

Increased understanding of tradeoffs between domain-independent and domain-dependent techniques

Integrated systems that draw on multiple domains and provide the user with with greater task flexibility

Possibly a notation system that allows formulation of models that are sensitive to domain constraints without having specific commitments to particular domains

Control Structure/Programming Methodology

23. The control hierarchy of robots sometimes implemented on multiple microprocessors has at most 5 levels now.

1. Servo control of joints

2. Coordinate transformation and coordinated joint motion.

3. Interpolated path planning for smooth motion paths.

Individual elements of progress (not all in any one offering) will be developed.

. Graphical layout of robotic cells and programming will be commercialized

. Hierarchical task-oriented interface languages designed for process planners will be developed .

Levels six and seven as defined in the previous column will permit domain-dependent , sensor-based intelligent robots. Many integration issues and advances to technology will still be open questions. Robotics will broaden in scope beyond manufacturing to limited-domain automatic devices in new areas.

Now In 5 Years In 10 Years

4. Simple subroutines, use of sensors, and lock-step coordination

5. Rudimentary operating system, structural language, complex sensor interface, hierarchical constructs

. Robot operating systems will do more for the user who uses sensors to permit task orientation

. Interfaces to other nonhomogeneous computers will broaden coordination beyond lock-step available now

. Multiple arm, dexterous hand, locomotive control, and other new mechanical advances

will define a sixth level of control and be available

. The incorporation of AI technology in the form of expert systems, natural-language front ends) and knowledge representation will define a seventh level of control.

. Data bases from CAD, CAM) and other sources will be incorporated to the language and control structure

REFERENCES

1.

National Bureau of Standards. 1980. Proceedings of NBS/Air Force ICAM Workshop on Robot Interfaces, June 4-6. NBSIR 80-2152.

2. Taylor, R. H., P. D. Summers, and J. M. Meyer. 1982. AML: A Manufacturing Language. International Journal of Robotics Research 1(3):19-41.

3. Birk, J. and R. Kelley, eds. 1980. Research Needed to Advance

the State of Knowledge in Robotics.
Kingston: Rhode Island

University.

4. Roth, B. Kinematic Design for Manipulation, in [3], pp. 110-118.

5. Dubowsky, S. Dynamics for Manipulation, in [3], pp. 119-128.

6. Houston, R. Compliance in Manipulation Links and Joints, in [3], pp. 129-145.

7. Paul, R. P. 1981. Robot Manipulators Mathematics Programming

and Control. Cambridge, Mass.: MIT Press.

8. Brady, M. and J. Hollerbach. 1982. Robot Motion: Planning and

Control. Cambridge, Mass.: MIT Press.

9. Toepperwein, L. L., M. T. Blackmon, R. Fukui, W. T. Park, and B. Pollard. 1980. ICAM Robotics Applications Guide. Vol. II. Technical Report AFWAL-TR-80-4042.

10. Salisbury, J. K. and J. Craig. 1982. Articulated Hands: Force

Control and Kinematic Issues. International Journal of Robotics

Research 1(1):4-17.

11. Hollerbach, J. M. 1982. Workshop on Dexterous Hands. MIT AI Memo.

12.

Orin, D. E. 1982. Supervisory Control of a Multilegged Robot. International Journal of Robotics Research 1(1):79-91.

13. Gleason, G. J. and G. Again. 1979. A Modular Vision System For Sensor Control Manipulation and Inspection. SRI Report, Project 4391. SRI International.

14. Lavin, M. A. and L. I. Lieberman. 1982. AML/V: An Industrial Machine Vision System. International Journal of Robotics Research 1(3):42-56.

15. Nagel, R. N., et al. 1979. Experiments in Part Acquisition

Using Robot Vision. SME Technical Paper MS 79-784.

16. Brady, M. 1982. Computational Approaches to Image Understanding.

Computing Surveys 14:4-71.

17. Nevins, J. L., et al. Exploratory Research in Industrial Assembly and Part

Mating. Report No. R-1276. Cambridge, Mass.:

Charles Stark Draper Laboratory. 193 pp.

18. Harmon, L. D. 1982. Automated Tactile Sensing. International Journal of Robotics Research 1(2):3-32.

19. Bejczy, A. K. 1979. Manipulator Control Automation Using Smart Sensors. Paper delivered at Electro/79 Conference, New York, April 24-26.

20. Raibert, M. H. and J. E. Tanner. 1982. Design and Analysis of a VLSI Tactile Sensor. International Journal of Robotics Research. 1(3):3-18.

21. Hillis, W. D. 1982. A High Resolution Image Touch Sensor. International Journal of Robotics Research. 1(2):33-44.

22. Albus, J. S., A. J. Barbera, M. L. Fitzgerald, R. N. Nagel, G. J.

VanderBrug, and T. E. Wheatley. 1980. Measurement and Control

Model for Adaptive Robots. Pp. 447-466 in Proceedings, 10th

International Symposium on Industrial Robots, Milan, Italy, March

5-7.

23.

Nagel, R. N., et al. 1982. Connecting the Puma Robot With the

MIC Vision System and Other Sensors. Pp.447-466 in Robot VI

Conference Proceedings, Detroit, March 2-4.

24. D. R. Brown, et al. 1982. R&D Plan for Army Applications of AI/Robotics. SRI Project 3736. SRI International. 324 pp.

25.

Nau, D. S. 1982. Expert Computer Systems and Their Applicability to Automated Manufacturing. NBSIR 81-2466.

26.

Charniak, E., and Y. Wilks, eds. 1976. Computational Semantics:

An Introduction to Artificial Intelligence and Natural Language

Comprehension. Amsterdam: North Holland Publishing Co.

27. Lehnert, W., and M. Ringle, eds. 1982. Strategies for Natural

Language Processing. Hillsdale, N.J.:
Lawrence Erlbaum

Associates.

28. Nilsson, N. J. 1971. Problem Solving
Methods in Artificial

Intelligence. New York: McGraw-Hill.

29.

Schank, R., and R. Abelson. 1977. Scripts,
Plans, Goals and Understanding. Hillsdale,
N.J.: Lawrence Erlbaum Associates.

30. Waltz, D. L. 1982. Artificial
Intelligence. Scientific American.

247(4):118-133.

31. Winston, P. H. 1977. Artificial
Intelligence. Reading, Pa.:

Addison Wesley.

32. Proceedings for the Conference on
Applied Natural Language Processing, Santa
Monica, Calif., February 1983.

33. Proceedings for the Association of
Artificial Intelligence Conference on
Artificial Intelligence (IJCAI 1969, 1973,
1975, 1977, 1979, 1981).

34. Ballard, D. H. and C. M. Brown. 1982. Computer Vision. Englewood Cliffs, N.J.: Prentice-Hall.
35. Rosenfeld, A. 1983. Picture Processing: 1982. Computer Science Technical Report. College Park: University of Maryland.
36. Dennicoff, M. 1982. Robotics in Japan. Washington, D.C.. Office of Naval Research.
37. Raibert, M., and J. Craig. 1981. Hybrid Controller. IEEE Systems Management Cybernetics.
38. Barr, A., and E. A. Feigenbaum, eds. 1981, 1982. Handbook of Artificial Intelligence, vols. I-III. Stanford, Calif.:
HeurisTech Press.
39. State of the Art of Vision in Japan, IEEE Computer Magazine (13)
1980.

GLOSSARY OF ACRONYMS

AFOSR Air Force Office of Scientific Research

AI artificial intelligence

AML manufacturing language developed at IBM

AMRDC U.S. Army Medical Research and Development Command

ASB Army Science Board

ASP Automated Ammunition Supply Point

ATE automatic test equipment

BITE built-in test equipment

C3I command, control, communication, and intelligence

CAD/CAM computer-aided design and manufacturing

CAI computer-aided instruction

CARP computer-aided robot programming

CMU Carnegie-Mellon University

CPU central processing unit

CRT cathode ray tube

DARPA Defense Advanced Research Projects Agency

DART expert system for the diagnosis of equipment failure

DEC Digital Equipment Corporation

DMA Defense Mapping Agency

ES expert system

FLIR forward-looking infrared

FMS flexible manufacturing system

GE General Electric Company

GM General Motors Corporation

Hawk-Missile CAI trainer at Fort Bliss Air Defense School

ICAM Integrated Computer-Aided Manufacturing program of the U.S. Air Force

IR industrial robot

IRCC instrumented remote center of compliance developed at Draper Laboratories

JPL Jet Propulsion Laboratory

MACSYMA symbolic mathematics expert system

MIC

MIT

MYCIN

NBC

NBS

NSF

ONR

Prospector

PUFF

P3I

RAIL

RAMS

R&D

REMBASS

RIA

RPI

SAR

SRI

VAL

VHF

VHSIC

VIMAD

VLSI

VTRONICS

computer language developed at McDonnell Douglas

Machine Intelligence Corporation Massachusetts Institute of Technology
production system for diagnosis and treatment

of infectious diseases nuclear, biological) and chemical National Bureau of
Standards National Science Foundation Office of Naval Research

expert system to aid in exploration for minerals

pulmonary function diagnosis expert system preplanned product improvement
Pascal-based second generation language by IBM reliability, availability,
maintainability)

and supportability research and development

remotely monitored battlefield sensor system Robot Institute of America

Rensselaer Polytechnic Institute synthetic aperture radar Stanford Research
Institute

language developed by Unimation for Puma robot very high frequency

Very High Speed Integrated Circuits Voice Interactive Maintenance Assistance

Development system (supported by DARPA) very large-scale integration

set of projects for onboard, embedded sensing of vehicular malfunctions with
built-in test equipment (BITE)