


Computer Science and Artificial Intelligence



Panel on Computer Science and Artificial Intelligence,
National Research Council

ISBN: 0-309-59085-X, 32 pages, 6 x 9, (1997)

This free PDF was downloaded from:
<http://www.nap.edu/catalog/5812.html>

Visit the [National Academies Press](http://www.nap.edu) online, the authoritative source for all books from the [National Academy of Sciences](http://www.nap.edu), the [National Academy of Engineering](http://www.nap.edu), the [Institute of Medicine](http://www.nap.edu), and the [National Research Council](http://www.nap.edu):

- Download hundreds of free books in PDF
- Read thousands of books online, free
- Sign up to be notified when new books are published
- Purchase printed books
- Purchase PDFs
- Explore with our innovative research tools

Thank you for downloading this free PDF. If you have comments, questions or just want more information about the books published by the National Academies Press, you may contact our customer service department toll-free at 888-624-8373, [visit us online](http://www.nap.edu), or send an email to comments@nap.edu.

This free book plus thousands more books are available at <http://www.nap.edu>.

Copyright © National Academy of Sciences. Permission is granted for this material to be shared for noncommercial, educational purposes, provided that this notice appears on the reproduced materials, the Web address of the online, full authoritative version is retained, and copies are not altered. To disseminate otherwise or to republish requires written permission from the National Academies Press.

Computer Science and Artificial Intelligence

NRL STRATEGIC SERIES

Panel on Computer Science and Artificial Intelligence
Naval Studies Board
Commission on Physical Sciences, Mathematics, and Applications
National Research Council

NATIONAL ACADEMY PRESS

Washington, D.C. 1997

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 panel responsible for this 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 Academy of Sciences is a private, nonprofit, self-perpetuating society of distinguished scholars engaged in scientific and engineering research, dedicated to the furtherance of science and technology and to their use for the general welfare. Upon the authority of the charter granted to it by the Congress in 1863, the Academy has a mandate that requires it to advise the federal government on scientific and technical matters. Dr. Bruce Alberts is president of the National Academy of Sciences.

The National Academy of Engineering was established in 1964, under the charter of the National Academy of Sciences, as a parallel organization of outstanding engineers. It is autonomous in its administration and in the selection of its members, sharing with the National Academy of Sciences the responsibility for advising the federal government. The National Academy of Engineering also sponsors engineering programs aimed at meeting national needs, encourages education and research, and recognizes the superior achievements of engineers. Dr. William A. Wulf is president of the National Academy of Engineering.

The Institute of Medicine was established in 1970 by the National Academy of Sciences to secure the services of eminent members of appropriate professions in the examination of policy matters pertaining to the health of the public. The Institute acts under the responsibility given to the National Academy of Sciences by its congressional charter to be an adviser to the federal government and, upon its own initiative, to identify issues of medical care, research, and education. Dr. Kenneth I. Shine is president of the Institute of Medicine.

The National Research Council was organized by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purposes of furthering knowledge and advising the federal government. Functioning in accordance with general policies determined by the Academy, the Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in providing services to the government, the public, and the scientific and engineering communities. The Council is administered jointly by both Academies and the Institute of Medicine. Dr. Bruce Alberts and Dr. William A. Wulf are chairman and vice chairman, respectively, of the National Research Council.

This work was performed under Department of Navy Contract N00014-93-C-0089 issued by the Office of Naval Research under contract authority NR 201-124. However, the content does not necessarily reflect the position or the policy of the Department of the Navy or the government, and no official endorsement should be inferred.

The United States Government has at least a royalty-free, nonexclusive, and irrevocable license throughout the world for government purposes to publish, translate, reproduce, deliver, perform, and dispose of all or any of this work, and to authorize others so to do. Copyright 1997 by the National Academy of Sciences. All rights reserved.

Copies available from:
Naval Studies Board
National Research Council
2101 Constitution Avenue, N.W.
Washington, D.C. 20418
Printed in the United States of America

PANEL ON COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE

Ruth M. Davis, Pymatuning Group, Inc., *Chair*
Walter R. Beam, Alexandria, Virginia
George Cybenko, Dartmouth College
Steven K. Feiner, Columbia University
W. Michael McCracken, Georgia Institute of Technology
Brian P. McCune, McCune & Associates
Raj Reddy, Carnegie Mellon University
Victor Vyssotsky, Digital Equipment Corporation

Navy Liaison Representatives

Paul G. Blatch, Office of the Chief of Naval Operations (N911T1)
Ronald N. Kostoff, Office of Naval Research

Consultant

Sidney G. Reed, Jr.

NAVAL STUDIES BOARD

David R. Heebner, Science Applications International Corporation (retired), *Chair*
George M. Whitesides, Harvard University, *Vice Chair*
Albert J. Baciocco, Jr., The Baciocco Group, Inc.
Alan Berman, Center for Naval Analyses
Norman E. Betaque, Logistics Management Institute
Norval L. Broome, Mitre Corporation
Gerald A. Cann, Raytheon Company
Seymour J. Deitchman, Institute for Defense Analyses (retired)
Anthony J. DeMaria, DeMaria ElectroOptics Systems, Inc.
John F. Egan, Lockheed Martin Corporation
Robert Hummel, Courant Institute of Mathematical Sciences, New York University
David W. McCall, AT&T Bell Laboratories (retired)
Robert J. Murray, Center for Naval Analyses
Robert B. Oakley, National Defense University
William J. Phillips, Northstar Associates
Mara G. Prentiss, Jefferson Laboratory, Harvard University
Herbert Rabin, University of Maryland
Julie JCH Ryan, Booz, Allen and Hamilton
Harrison Shull, Naval Postgraduate School (retired)
Keith A. Smith, U.S. Marine Corps (retired)
Robert C. Spindel, Applied Physics Laboratory, University of Washington
David L. Stanford, Science Applications International Corporation
H. Gregory Tornatore, Applied Physics Laboratory, Johns Hopkins University
J. Pace VanDevender, Prosperity Institute
Vincent Vitto, Lincoln Laboratory, Massachusetts Institute of Technology
Bruce Wald, Arlington Education Consultants

Navy Liaison Representatives

Paul G. Blatch, Office of the Chief of Naval Operations (N911T1)
Ronald N. Kostoff, Office of Naval Research

Staff

Ronald D. Taylor, Director
Peter W. Rooney, Program Officer
Susan G. Campbell, Administrative Assistant
Mary (Dixie) Gordon, Information Officer
Christopher A. Hanna, Project Assistant

COMMISSION ON PHYSICAL SCIENCES, MATHEMATICS, AND APPLICATIONS

Robert J. Hermann, United Technologies Corporation, *Co-Chair*

W. Carl Lineberger, University of Colorado, *Co-Chair*

Peter M. Banks, Environmental Research Institute of Michigan

Lawrence D. Brown, University of Pennsylvania

Ronald G. Douglas, Texas A&M University

John E. Estes, University of California at Santa Barbara

L. Louis Hegedus, Elf Atochem North America, Inc.

John E. Hopcroft, Cornell University

Rhonda J. Hughes, Bryn Mawr College

Shirley A. Jackson, U.S. Nuclear Regulatory Commission

Kenneth H. Keller, University of Minnesota

Kenneth I. Kellermann, National Radio Astronomy Observatory

Margaret G. Kivelson, University of California at Los Angeles

Daniel Kleppner, Massachusetts Institute of Technology

John Kreick, Sanders, a Lockheed Martin Company

Marsha I. Lester, University of Pennsylvania

Thomas A. Prince, California Institute of Technology

Nicholas P. Samios, Brookhaven National Laboratory

L.E. Scriven, University of Minnesota

Shmuel Winograd, IBM T.J. Watson Research Center

Charles A. Zraket, Mitre Corporation (retired)

Norman Metzger, Executive Director

About this PDF file: This new digital representation of the original work has been recomposed from XML files created from the original paper book, not from the original typesetting files. Page breaks are true to the original; line lengths, word breaks, heading styles, and other typesetting-specific formatting, however, cannot be retained, and some typographic errors may have been accidentally inserted. Please use the print version of this publication as the authoritative version for attribution.

Preface

The Panel on Computer Science and Artificial Intelligence was established in early 1992 by the Naval Studies Board (NSB) of the National Research Council in response to a request from the Naval Research Laboratory (NRL). The NRL called on the NSB to establish a panel of experts in the fields of computer science and artificial intelligence to interact informally with the laboratory's research staff regarding plans, facilities, capabilities, prospects, and problems faced by the represented communities (government, academia, and industry) and to address the tasks contained in the formal terms of reference. Those tasks delineated the following priority topics:

- *Software production.* Provide a critical examination of the scientific issues that could be pursued (as opposed to commercial developments that are likely to evolve without fundamental R&D activities).
- *Adaptive software (machine learning).* What opportunities are envisioned in this area?
- *Interface technology.* What are the scientific issues that are to be examined over the next decade and which hold the potential for providing an improved base for sound developments in this field? What centers of activity worldwide are currently active leaders in such activities?
- *Speech synthesis/recognition.* What fundamental pursuits are likely to lead to further principles in this field (as opposed to commercial developments, which will indeed provide more sophisticated systems under many conditions)?
- *Neural networks.* Which scientific pursuits are required to place the behavior of these systems on a firm basis? Where are the leadership roles associated with this set of issues?
- *Facilities.* What facilities are the highest priority to emphasize in furthering the unique strengths that a government laboratory brings to this field? Which facilities are most appropriate at a university?

In addition, communications between Timothy Coffey, Director of Research at the Naval Research Laboratory, and the panel's chair, Ruth M. Davis, provided further clarification concerning the terms of reference and the objectives of the panel: The panel was to provide an outside perspective on key scientific and technical topics and to highlight technical opportunities for NRL. In particular, the panel was asked for such perspectives in the areas of artificial intelligence (AI) and human-computer interface (HCI). The panel was also requested to assess industrial interest in recommended opportunity areas.

The panel was not asked to conduct a critical review of NRL's current research program in computer science and AI/HCI and did not do so.

The panel performed its task in the course of four meetings: March 26-27, 1992, at which discussions were held with Timothy Coffey and presentations were made by NRL's Information Technology Division and Human-Computer Interface Laboratory; June 24-25 and September 28-29, 1992, at which presentations were made by government and academic experts on computer science and AI research programs; and October 21-22, 1992, which included further discussions with Timothy Coffey.

NRL has a considerable investment in computer science and is a node of the U.S. High Performance Computing Consortium. Its AI research laboratory was established some 20 years ago and is the only government laboratory of its kind.

About this PDF file: This new digital representation of the original work has been recomposed from XML files created from the original paper book, not from the original typesetting files. Page breaks are true to the original; line lengths, word breaks, heading styles, and other typesetting-specific formatting, however, cannot be retained, and some typographic errors may have been accidentally inserted. Please use the print version of this publication as the authoritative version for attribution.

Contents

Chapter 1-	Introduction	1
	Content	1
	Selected Definitions	1
	Grand Challenge Areas	2
Chapter 2-	Artificial Intelligence and Human-Computer Interface	4
	Importance of AI and HCI to the Department of Defense	4
	Types of Investment Required for AI and HCI	4
	General Recommendations Regarding AI and HCI	5
Chapter 3-	Grand Challenge Areas	7
	Representation and Modeling of Complex Systems	7
	Collaborative Problem Solving	8
	Machine Learning and Adaptive Systems	11
	Reasoning Under Uncertainty	12
	Virtual Worlds (Reality)	15
	Neurophysiological Models of Cognition	17
	Industrial Interest in Grand Challenge Areas	18

About this PDF file: This new digital representation of the original work has been recomposed from XML files created from the original paper book, not from the original typesetting files. Page breaks are true to the original; line lengths, word breaks, heading styles, and other typesetting-specific formatting, however, cannot be retained, and some typographic errors may have been accidentally inserted. Please use the print version of this publication as the authoritative version for attribution.

Chapter 1

Introduction

CONTENT

The focus of this report is on artificial intelligence (AI) and human-computer interface (HCI) technology. Observations, conclusions, and recommendations regarding AI and HCI are presented in terms of six grand challenge areas which serve to identify key scientific and engineering issues and opportunities. [Chapter 1](#) presents the panel's definitions of these and related terms. [Chapter 2](#) presents the panel's general observations and recommendations regarding AI and HCI. Finally, [Chapter 3](#) discusses computer science, AI, and HCI in terms of the six selected "grand challenge" areas and three time horizons, that is, short term (within the next 2 years), midterm (2 to 6 years), and long term (more than 6 years from now) and presents additional recommendations in these areas.

SELECTED DEFINITIONS

The scientific and engineering disciplines that make up or are closely allied with AI and HCI are of relatively recent origin. They date back to the 1950s at the earliest, with the first encyclopedia article on AI appearing in 1967. AI and HCI have been rapidly changing and expanding both in intellectual content and in application, especially in the last decade. The recent accelerated pace of change has been due in no small part to the almost breathtaking innovations and product developments in the supporting microelectronics, electrooptics, and display technologies—all hardware-intensive areas.

Growth in the amount of new terminology and technical jargon always accompanies such changes and advances in technology. Recognizing the futility of attempts to craft comprehensive definitions in light of these rapid changes, the panel has opted to provide brief descriptions of four frequently used terms: artificial intelligence, human-computer interface, virtual worlds, and synthetic environments.

Artificial Intelligence

Artificial intelligence is the collection of computations that at any time make it possible to assist users to perceive, reason, and act. Since it is computations that make up AI, the functions of perceiving, reasoning, and acting can be accomplished under the control of the computational device (e.g., computers or robotics) in question.¹

AI at a minimum includes

- Representations of "reality," cognition, and information, along with associated methods of representation;
- Machine learning;
- Representations of vision and language;
- Robotics; and
- Virtual reality (defined below).

Human-Computer Interface

Human-computer interface (HCI) consists of the following:

- The machine integration and interpretation of data and their presentation in a form convenient to the human operator or user (i.e., displays, human intelligence emulated in computational devices, and simulation and synthetic environments).²
- The bidirectional communication of information between two powerful

¹ Modification of the definition presented in Patrick Winston, *Artificial Intelligence*, 3rd Edition, Addison-Wesley, Reading, Mass., 1992, p. 5.

² *Defense Science and Technology Strategy*, Office of the Secretary of Defense, Defense Research and Engineering, Washington, D.C., July 1992 (extracts).

information processors: people and computers. Information can be in the form of data, symbolic knowledge, or control specifics.³

Virtual Worlds

A virtual world, or virtual reality, is a precise re-creation of a real-world environment via multisensory data and computer graphics that allows interaction between humans and synthesized objects. It consists of a set of multisensory devices employed as both actuators and effectors. "Virtual" is often used synonymously with computer-generated or synthetic.

Synthetic Environments

A synthetic environment is a reconstructed multipurpose environment with a mix of real and computer-synthesized (simulated) objects under computer control. It allows interaction between combinations of real and synthesized objects. A synthetic environment consists of a digital and analog representation of a physical environment with specified fidelity and complexity and is scalable to any size and degree of complexity.

GRAND CHALLENGE AREAS

The grand challenge areas selected by the panel should be understood in the context of the usage of that term in the High Performance Computing and Communications (HPCC) 1991 Initiative of Alan Bromley, President Bush's Science Advisor, stated as follows: "The HPCC Program is driven by the recognition that unprecedented computational power and capability is needed to investigate and understand a wide range of scientific and engineering grand challenge problems. These are fundamental problems whose solution is critical to national needs."⁴

The panel tailored this concept to its work in the following manner: Grand challenge areas are those fundamental problem areas to which the application of scientific and engineering resources will yield much-needed improvements in capabilities and performance. They also serve to identify key scientific and engineering issues and opportunities.

In selecting the grand challenge areas, the panel applied the following two constraints:

- Only leading-edge technologies were considered, and
- Research and technology application communities had to believe that the grand challenges were susceptible to resolution and that their resolution would provide demonstrable value-added to nontrivial user groups.

The panel chose the following six grand challenge areas:

- Representation and modeling of complex systems,
- Collaborative problem solving,
- Machine learning and adaptive systems,
- Reasoning under uncertainty,
- Virtual worlds (reality), and
- Neurophysiological models of cognition.

Panel members concurred that organizing their observations and findings concerning the NRL priority topics in the terms of reference (see the preface) according to the grand challenge areas selected would bring focus to their work and would permit easy presentation of highlights, issues, and high-value applications. [Table 1.1](#) summarizes the relationship of the NRL priority topics to the grand challenges.

[Table 1.1](#) can be read as follows: Software production, particularly for increasingly complex systems, presents challenges in the extensive collaborative and sometimes

³ NRL Human Computer Interaction Laboratory information brochure, 1991.

⁴ Committee on Physical, Mathematical and Engineering Sciences, Federal Coordinating Council for Science, Engineering and Technology, *Grand Challenges: High Performance Computing and Communications*, supplement to the President's Fiscal Year 1992 Budget.

concurrent problem-solving effort involved, and in realistic representation and modeling.⁵ Software for complex systems also presents challenges in reasoning under uncertainty.⁶ Adaptive software presents challenges in machine learning and reasoning under uncertainty. Interface technology develops collaborative (with the user) problem solving and virtual reality. Speech synthesis and recognition present challenges in collaborative problem solving, machine learning/adaptive systems, reasoning under uncertainty, and virtual worlds. Neural networks present challenges in machine learning, reasoning under uncertainty, and neurophysiological models of cognition.

Table 1.1 Linkages Between Grand Challenge Areas and NRL Priority Topics.

GRAND CHALLENGE AREAS	NRL PRIORITY TOPICS
Representation and modeling of complex systems	Software production
Collaborative problem solving	Interface technology, speech synthesis/recognition
Machine learning and adaptive systems	Adaptive software, speech synthesis/recognition, neural networks
Reasoning under uncertainty	Adaptive software, speech synthesis/recognition, neural networks
Virtual worlds (reality)	Interface technology, speech synthesis/recognition
Neurophysiological models of cognition	Neural networks

The sixth topic in the terms of reference was facilities, and the related questions were "What facilities are the highest priority to emphasize in furthering the unique strengths that a government laboratory brings to this field? Which facilities are the most appropriate at a university?" This topic is addressed in the presentation of each grand challenge area to the extent possible on the basis of the panel's collective experience.

⁵ Computer Science and Technology Board, National Research Council, *Scaling Up: A Research Agenda for Software Engineering*, National Academy Press, Washington, D.C., 1989, p. 4.

⁶ John E. Hopkins and Kenneth Kennedy, *Computer Science, Achievements and Opportunities*, Society for Industrial and Applied Mathematics, Philadelphia, Pa., 1989, p. 73.

Chapter 2

Artificial Intelligence and Human-Computer Interface

IMPORTANCE OF AI AND HCI TO THE DEPARTMENT OF DEFENSE

From the time of the first embryonic computer developments in the late 1940s and early 1950s, DOD has encouraged and sponsored AI and HCI R&D in the laboratories of industry, academia, and government. Advancements in AI and HCI have allowed DOD to accelerate work toward its goal of improving national security while reducing the risk faced by individuals in hostile environments.

This goal has manifested itself in strong support for and pressure to advance technologies that

- Permit total replacement of humans in hostile environments by "satisfactory" perceiving, reasoning, and acting mechanical-electronic surrogates;
- Permit combinations of robots and humans to carry out needed functions in hostile environments with minimal or much reduced risk to the human;
- Permit remote control of robots or human-robot systems in hostile environments to substantially reduce the risk to the human;
- Improve dramatically via education and training the capability of humans to perform satisfactorily in hostile environments, principally via simulation, emulation, and modeling processes and procedures; and
- Expedite the design, development, and manufacture of products, systems, and platforms of high quality in a timely manner and at least cost via concurrent engineering, computer-aided design (CAD), computer-aided manufacturing (CAM), flexible computer-integrated manufacturing (FCIM), and other automated design and production technologies.

The hardware, software, reasoning, and representation technologies that make up AI and HCI have played a significant role in approaching this strategic DOD goal. Relevant technology advances include

- Autonomous, unmanned vehicles (AUVs),
- Expert systems,
- Distributed product design,
- "Smart" weapons and platforms,
- "Lights out" factories,
- Real-time space asset command and control (C^2),
- SIMNET and other performance-enhancing simulation, and
- Real-time remote sensor, weapons, and platform management.

With the termination of the Cold War, and the resulting reduction in budget, force levels, and manpower, DOD has begun placing greater emphasis on technologies such as those just mentioned in order to leverage remaining resources.

TYPES OF INVESTMENT REQUIRED FOR AI AND HCI

NRL should be aware that two factors will influence the type of investment required for AI and HCI, as follows:

- AI research is more labor intensive than capital intensive, and HCI research is

more capital intensive than labor intensive.

- AI is dependent upon high-performance computer systems, which fortunately are rapidly decreasing in price and operating costs. Industry competition should continue to drive these costs down.

Industrial interest in AI and HCI can be measured in terms of their use in industrial operations and the extent of external (usually government) funding of industrial programs.

AI practitioners and users would concur that AI has an apparently solid future in academia and a bright, expanding future in applications. Today AI is pervasive in academic institutions, although its applications are primarily in industry and government. Well-known and widespread examples of AI applications include the following: expert systems, multisystem analysis, natural language processing, pattern recognition, robotic control, and computer vision.

HCI is rapidly acquiring all of the attributes of a scientific specialty area of computer science and, according to many experts, will shortly be perceived as such. The hardware resource base for HCI is shared fairly equally between academia and industry.

The skill base for AI and HCI is currently shared fairly equally between academia and industry. Government is primarily an HCI customer rather than a research source. The diversity of the required skill base for AI and HCI makes it difficult for a single research organization to sustain the minimum skill requirements. One attractive option is for NRL to form long-term arrangements with other research groups to pool skill bases in order to meet programmatic needs. This pooling is a cost-effective means for obtaining more current skills than are often available in a government laboratory with a stable professional population and a low attrition rate. Fortunately, the growing success of multimedia networking has essentially negated any justification for "owning" all needed skills at one facility. NRL is encouraged to enter into such strategic alliances via multimedia networking to maintain the necessary diversity in its AI and HCI skill base. It is further encouraged to establish joint projects in as many research or application areas as is feasible. However, the laboratory is cautioned that cooperative networking programs need to be explicitly defined, with formal, agreed-on objectives.

GENERAL RECOMMENDATIONS REGARDING AI AND HCI

The panel made a number of observations and reached a number of general recommendations regarding AI and HCI:

- NRL's research program in AI and HCI should be goal-oriented, with specific, substantive content, and not simply aimed at maintaining an "awareness" of the AI and HCI fields. The AI and HCI strategic planning process therefore would be more akin to a corporate model for planning than to an academic model.
- In support of the Navy's many unique and varied operational requirements, NRL should support and promote AI as a strong research component. Linking the strategic planning process to known Navy needs for AI and HCI technology, particularly during this period of restructuring and downsizing when operational leverage is paramount, will allow NRL to better assess its research and programmatic strengths alongside those of other defense research centers and laboratories.
- NRL management should manage and support AI as a recognized scientific specialty area within computer science. In doing so, it should be recognized that AI is still immature and is more "science" than engineering in nature.
- NRL should accelerate its HCI programs and support capital investment in HCI as a key research area for a defense laboratory.

- HCI is dependent upon customized electronic devices and components such as high-resolution displays and interactive electronics and often requires teleoperating and video devices. Customized HCI equipment rather than high computer power epitomizes HCI needs. Budgets for both AI and HCI should reflect the need for periodic equipment replacement and modernization due to the rapid obsolescence of supporting equipment.

Chapter 3

Grand Challenge Areas

After considerable discussion and debate, the following grand challenge areas were selected:

- Representation and modeling of complex systems,
- Collaborative problem solving,
- Machine learning and adaptive systems,
- Reasoning under uncertainty,
- Virtual worlds (reality), and
- Neurophysiological models of cognition.

Each of the six grand challenge areas is discussed in turn in this chapter. For the most part, these discussions include a description of the grand challenge area; research objectives over the short term (within the next 2 years), midterm (2 to 6 years), and long term (more than 6 years from now); ongoing leading-edge activities and organizations; needed resources delineated in terms of research skills, facilities, cooperative arrangements and opportunities, and level of effort; and opportunities for NRL.

Some subjects listed above may not be included in all discussions. The reader is referred to [Table 1.1](#) and its accompanying discussion in [Chapter 1](#) for an explanation of the cross-linkage of the grand challenge areas to the NRL priority topics.

The panel recommends that NRL give highest priority to investment and applications in the six grand challenge areas.

The panel's informal assessment of industrial interest in the six grand challenge areas is given at the end of this chapter.

REPRESENTATION AND MODELING OF COMPLEX SYSTEMS

Description

Prior to the advent of the computer, traditional engineered systems (e.g., bridges, airplanes, and ships) were modeled by a mathematical formalism combined with a set of engineering heuristics for how to apply the theory to given situations. Today, we are faced with the design, construction, maintenance, and use of much more complex systems. They require representations of the problem that go beyond systems of equations and modeling techniques and require more than closed-form solving, numerical approximation, or numerical simulation. The dimensions of this complexity can be exemplified as follows:

- A system today is a heterogeneous system of subsystems, so that no one representation and modeling paradigm suffices.
- A system is typically controlled by a software-based supervisory system that is usually the single most complex subsystem and is best modeled in detail by discrete, logical models rather than continuous, physical ones.
- One or more human users are part of the overall system and require interactions with the other subsystems to allow for monitoring and control.

The grand challenge is to be able to model entirely in software a range of complex systems. Two specific application milestones that would validate this achievement are the following:

- The ability to design and evaluate in software a sophisticated new weapon system, including carrying platform, sensors, weapons, communications, control systems, and human decision makers. This "prototype before build" capability (smaller-scale examples of which include Boeing's use of computer-aided design (CAD), computer-aided software engineering

(CASE), and computer-aided manufacturing (CAM) in the design of the 777) would help the Navy to stay even with rapidly changing threats, technologies, priorities, and budgets.

- The building of a fully autonomous undersea or air vehicle to perform one mission with reasonable generality.

Research Objectives

Desirable research objectives in this area are as follows:

- *Short term*—Integration of discrete, continuous, and symbolic representations of models.
- *Midterm*—Validation of internal consistency and completeness and against external specifications; representation and explanation of complex models to humans; and computation-constrained, satisfactorily approximate solutions based on dynamic, variable-resolution models controlled by metamodels.
- *Long term*—Global optimization across complex models; and very high level languages for system design (including decision-support software) and theory of model design.

Ongoing Leading-Edge Activities and Organizations

Types of ongoing research activities in this area and organizations involved are as follows:

- Knowledge representation activities—Stanford University and Carnegie Mellon University,
- Simulation/modeling activities—Bolt, Beranek and Newman Laboratories,
- Software engineering activities—Defense Advanced Research Projects Agency (DARPA) knowledge-based systems application (KBSA) centers, Kestrel Institute, University of Southern California (USC)/ISI, Stanford University, Carnegie Mellon University, and Harvard University, and
- Hardware computer-aided design (CAD) activities—DARPA centers, Stanford University, and Carnegie Mellon University.

Needed Resources

The necessary ingredients for research in this area are summarized below:

- *Skills*—Expert systems, simulation and modeling, automated software design and synthesis, computer-aided design, and human physiology.
- *Facilities*—Major computational resources to execute and validate complex models: distributed workstations plus parallel computer available over the network could suffice.
- *Cooperation opportunities*—Multiple groups must collaborate to model complex systems (e.g., a ship) because the expertise never resides totally in one organization. Two or more research groups could collaborate by modeling different subsystems of a problem.
- *Level of effort*—For small, theoretical tasks, one person could make progress; large, demonstration-oriented tasks would require a group of at least five, for example, one engineering system researcher, one simulation/modeling researcher, one application expert, one knowledge engineer, and one programmer.

COLLABORATIVE PROBLEM SOLVING

Description

Collaboration or group work is generally described as a social process used to more effectively perform tasks. These tasks include problem solving, idea generation, decision making, and conflict resolution. Collaboration includes both formal and informal methods. Formal methods can be characterized as structured collaboration sessions with

predetermined agendas and anticipated results. In contrast, informal methods are generally ad hoc in nature and constitute much of the daily practice of performing tasks in organizations (e.g., one seeks help from a colleague on an attribute of a problem one is trying to solve). Many researchers are emphasizing informal collaboration as a model of problem solving in organizations. Quite often, formal methods are used because of the size of the constituency needed to solve the problem, the diversity of opinions on its potential solution, and the anticipation of adversarial interactions. Although formal and informal methods are very different models of interaction, in actual practice in organizations the methods are often intertwined.

Understanding how people solve problems is requisite to determining whether or not technologies can aid in the process. A number of studies comparing problem solving with and without supporting technologies have been made over the past several years. Many of them have found that formal meetings using group decision support technologies improved the productivity of groups, shortened the time to final decisions, and in some cases quantitatively improved the decisions themselves. These studies have mainly considered formal meetings, where participants were schooled in the technologies, and therefore some researchers have questioned the validity of these studies. Sociologists, psychologists, anthropologists, and cognitive scientists have been characterizing the tasks of groups, but several issues have not been studied in depth. For example, most task characterizations do not include group activities that are important to the conduct of tasks but are not directly related to the tasks, such as learning and interpersonal relationship building. In addition, there has been little study of the impact of collaboration technologies on the problem solving task itself (primarily owing to the current state of the technology).

Research Objectives

The grand challenge in collaborative problem solving is to provide

- More accurate descriptions of tasks and their interrelationships,
- A broader tool framework to support multiple decision-making models and dimensions,
- Understanding of the impact of group decision support system (GDSS) tools on the decision-making process,
- Understanding of the use of intelligent agents in the decision-making process, and
- Incorporation of corporate memory into the decision-making process.

Desirable research objectives can be summarized as follows:

- *Short term*—An example of GDSS is Ventana Corporation's Group Systems Tools, originally developed at the University of Arizona and now a commercial product. Numerous other commercial GDSS tools are in use in industry and government. Most of these commercial tools have been developed by using the literature as a guide, even though some of them were originally developed in research environments. At this stage of the technology, most of the systems are based on portrayals of commonly understood activities in organizations. Elements contributing to decision making, such as brainstorming and voting, are encapsulated in these tools, with the primary objective to preserve information and to force structure on the decision-making process.

Decision making has been characterized as having three dimensions: time (synchronous or asynchronous), locale (collocated or distributed), and group size. Research needs to be conducted to verify that these are sufficient parameters, or to add others in order to more accurately describe the impact of these dimensions on decision-making processes and then reflect that research in the technologies.

For example, does a set of tools designed for collocated groups function similarly for distributed groups with the addition of video and audio support, and do those additions properly rectify the lack of collocation?

- *Midterm*—Midterm research is needed to more properly formulate cognitive models of decision making in terms of the true interactions of humans and organizations, including studying the impact of the technologies themselves on the process. The question of the utility of formal processes also needs to be addressed. Researchers in organizational behavior and management should understand the real needs of formal processes of decision making, since these processes were developed to preserve information, force rigor on the process, and collocate stakeholders to ensure unanimity in the decision. The impact of decentralized organizations and the dissolution of hierarchical management techniques will have significant impacts on the technologies that were designed for hierarchical decision processes.
- *Long term*—Long-term research implications are anticipated to be quite complex. Questions that could be addressed in long-term research include the following: If decision making is supported by intelligent agents incorporated into the tools, what is the impact on the process (e.g., in biasing of decisions or change of resolution time)? If every element of the process and the results of decisions is preserved (including analysis of all interactions), will decisions ever be made, and if so, will they be more conservative? If elements of interaction can be appropriately modeled and described, will decision making revert to the highest levels of an organization, or, to put it another way, if decisions can be made by machine, will empowerment of the employees become a moot point?

Ongoing Leading-Edge Activities and Organizations

Current leading-edge research activities in this area are being carried out by the following organizations:

- Georgia Tech and the University of Arizona, Centers for Information Management Research;
- Xerox Palo Alto Research Center;
- Lotus Development Corporation;
- Aarhus University, Denmark;
- MCC;
- University of Michigan;
- Hewlett Packard Laboratories; and
- Indiana University.

A partial list of current major users of collaborative problem solving technology includes the following:

- Army, Information Systems Command;
- Many major corporations, including IBM, BellSouth, and Boeing; and
- Numerous universities, including the University of Arizona, Georgia Tech, Indiana University, and the University of Georgia.

Needed Resources

An effective research program in collaborative problem solving would require skills in the following areas: cognitive science, sociology, anthropology, computer science, organizational behavior, and management information systems.

Opportunities for NRL

A number of highly leverageable technologies are applicable to collaborative problem solving, for example, in meetings. Traditional meetings are often frustrating and unproductive because of factors such as lack of group consensus, poor notetaking, and low group participation. Many times, groups become sidetracked. In addition, some participants may have a hidden agenda, while others may be apprehensive about participating

because they are uncertain how their ideas or comments will be received. It has been found that if GDSS, also referred to as Electronic Meeting Systems, is used to apply information technology to the meeting environment, group productivity, efficiency, and satisfaction are improved.⁷ When this technology is used efficiently, meetings ultimately culminate in an effective decision-making process by the given group.

This technology can be used in any type of decision-making process, including procurement decisions, strategic planning activities, planning and implementing of total quality management (TQM) principles, system and software requirement definition and representation, and war gaming (both adversarial and cooperative).

To take best advantage of such technology opportunities, NRL would benefit from having or developing the following critical resources:

- Ability to conduct large-scale, long-term case studies;
- Ability to incorporate AI research results into collaborative problem-solving technologies; and
- Connections with major research universities working in the area (e.g., through cooperative shared agreements).

MACHINE LEARNING AND ADAPTIVE SYSTEMS

Description

A grand challenge for engineering is to design systems that can adapt to and learn from their operating environments. While learning and adaptation are basic capabilities in biological systems, they have been extremely difficult to implement in synthetic, engineered systems. Only recently have the memory and computational requirements necessary for minimal learning and adaptation become available. These advances have resulted in an explosion of ideas and research in machine learning.

Specific successful implementations of learning and adaptive systems exist in the areas of supervised learning (e.g., the ability to generalize from empirical input/output data), unsupervised learning (e.g., the ability to categorize unlabeled data using generic clustering concepts), and adaptive filtering and control (e.g., where plant structural models are known and plant dynamics need to be estimated).

In spite of successful applications and a growing body of theory, machine learning and adaptive systems suffer from having poor models of the underlying problems being solved. There are no universal criteria for characterizing one learning problem as being more difficult than another (e.g., in what quantitative measure is learning Japanese easier or more difficult than learning to discriminate sonar signals?). While some measures such as metric entropy have been proposed, they are theoretical constructs that cannot be applied directly to practical problems. The present situation, in which many methods such as neural networks, case-based reasoning, and more traditional statistical methods compete on the basis of hyperbole and performance on toy problems, cannot continue.

The grand challenge in machine learning and adaptive systems is to identify a scientific methodology and theory that characterize classes of practical learning and adaptation problems. These characterizations would then be used to evaluate performance, select solution methods, and eliminate heuristics as much as possible.

Progress in this area would result in better validation techniques and a guarantee of performance in the numerous autonomous applications for which learning and adaptation are currently proposed.

⁷ Nunamaker, J.F., Alan R. Dennis, Joseph S. Valacich, Douglas R. Vogel, and Joey F. George, Electronic Meeting Systems to Support Group Work, *Communications of the ACM*, 34(7), June 1991.

Research Objectives

In the three time frames of interest, several research objectives can be suggested:

- *Short term*—A number of Navy applications are suitable for solution by machine learning and/or adaptive methods. Apply various competing existing learning and adaptation techniques to these applications and build a database of performance results.
- *Midterm*—Study the performance database and empirically categorize which applications are best solved by which methods. Develop models for learning and adaptation problems that explain the database results. Develop the appropriate mathematical theory to support the models.
- *Long term*—Apply and correct the mathematical models that characterize learning problems on new applications to verify the theory.

Ongoing Leading-Edge Activities and Organizations

Types of leading-edge research activity and organizations involved include the following:

- Empirical evaluation of classification techniques (Lincoln Laboratories, R. Lippmann),
- Statistical interpretations of neural networks (Brown University, S. Geman),
- Statistical experiments on neural networks (Oxford University, D. Ripley),
- Structural properties of learning problems (AT&T Bell Laboratories, Holmdel, V. Vapnik), and
- Minimal description length learning criterion (IBM San Jose, J. Rissanen).

Needed Resources

For any effective research program in machine learning and adaptive systems, the types and combinations of resources required can be summarized as follows:

- *Skills*—Mathematicians, computer scientists, statisticians, engineers, and cognitive scientists need to cooperate in this effort.
- *Facilities*—Access to Navy-relevant applications is essential. Significant computer power is needed to solve realistic applications and perform simulations. Such facilities already exist at NRL.
- *Cooperative arrangements*—Arrangements to share data and approaches with other researchers working on similar empirical investigations are essential. Empirical data can be collected at different sites, providing efforts are coordinated.
- *Level of effort*—About three senior researchers and six assistants would be required.

Opportunities for NRL

NRL already has a number of efforts in machine learning and adaptive systems. Focusing those activities with the goal of discovering models for learning requires better coordination but not necessarily more resources. Overall, the field of machine learning is full of techniques but short on methodology. NRL can have significant impact on the scientific application of existing techniques on real applications problems.

REASONING UNDER UNCERTAINTY

Description

Modern problems, including many that face Navy planners, are increasingly complex. These complexities can be characterized as follows:

- Problems have broad, possibly global, scope;
- There are subtle, time-varying multiple objectives, and different constituencies

in the United States and abroad often have different sets of objectives;

- Some of the simplest and most desirable solutions are not attainable within the constraints that will exist; and
- Effective solutions may require involvement of forces and organizations outside of those that have traditionally cooperated.

Methods are available for computer-aided solution of even highly complex problems, if (1) they are governed by stable constraints or rules; (2) they are solvable using well-understood processes and resources, under control of the responsible manager; (3) they are analogous to similar situations within the knowledge or experience of the manager and organization; and (4) they are subject to little or no uncertainty in knowledge of the circumstances.

The grand challenges address aspects of the solving of complex problems, through representation and modeling, collaborative solution, machine learning, and virtual worlds. While few complex problems are without significant uncertainties, those uncertainties are commonly suppressed in the interest of making the problem easier, or, at best, the problem is solved for a range of parameters that are presumed to include the uncertainties to a satisfactory degree. While such methods are suitable when the nature of the solution is invariant to the uncertainty, they are *not* satisfying when uncertainty is so great that totally different approaches to solution could be required to meet the potential demands.

Uncertainty takes many forms in solving complex military problems. Usually, knowledge is fragmentary. One may know the classes and counts of ships constituting an enemy task force, but not the particular vessels, their state of readiness, ammunition supplies, and other critical information. On the other hand, even the information one supposedly has in hand regarding enemy resources may be erroneous. Erroneous assessments of one's own resources or of enemy resources have historically been common in military actions. Union General McClellan's belief that Confederate troops near Richmond in 1862 outnumbered his own, though quite incorrect, led to a reticence to attack and allowed Lee's much smaller army to chase him away from threatening Richmond. Further uncertainty may be imposed by deliberately deceptive actions—anything from decoy weapons to radio reporting of nonexistent activity. Overall, the battle manager must deal with uncertainties in everything from the capabilities of the opposable forces to the objectives of the enemy (or even of some of the forces reporting to him).

Not all of the uncertainty needs to be dealt with at the same time. In military activity, uncertainty prior to action often occurs in understanding the circumstances and objectives of an opponent or in assessing one's own capabilities and limitations. During military actions, a high level of confusion in reports of sightings and actions is not uncommon. The top-level realities of the drama are seldom perceived correctly by the individual players, and even the sequence of events is sometimes unclear from first-hand reports.

Preaction knowledge of the enemy is based on surveillance or military intelligence. Knowledge of one's own resources relies on complete, unbiased reporting of readiness by one's own commanders. During military action, new information arrives through the agency of command, control, and intelligence (C²I) systems.

It is probably easiest to find examples of uncertainty in intelligence activities. Examples abound from earlier wars; for example, the misreporting of cruisers as aircraft carriers (or the converse) was not uncommon in World War II. During the Cold War, continuing intelligence interest in the numbers and capabilities of Soviet intercontinental ballistic missiles (ICBMs), sea-launched ballistic missiles (SLBMs), and other high-technology weapons was thwarted by the secrecy of the Soviet Union.

It has already been established in DOD long-term intelligence activities that rule-based expert systems have utility for weapon system identification using familiar modes of observation. There is much less experience and

little confidence in approaches for reasoning under uncertainty in the near-real-time decisions of military tactical management. There, both the situation itself and reports about it are changing. If battle managers are to make rational decisions, they must reason under uncertainty and with time constraints.

An idealistic way of phrasing the grand challenge in this area is to provide means whereby a military commander can make optimal use of all available valid and timely information. Implied in this objective is not only the ability to recognize uncertainty, but also the ability to use redundancy in data from reliable independent sources to reduce that uncertainty.

This is a multifaceted problem area. The nature of its solutions may, likewise, be manifold. An approach in which a powerful computer suggests a specific action and adds "trust me!" is clearly not acceptable. The knowledge and wisdom of human decision makers must be incorporated in the outcomes. If there are many human experts involved, some form of "calibration" of their opinions and judgment is desirable.

Research Objectives

- *Short term*—In the short term, it is difficult to identify specific objectives. For one thing, the scope of the set of problems is so great that subsets of interest to the corporate Navy need to be defined. One of the possible subareas would be the application of *belief theory* (of which the Dempster-Shaffer approach may be the best known) to characterize the confidence of estimates. This could be used, for example, to calibrate intelligence estimates based on (belief) characteristics of the estimators. It could also be used by a decision maker to determine, in comparison to contemporary or historical norms, if his or her decision is aggressive or nonaggressive, timely (or precipitate or slow), and where its weakest aspects may lie.

For intelligence applications, applied research could address reduction of uncertainty in both automated and human reports. If, for example, the limitations of sensors are properly incorporated in fusing their data, more accurate assessment should be possible.

We have traditionally relied on human interpreters to scan sensor imagery. This too is a case of reasoning under uncertainty, since the objects of interest may be hidden or camouflaged and the image quality may be naturally limited by range, bandwidth, and/or atmospheric. Large differences are seen in human ability to recognize militarily significant objects in infrared imagery. However, whether the next logical step would be to provide automation aids for human or interpreters or to scan and analyze electronically is not clear. Although accuracy is uncertainty-limited in this application, credible research must deal with realistic constraints and objects.

A possible set of objectives can be delineated as follows for the midterm and long term:

- *Midterm*—Devise metrics appropriate to comparison of traditionally incommensurate information and devise computer-applicable representation algebra(s) that will permit manipulation of currently incommensurate sets of related data, which can probably reduce uncertainty in individual data, preferably using methods faster than enumeration.
- *Long term*—Devise computational methods that permit the combination of incomplete and/or incorrect data from multiple contemporary observations, context description, and historical experience to demonstrably reduce the

uncertainty in situations or events described by current data.

Ongoing Leading-Edge Activities and Organizations

Types of activities in this area, and organizations and individuals involved, include the following:

- Information fusion—TRW, and
- Evidence theory—George Mason University, D. Schum.

Needed Resources

For an effective research program, the following kinds of resources are needed:

- Mathematics expertise in representational and computational algebras,
- Computer application expertise in numerical analyses and computer problem solving,
- Military strategic or tactical data streams as well as business data suitable for use as source, and
- High-performance computing power.

These resources are available or accessible to NRL by cooperative arrangements.

VIRTUAL WORLDS (REALITY)

Description

The scientific and engineering area is concerned with the development of knowledge-based multimedia interfaces that support real-time intersection with true three-dimensional input and output devices.

The current state of the art in commercial user interfaces emphasizes the use of two-dimensional windows, with which users interact by using two-dimensional devices such as "mice." Even users of three-dimensional graphics workstations usually view graphics whose projections appear in two-dimensional windows and are manipulated under mouse control. Research in three-dimensional user interfaces addresses the use of interactive three-dimensional graphics, coupled with true three-dimensional stereo displays and three-dimensional interaction devices that monitor the user's movements in three-space, as well as three-dimensional audio and haptic feedback. The goal is to harness the physiological capabilities and training that enable us to perform physical tasks effectively in three dimensions and apply them to develop effective user interfaces for computer-based tasks.

For an information-intensive user interface to be effective for a wide range of situations and users, it must be able to

- Design and present information to people on-the-fly, using multiple output media and
- Understand user input couched in multiple input media. To meet these requirements, a system must
- Generate technical material in real time in each individual output medium (written text, speech, static graphics, and animation),
- Understand technical material in real time in each individual input medium with performance equal to that of a human expert, and
- Coordinate real-time generation and understanding of multimedia interactions with humans, combining multiple input and output media, with performance equal to that of a human expert.

Research Objectives

The set of suggested short-term, midterm, and long-term objectives considered in totality is as follows:

- Develop real-time operating system support for highly parallel asynchronous input (from large numbers of three-dimensional trackers) and output (to multiple display modalities).

- Build effective "augmented realities" that enrich the user's existing environment with additional information, merging synthesized material with what the user normally sees, hears, and feels—overlying or replacing it, as appropriate.
- Develop better hardware that produces high-quality, high-resolution, wide-field displays (graphics, sound, haptic, and temperature) and tracking (hand, body, and eye). For example, there is a key need for "see-through" displays whose images are overlaid on the environment to build "augmented realities." A general-purpose visual display technology must allow differential visual accommodation, corresponding to real and virtual objects at different distances in the same image. It must also be able to perform full visible-surface determination with all objects, both real and virtual. Virtual objects should be able to occlude real objects, and real objects should be able to occlude virtual ones.
- Determine how to map abstract task domains effectively to a three-dimensional environment in which it will be possible to visualize and manipulate objects in the domain.
- Determine how to take advantage of the richness of three-dimensional gesture to reduce reliance on icons to express actions in current user interfaces. For example, rather than moving an item to the "trash can," it may be disposed of by using an appropriate gesture.
- Determine how to ensure that three-dimensional user interfaces will be usable, especially in an environment that supports end-user programming and customization. The problem is that in a world of whole-body computer interaction, there may no longer be any distinction between human factors (as usually understood) and the human factors of computer interfaces. The existing hardware that limits capabilities (and that also limits mistakes) will be gone.
- Apply AI techniques (e.g., interactive knowledge-based generation and understanding) to design virtual worlds for visualization automatically.
- Design high-quality multimedia systems, but only after designing systems that function well in a single medium. Proceeding in this fashion is particularly important with regard to knowledge-based information presentation systems. Improvements are needed in the ability to perform high-quality generation and understanding in individual media. There is much work to be done in generation and understanding of individual media, ranging from those media that have long been explored by AI researchers (e.g., written text and speech) to less well-charted terrain (e.g., graphics, audio, and haptics).
- Develop methods to predict and evaluate presentation quality. The system should be able to predict the quality of a presentation in the course of designing it (and, on the basis of these predictions, to refine the presentation until it is adequate). This requires the ability to evaluate the presentation, estimating how it "will" affect the user (and evaluating the user's response, estimating how it "has" affected the user). The ability to evaluate the presentation makes possible time-quality tradeoffs. For example, in a crisis situation, a timely rush job might be preferred over a later, higher-quality presentation.
- Develop generation and understanding capabilities for temporal media (media in which information context is presented over time in a way that is controlled explicitly by the producer) in animation, speech, and audio. Issues here include how to phrase information (e.g., for maximal comprehension). For example, the ability must be developed

- to generate output and understand input that communicates complex temporal relations.
- Develop facilities for coordinated generation and understanding of multiple media. The key challenge is to ensure that material in different media reinforce rather than interfere with each other.

Ongoing Leading-Edge Activities and Organizations

Leading groups and their area of activities include the following:

- Brown University—three-dimensional user interfaces (visual displays);
- Columbia University—three-dimensional UI, virtual worlds, multimedia-VI (MM-VI), MM-UI;
- Deutsche Forschung Künstlichen Intelligenz (Saarbrücken)—MM-UI;
- IBM T.J. Watson—virtual worlds;
- NASA Ames—virtual worlds;
- University of North Carolina—virtual worlds, three—dimensional UI;
- University of Pennsylvania—three-dimensional UI, virtual worlds;
- University of Washington—virtual worlds; and
- Xerox Palo Alto Research Center—three-dimensional UI.

Needed Resources

In order to do effective research in this area, the following types of resources are required:

- *Skill base*—Computer scientists, cognitive scientists, electronic engineers, optoelectronic engineers, and application area specialists.
- *Equipment*—High-performance, three-dimensional workstations and three-dimensional interaction and display devices (e.g., graphics, sound, and touch).

These resources exist within or are accessible to NRL through cooperative arrangements.

NEUROPHYSIOLOGICAL MODELS OF COGNITION

Description

Current approaches to human-computer interfaces are largely based on traditional sensory capabilities, that is, sight, sound, and touch. There is much commercial and research activity exploring those interfaces. This grand challenge goes beyond those ideas and proposes to explore internal neurophysiological representations of knowledge with the ultimate goal of using such representations for direct low-level computer interactions with the human nervous system. By the same token, understanding the internal representation of knowledge will be likely to result in better computer implementation of learning, cognitive, and other intelligent functions.

This challenge is timely. The past two decades have witnessed significant progress in our understanding of the biological mechanisms for memory, learning, and sensory processing. Most of that progress has been at the low-level, neuronal level, whereas correlations with higher-level functions and representations useful for cognition and intelligence are not yet understood. Accelerating the study of those correlations will allow more direct human-computer interfaces to be implemented.

Notable examples of these ideas are already being explored. Preliminary studies of semiconductor chip implants into neuronal tissue are being studied for motor control interfaces (at Dartmouth Medical Center and Stanford University). Moreover, the use of the electroencephalographic (EEG)-type readings of brain activity has allowed researchers to interface thought patterns directly with computer input (at Fujitsu Laboratories in Japan). These are but two examples of researchers attempting to bridge the gap between low-level neural activity and higher-level functionality.

Continued progress will require collaboration among neurophysiologists,

computer scientists, electrical engineers, mathematicians, and psychologists. Success will enable the development of more efficient human-computer interfaces that occur at a lower level and the design of better-performing artificial intelligence systems.

Research Objectives

The following objectives are suggested for research in this area:

- *Short term*—Organize an interdisciplinary NRL team to select a specific cognitive function and/or knowledge representation. Models for that phenomenon are formulated. Experiments and equipment to test those models are designed.
- *Midterm*—Perform experiments to validate hypothesized models. Models are modified and retested. Development of novel interface technology coincides with experimentation.
- *Long term*—Implement technology to use verified models for low-level human-computer interfaces in Navy applications.

Ongoing Leading-Edge Activities and Organizations

Types of research and institutions and individuals involved include the following:

- Neurophysiological models of the ear (N. Kiang, Massachusetts General Hospital),
- Neural chip implants (J. Rosen, Dartmouth Medical School and White River Junction VA Hospital; G. Kovaks, Stanford University),
- Silicon retina and ear (C. Mead, California Institute of Technology),
- Brain wave-computer interface (Fujitsu Laboratories, Japan), and
- Functional electrical stimulation and cardiomyoplasty (various engineering and medical schools).

Needed Resources

For effective research in this area, the following kinds of resources are needed:

- *Skills*—This research is best carried out by teams whose members are versant in neuroscience, physiology, signal processing, control theory, computer science, and mathematical modeling and instrumentation.
- *Facilities*—An effort encompassing all aspects of the challenge would most likely consist of team members who have access to dedicated facilities. Wet laboratories for experimenting with live tissue might be best located in hospitals or medical schools. Modeling and computing facilities can be off-site. Fabrication of interface instruments will probably require machine room capabilities typically found at NRL and other engineering research laboratories.
- *Cooperative arrangements*—Arrangements to share data and approaches with other researchers working on similar empirical investigations are essential. Empirical data can be collected at different sites, providing coordination is made.
- *Level of effort*—Five senior scientists distributed over the technical areas, ten research associates, and five technicians are required for a sustained multiyear effort.

Opportunities for NRL

NRL has significant in-house expertise in neural networks, control, cognitive science, and instrumentation. With unique access to Navy applications, the laboratory can develop its research program in this direction and play a leading role in future opportunities as the field opens up.

INDUSTRIAL INTEREST IN GRAND CHALLENGE AREAS

The panel's informal assessment of

industrial interest in the six grand challenge areas is as follows:

- *Representation and modeling of complex systems*—considerable and increasing interest by industry,
- *Collaborative problem solving*—considerable and increasing interest by industry,
- *Machine learning and adaptive systems*—significant interest and investment by industry,
- *Reasoning under uncertainty*—some interest by industry but no trend apparent in industry,
- *Virtual Worlds (reality)*—significant interest and investment by industry and a growing consumer and commercial market envisioned by industry, and
- *Neurophysiological models of cognition*—little interest by industry.