-- Diagnose need for change in patient ventilation;

-- Identify indications for proceeding forward or back in the process of weaning from a ventilator; and

-- Make interpretations in light of measured test results, patient history, record of therapies and results of therapies and observations.

4. Implement rules for interpretation of respiratory insufficiency data with a new heuristic interpretation system including the following major features:

   -- Forms time-dependent hypotheses about the patient state;

   -- Infers desired courses of action based on measured patient state, observations, and expectations of future course.

   -- Uses models, both heuristic and mathematical, for generating an expectation of the immediate patient course;

5. Create an advisory committee, including outside experts in clinical medicine and computer science, to review the progress to date. They will review conceptual formulations, system design, scope and detail of the clinical knowledge and system operation. The advisory group will be asked to help to identify additional important considerations for the clinical knowledge base and the computer implementation, suggest improved ways to conceptualize or implement problems, and evaluate the soundness of the results to date.

4. Significance

Science advances by quantitation and development of general theories. The practice of medicine advances along one path by integrating quantitative measurements and general theories into the routine of existing clinical practice. The world of clinical medicine includes a complicated interaction among human patients, complex physiology, and proud, human clinical staff. This project is based on the assertion that good, quantitative measurements of physiological state are useful if effectively related to the human and physiological complexities of the clinical world. The best possibility we see for making new quantitative measurements far more generally useful in clinical medicine lies in knowledge-based interpretation of well understood physiologically relevant measurements. The improved care of the sick patient is our objective. This project, if successful, will directly improve the ability of the clinical staff to properly diagnose and manage the patient with respiratory insufficiency. It will lay the foundation for extension of successful methodologies of interpretation of the general problem of interpreting measurements of physiological state.

Appendix I

OVERVIEW OF ARTIFICIAL INTELLIGENCE RESEARCH

ARTIFICIAL INTELLIGENCE RESEARCH

What is it? What has it achieved? Where is it going?

Excerpt from a report by
Professor Edward A. Feigenbaum
Stanford University

May 1975

INTRODUCTION

In this briefing, these questions will be discussed as succinctly as possible:

I. What is the scientific field of artificial intelligence research, as seen from various viewpoints? What are the general goals of the field?

II. What are its practical working goals? What are some achievements relative to these goals (circa 1973)?

III. What steps (new goals, problems, potential achievements) seem to lie ahead, within a five year horizon?

ARTIFICIAL INTELLIGENCE (alias INTELLIGENT COMPUTER SYSTEMS):

General View;

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

Some scientists view the performance of complex symbolic reasoning acts by computer programs as the sine qua non for artificial intelligence programs, but this is necessarily a limited view.

Yet another view unifies AI research with the rest of Computer Science. It is an oversimplified view, but worthy of consideration. The potential uses of computers by people to accomplish tasks can be "one-dimensionalized" into a spectrum representing the nature of instruction that must be given the computer to do its job. Call it the WHAT-TO-HOW spectrum. At one extreme of the spectrum, the user supplies his intelligence to instruct the machine with precision exactly HOW to do his job, step-by-step. Progress in Computer Science can be seen as steps away from that extreme "HOW" point on the spectrum: the familiar panoply of assembly languages, subroutine libraries, compilers, extensible languages, etc. At the other extreme of the spectrum is the user with his real problem (WHAT he wishes the computer, as his instrument, to do for him). He aspires to communicate WHAT he wants done in a language that is comfortable to him (perhaps English); via communication modes that are convenient for him (including perhaps, speech or pictures); with some generality, some abstractness, perhaps some vagueness, imprecision, even error; without having to lay out in detail all necessary subgoals for adequate performance - with reasonable assurance that he is addressing an intelligent agent that is using knowledge of his world to understand his intent, to fill in his vagueness, to make specific his abstractions, to correct his errors, to discover appropriate subgoals, and ultimately to translate WHAT he really wants done into processing steps that define HOW it shall be done by a real computer. The research activity aimed at creating computer programs that act as "intelligent agents" near the WHAT end of the WHAT-TO-HOW spectrum can be viewed as the long-range goal of AI research. Historically, AI research has always been the primary vehicle for progress toward this end, though science as a whole is largely unaware of the role, the goals, and the progress.

HISTORICAL TRACE

The Working Goals of the Science;
Progress toward those goals;

The root concepts of AI as a science are 1) the conception of the digital computer as a symbol-processing device (rather than as merely a number calculator); 2) the conception that all intelligent activity can be precisely described as symbol-manipulation. (The latter is the fundamental working hypothesis of the AI field, but is controversial outside of the field.) The first inference to be drawn therefrom is that the symbol-manipulations which constitute intelligent activity can be modeled in the medium of the symbol-processing capabilities of the digital computer.

This intellectual advance--which gives realization in a physical system, the digital computer, to the complex symbolic processes of intelligent action and decision--with detailed case studies of how the realization can be accomplished, and with bodies of methods and techniques for creating new demonstrations--ranks as one of the great intellectual achievements of Science, allowing us finally to understand how a physical system can also embody mind. The fact that large segments of the intellectual community do not yet understand that this advance has been made does not change its truth or its fundamental nature.

Three global "working goals" have dominated the AI field for the 17 years of its existence. These are:

1. Understanding heuristic search as a processing scheme sufficient to account for much intelligent problem solving behavior; and exploring the scope and pervasiveness of heuristic problem solving.

2. Semantic information processing: developing precise formulations of "understanding" by programs, and "meaning" of symbols that are input or stored; the acquisition, storage, and deployment of knowledge of the world in the service of symbolic problem solving.

3. Information Processing Psychology: developing precise models of human behavior in symbolic-processing tasks.

The first two goals represent the fundamental paradigms that have dominated the field. The third cuts across these orthogonally, and involves intense interdisciplinary contact with Psychology, and Linguistics.

GOAL 1. HEURISTIC SEARCH, HEURISTIC PROGRAMMING, SYMBOLIC PROBLEM SOLVING PROGRAMS

In the first decade, the dominant paradigm of AI research was heuristic search. In this paradigm, problem solving is conceived as follows: A tree of "tries" (aliases: subproblems, reductions, candidates, solution attempts, alternatives-and-consequences, etc.) is sprouted (or sproutable) by a generator. Solutions (variously defined) exist at particular (unknown) depths along particular (unknown) paths. To find one is a "problem". For any task regarded as nontrivial, the search space is very large. Rules and procedures called heuristics are applied to direct search, to limit search, to constrain the sprouting of the tree, etc. While some of this tree-searching machinery is entirely task-specific, other parts can be made quite general over the domain of designs employing the heuristic search paradigm. Two notions are critical. The first is that problem solvers generally face a "maze" of alternative courses of decision and action that is huge compared with their processing resources. The second is the use of heuristic knowledge to steer carefully through large mazes toward a solution seeking the plausible and potentially fruitful avenues, avoiding the absurdities and the high-risk paths. Heuristic knowledge is usually informal knowledge--to be distinguished from formal knowledge that is assertable with the rigor of proof. Polya, the famous mathematician who wrote Patterns of Plausible Inference and other books on problem solving, calls heuristic reasoning "the art of good guessing." Heuristic knowledge is often "common sense" knowledge of the world, rules-of-thumb for generally acceptable performance, or rules of good practice in specific situations. When we speak of the "expertise" of an expert, and the "good judgment" he brings to bear on complex problems in his domain, we often are speaking of the heuristics he has developed to search effectively.

Provocative essays by Polya notwithstanding, the first serious and detailed studies of heuristic problem solving ever done by Science were done as AI research in its first decade. As with any other science, progress came by the detailed examination of specific cases, from which gradually emerged both a broad picture of the nature of the phenomena being studied and, within this, more formal theories for specific parts.

Three sub-goals of heuristic programming are discernable.

SUBGOAL 1A. Demonstrate sufficiency of heuristic search for tasks of intellectual difficulty.

These heuristic programming efforts dealt with almost "pure" symbolic reasoning tasks (i.e., tasks not requiring much coupling to real-world knowledge), and used inference schemes that were either ad-hoc or of limited scope. Notable successes during this "prove-the concept" phase were: the Logic Theory Program, that proved theorems in Whitehead & Russel's propositional calculus; the Geometry Theorem Proving program, that proved theorems in Euclidean geometry at a level of competence exceeding that of the excellent high school geometry student; the Symbolic Integration program, that solved college freshman symbolic integration problems about as well as MIT freshmen; chess-playing programs that play respectable "club player" C or B Class chess; a checker playing program that was virtually unbeatable, except by the country's top few players (notable also for remarkable self-improvement in performance by analysis of its own play and "book-move" good play); and a number of competent management science applications (assembly-line balancing, warehouse location, job-shop scheduling, etc.).

To recapitulate briefly: the key concepts are: search in problem solving; and the use of generally informal knowledge to guide search effectively. The AI community was the first to devote serious scientific effort to developing the idea of the use of informal knowledge in problem solving, with notable successes. Few in Science recognize that this achievement has been made and is ready for exploitation.

SUBGOAL 1B. Generality in Problem Solving Programs

Generality here means the use of a small set of problem solving methods of wide applicability to solve problems of many different types. Each of the problems posed is stated to the program in a particular representation (or framework) with which the set of methods is constructed to handle.

The subgoal of generality arises first as a reaction to the array of "specialty" programs mentioned above; second, from the general observation that the ability to do a wide range of tasks is a special touchstone of intelligence; third, from a direct assessment that as the diversity and heterogeneity of the tasks handled by an agent increases, the likelihood that it can do them all without intelligent action decreases; and fourth, from the argument that any ultimate intelligent agent must have wide generality, since it must take the world and its problems as they come without any intermediary, making generality an important independent desideratum.

This subgoal was pursued with vigor for ten years in a number of projects, was important for its feedback value in clarifying issues for the AI field, and has temporarily (at least) been put back on the shelf as the field begins to explore knowledge-based problem solvers and issues in the representation of knowledge.

There were two discernable subthemes. The first was an attempt to create abstract heuristic search methods that were divorced from any particular content. Examples were: the General Problem Solver, which used a variant of heuristic search known as means-ends analysis; MULTIPLE, which introduced adaptivity in the selection of what subproblem to choose "next" in a search; and REF-ARF, which extended the generality of ordinary procedural programming languages to include the embedding of non-procedural problems of constraint satisfaction.

The second subtheme was the construction of theorem provers that take problems expressed as theorems to be proved in the first-order predicate calculus. This line of work was motivated by the (correct) observation that the scope for representing real-world facts and situations in first-order predicate calculus is very great; and by the invention of the resolution method, a computational method for finding proofs for theorems in this calculus. There has been continuous improvement on the basic method, taking the form of proposing more powerful inference techniques, rather than the form of specific ways for programs to adapt to particular problems. The very strength of the formulation in terms of generality, namely its complete homogenization of the particular task (all tasks are seen and dealt with in the same logical formalism) turns effort away from how to exploit the particularities of special classes of tasks. But it appears that only by exploiting the particularities can significant reduction in search be achieved. From a practical point of view the only proofs produced by such problem solvers were "shallow" proofs.

Much of this line of research has been temporarily "shelved", awaiting further knowledge on how best to represent knowledge for computer processing. Problems that are essentially simple when represented in their "natural" representation appear extraordinarily complicated when translated into first-order predicate calculus. The current search for theorem provers using higher-order logics is based not on the attempt to increase the raw expressive power, so to speak, of first-order logic, but on the belief that naturalness of expression will ultimately pay off.

SUBGOAL 1C: High-Performance Programs that perform at near-human level in specialized areas

As the heuristic programming area matured to the point where the practitioners felt comfortable with their tools, and adventuresome in their use; as the need to explore the varieties of problems posed by the real-world was more keenly felt; and as the concern with knowledge-driven programs (to be discussed later) intensified, specific projects arose which aimed at and achieved levels of problem solving performance that equalled, and in some cases exceeded, the best human performance in the tasks being studied. The example of such a program most often cited in the Heuristic DENDRAL program, which solves the scientific induction problem of analyzing the mass spectrum of an organic molecule to produce a hypothesis about the molecule's total structure. This is a serious and difficult problem in a relatively new area of analytical chemistry. The program's performance has been generally very competent and in "world's champion" class for certain specialized families of molecules. Similar levels of successful performance have been achieved by some of the MATHLAB programs that assist scientists in doing symbolic mathematics. The effectiveness of MATHLAB's procedures for doing symbolic integration in calculus is virtually unexcelled. Yet another example, with great potential economic significance, involves a program for planning complex organic chemical syntheses from substances available in chemical catalogs. The program is currently being used as an "intelligent assistant" in a new and complex organic synthesis.

GOAL 2. SEMANTIC INFORMATION PROCESSING (S.I.P.)

The use of the term "semantic" above is intended to connote, in familiar terms, something like: "What is the meaning of..." or "How is that to be understood..." or "What knowledge about the world must be brought to bear to solve the particular problem that has just come up?" The research deals with the problem of extracting the meaning of: utterances in English; spoken versions of these; visual scenes; and other real-world symbolic and signal data. It aims toward the computer understanding of these as evidenced by the computer's subsequent linguistic, decision-making, question-answering, or motor behavior.

Thus, for example, we will know that our "intelligent agent" understood the meaning of the English command we spoke to it if: a) the command was in itself ambiguous; b) but was not ambiguous in context; and c) the agent performed under the appropriate interpretation and ignored the interpretation that was irrelevant in context.

In this goal of AI research, there are foci upon the encoding of knowledge about the world in symbolic expressions so that this knowledge can be manipulated by programs; and the retrieval of these symbolic expressions, as appropriate, in response to demands of various tasks. S.I.P. has sometimes been called "applied epistemology" or "knowledge engineering".

To summarize: the AI field has come increasingly to view as its main line of endeavor: knowledge representation and use, and an exploration of understanding (how symbols inside a computer, which are in themselves essentially abstract and contentless, come to acquire a meaning).

To classify all of the current work into a relatively simple set of subgoals is a formidable and hazardous undertaking. Nevertheless, here is one rough cut (stated for convenience as questions).

A. How is the knowledge acquired, that is needed for understanding and problem solving; and how can it be most effectively used?

B. How is knowledge of the world to be represented symbolically in the memory of a computer?

B1. What symbolic data structures in memory make the retrieval of this information in response to task demands easy?

C. How is knowledge to be put at the service of programs for understanding English?

D. How is sensory knowledge, particularly visual and speech, to be acquired and understood? How is knowledge to be applied to intelligent action of effectors, such as arms, wheels, instrument controls, etc.

Significant advances on all of these fronts have been made in the last decade. The area has a rather remarkable coherence--with individual projects threading through a number of the goals stated above (this makes excellent science and difficult exposition!)

GOAL 2A. Knowledge Acquisition and Deployment for Understanding and Problem Solving

The paradigm for this goal is, very generally sketched, as follows:

a. a situation is to be described or understood; a signal input is to be interpreted; or a decision in a problem-solution path is to be made.

Examples: A speech signal is received and the question is, "What was said?" The TV camera system sends a quarter-million bits to the computer and the question is, "What is out there on that table and in what configuration?" The molecule structure-generator must choose a chemical functional group for the "active center" of the molecular structure it is trying to hypothesize, and the question is, "What does the mass spectrum indicate is the ´best guess´?"

b. Specialized collections of facts about the various particular task domains, suitably represented in the computer memory (call these Experts) can recognize situations, analyze situations, and make decisions or take actions within the domain of their specialized knowledge.

Examples: In the CMU Hear-Say Speech Understanding System, currently the Experts that contribute to the Current Best Hypothesis are an Acoustic-Phonetic Expert, a Grammar Expert, and a Chess Expert (since chess-playing is the semantic domain of discourse). In Heuristic DENDRAL, the Experts are those that know about stability of organic molecules in general, mass spectrometer fragmentation processes in particular, nuclear magnetic resonance phenomena, etc.

For each of the sources of knowledge that can be delineated, schemes must be created for bringing that knowledge to bear at some place in the on-going analysis or understanding process. The view is held that programs should take advantage of a wide range of knowledge, creating islands of certainty as targets of opportunity arise, and using these as anchors for further uncertainty reduction. It is an expectation that always some different aspect provides the toe-hold for making headway--that is , that unless a rather large amount of knowledge is available and ready for application, this paradigmatic scheme will not work at all.

Within this paradigm lie a number of important problems to which AI research has addressed itself:

a. Since it is now widely recognized that detailed specific knowledge of task domains is necessary for power in problem solving programs, how is this knowledge to be imparted to, or acquired by, the programs?

a1. By interaction between human expert and program, made ever more smooth by careful design of interaction techniques, languages "tuned" to the task domain, flexible internal representations. The considerable effort invested by the AI community on interactive time-sharing and interactive graphic display was aimed toward this end. So is the current work on situation-action tableaus (production systems) for flexibly transmitting from expert to machine details of a body of knowledge.

a2. "Custom-crafting" the knowledge in a field by the painstaking day-after-day process of an AI scientist working together With an expert in another field, eliciting from that expert the theories, facts, rules, and heuristics applicable to reasoning in his field. This was the process by which Heuristic DENDRAL's "Expert" knowledge was built. It is being successfully used in AI application programs to: diagnosis of glaucoma eye disease, to treatment planning for infectious disease using antibiotics, to protein structure determination using X-ray crystallography, to organic chemical synthesis planning, to a military application involving sonar signals, perhaps to other areas, and of course to chess.

a3. By inductive inference done by programs to extract facts, regularities, and good heuristics directly from naturally-occurring data. This is obviously the path to pursue if AI research is not to spend all of its effort, well into the 21st Century, building knowledge-bases in the various fields of human endeavor in the custom-crafted manner referred to above. The most notable successes in this area have been:

...the Meta-DENDRAL program which, for example, has discovered the mass spectrum fragmentation rules for aromatic acids from observation of numerous spectra of these molecules--rules previously not explicated by the DENDRAL chemists.

...a draw-poker playing program that inferred the heuristics of good play in the game by induction (as well as by other modes, including the aforementioned interaction with experts).

a4. By processes of analogical reasoning, by which knowledge acquired about one area can be used to solve problems in a another area if a suitable analogy can be drawn. Our human experience tells us that this approach is rich in possibilities. One successful project can be cited (and that is a limited success); a program that discovers an analogy (in full-blown detail) between a theorem-to-be-proved in modern algebra and another theorem in algebra whose proof is known. The analogy is used to pinpoint from a large set of facts those few that will indeed be relevant to proving the new theorem.

GOAL 2B. Representation of Knowledge

The problem of representation of knowledge for AI systems is this: if the user has a fact about the world, or a problem to be stated, in what form does this become represented symbolically in the computer for immediate or later use? Three approaches are being pursued:

B1. the approach via formal logic. As mentioned before, first-order predicate calculus was tried, but was found to be too cumbersome to represent ordinary situations and common-sense knowledge. Set theory and higher-order logics are currently under examination as better candidates to be a medium for homogeneous representation.

B2. The ad-hoc approach. Most problem domains have a "natural" representation that human experts use when operating in the domain. Translate that representation fairly directly for the computer, and tailor the information processes to work with it. This is the approach commonly taken, in DENDRAL, MATHLAB, in chess playing programs, visual scene analysis, and so on (almost everywhere). Though it gets the job done, it creates serious problems for the cumulation of knowledge, techniques, and programs in the science because of the inhomogeneity that arises therefrom throughout the collection of AI projects undertaken. One way out of the dilemma is to do research on the problem of translation (by program) from one ad-hoc representation into another (the so-called "shift of representation" problem). Little work has been done on this problem, except one excellent "pencil-and-paper" exercise in connection with a simple puzzle, and one subprogram in DENDRAL (the Planning Rule Generator, that translates mass spectral knowledge from its form as fragmentation processes to a form useful for pattern matching).

B3. the approach via a "computable" semantic theory. In this approach, computational linguists attempt to analyze the full range of actions, actors, objects, and their relations, of which the common-sense world is composed; then refine and formalize these into a useable computational theory for representing facts, utterances, problems, etc. The most successful of these efforts is the Conceptual Parser (and its follow-on, MARGIE, which successfully accomplishes English paraphrase and common-sense inference).

In lieu of a tight, parsimonious computable semantic theory, other more ad-hoc systems, known as semantic-net-memory models, have developed experimenting with various sorts of actor-action-object-relation data structures. Semantic-net-memory models have a ten year history relating particularly to intelligent question-answering. Perhaps most successful of these is the HAM program which combines ideas from semantic theory, semantic-net-memory structures, and more traditional linguistic analysis (all in the context of a rather good model of human sentence comprehension, validated with dozens of careful laboratory experiments).

GOAL 2C. Programs for Understanding English

One can readily observe that it will be almost impossible to disentangle the skein of research on understanding natural language (English) from the coordinate efforts in representation and deployment of knowledge. Most of the state-of-the-art programs for understanding English employ, in one form or another, the basic S.I.P. paradigm outlined previously. These systems have substantial linguistic components that are highly sophisticated compared with anything done in the past. All of them incorporate linguistic theory that has an intimate and continuous tie-in between grammar "Experts" and domain-dependent "Experts". Although the domains about which they admit discourse are still modest and discrete, they are many times richer than anything done previously. The state-of-the-art is represented by the SHRDLU program for conducting a dialogue with a simulated robot about a world of blocks, boxes, and pyramids on a table; and the Lunar Rocks program for conducting a dialogue about properties of and transformations upon NASA moon-rock samples. The SHRDLU program, for example, will carry out commands, answer questions, and generally be aware of what it was doing, so as to answer "how" and "why" questions about its behavior.

The internal structure of these systems exhibits an interesting evolution over the semantic-net-memory systems, and they appear to be a long way from the heuristic search schemes mentioned earlier. They are essentially large programs written within a programming system that provides search and matching capability. There is no factorization between a data base (i.e., semantic net) and a small set of methods that process the data base. Rather, the entire system appears to be a large collection of special purpose programs for dealing with a multitude of special cases. They give the appearance of being a highly distributed system, in which the intelligent action resides throughout the entire program.

GOAL 2D. Acquiring and Understanding Sensory Data.

The goal here is to discover broadly applicable methods for extracting from sensory data (chiefly visual and aural) the information that is specifically responsive to users' needs. Two classes of needs may be noted: the need to facilitate communication between man and machine; and the need to apply computers to intrinsically perceptual tasks. The former is exemplified by the desire to talk, rather than type, to computers; the latter is illustrated by the task of automatically guiding an effector on the basis of visual data. To satisfy either (or both) of these needs, it is necessary to move from well-understood problems of sensing data to much more difficult problems of interpretation.

SUBGOAL 2D1. Visual Scene Analysis. Computer-based analysis of visual scenes has its roots in work on optical character recognition (early to mid-Fifties) and by work in automatic photoreconnaissance. These tasks are essentially two-dimensional. Little is lost by disregarding dimensions of objects in a direction orthogonal to the picture plane.

AI research on scene analysis began in the early sixties with the work of Roberts on pictures of polyhedra. This work (and its intellectual descendants) differs from the earlier two-dimensional work in two major respects: first, it explicitly considers, and capitalizes on, the three-dimensional properties of objects and their perspective representations; second, it utilizes a variety of special processing steps and decision-making criteria, in contrast to the earlier template-match/classify paradigm.

Robert's work spawned five years of intensive research on pictures of collections of polyhedra. One theme, centered on the archetypical question "Is an edge present in a given (small) region of the picture?", led to the development of edge detecting, contour following, and region finding programs. A second theme, centered on teasing out the properties of polyhedra and their representations, led to an elegant theory of permissible representations of edges and vertices, and their relations to three-dimensional polyhedra - a theory not previously discovered by projective or descriptive geometers.

Work in the polyhedral objects domain culminated in several programs capable of describing, in more or less complete detail, pictures of complicated collections of polyhedra, even taking into account shadows cast by these objects. At the same time, more complicated types of scenes began to be seriously studied. This has led to current interest in the use of color, texture, and range data, and has stimulated interest in program organizations capable of capitalizing on these multiple perceptual modalities. For example, in one paradigm perception is viewed as a problem-solving process that uses many varieties of knowledge to select perceptual operators, to guide their application to sensory data, and to evaluate the results obtained therefrom.

SUBGOAL 2D2. Speech Understanding. Research on computer recognition of speech signal data began in the Fifties with work on the recognition of isolated words. Some observations will be made here on the relation between speech understanding research and the on-going body of AI research.

The fundamental idea driving research on speech understanding is that "recognition" is impossible (in flowing natural speech) without understanding, and that understanding is impossible without extensive knowledge about the domain of discourse. This view arises in part from the observation that ambiguities and omissions at both the acoustic and semantic level do not arise as bizarre or pathological exceptions but instead are commonplace events. Speech understanding research thus relies heavily on progress in the basic AI research problems of knowledge acquisition, representation, and deployment. This situation in unlikely to change regardless of advances in processing acoustic signals.

GOAL 2E. Intelligent control of Effectors

This goal concerns the creation of devices and control programs for bringing about specified changes in the physical world. The effectors that have attracted the most attention have been mechanical manipulators and mobile vehicles but this has been largely a matter of experimental convenience. In principle, they could as easily have been subsystems of spaceships or manufacturing tools.

Early work in "intelligent" effectors dates back two decades, but systematic work did not begin until about 1966, at which time some progress had already been made in developing symbolic problem solving programs to control effectors. Since then there has been considerable interest in computer-controlled effectors because problems of effector control excite a set of important issues for AI research. The following is a rough characterization of the subgoals of work on effector control:

SUBGOAL E1. Monitoring Real-World Execution of Problem Solutions: The special touchstone of effector control research is that a problem is never "solved" until the real, physical world has been altered in a fashion that satisfies the task specification (in contrast to other problem solving programs whose responsibility ends with the symbolic presentation of a good solution). Thus, an effector control program should ideally be prepared to deal with any eventuality that affects the execution of a theoretically correct solution, be it initial misinformation, accidental dynamic effects, etc. These demands strongly influence all levels of program organization and strategy. Problem solving and execution monitoring must be made to interact intimately. The most advanced work of this type is probably the STRIPS-PLANEX system (for the control of a mobile vehicle) that can detect and gracefully recover from a wide variety of execution difficulties.

SUBGOAL E2. Modelling "Everyday" worlds: To control effectors by computer requires that the computer have adequate models of everyday situations. It has become important to model occlusion, obstruction, relative location, etc., and this has been done to the extent necessary to handle various simple manipulation and locomotion problems.

SUBGOAL E3. Planning in the face of uncertainty. Problem-solving programs for the control of effectors that operate on the physical world must be able to work routinely with incomplete and inaccurate information. This creates a need to do research on programs that can form contingency plans, can plan to acquire information, can decide when to execute actions in the physical world, even if the plan is incomplete, and so forth. Some research of this type has been done.

SUBGOAL E4. Low-Level Control. By low-level control is meant: programs that interact more-or-less directly with the effector mechanism, and that do not engage in global planning or problem solving. Research on this topic is producing a new and potentially important branch of classical automatic control. Although little has been formalized to date, enough experience has been acquired to permit the construction of interesting demonstrations. Among the most impressive of these is an arm control program that can drive the arm in partially constrained ways; for example, the arm can be made to turn a crank by dynamically constraining the necessary degrees of freedom.

SUBGOAL E5. Hardware Development. The manipulators available in 1966, whether based on prosthetic limbs or industrial put-and-take machinery, were generally too primitive to be of long-term value for AI research. This situation fostered a fairly significant hardware development effort that produced a useful arm-hand device. Similarly, sensing devices received some development efforts. Examples of this work are newly developed optical range finders, and special tactile, force, and torque sensors.

GOAL 3. INFORMATION PROCESSING PSYCHOLOGY: DEVELOPING DETAILED SCIENTIFIC MODELS OF HUMAN SYMBOLIC PROCESSING BEHAVIOR.

Since its inception, one focus of AI research has been the study of the symbol manipulation processes capable of explaining and predicting human behavior in a wide range of cognitive tasks. As science, the endeavor is entirely classical in intent and method, employing model construction and validation. Empirical data from well-controlled laboratory experiments is obtained from psychologists or generated by the researchers in their own laboratories. Induction from this data leads to the formulation of a symbol-processing model which purports to explain the observed phenomena. This model is given a precise form as computer programs and data structures (since the computer as a general symbol-processing device is capable of carrying out any precisely specified symbol-manipulation process; this step is entirely analogous to the model-implementation step taken by the physicist when he translates his physical model into the form of a set of differential equations). A computer is then used to generate the complex and remote consequences of the symbol-processing postulates of the model for the particular laboratory situations and stimuli being studied. These consequences and predictions are tested against empirical data; differences are noted and analyzed; the model is refined and run again; iterations continue until a satisfactory state of agreement between model's predictions and empirical data is achieved.

From one point of view, the endeavor is to be seen as Theoretical Psychology. From another point of view, it can be seen as a systematic attempt by AI research to understand intellectual activity as it occurs in nature (i.e., in humans) so that artifacts capable of performing such intellectual activity can be constructed upon the principles discovered. The interplay between these two views has been very strong.

Information Processing Psychologists have usually chosen their problems in areas that have been of "classical" concern to Psychology, though some of these areas have been reopened to serious investigation because of the successes of the information processing approaches. The following are brief sketches of some subgoals of the effort in Information Processing Psychology.

SUBGOAL 3A. Functional reasoning. Analysis and modeling has been done for human behavior in solving logic problems, complex crypt-arithmetic puzzles, and chess-play problems. The models, and the predictions derived from them, are so detailed that no comparison with previous work on the psychology of problem solving is meaningful. The work is a scientific revolution, and has had a great paradigmatic and methodological impact upon Psychology. The principal innovators, Newell and Simon, have had their contributions recognized by election to the National Academy of Science; Simon was awarded the Distinguished Scientific Contribution Award of the American Psychological Association, more or less the "Nobel Prize" of Psychology.

SUBGOAL 3B. Rote Memory and Short-term Memory phenomena: Storage and retrieval processes for short-term memory. Rote memorization effects. Discrimination and association learning for verbal materials. These and related phenomena of verbal learning and memory have been studied intensely by experimental psychologists in this century. A few dozen solid empirical generalizations are known. A set of closely related information processing models is capable of explaining many of these (roughly speaking, 15-20 of the "classical" phenomena).

SUBGOAL 3C. Long-term associative memory: Associative retrieval from associative memory nets of several hundred to a few thousand symbols. Interaction of English sentence processing and memory. The symbolic representation of knowledge (i.e., facts about the world) in memory. The work is currently very active, highly promising, and is causing a mini-revolution in thinking among psychologists who study memory.

SUBGOAL 3D. Pattern induction/concept formation. Induction of models of pattern regularities in strings of symbols. Induction of the "generating rule" from the exhibition of instances of the rule.

SUBGOAL 3E. Phenomena of neurosis. The behavior studied is neurotic symbol-processing behavior, viewed as processing distortions of otherwise "normal" linguistic and problem-solving processes. A highly successful model of paranoid behavior has been developed, incorporating some English language processing.

These examples are but pieces of a bigger picture, which looks something like this:

1. It is no surprise that Psychology has been strongly affected by the information processing concepts and tools of AI research since both sciences are concerned with the study of cognition. The magnitude of the impact is the big surprise. It is probably fair to say that the dominant paradigm currently structuring Experimental Psychology in this country is the information processing paradigm. Upon no other area of science has AI research had such a strong impact.

2. The scientific study of human thought has been accelerated greatly during the last fifteen years because of the AI impact. It is not much of an overstatement to say that the AI impact has revitalized the study of thinking by Psychology, making this scientific enterprise tractable, fruitful, and respectable.

VIEW OF THE FUTURE: What lies within a five year horizon?

An extrapolation of the research directions previously described into the future faces at least two problems. First, there are the usual uncertainties that loom because of unpredictable advances and wishful thinking. Second, the imposition by ARPA of research priorities upon the course of events that would "normally" ensue will have a large effect. Thus, the question of "what should happen" is as big a question as "what will happen."

This exposition is made difficult by the fact that the structure of the field, as outlined above in terms of Goals, will show strong confluences during the future period. Any simple presentation goal-by-goal would be misleading, and was not attempted. Instead, each identifiable focus is stated and then given an extended discussion.

The main thrusts of the Artificial Intelligence community in the next five year period will be:

1. Development of applications programs that represent and use knowledge of carefully delimited portions of the real-world for high-performance problem solving, hypothesis induction, and signal data interpretation.

   The next period is likely to be a period of consolidation of AI's previous gains into meaningful real-world applications. High levels of competence in the performance of difficult tasks will be the hallmark. In addition to growing attitudes toward becoming more relevant, the AI community's current major interest in knowledge structuring and use will naturally lead it to bodies of real-world knowledge that are rich in structure and challenges. An extrapolation indicates applications to domains in science (much as the DENDRAL and MATHLAB programs were developed); and in medicine (current activity includes programs that deal with Infectious Diseases and with Glaucoma); perhaps more routine aspects of architecture (e.g., space layout and design); perhaps design in electronics (e.g., layout of IC and PC electronics, actual circuit design to functional specs); management science applications (e.g., logistics management and control, crew scheduling for aircraft fleets). The most significant application will be to computer science itself, namely the automation of many programming functions (to be discussed later). Application to some of the less routine aspects of office document processing is a likely event (discussed later). With appropriate stimulus from ARPA, or other service agencies, these application priorities could be shifted toward defense problems, particularly those related to signal processing (e.g., application to seismic or sonar signal interpretation). In such applications, interpretation of what the signal means is made in terms of knowledge about the signal-generating source and the environment in which the signal occurs. All of these applications will be characterized by careful choice of domain, careful delimiting of the extent of knowledge necessary to do the job, and close coupling with human experts to gain the knowledge necessary. None of these programs will be "general problem solvers" of the old genre. Characteristic of some of these applications will be one-line interaction with human experts, not only to "tune" the knowledge used by the program, but also to intervene in decisions for which human expertise dominates that of the program, or where the relevant knowledge has not been made explicit and formalized for computer processing.

2. The development, in particular, of that area of application involving the synthesis of computer programs (the so-called "automatic programming" problem).

The particular application of AI techniques to the task of synthesizing computer programs from imprecise and non-procedural descriptions of what a user wants a computer to do for him is the AI problem area whose time has come. This area will be the subject of a separate and detailed program plan. It is an AI application of tremendous economic, and industrial importance, since computer programming is today a major bottleneck in the application of computers to technological and business problems. What is worse, virtually no advances of substantial impact upon this problem have been made in the last decade in other areas of computer science (with the possible exception of the interactive editing, debugging, and running of programs). The automatic programming problem is, furthermore, the quintessential problem that fits the WHAT-TO-HOW characterization of the nature of the science of Artificial Intelligence. It is the meeting ground of many of the tributaries of AI research: problem solving, theorem proving, heuristic knowledge and search, understanding of English (perhaps even speech), and advanced systems work. It is an ideal problem from the viewpoint of knowledge-based systems--the main line of current AI research. The essential activity in building such systems is the extraction and formalization of knowledge of the specific task domain. In the art of programming, computer scientists are their own best experts, and for years have been engaged in formalizing what is known about programming, mathematically and in other forms. Following this line of reasoning, the programming task that may be best suited is systems programming. An example of a specific systems programming task that may be accomplishable within the period is: development of an automatic programming system that will produce operating system code for a minicomputer like the PDP11/45, in response to functional specifications for instrument control and data-handling, where the specs are given in functional terms by a scientist putting together the instrument-computer package, not his (until now inevitable) programmer.

3. The extension of current ideas about the processing and understanding of English to more extensive domains of discourse and with greater flexibility, to the point of practical front-end processors for large applications programs.

In the coming period, programs for understanding English in limited "universes of discourse" will achieve practicality, and will be made available as the linguistic interaction vehicle in some of the larger AI applications programs, e.g., the automatic programming systems mentioned above. Since these applications programs will be domain-limited anyway, it will not be an extraordinarily difficult task to construct for them front-end processors that understand English in that domain. Since currently the field has only "demonstration programs" that exhibit (limited) understanding of English, much more research will be undertaken in these directions: examining how well current techniques extrapolate to broader domains of knowledge; developing techniques for establishing context of an interaction and maintaining that context throughout the conversation; and extending methods for drawing inferences from the continually updated context. Research on semantic theory, previously mentioned in connection with representation of knowledge, will be applied to specific problems of linguistic interaction involving actors, actions, objects, and common-sense knowledge. The area of language understanding is so rich in possibilities and implications that it is not unreasonable to consider developing a separate program plan for it within the next two years.

4. Initial exploration of office-work tasks as an area of development and application; the careful choosing and shaping of specific tasks in this enormous arena of human endeavor; and some limited applications progress on these tasks.

The AI research community has been searching for problem domains of significance to science, technology, or industry that would provide an integrating theme for the various subareas of AI work. These subareas have a considerable coherence of concepts and techniques, but the centripetal force of a real-world theme is necessary to make this coherence a practical reality. Production assembly by combinations of vision, manipulation and problem solving programs is an attempt to establish such a theme. Increasingly the feeling is growing in the AI community that the development of "intelligent assistant" programs for ordinary office work is a useful and important focus. There are two reasons for this. First, much of current AI research fits the task area well (e.g., semantic-net-memory structures, question answering programs, natural language understanding, "intelligent assistant" interaction programs, etc.). Second, the explosion of use of the ARPA network for "office work" tasks quite apart from computation (uses such as message processing, message and document filing, information retrieval from large data bases, composing and editing of documents, etc.) provides an excellent medium in which to do the work. The AI community, perhaps with a push from ARPA, has the capability to do significant work on the office automation problem in the next period. A carefully thought-through program plan will probably be the first output of the field in this area (should be organized and completed within the next two years), followed by initial exploratory ventures along the lines laid down in the plan.

Again, as with all the knowledge-based systems of this decade, the specific tasks worked upon will of necessity be carefully delimited. The general "intelligent office assistant" is well beyond the horizon, but specific assistant-programs for handling some of the office-work flow of information on the ARPA network can be realized within five years.

5. Intensive developmental work on the speech-understanding problem.

6. Expansion of computer vision research to: knowledge-based program organizations; development of a repertoire of low-level perceptual operators for color, range and texture, and exploitation of these modalities; first practical applications of scene analysis to selected tasks in industrial and biomedical settings; and use of interactive scene analysis for both research and application purposes.

Scene analysis programs consist of a combination of sensing-and-measuring primitive perceptual operators (like line-finders) and higher-level knowledge-based procedures (like line-proposers). Because of general awareness of the limitations of current primitive operators (at least as they are applied to monochrome pictures), the research will place increased emphasis on the acquisition and low-level analysis of color and range data. Higher level procedures will use knowledge of: three-dimensional properties of objects other than polyhedral objects; perceptual properties of objects; many varieties of contextual constraints among objects; and properties of the primitive operators (like computational cost, reliability, and domain of applicability). Practical applications will probably focus on industrial tasks like work-piece identification and location, inspection, and manipulator control. The scene analysis research issues in these applications may turn out to be pedestrian, but concerns about cost, reliability, and reprogrammability will become prominent. Biomedical scene analysis problems will continue to stimulate research; application to medical mass-screening tasks may occur. Interactive scene analysis will be an important focus. In research settings, interactive scene analysis will be used to construct large scene-analysis systems through the incremental accumulation of knowledge; in application settings it will be used to achieve flexible scene analysis systems that can be easily "re-programmed" by users who are not computer scientists.

7. Expansion of arm-hand effector technology and associated program control, with some practical applications of simple forms of this technology in industrial settings.

There will be considerable activity in the transfer of ARPA-initiated work on effector control to industrial settings. Hardware realizations of a rich variety of mechanical effectors, with their tactile, force, and torque sensors, will appear. Visual feedback in controlling effectors will be a feature of many of the applications. Basic research on the hardware and software technology of effector control will continue, if support from ARPA or other agencies is forthcoming. More broadly-based research on effector control is likely to be stimulated by the appearance of relatively inexpensive experimental hardware. Researchers who are currently unable to develop one-of-a-kind devices because of their cost will enter the field.

8. Expanded basic research on acquisition, deployment and representation of knowledge to support knowledge-based systems development.

Though the main thrust of AI research is in the direction of knowledge-based programs, the fundamental research support for this thrust is currently thin. This is a critical "bottleneck" area of the science, since (as was pointed out earlier) it is inconceivable that the AI field will proceed from one knowledge-based program to the next painstakingly custom-crafting the knowledge/expertise necessary for high levels of performance by the programs. In the next period, the following kinds of fundamental explorations must be pursued and strongly encouraged:

a. Additional case-study programs of hypothesis discovery and theory formation (i.e., induction programs) in domains of knowledge that are reasonably rich and complex. It is essential for the science to see some more examples that discover regularities in empirical data, and generalize over these to form sets of rules that can explain the data and predict future states. It is likely that only after more case-studies are available will AI researchers be able to organize, unify and refine their ideas concerning computer-assisted induction of knowledge.

b. Development of interactive interrogative techniques, coupling a program to a human expert, by means of which the program systematically elicits from the expert particular facts, useful heuristics, and generalizations (or models) in the domain of the human's expertise. Again, specific case-studies are desirable. Their development need not await the arrival of English language understanding programs to facilitate the interaction and interrogation. (Stylized languages designed for the specific case-study domains will serve for now.)

c. Exploration of a variety of methods for bringing together disparate bodies of knowledge held by a program to assist in the solution of specific problems that the program is called upon to solve. The nature of this problem was discussed earlier under Goal 2A. If there are to be a number of Experts (i.e., specialized knowledge bases) interacting in the solution of a problem, how should their interaction be arranged? Is there an Executive Program "in charge" of sequencing the activity of the Experts? If so, what is the nature of the Executive Program's knowledge about each Expert, and the appropriateness of calling that Expert to assist at a specific point in the process? Should the Experts be relatively independent, each with its own situation-recognizer to trigger its activity? These particular questions are posed here not in an effort to characterize the problem completely (or even adequately), but to give the flavor of the experimental inquiry that needs to be pursued in the coming period - a period in which major AI programming efforts will be directed toward knowledge-based systems with multiple sources of knowledge.

d. Theoretical and experimental studies of representation of knowledge. This basic and difficult problem is not one that is likely to have a "solution" in a five year period. Theoretical studies will continue to search for a logical calculus in terms of which to formalize and store knowledge in a fairly "natural" way, and for logical processors that will compute efficiently within this formalism. Experimental studies will attempt to deal with the usual nonhomogeneity of representation among different bodies of knowledge directly, by programming translations of representations from one "natural" representation to another as necessary in those situations requiring communication between Experts for joint problem-solving.

9. Continuing basic research on various mathematical-logical problems such as formal models for heuristic search, theorem proving methods, and mathematical theory of computation.

Because heuristic search has been a central theme of AI problem solving research, it is likely that attempts at mathematical formulation and analysis of heuristic search methods will continue. No existing research thrusts indicate that this work should have high priority at this time. However, the situation is unstable in the sense that a few key results (e.g., new theorems or, more likely, new formulations of heuristic search) might cause a rush of activity along lines of formal analysis.

A similar situation attends theorem-proving research. There are currently no critical ideas acting as a forcing function, but nonetheless the problem appears to some scientists to be central for progress in the long run. In their view, to state that a computer can be used as a "symbolic inference engine" is equivalent to saying that it is a "logic engine"; and what makes a "logic engine" turn over is a theorem prover over the domain of some logical calculus. The search for appropriate logical calculi and associated theorem provers will therefore continue.

The work in mathematical theory of computation has been peripheral to the AI mainstream, but recently has been gaining momentum and importance, and will enter the mainstream as basic research for automatic programming efforts. To write programs capable of synthesizing programs obviously requires a thorough understanding of the nature of programs. One kind of understanding is gained by formal description and mathematical analysis (the kind of understanding we take so much for granted in some physical sciences and engineering). To the extent that useful formal descriptions of how programs are put together and what programs do can be discovered; and to the extent that powerful theorems can be proved within the formalism; the work on mathematical theory of computation could aid significantly in the practical work of constructing automatic program synthesizers and verifiers. Thus, there are noteworthy "breakthrough" possibilities in this area.

A prediction of the most likely course of events in these tasks of formal analysis is that they will be low-key, low cost, high risk/high payoff.

10. Continuing research on modeling of human cognitive processes using information processing techniques.

At the interface between AI and the psychology of human perception and thought, the research tempo has been increasing for some time. In the coming period it is likely that new methodology, new conceptual insights, and new models will have a continuing dramatic impact on Psychology. The feedback to on-going AI research will continue to be important, particularly in the areas of perception and memory. The principal developments are likely to be these:

a. Methodological: analysis by program of the thinking-aloud protocols of humans solving complex problems (i.e., "data reduction" that requires some language understanding and complex inductive inference), resulting in a speed-up in this critical empirical procedure of perhaps a factor of 100. A typical complete protocol analysis of human data in a puzzle-solving task currently takes, without computer assistance, 100 hours.

b. Short-term memory. The processes of human short-term memory will be so well modeled and understood as a result of research in this period that the topic will cease to be of major theoretical interest to psychologists.

c. Long-term memory. A very good model of human long-term associative memory will be developed. The program which realizes this model will be given a great deal of "garden variety" knowledge of the everyday world, as the basis for empirical testing. Such a model will undoubtedly prove to be an important subsystem in larger programs that attempt language understanding in contexts involving common-sense knowledge. Only the beginnings of such a memory model exist today.

d. Visual perception. The most important impact of AI on Psychology in the coming period may be the initial formulation on an information processing theory of human visual perception of common 3-D forms, along the lines of the visual processing concepts and operators developed by AI vision research. AI vision research stands on the threshold of Psychology awaiting an intellectual push like the one given to problem solving in late Fifties. If the push is made, and is successful, it will noticably dent the theory of visual perception in five years and totally capture it within ten years.

Appendix II

AI HANDBOOK OUTLINE

NOTE:
The following material is a tentative outline of a handbook on artificial intelligence planned for publication. It is not to be cited or quoted out of the context of this report without the express permission of Professor E. A. Feigenbaum of Stanford University. This handbook is intended for two kinds of audience; computer science students interested in learning more about artificial intelligence, and engineers in search of techniques and ideas that might prove useful in applications programs. Articles in the first seven sections are expected to appear in the first volume to be published in preliminary form by September 1977. The remaining articles are expected to appear in the second volume to be published in preliminary form by June 1978.

The following is a brief checklist that was used to guide the computer science students engaged in writing articles for the handbook. It is, of course, only a suggested list.

i) Start with 1-2 paragraphs on the central idea or concept of the article. Answer the question "what is the key idea?"

ii) Give a brief history of the invention of the idea, and its use in A.I.

iii) Give a more detailed technical description of the idea, its implementations in the past, and the results of any experiments with it. Try to answer the question "How to do it?.

iv) Make tentative conclusions about the utility and limitations of the idea if appropriate.

v) Give a list of suitable references.

vi) Give a small set of pointers to related concepts (general/overview articles, specific applications, etc.)

vii) When referring in the text of an article to a term which is the subject of another handbook article, surround the term by +'s; e.g. +Production Systems+.

AI Handbook Articles

I. INTRODUCTION
   A. Philosophy
   B. Relationship to Society
   C. History
   D. Conferences and Publications