


TP11
G865 7. 外借 8861112

AI in the 1980s and Beyond

An MIT Survey

Edited by W. Eric L. Grimson
and Ramesh S. Patil



E8861112

The MIT Press
Cambridge, Massachusetts
London, England

511104

PUBLISHER'S NOTE

This format is intended to reduce the cost of publishing certain works in book form and to shorten the gap between editorial preparation and final publication. Detailed editing and composition have been avoided by photographing the text of this book directly from the author's prepared copy.

© 1987 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

This book was printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

AI in the 1980s and beyond.

(The MIT Press series in artificial intelligence)

Includes index.

1. Artificial intelligence. I. Grimson, William
Eric Leifur. II. Patil, Ramesh S. III. Massachusetts
Institute of Technology. IV. Series.

Q335.A42 1987 006.3 87-3241

ISBN 0-262-07106-1

AI in the 1980s and Beyond

The MIT Press Series in Artificial Intelligence

Edited by Patrick Henry Winston and Michael Brady

Artificial Intelligence: An MIT Perspective, Volume I: Expert Problem Solving, Natural Language Understanding, Intelligent Computer Coaches, Representation and Learning edited by Patrick Henry Winston and Richard Henry Brown, 1979

Artificial Intelligence: An MIT Perspective, Volume II: Understanding Vision, Manipulation, Computer Design, Symbol Manipulation edited by Patrick Henry Winston and Richard Henry Brown, 1979

NETL: A System for Representing and Using Real-World Knowledge by Scott Fahlman, 1979

The Interpretation of Visual Motion by Shimon Ullman, 1979

A Theory of Syntactic Recognition for Natural Language by Mitchell P. Marcus, 1980

Turtle Geometry: The Computer as a Medium for Exploring Mathematics by Harold Abelson and Andrea diSessa, 1981

From Images to Surfaces: A Computational Study of the Human Early Visual System by William Eric Leifur Grimson, 1981

Robot Manipulators: Mathematics, Programming and Control by Richard P. Paul, 1981

Computational Models of Discourse edited by Michael Brady and Robert C. Berwick, 1982

Robot Motion: Planning and Control edited by Michael Brady, John M. Hollerbach, Timothy Johnson, Tomás Lozano-Pérez, and Matthew T. Mason, 1982

In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension by Michael G. Dyer, 1983

Robotics Research: The First International Symposium edited by Michael Brady and Richard Paul, 1984

Robotics Research: The Second International Symposium edited by Hideo Hanafusa and Hirochika Inoue, 1985

Robot Hands and the Mechanics of Manipulation by Matthew T. Mason and J. Kenneth Salisbury, Jr., 1985

The Acquisition of Syntactic Knowledge by Robert C. Berwick, 1985

The Connection Machine by W. Daniel Hillis, 1985

Legged Robots that Balance by Marc H. Raibert, 1986

Robotics Research: The Third International Symposium edited by O. D. Faugeras and Georges Giralt, 1986

Machine Interpretation of Line Drawings by Kokichi Sugihara, 1986

ACTORS: A Model of Concurrent Computation in Distributed Systems by Gul A. Agha, 1986

Knowledge-Based Tutoring: The GUIDON Program by William Clancey, 1987

Visual Reconstruction by Andrew Blake and Andrew Zisserman, 1987

AI in the 1980s and Beyond: An MIT Survey edited by W. Eric L. Grimson and Ramesh S. Patil, 1987

Series Foreword

Artificial intelligence is the study of intelligence using the ideas and methods of computation. Unfortunately, a definition of intelligence seems impossible at the moment because intelligence appears to be an amalgam of so many information-processing and information-representation abilities.

Of course psychology, philosophy, linguistics, and related disciplines offer various perspectives and methodologies for studying intelligence. For the most part, however, the theories proposed in these fields are too incomplete and too vaguely stated to be realized in computational terms. Something more is needed, even though valuable ideas, relationships, and constraints can be gleaned from traditional studies of what are, after all, impressive existence proofs that intelligence is in fact possible.

Artificial intelligence offers a new perspective and a new methodology. Its central goal is to make computers intelligent, both to make them more useful and to understand the principles that make intelligence possible. That intelligent computers will be extremely useful is obvious. The more profound point is that artificial intelligence aims to understand intelligence using the ideas and methods of computation, thus offering a radically new and different basis for theory formation. Most of the people doing artificial intelligence believe that these theories will apply to any intelligent information processor, whether biological or solid state.

There are side effects that deserve attention, too. Any program that will successfully model even a small part of intelligence will be inherently massive and complex. Consequently, artificial intelligence continually confronts the limits of computer science technology. The problems encountered have been hard enough and interesting enough to seduce artificial intelligence people into working on them with enthusiasm. It is natural, then, that there has been a steady flow of ideas from artificial intelligence to computer science, and the flow shows no sign of abating.

The purpose of this MIT Press Series in Artificial Intelligence is to provide people in many areas, both professionals and students, with timely, detailed information about what is happening on the frontiers in research centers all over the world.

Patrick Henry Winston

Michael Brady

Preface

In 1959, the first memoranda from the Artificial Intelligence Project, a joint project between MIT's Research Laboratory of Electronics and the Computation Center and the precursor to the present-day Artificial Intelligence Laboratory, were published. Scanning through these first memos provides a fascinating glimpse into the early history of a fledgling field. While many of the first memos dealt with aspects of the LISP programming language, which was being developed explicitly to support the AI Project, and which as a consequence is the second oldest active programming language after FORTRAN, one can already see the seeds of today's AI programs begin planted. For example, Memo 10 deals with a symbolic differentiation program – a problem that present day freshmen at MIT see in the second week of the first class in computer science. Memo 16 describes a Question-Answering System for English, and Memo 17 is entitled "Programs with Common Sense" perhaps heralding today's expert systems. Memo 41 describes several years work on a chess program, and memo 53 describes a system to recognize hand drawn characters. A similar glance at the early history of other pioneering AI centers shows the same thing. For example, a 1954 memo from Carnegie-Mellon University describes their chess machine, and a 1958 memo describes a heuristic program to do assembly line balancing in industrial plants.

If one can already find documents on natural language processing, game playing, vision, reasoning and learning in the late 50's and early 60's, critics of AI might well wonder what, if any, progress has been made in the last twenty-five years. In many cases, these early efforts looked at limited slices of the problem, and proposed initial solutions which have since undergone significant refinement, extension or replacement. Today's AI programs, while addressing many of the same problems, are much more sophisticated than their early predecessors, and have a much broader scope of application. As a consequence, AI is graduating from the university research lab to the marketplace. This has led to a broad based interest in AI, as evidenced by the recent spate of articles appearing in national news magazines and television documentaries, and job advertisements from virtually every major corporation for AI experts.


Given the surge of public interest in AI and its emergence into the commercial forum, Christine Simonsen of MIT's Seminar Office concluded

AI in the 1980s and Beyond

8861112

Contents

Series Foreword	vii
Preface	ix
Artificial Intelligence: A Perspective Patrick H. Winston	1
Knowledge-Based Systems: The View in 1986 Randall Davis	13
Expert Systems Tools and Techniques: Past, Present and Future Peter Szolovits	43
A Case Study on Evolution of System Building Expertise: Medical Diagnosis Ramesh S. Patil	75
Artificial Intelligence and Software Engineering Charles Rich and Richard C. Waters	109



Intelligent Natural Language Processing: Current Trends and Future Prospects Robert C. Berwick	155
Automatic Speech Recognition and Understanding Victor W. Zue	185
Intelligent Vision Michael Brady	201
Making Robots See and Feel W. Eric L. Grimson	243
Robot Programming and Artificial Intelligence Tomás Lozano-Pérez	281
Robot Hands and Tactile Sensing John M. Hollerbach	317
Autonomous Mobile Robots Rodney A. Brooks	343
Index	365

Artificial Intelligence: A Perspective¹

Patrick H. Winston

The primary goal of Artificial Intelligence is to make machines smarter. The secondary goals of Artificial Intelligence are to understand what intelligence is (the Nobel laureate purpose) and to make machines more useful (the entrepreneurial purpose). Defining intelligence usually takes a semester-long struggle, and even after that I am not sure we ever get a definition really nailed down. But operationally speaking, we want to make machines smart.

The typical big-league, artificial-intelligence laboratory, and there are many of them now, will be involved in work like that shown in Figure 1. We at the MIT Artificial Intelligence Laboratory work in robotics, a field spanning manipulation, reasoning, and sensing. We do research in learning, language, and what some people call expert systems, something that I prefer to call design-and-analysis systems, by virtue of the common misuse of the term *expert systems*. We are also involved in issues basic to Computer Science, such as programming and computer architecture.

The Past: Six Ages

The history of Artificial Intelligence can be divided into a variety of ages, as shown in Figure 2. First is the prehistoric time, starting in 1842 when

¹This paper is the introduction to *The AI Business: Commercial Uses of Artificial Intelligence*, Patrick H. Winston and Karen A. Prendergast, editors, MIT Press, Cambridge, MA, 1984. It is reprinted, with minor revisions, by permission of the publisher.

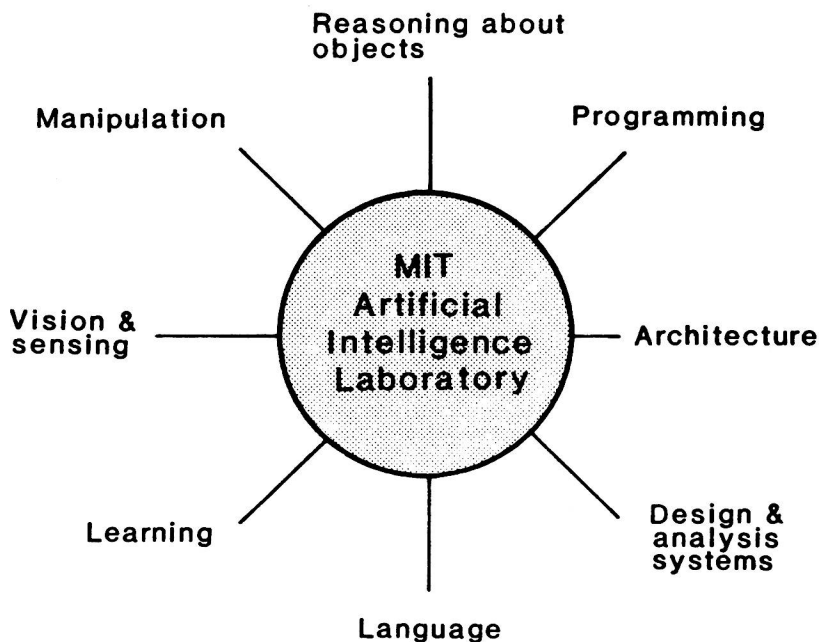


Figure 1. Subfields of Artificial Intelligence.

Charles Babbage first tinkered with his machines. Lady Lovelace, for whom the ADA programming language is named, was Babbage's main sponsor. She was besieged by the press, wondering if Babbage's machines would ever be as smart as people. At that time, she intelligently denied it would ever be possible. After all, if you have to wait for a hundred years or so for it to happen, it is best not to get involved.

The prehistoric times extended to about 1960 because the people who wanted to work on the computational approach to understanding intelligence had no computers. Still, people like Claude Shannon and John von Neumann made many speculations.

Around 1960 we start to speak of the Dawn Age, a period in which some said, "In ten years, they will be as smart as we are." That turned out to be a hopelessly romantic prediction. It was romantic for interesting reasons, however. If we look carefully at the early predictions about Artificial Intelligence, we discover that the people making the predictions were not lunatics, but conscientious scientists talking about real possibilities.




Figure 2. Ages of Artificial Intelligence.

They were simply trying to fulfill their public duty to prepare people for something that seemed quite plausible at the time.

The Dawn Age was sparked by certain successes. A program for solving geometric analogy problems like those that appear on intelligence tests was developed. Another was a program that did symbolic integration, spawning today's MACSYMA and other mathematics manipulation systems. These two examples, integration and analogy, are particularly worth noting because they introduced ideas that have become extraordinarily popular in the creation of expert systems. Retrospectively, the analogy program was based on the paradigm of describe-and-match, and the inte-

gration program was based on the paradigm of if-then rules.

I call the next period the Dark Period because little happened. There was a dry spell because the tremendous enthusiasm generated by the Dawn Age made everyone think that the enterprise of creating intelligent computers would be too simple. Everyone searched for a kind of philosopher's stone, a mechanism that when placed in a computer would require only data to become truly intelligent. The Dark Age was largely fueled by over-expectation.

Then we had a Renaissance. During this Renaissance people doing Artificial Intelligence began to make systems that caught people's eyes. MYCIN and other systems developed during this period are the harbingers of today's excitement.

The Renaissance was followed by the Age of Partnerships, a period when researchers in Artificial Intelligence began to admit that there were other researchers, particularly linguists and psychologists, with whom people working in Artificial Intelligence can form important liaisons.

I like to call our present age the Age of the Entrepreneur.

If there were substantial ideas about how to do impressive things as early as 1960, why has it taken until the 1980s to talk about how Artificial Intelligence might be commercialized?

The Successes

Let us agree that something has to be well known and in daily use to be successful. By this definition, there are only a handful of successful systems clearly containing artificial-intelligence technology.

One of the most conspicuous successes is the XCON system (also known as R1) developed by Digital Equipment Corporation and Carnegie-Mellon University for doing computer configuration. Others are DENDRAL and PUFF, products of Stanford University, developed for analyzing mass spectrograms and for dealing with certain lung problems. Still others include General Motors' CONSIGHT system and Automatix's AUTOVISION^R II, both of which endow increasingly intelligent robots with a limited but important ability to see.

Other successes are less domain specific. One, a product of Artificial Intelligence Corporation, is INTELLECT, a natural language interface system. Another is MACSYMA, a giant system for symbolic mathematics

developed at the Massachusetts Institute of Technology and marketed by Symbolics, Incorporated.

As I recently went over this list of successes with some friends, one pointed out that I had left out some of the most dramatic developments of Artificial Intelligence. One is the LISP programming language, a serious by-product of Artificial Intelligence. It is not surprising that the first major spinoffs of the MIT Artificial Intelligence Laboratory were two LISP Machine companies, Symbolics, Incorporated, and LISP Machine, Incorporated. If we go even further back, there are those who would argue that time-sharing was a major development that came out of Artificial Intelligence. Time-sharing is not Artificial Intelligence, but Artificial Intelligence demanded it.

Expert Systems

Human experts specialize in relatively narrow problem-solving tasks. Typically, but not always, human experts have characteristics such as the following: Human experts solve simple problems easily. They explain what they do. They judge the reliability of their own conclusions. They know when they are stumped. They communicate smoothly with other experts. They learn from experience. They change their points of view to suit a problem. They transfer knowledge from one domain to another. They reason on many levels, using tools such as rules of thumb, mathematical models, and detailed simulations.

An expert system is a computer program that behaves like a human expert in some useful ways. Today's state of the art is such that expert systems solve simple problems easily, occasionally explain their work, and occasionally say something about reliability.

Some expert systems do synthesis. XCON configures computers, for example. Other rule-based expert systems do analysis. MYCIN diagnoses infectious diseases, and the DIPMETER ADVISOR interprets oil well logs.

Currently, there are a dozen or two serious expert systems whose authors have commercial aspirations. By dropping the qualifier *serious*, the number grows to a few thousand. The reason is that creating a simple, illustrative expert system is now a classroom exercise in advanced artificial-intelligence subjects. Soon expert systems will be created in elementary courses in computing at the early undergraduate level.

All of this activity has attracted top-management interest, aroused the entrepreneurial spirit, and stimulated investor curiosity. Are the interest, the spirit, and the curiosity misadvised? It is too soon to be sure because few projects have had time to succeed and none has had time to fail.

Nevertheless there are some questions that can be answered, or at least debated. The list includes the following:

- Can today's technology revolutionize whole industries, or can it just deal with isolated, albeit important, targets of opportunity?
- Where are the most susceptible problems: engineering design, equipment maintenance, medicine, oil, finance?
- What are the obstacles to introducing expert systems: finding the right people, working with the existing human experts, getting snared by technically exciting but off-the-mark ideas?
- How hard will it be to build systems that exhibit more of the talents of real human experts?

Work and Play

A work station is a computer system that can be an exciting, productive partner in work or play. To be a good work station, a computer system must offer many features. First, we must be able to talk to the computer system in our own language. For some systems that language must be English or another natural language; for other systems the language must be that of transistors and gates, or procedures and algorithms, or notes and scales. Second, we must be able to work with the computer system the way we want to, not necessarily the way dogma dictates. In engineering design, for example, some people work bottom up; others prefer to work top down; still others work middle out or back and forth. All should be accommodated. Third, the computer system must constitute a total environment. Everything we need should be smoothly accessible through the system, including all the necessary computational tools, historical records, and system documentation. And fourth, the computer system's hardware must be muscular and the graphics excellent.

Some existing work-station products, like Daisy Systems Corporation's LOGICIAN and GATE MASTER, are extraordinarily important in the design of extremely complicated integrated circuits, often containing tens of thousands of transistors. Another work-station-oriented product, Artificial

Intelligence Corporation's INTELLECT, is not so domain oriented. INTELLECT is designed to be a powerful interface between decision makers and whatever data bases they need to work with. While INTELLECT began as a natural language interface, it is becoming the hub of a multitool, multifile information system, with much of the power residing in the parts having no direct concern with English input and output.

Companies such as Daisy Systems Corporation and Artificial Intelligence Corporation may be merely among the first flow of a potential cornucopia. People are developing work stations for such diverse activities as tax planning, chemical synthesis, robot assembly, musical composition, expository writing, and entertainment.

Where are the likely early successes? Key questions in determining this include the following:

- How important is natural language interaction? What does it take to get natural language interaction?
- What constitutes a minimally muscular computer and minimally excellent graphics?
- How important is it for work-station modules to be able to explain what they do? How important is it for users to be able to intervene whenever they want?
- Who can design and build work stations with human-like intelligence? A dozen people? Any computer engineer willing to learn?

Robotics

An intelligent robot is a system that flexibly connects perception to action. Humans are examples of intelligent robots for the following reasons. First, we can see and feel forces. Consequently we can cope with uncertain positions and changing environments. Second, we have graceful arms capable of fast, accurate motion, together with fantastic hands capable of grasping all sorts of objects. Third, we think about what we do. We note and avoid unexpected obstacles. We select tools, design jigs, and place sensors. We plan how to fit things together, succeeding even when the geometries are awkward and the fits tight. We recover from errors and accidents.

In contrast, most of today's industrial robots are clumsy and stupid. For the most part they cannot see, feel, move gracefully, or grasp flexibly, and they cannot think at all. Most of today's industrial robots move repetitively through boring sequences, gripping, welding, or spraying paint at