

ADVANCED AUTOMATION FOR SPACE MISSIONS

Robert A. Freitas Jr., Space Initiative
100 Buckingham Drive, Suite 253A, Santa Clara, California 95051

Timothy J. Healy, Dept. Electrical Engineering & Computer Science
University of Santa Clara, Santa Clara, California 95053

James E. Long, Jet Propulsion Laboratory
4800 Oak Grove Drive, Pasadena, California 91103

ABSTRACT

A NASA/ASEE Summer Study conducted at the University of Santa Clara in 1980 examined the feasibility of using advanced artificial intelligence and automation technologies in future NASA space missions. Four candidate applications missions were considered. (1) An intelligent Earth-sensing information system, (2) an autonomous space exploration system, (3) an automated space manufacturing facility, and (4) a self-replicating, growing lunar factory. The study assessed the various AI and machine technologies which must be developed if such sophisticated missions are to become feasible by century's end.

1. INTRODUCTION

During the summer of 1980 eighteen educators from throughout the United States joined with fifteen NASA program engineers from six NASA centers to examine ways in which advanced automation, including artificial intelligence (AI) and robotics, might be used in space missions in the next several decades. The study was supported by the National Aeronautics and Space Administration (NASA) and the American Society for Engineering Education (ASER) because of an increasing realization that advanced automatic and robotic devices must play a major role in space exploration and utilization, and may provide enormously beneficial capabilities at affordable cost.

The 10-week, 10,000-man-hour project was hosted by the University of Santa Clara and co-directed by James Long and Timothy Healy. Team members were given the task of selecting and defining a number of representative space ventures which would benefit from or even require extensive application of machine intelligence and automation, and assessing existing and foreseeable AI technologies necessary to accomplish the proposed missions (Long and Healy, 1980). Major support in the field of artificial intelligence was provided by SRI International, and a number of industrial concerns made contributions in the area of system implementation. Further, the study received firm backing from NASA Headquarters personnel, including two unprecedented on-site personal visits by the agency administrator, Dr. Robert A. Frosch. The strong NASA support signals a new perception of the tremendous potential of artificial intelligence in future space missions.

For purposes of the summer study several specific areas within AI were taken as representative of highly sophisticated computer-based systems which may be required in future space applications, including: Planning and problem-solving; perception;

natural languages; expert systems; automation, teleoperation, and robotics; distributed data management; and cognition and learning (AAAI, 1980; Arden, 1980; Winston and Brown, 1979). The latter category includes the concept of an adaptive machine able to study a new situation or environment, formulate hypotheses about the environment, test those hypotheses with additional data, and decide whether or not a new hypothesis should be added to the machine's existing model of the environment. The 1980 study was predicated on the assumption that such machine capabilities will become available by the end of the present century.

NASA has a long record of using automation and computers in its space missions. However, in 1979 agency activities were examined by an ad hoc advisory committee and criticized for being "5 to 15 years behind the leading edge in computer science and technology" (Sagan, 1980). The committee also found that "the advances and developments in machine intelligence and robotics needed to make future space missions economical and feasible will not happen without a major long-term commitment and centralized, coordinated support." A part of NASA's response to this criticism was to commission the University of Santa Clara study, the results of which are briefly described herein.

2. MISSION DESCRIPTIONS

The study group divided into four teams in order best to focus its examination of possible space applications of artificial intelligence. These teams were: Terrestrial Applications, Space Exploration, Non-Terrestrial Utilization of Materials, and Replicating Systems Concepts. Each team chose a challenging venture that could be used to identify critical technology needs for future research and development. A brief description of each mission is given here. A far more complete discussion may be found in an extensive Technical Summary of the summer study (Long and Healy, 1980). Major criteria for selection were that missions be implementable within a few decades and that they demonstrate a significant machine intelligence requirement. Novelty was not considered to be a factor of primary importance.

2.1 Terrestrial Applications: An Intelligent Earth-Sensing Information System

The Terrestrial Applications team concluded that artificial intelligence techniques would be most useful in near-Earth missions which generate data

at very high rates - such as the Tiros weather and Landsat terrestrial imaging satellites. One of the weaknesses of such systems is that they tend to return far more data to Earth than can ever be effectively used. The problem of culling a small amount of valuable information from a large volume of useless data is a task now left largely to users. The team proposed a new, highly versatile Intelligent Earth Sensing Information System (IESIS), able to perform substantial amounts of selection and interpretation of incoming sensor data and provide more useful information tailored to the needs of the individual user. A new philosophy of goal-oriented data collection, in which information is gathered to meet specific objectives (e.g., site, observation timing, sensor sets), was articulated by the team as the cornerstone of the proposed mission.

IESIS has the following major features: (1) An intelligent satellite network which gathers data in a goal-directed manner, based on specific requests for observation (such as a farmer requesting once-a-week surveillance of his cornfield) and on prior knowledge contained in a detailed, self-correcting "world model"; (2) a user-oriented "natural Language" interface which permits requests to be satisfied in plain English, without additional human intervention, using information retrieved from the system library or from direct observations made by a member satellite within the network; (3) a medium-level on-board decisionmaking capability that optimizes sensor utilization without compromising user requests; and (4) a library of stored information which provides a detailed set of all significant planetary features and resources, adjustable for seasonal and other identifiable variations, and accessible through a comprehensive cross-referencing system.

The heart of IESIS is, however, the world model, a self-correcting description of the environment under observation to any desired level of detail. This eliminates the need for acquiring and storing large quantities of redundant information by making use of two important AI elements: First, a "state component," which defines the physical state of the world to a predetermined accuracy and completeness at some specified time; and, second, a "theory component," which permits derivation of parameters of the world state not explicitly stored in the state component and allows forecasts of the time evolution of the state of the world. IESIS retains the complete Earth model in a ground-based central systems computer and an appropriate subset thereof on-board the main satellite. Orbiter sensors still collect extensive data, but only changes in the world model are downlinked, rather than the entire data stream. The result is an effective data compression system which removes redundancy over time.

2.2 Space Exploration: Titan Demonstration of a General-Purpose Exploratory System

NASA has been exploring the Solar System for more than twenty years. Such missions no doubt will continue in the decades ahead and may include visits to the more distant planets, moons, asteroids, comets and numerous yet-undiscovered bodies. It is conceivable that by the 21st century spacecraft may exit the Solar System in search of other planetary systems many light-years from Earth. The Space Ex-

ploration team thus identified interstellar navigation by automatic probe as an ultimate goal, but defined their overall study concept in terms of an autonomous general-purpose space exploration system.

One of the major problems with present interplanetary exploration strategies is that they typically require three distinct stages: Initial reconnaissance, exploration, and intensive study. Especially in the case of relatively distant bodies, the sequential character of the examination leads to inordinately lengthy total investigation times. The team concluded that the three stages can be telescoped into a single mission by incorporating advanced machine intelligence to produce a single integrated scientific phase of discovery. On-board AI systems are required to make certain initial decisions about sites to be explored in detail, the nature of the exploration, and the best ways to conduct intensive studies.

As a preliminary shakedown voyage for this new technology which could help pave the way for more ambitious exploratory ventures both within and beyond the Solar System, the team proposed a demonstration mission to Titan (the largest natural satellite of Saturn). This would be capable of independent operation starting from launch in low Earth orbit; navigation and propulsion system control during interplanetary transfer to Saturn; rendezvous with Titan and orbital insertion; automatic landing site decision-making; deployment of various subsatellites, landers, and fliers on and about Titan; and subsequent monitoring and control of atmospheric and surface exploration.

Of course, decisions about succeeding steps in the exploration of Titan could well be made directly by earth bound scientists since transmission delay time is only about one hour. However, when explorer craft are sent to other star systems the delay time will stretch to years, and decisions concerning successive stages of investigation must be made on-board the spacecraft. The purposes of the Titan mission arc to enhance the capabilities of semi-autonomous vehicles in the short-term, and to refine and demonstrate the effectiveness and versatility of fully autonomous exploration in the long-term.

A major finding of the study team was that automated hypothesis formation is highly desirable for sophisticated interplanetary missions within the Solar System but is absolutely essential for interstellar exploration. Machine intelligences capable of unassisted scientific and operational hypothesis formation must be able to handle three distinct classes of inferential thinking: (1) Analytic inference (application of existing scientific classification schemes), (2) inductive inference (logical processes for generating universal statements about an entire domain based on quantitative or symbolic information from a restricted part of that domain), and (3) abductive inference (a method for evolving new information classification schemes using old theories, old schemes, old predictions, and novel contradictory data as inputs). An important feature of the Titan spacecraft is that it would carry a world model of the target for exploration, the best available record of all pertinent features of the body in view of the research to be conducted. The probe would use its sensors to accumulate data about Titan, generate hypotheses about the sensed environ-

ment, test these hypotheses using new data, then update the scientific model as required.

2.3 Non-Terrestrial Utilization of Materials: Automated Space Manufacturing Facility

The Non-Terrestrial Utilization of Materials team studied the concept of a space manufacturing facility, initially located in low Earth orbit and using terrestrially-provided raw materials but constantly evolving towards ever-greater independence from Earth resupply (Miller and Smith, 1979; SAI, 1978). This would be a permanent, fully automated or tele-operated facility ultimately for the utilization of non-terrestrial matter retrieved from asteroids, the Moon, and other planets. Extensive use of robotics and teleoperation techniques are required for efficient manufacturing, repair, and for building new generations of equipment for expansion and diversification. Automatic and robotic devices must be used to control processors, manipulate and transport components, organize fabrication, and generally perform a wide variety of typical manufacturing tasks. Sophisticated new "telepresence" (Minsky, 1979, 1980) technology - full sensory feedback and effector control - might permit, say, an Earth-based excavation or construction worker to drive a tractor on the Moon from a comfortable terrestrial ground station or a command platform located in low Earth orbit (LEO).

Such a mission requires important advances in visual, tactile, and force sensors, machine decision-making, adaptability, mobility, and many other areas of AI technology. Rapid advancements now are being made in many of these fields in connection with the automation of factories on Earth. It is, however, in NASA's interest to promote additional directed research into problems of manufacturing unique to the space environment.

2.4 Replicating Systems Concepts: Self-Replicating Lunar Factory and Demonstration

Someday it may become necessary to build very large, massive space structures - perhaps to erect huge observational platforms, mine and utilize significant quantities of lunar or asteroidal resources, terraform satellite or planetary environments, or harness the energy of the Sun on a grand scale never before imagined. One way to build the large number of space factories needed to construct such systems rapidly is by the use of replicating systems. Such systems may grow exponentially, producing great masses of organized matter in a relatively short period of time (Freitas, 1980, 1981; von Tiesenhausen and Darbro, 1980).

The Replicating Systems Concepts team defined, as an ultimate challenge for advanced artificial intelligence and automation, a factory on the Moon which completely replicates itself using only lunar materials and solar energy. The basic concept of machine self-reproduction was shown theoretically feasible decades ago by John von Neumann (1966), but actual implementation will be extremely difficult. To arrive at a system capable of building all of its own components and then assembling them into an exact duplicate will require major advances in automated materials processing, computer-aided manufacturing and parts fabrication (CAD/CAM) technology),

robot assembly techniques, storage and inventory maintenance, inspection and repair capabilities, scheduling and other aspects of general factory management requiring very sophisticated AI techniques.

The central theoretical issue is closure. Can a real machine system itself produce and assemble all the kinds of parts of which it is comprised? In a generalized terrestrial industrialized economy manned by humans the answer clearly is yes (e.g., American industry), since the set of machines which make all other machines is a subset of the set of all machines. In space a few percent of total system mass - in particular those items most difficult to produce such as ball bearings, motors, or integrated circuits - could feasibly be supplied from Earth-based manufacturers as "vitamin parts." Alternatively, the system could be designed with components of very limited complexity (Heer, 1980). The minimum size of a self-sufficient "machine economy" remains unknown.

3. ARTIFICIAL INTELLIGENCE TECHNOLOGY ASSESSMENT

A principle goal of the summer study was to identify advanced AI, robotics, and automation technology needs for mission capabilities representative of desired NASA programs in the 2000-2010 time period. Six major classes of technology requirements derived during the mission definition phase of the project were identified as having maximum importance and urgency. These general classes of requirements individually were assessed by considering in each case the current state of the relevant technology; the specific technological goals to be achieved; and the developments needed to achieve these goals.

3.1 Autonomous world Model Generation

Perhaps the most important immediate requirement is the development of the technology necessary autonomously to map, manage, and re-instruct a world model based information system, at least a part of which is operating in space. Research is needed generally in the following areas:

- Techniques for autonomous management of an intelligent space system.
- Mapping and modeling criteria for creation of compact world models.
- Autonomous mapping from orbital imager.
- Efficient, rapid image processing based on comparisons with world model information.
- Advanced pattern recognition, signature analysis algorithms and multisensor data/knowledge fusion.
- Explicit models of system users.
- Fast, high-density computer systems suitable for space application of world model computations and processing.

3.2 Machine Learning and Hypothesis Formation

Learning, the process at the heart of the problem of developing "intelligent machines", is afield of AI requiring major research for space applications. More work is needed to develop the theoretical basis of machine learning, to identify which operations are essential to learning, and to deter-

mine how these functions might be implemented in hardware. A machine intelligence which learns must be capable of (1) formulating hypotheses which apply existing concepts to events and processes found in a new environment (analytic inference), and of (2) hypothesizing new concepts, laws and theories whenever existing ones are inadequate (inductive and abductive inferences).

Analytic inferences have received the most complete treatment within the AI research community. For example, rule-based expert systems can apply detailed diagnostic classification schemes to data on events and processes in some given domain and produce appropriate identifications. However, these consist solely of complex diagnostic rules describing the phenomena in some domain. They do not include models of the underlying physical processes of these phenomena. In general, state-of-the-art AI treatments of analytic inference fail to link detailed classification schemes with fundamental models required to deploy this detailed knowledge with maximal efficiency.

Inductive inferences receive a less complete treatment than analytic inferences, although some significant advances have been made. For instance, a group at the Czechoslovak Academy of Sciences has developed formal techniques for moving from data about a restricted number of members of a domain, to observation statement(s) which summarize the main features or trends of this data, to a theoretical statement which asserts that an abstractive feature or mathematical function holds for all members of the domain (Hajek and Havranek, 1978). Another research effort attempts to integrate fundamental models with specific abstractive, or generalizing, techniques. However, this work is at the level of theory development - a working system has yet to be implemented in hardware.

Abductive inference has scarcely been touched by the AI community. Tentative first steps have been taken, as for example current efforts in "non-monotonic logic" presented at a recent AI conference held at Stanford University (Israel, 1980; Weyrauch et al., 1980). These attempts to deal with the invention of new or revised knowledge structures are hampered (and finally undermined) by their lack of a general theory of abductive inference. One notable exception is the recent work of Frederick Hayes-Roth (1980), who takes a theory of abduction developed by Imre Lakatos for mathematical discovery and operationalizes two low-level members of Lakatos' family of abductive inferential types. Still, this work is but a preliminary step toward the implementation of workable systems of mechanized abduction.

3.3 Natural language and Other Man-Machine Communication

If Earth-sensing systems and remotely operated devices are to be used by a wide range of relatively unsophisticated users, substantial improvements are necessary in man-machine communication. This includes the development of high-level natural languages, machine generation and recognition of speech, visual perception and image generation.

In those instances in which the environment is highly restricted with respect to both the domain of discourse (semantics) and the form of appropriate statements (syntax), serviceable interfaces are

just possible with state-of-the-art techniques. However, any significant relaxation of semantic and syntactic constraints produces very difficult problems in AI. For general use the following capabilities are highly desirable, and probably necessary, for efficient and effective communication: Domain model, user model, (general, idiosyncratic, contextual), dialogue model, explanatory capability, and reasonable default assumptions.

Recognition and understanding of fluent spoken language adds further complexity to that of ordinary keyed language/phoneme ambiguity. In noise-free environments involving restricted vocabularies, it is possible to achieve relatively high recognition accuracy, though presently not in real time. In more realistic operating scenarios, oral fluency and recognition divorced from semantic understanding is not likely to succeed. The critical need is the coupling of a linguistic understanding system to the spoken natural language recognition process. On a related research front, the physical aspects of machine speech generation are ready for applications, although some additional "cosmetic" work is still necessary for general use.

Some motor-oriented transfer of information from humans to machines already has found limited use, such as lightpens, joysticks, and head-eye position detectors employed for military target acquisition. An interesting alternative for space applications is the "show and tell" approach. In this method a human manipulates an iconic model of the real environment. A robot "watches" these actions (perhaps complemented by some further information spoken by the human operator), then duplicates them in the real environment. Robot action need not be real time with respect to human operator action - the machine may analyze the overall plan, ask questions, and cooperatively optimize the original course of action. The operator plays the role of "editor" of the evolving robot program. Show and tell tasks can be constructed piecemeal, thus allowing a job to be described to the machine which requires many simultaneous and coordinated events. Finally, the fidelity of robot actions to the human example may vary in significant ways (e.g., size scale, mass scale, or speed of performance), allowing the machine to optimize the task in a manner alien to human thinking.

3.4 Space Manufacturing Automation

To achieve the goals of space manufacturing and replicative automation, space processing likely must progress from terrestrial simulations to low Earth orbit experimentation with space production techniques, and ultimately to processing lunar materials and other non-terrestrial resources into feedstock for more basic product development. There are four major components of any "universal" manufacturing system: (1) Extraction and purification of useful materials from raw resources, (2) fabrication of product components, (3) product component assembly, and (4) system control.

Extraction and purification technologies for processing raw materials on the lunar surface or elsewhere lie far beyond state-of-the-art. Sophisticated, highly automated chemical, electrical, and crystallization techniques must be developed yield-

ing a far broader range of elements and materials than is presently possible. Product component fabrication involves primary shaping and finishing operations. Shaping technologies of greatest utility for fully automatic space manufacturing are casting and powder metallurgy, both of which can produce parts ready for use without further processing. Elimination of manual mold preparation should be sought, possibly through the use of computer-controlled containerless forming. Laser and electron-beam techniques appear promising for highly automated finishing. Product component assembly requires robot/teleoperator vision and end-effectors which are smart, self-preserving, and dexterous. Placement accuracy of 1/1000th inch and repeatability of 5/10,000ths inch are desirable in electronics assembly tasks. High-capacity arms and multi-arm coordination must be developed.

Control of a large-scale space manufacturing or replicating system demands a distributed, hierarchical, dynamic, machine-intelligent information system. For inventory control, an automated storage and retrieval system well-suited to the space environment is necessary. The ability to gauge and measure products - inspection or quality control - is essential, and a general-purpose high-resolution AI vision module is needed for quality control of complex products and components. Advances in artificial intelligence should also include the embodiment of managerial and repair skills in an autonomous, adaptive-control expert system.

3.5 Teleoperators and Robot Systems

Teleoperators permit action or observation at a distance by human operators. Teleoperation represents a technology intermediate between fully manned and autonomous robot activity, having motor functions (commanded by the man) with many possible capabilities and sensors (possibly multiple, special-purpose) to supply information. The human being controls and supervises through a mechanical or computer interface. As technology advances and new requirements emerge, more and more of command/control functions must reside in the computer with the man assuming a more supervisory role. With the development of superior AI methods, computers eventually may perform "mental" functions of far greater complexity, eventually becoming virtually autonomous.

An important goal of teleoperator development must be to give the operator the ability to sense remote environments as realistically as possible, an effect termed "telepresence" by Minsky (1979, 1980). The capacity closely to relate action and reaction at the remote site and at the control room requires major advances in manipulators (coordination and cooperation of multiple manipulator arms and hands); force reflection and servoing; visual, audio, tactile, radar/proximity and other sensors; comprehension of variable conditions of scene illumination, wide or narrow viewing fields, and three-dimensional information via stereo displays, planar beams, or holograms; and systems to circumvent the disorienting effects of communication time delays in sensor/effector feedback loops.

Two other distinct classes of teleoperators may be required for complex, large-scale space operations such as a manufacturing facility or replicating factory. First is a free-flying system combin-

ing the technology of the Man Maneuvering Unit with the safety and versatility of remote manipulation. The free-flying teleoperator can be used for satellite servicing, stockpiling and handling materials, operations requiring autonomous rendezvous, station-keeping, and docking capabilities. Second, mobile or walking teleoperators may be useful in various manufacturing processes and for handling hazardous materials. The device would automatically move to a designated internal or external work site and perform either preprogrammed or remotely controlled functions. For manufacturing and repair such a system could transport astronauts to the site. Manipulators could be locally controlled for view/clamp/tool operations or as a mobile workbench.

3.6 Computer Science and Technology

NASA's role, both now and in the future, is fundamentally one of information, acquisition, processing, analysis, and dissemination. This requires a strong institutional expertise in computer science and technology. General computer science research avenues and capabilities required to implement the types of missions proposed by the study teams include: computer systems, software, management services, and computer systems engineering. State-of-the-art technology already is a part of NASA programs in the natural sciences, engineering, simulation and modeling. A substantial commitment to research in machine intelligence, real time systems, information retrieval, supervisory and computer systems is required.

4. CONCLUSIONS AND RECOMMENDATIONS

Advanced machine technology is essential in realizing a major space program capability for extra-terrestrial exploration and resource utilization within realistic temporal and economic limits. To this end, the study group reached the following general conclusions and recommendations:

- (1) Machine intelligence systems with automatic hypothesis formation capacity are necessary for autonomous examination of unknown environments. This capability is highly desirable for efficient exploration of the Solar System and is essential for the ultimate investigation of other star systems.
- (2) The development of efficient models of Earth phenomena and their incorporation into a world model based information system are required for a practical, user-oriented, Earth resource observation network.
- (3) A permanent manned facility in low Earth orbit is an important element of a future space program. Planning for such a facility should provide for a significant automated space manufacturing capability.
- (4) New, automated space materials processing techniques must be developed to provide long-term space manufacturing capacity without major continuing dependence on Earth resupply.
- (5) Replication of complex space manufacturing facilities is a long-range need for ultimate large-scale space utilization. A program to develop

and demonstrate major elements of this capability should be undertaken.

- (6) General and special purpose teleoperator/robot systems are required for a number of manufacturing, assembly, inspection and repair tasks.
- (7) An aggressive NASA development commitment in computer science is fundamental to the acquisition of machine intelligence/automation expertise and technology required for the mission capabilities described earlier. This should include a program for increasing the number of people trained in the relevant fields of computer science and artificial intelligence.

REFERENCES

American Association for Artificial Intelligence (AAAI), Proceedings of the First Annual National Conference on Artificial Intelligence, Stanford University, 18-21, August 1980.

Arden, B. W., Ed., What Can Be Automated?, MIT Press, Cambridge, Massachusetts, 1980.

Freitas, Robert A., Jr., "A Self-Reproducing Interstellar Probe," Journal of the British Interplanetary Society, Vol. 33, 1980, pp. 251-264.

——, "A Self-Replicating, Growing Lunar Factory," paper submitted to the 5th Princeton/AIAA/SSI Conference on Space Manufacturing, 18-21, May 1981, Princeton University, 1981a, 35pp.

Hajek, P. and Havranek, T., Mechanizing Hypothesis Formation: Mathematical Foundations for a General Theory, Springer-Verlag, Berlin, 1978.

Hayes-Roth, F., "Theory-Driven Learning: Proofs and Refutations as a Basis for Concept Discovery," Workshop on Machine Learning, Carnegie-Mellon University, July, 1980.

Heer, Ewald, Ed., Proceedings of the Pajaro Dunes Goal-Setting Workshop, unpublished draft notes, June 1980.

Israel, David J., "What's Wrong with Non-Monotonic Logic?" paper delivered at the First Annual National Conference on Artificial Intelligence, The American Association for Artificial Intelligence, Stanford University, 18-21, August 1980.

Long, James E. and Healy, Timothy J., Advanced Automation for Space Missions: Technical Summary, A Report of the 1980 NASA/ASEE Summer Study on the Feasibility of Using Machine Intelligence in Space Applications, University of Santa Clara, 15 September 1980. (Copies available from Tim Healy, EECS, University of Santa Clara, Santa Clara, CA 95053.)

Miller, M., "Toward a Remotely-Manned Energy and Production Economy," MIT Artificial Intelligence Laboratory, A.I. Memo No. 544, September 1979. 19pp. See also "Telepresence," OMNI, Vol. 2, June 1980, pp. 44-52.

Sagan, C., Chmn., Machine Intelligence and Robotics: Report of the NASA Study Group, NASA JPL Report No.

715-32, March 1980.

Science Applications, Inc. (SAI), Space Industrialization, Volumes 1-4, NASA Contract NAS-8-32197, SAI-79-(602-605)-HU, Huntsville, Alabama, 15 April 1978.

Von Neumann, J., Theory of Self-Reproducing Automata, University of Illinois Press, Urbana, Illinois, 1966. Compiled and edited by A. W. Burks.

Von Tiesenhausen, Georg and Darbro, Wesley A., Self-Replicating Systems - A Systems Engineering Approach, NASA TM-78304, July 1980.

Weyrauch, R., Doyle, J., Hewitt, C., McCarthy, J., McDermott, D., Reiter, R. and Thompson, A., "Non-Monotonic Logic Panel," panel chaired by Richard Weyrauch, First Annual National Conference on Artificial Intelligence, The American Association for Artificial Intelligence, Stanford University, 18-21, August 1980.

Winston, P. H. and Brown, R. H., eds., Artificial Intelligence: An MIT Perspective, MIT Press, Cambridge, Massachusetts, 1979. Volumes I and II.