FORECASTING AND ASSESSING THE IMPACT OF ARTIFICIAL INTELLIGENCE ON SOCIETY*

Oscar Firschein
Martin A, Fischler
Lockheed Research Laboratory
Palo Alto, California

L. Stephen Coles
Jay M. Tenenbaum
Stanford Research Institute
Menlo Park, California

Abstract

At the present stage of research in artificial intelligence, machines are still remote from achieving a level of intelligence comparable in complexity to human thought. As computer applications become more sophisticated, however, and thus more influential in human affairs, it becomes increasingly important to understand both the capabilities and limitations of machine Intelligence and its potential impact on society. To this end, the artificial intelligence field was examined in a systematic manner. The study was divided into two parts:

(1) Delineation of areas of artificial intelligence, and postulatio" of hypothetical products resulting from progress in the field, and

(2) A judgmental portion, which involved applications and implications of the products to society.

For the latter purpose, a Delphi study was conducted among experts in the artificial intelligence field to solicit their opinion concerning prototype and commercial dates for the products, and the possibility and desirability of their applications and implications.

1. Introduction

In the near future, intelligent machines will replace or enhance human capabilities in many areas previously considered strictly within the human domain. Such machines will affect the everyday lives of people in ways unique in the history of mankind, and will change the face of society as we now know it (c.f., Weizenbaum, 1972). Although many philosophers, science fiction writers, and others have conjectured about robots and other products related to artificial intelligence, as far as we know there has been little effort to examine the AI field as a whole (in its relation to society) in an organized manner. At the request of the Institute of Electronic and Electrical Engineers (IEEE) to conduct a technical forecasting and assessment study in one of its various fields of interest, the San Francisco Chapter of tho IEEE Systems, Man, and Cybernetics Society selected the field of artificial intelligence; this paper reports on the results of this study, After an initial framework was established, the cooperation of other organizations, such as the ACM Special Interest Group on Artificial Intelligence (SIGART), the World Future Society, and the Institute for the Future, was obtained.

The present paper describes the derivation of the Delphi questionnaire which forms an important part of the study, an associated detailed examination of the AI field, the Delphi study itself, the results obtained in the study, and finally, the implications of the study.

Forecasting and Assessing the AI Field

Why bother to make an effort to predict future developments in AI? On the one hand it hardly seems worthwhile to plan further ahead than ten years. Things are changing so rapidly that any attempt to predict conditions more than a decade hence are likely to be seriously wrong in at least several important respects. The combinatorial power of the exponential curve tends to make short-term forecasts overly optimistic while long-term forecasts are observed to be overly pessimistic. Moreover, the unanticipated appearance of saturavion effects can radically alter the basis of a forecast.

On the other hand, we tend to make decisions practically every day that must reach beyond a decade, or occasionally even a centuryl For example, new buildings are normally designed for an occupancy of one hundred years. Because computer systems also have the potential for long term effects on human life and society, It. is important that the implications of such systems be understood. It is our hope that this study can be useful to those who are given responsibi1ity for future-based decisions.

To put. this study into perspective, it is useful to indicate the present and predicted state of the art in computer hardware, and to indicate the AI-like devices that are currently on the market. It has been recently indicated by Foster (1972), that before 1980 a general-purpose microcomputer complete with central processor and internal working memory (but no peripherals) will be available on a single chip in small quantities for a cost of between \$1.00 and \$10.00. About this time frame, magnetic bubble technology is expected to bring the cost of memory from the present \$0.01 a bit to \$0.001 a bit. In a recent article, Broers and liatzakls (1972) predict 100,000 transistors and similar devices on a silicon chip a few milimeters square. In 25 years it is predicted that a single chip computer will be available capable of 20 million instructions a second with 65K of internal memory selling for about \$1.001 Even if this estimate were off by an order of magnitude, the social significance is enormous.

* The research reported here was sponsored in part by a grant from the National Science Foundation (GJ-37696) to the Institute of Electrical and Electronics Engineers.

Table 1 indicates some AI products which have been reported in the trade press. Although some of these products are relatively unsophisticated compared to the corresponding AI products postulated in our study, they do indicate current trends and the state of the art.

2. Setting Up The Study

In assessing the effects of the artificial intelligence field on society, we found it convenient to break the effort into two parts:

(1) A technical portion in which the AI field is examined, capabilities derived, and products postulated.

(2) A more Judgmental portion in which applications and implications for society are postulated, and their effects on society are Judged,

It is in the latter portion of the study that the Delphi Forecasting technique is most useful.

The overall methodology used in the study is shown in Fig. 1.

(1) The topic and subtopic areas of AI were first derived using such sources as the ACM Computing Reviews Index, the IJCAI Call-For-Papers, and through discussions with experts in the field.

(2) These topic areas were examined individually to determine what capabilities might be available if success were achieved in preeent-day efforts.

(3) From the capabilities available, products were postulated based on these capabilities.

Steps (2) and (3) were iterated several times, since conceptualizing products often leads one to a closer investigation of the availability of capabilities for that product.

(4) Examining the product list, we organized the items according to the social areas on which the products will have a major impact.

(5) The applications and implications of the AI products on society were postulated.

(6) A Delphi study was performed to solicit expert opinion concerning the items postulated.

(7) The responses to the Delphi questionnaire were analyzed.

We then iterated back to step (3) as required. The results of the more technical portion of the study are given in Section 3, while the Judgmental portions are given in Section 4.

3. Current AI Capabilities and Derived Products

In order to postulate products that could result from advances in the artificial intelligence field, we first partitioned the field into basic subject categories, and then analyzed the capabilities currently available in each category. The following categories and the capabilities associated with each category were used in the study:

1—Language understanding. The ability to "understand" natural language and to respond in natural language. Ability to translate from spoken language to a written form. Ability to translate from one natural language to another.

2—Problem solving. Ability to formulate a problem in a suitable representation, to plan for its solution, and to know when new information is needed and how to obtain it.

3--Perception (visual), The ability to analyze a sensed scene by relating it to an internal model which represents the perceiving organism's "knowledge of the world." The result of this analysis is a structured set of relationships between entities in the scene.

4—Modeling. The ability to develop an internal representation and set of transformation rules which can be used to predict the behavior and relationship between some set of real-world objects or entities.

5—Learning and adaptive systems. The ability to adapt behavior baged on previous experience, and to develop general rules concerning the world based on such experience,

6--Robots. A combination of most or all of the above abilities with the ability to move over terrain and manipulate objects.

7--Games. The ability to accept a formal set of rules for games such as Chess, Go, Kalah, Checkers, etc., and to translate these rules into a representation or structure which allows problem-solving and learning abilities to be used in reaching an adequate level of performance.

The areas of artificial intelligence and their subtopics are shown in Table 2. A representative collection of early work in A.I is given in Feigenbaum and Feldman (1963), and in MInsky (1968). We have focused our attention on current and recent research projects in AI and have summarized in Table 3 those efforts which we feel indicate currently obtainable levels of competence for each of the AI categories postulated.

By examining these AI capabilities, we were able to postulate a set of hypothetical products that might impact on society. The products were discussed with members of the Ar community, and eventually a list of twenty-one items was compiled, as given in Table 4.

* Note that the postulated products were based on what could, not what would or should be developed. The would and should aspects are dealt with in the Delphi portion of the study.

Table 1

SOME COMMERCIAL PRODUCTS USING AI TECHNOLOGY

Description	Name of Product	Company or Source	Function
Unique personal identification	Identimat 2000	Identimation Marketing, Northvale, N.J.	Device compares hand measurements with data coded on user's ID card.
Automatic voice verification		Bell Telephone Labs	Will be used to identify credit card users by having user repeat previously recorded phrase.
Medical monitoring	IMBLMS	Lockheed, Sunnyvale, Calif.	A "biobelt" which allows transmission of user's bio-activity to remote console for analysis.
Industrial robot (first generation)	Versatran	American Machine and Foundry, New Rochelle, N.Y.	"Blind" assembly line operation
	Unimate	Unimation, Inc. Danbury, Conn.	
Industrial robot (1-1/2 generation)	HIVIP, Mark I	Hitachi Ltd., Japan	Robot with TV vision can pick up objects from a moving conveyor belt.
Automated warehouse		Mobility Systems, Santa Clara, Calif.	Computer-controlled positioning of loading devices to correct storage bin.
Voice operated (a) supermarket checkout (b) baggage routing	VIP-100	Threshold Technology, Cinaminson, N.J.	(a) clerks dictate the prices of the items (b) baggage handlers dictate flight number or destination
Voice response system		Quantel Corp.	16,000 word vocabulary, "complex" sentence response.

Table 2

AREAS OF ARTIFICIAL INTELLIGENCE

1. Language Understanding
   - 1.1 Speech Understanding
   - 1.2 Semantic Information Processing (Computational Linguistics)
   - 1.3 Question Answering
   - 1.4 Information Retrieval
   - 1.5 Language Translation

2. Problem Solving
   - 2.1 Inference (Resolution-Based Theorem Proving, Plausible Inference, and Inductive Inference)
   - 2.2 Interactive Problem Solving
   - 2.3 Automatic Program Writing
   - 2.4 Heuristic Search

3. Perception
   - 3.1 Pattern Recognition
   - 3.2 Scene Analysis

4. Modeling
   - 4.1 The Representation Problem for Problem-Solving Systems
   - 4.2 Modeling Natural Systems (Economic, Sociological, Ecological, Biological, etc.)
   - 4.3 Hobot World Modeling (Perceptual and Functional Representations)

5, Learning and Adaptive Systems
   - 5.1 Cybernetics
   - 5.2 Concept Formation

6. Robots
   - 6.1 Exploration
   - 6.2 Transportation/Navigation
   - 6.3 Industrial Automation (e.g., Process Control, Assembly Tasks, Executive Tasks)
   - 6.4 Security
   - 6.5 Other (Agriculture, Fishing, Mining, Sanitation, Construction, etc.)
   - 6.6 Military
   - 6.7 Household

7, Games
   - 7.1 Particular Games (Chess, Go, Bridge, etc.)

Table 3

RESEARCH REPRESENTATIVE OF CURRENT COMPETENCE IN AI

1, Language Understanding

Winograd [I97I1: A system that understands reasonably complex declarative and imperative sentences about a limited (simulated robot) environment,

Badre LI972]: Charniak [I973]: Comprehension of children's stories in a reading text.

Thompson [I969]: Woods [I970]: Schank [I972]; Coles [I972]; Question-answering systems which can respond to queries posed in an impressive subset of! natural English about information stored in a relational data base.

Reddy-Vicens [I969]: Connected speech recognition with small vocabulary and limited number of speakers.

Reddy [I972]; Chess program to which moves are provided in the form of spoken chess notation. Uses knowledge of chess to decide phonetic ambiguities in favor of the more probable move.

2, Problem Solving—Automatic Programming

Strips, Fikes, et al. [I97I]: Uses resolution theorem prover to select subgoals for GPS-type, means-ends analysis. Used in SRI robot to specify a sequence of operations {Push, Go To, etc.) that accomplish a desired end-result,

Planner, Hewitt [I97I]: QA4, Rulifson, et al. [I972]: Conniver, Sussman, et al. [I972]: Procedures that accomplish a desired goal are retrieved by pattern matching. Planner provides an effective way nf organizing ad hoc knowledge about a problem domain.

Hacker, Sussman [I972]; Hacker is a programming system intended to emulate the programming style (try and debug) of its author. This ongoing work has thus far written successful programs for block-stacking with a simulated robot.

3, Perception

MIT, Minsky-Papert [I972]: Stanford, Falk [I972] hand/eye systems: Can understand scenes consisting of white polyhedral blocks, arbitrarily stocked on a black table top.

Keily [I970]; Fischler [1973]; Can find human faces in a cluttered room environment and recognize a limited number of those previously seen.

SRI mobile robot project, Hart, et al. [1973]: Current research intended to understand scenes depicting real-world office environments well enough for a mobile robot to be able to fetch specified objects.

4, Modeling

Forrester [I969] [I97I]; Urban and world models which indicate effects of present population trends.

McCarthy and Hayes [1969]: Hayes [I97I]; The problem of epistemology for a robot.

5, Learning and Generalization

Samuel [1959-72]: Parameter learning for move-evaluation in the game of checkers.

Winston [I970]: Learns the distinguishing attributes for concepts in the block world (e.g., an arch or tower) by forming structural descriptions of well-chosen examples.

Generalization of Strips Plans, Fikes, et al. [I972]: Generalizes a solution to a specific problem so that it can be used in whole or in part to subsequently solve similar problems (e.g., where the room names or objects differ).

Table 3 (Concluded)

6. Robotics (Integrated Problem-Solving and Perceptual Systems)

Stanford, Feldman [1971]: MIT hand/eye projects, Winston [I973]; Disassemble a tower of blocks and reassemble it elsewhere.

Stack blocks, Winston [1973]-and thread bolt, Paul [I973] using visual servoing.

SRI industrial automation project; Pick up selected objects from a moving conveyor belt.

Current work at each of the above projects involves the assembly and/or repair of simple structures (pumps, electronic circuit cards, etc.).

SRI, Nilsson [I969], robot project: Mobile vehicle which can navigate through a simple world of rooms and doorways to fetch large polyhedral objects. Currently attempting to extend these capabilities to enable the robot to do useful "office boy" tasks in a real, office environment.

University of Edinburgh, Ambler, et al.,[1973]: Stationary robot with mobile surroundings. Assemble simple toys (e.g. , cars) given a kit of pares and high level instructions about how to recognize and manipulate the various parts.

Japan: Assembly robot that can "read" blueprints, choose parts, and assemble them, Ejiri, et al.[1972]; "walking machine," Waseda University; mobile delivery robot, Toshiba Research Laboratory.

7. Game Playing

Chess program, Greenblatt [I969]: Plays Class B chess, master checker player, Samuel [I959], muster Kalah. Slagle [I97I], trivial GO, Ryder [I97I]

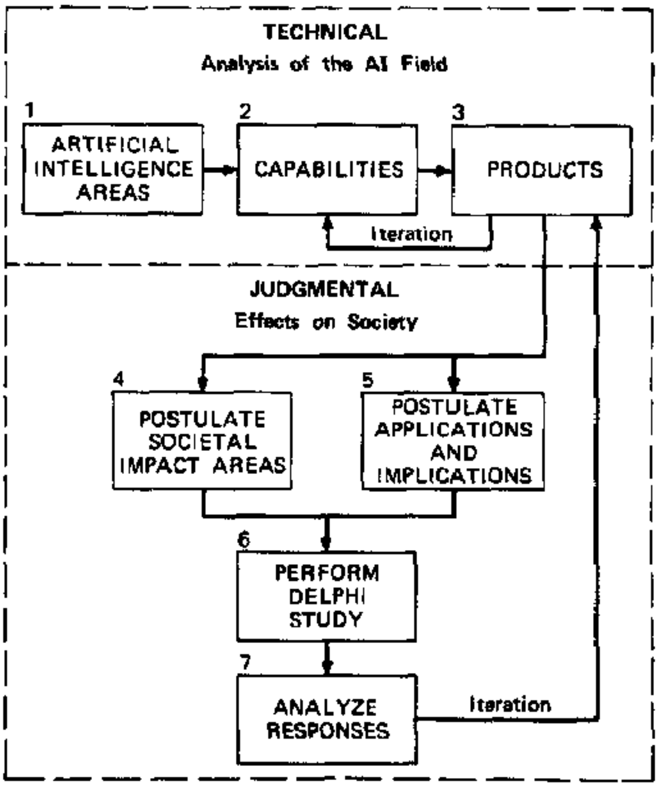

FIGURE 1 OVERALL METHODOLOGY USED IN THE STUDY

Table 4

PRODUCTS USED IN THE STUDY

1. Automated Inquiry System; An automatic information-retrieval system, using a man/machine dialog to determine user needs, which can search its data base to present the user with specific information or "facts," rather than references to other sources.
2. Automuted Intelligence System: General augmenter of human intelligence, capable of automatically monitoring ongoing streams of input data, coordinating facts, and making logical inferences to obtain insights and alert the human as appropriate.
3. Talking Typewriter; Voice typewriter, capable of converting spoken natural language into typewritten form in essentially real time, with an error rate equal to a human.
4. Automatic Language Translator: Language translating device capable of high-quality translation of text in one foreign language to another. (Both technical and commercial material.)
5. Automatic Identification System; System for automatically determining a person's identity by recognizing his voice, fingerprints, face, etc.
6. Mobile Robot: A military or police robot that can senrch for, locate, and deal with a prespecified subject or class of behavior patterns.
7. Machine/Animal Symblont: Tapping the brain of a living animal such as a bird or rut, to obtain preprocessed sensory inputs (visual, auditory, olfactory, etc.) to augment the capabilities of a mechanical device.
8. Automatic Diagnostician: System capable of interactive and/or automatic medical diagnosis based on querying the patient, examination of biological tests, etc.
9. Personal Biological Model: System periodically monitors patient's blood circulation, lung and heart function, muscle action, etc., to provide inputs to a personal biological model so that current status can be evaluated and the effects of medication and treatment can be simulated.
10. Computer Psychiatrist: A system in which the patient's written or verbal input is sufficiently understood that the system can legitimately counsel and advise the user for commonly encountered problems.
11. Computer-Controlled Artificial Organs: Capable ol replacing natural organs of the body (such as arms, legs, eyes, kidneys, heart, lungs) while still preserving homeostatic balance,
12. Computer Arbiter: A system in which adversaries enter their mutual complaints and the interactive system uses a data base of precedents and value rules to deliver either advice or a verdict.
13. Industrial Robot; Autonomous industrial robot capable of product inspection and assembly in an automated factory, using both visual and manipulative skills.
14. Voice-ResponBe Order-Taker; Capable of handling inquiries on order status, forming new orders, looking up data in a catalog, and indicating the result of its actions verbally.
15. Insightful Weather-Analysis System: A system which combines strict numerical analysis using measurements, with the type of insight used by a weather expert, e.g., analysis of two-dimensional weather maps, patterns, review of previous records, etc.
16. Insightful Economic Model: System which combines straightforward business statistics and input/output analysis with heuristic insights concerning the economy (unemployment, interest rates, monetary and fiscal policy, etc.) to predict the effects of economic policies.
17. Robot Chauffeur: Robot cars capable of operation on standard city streets and country highways, using visual sensors.
18. Robot Tutor; A system which can accept verbal and biosensed responses from the student and tailor the course of instruction to suit the pupil's needs. This would not be a simple-minded CAI-programmed learning device, but on a level with a good human teacher.
19. Universal Game Player: A system capable of playing Chess, Checkers, Kalah, Go, Bridge, Scrabble, Monopoly, etc., at a controllable level of proficiency, from master level to novice.
20. General Factotum: A robot servant, capable of smooth verbal interchange and versatile perception and manipulative activity in a household environment.
21. Creation and Valuation System; Capable of creative work in such areas as music, art (painting, sculpture, architecture), literature (essays, novels, poetry), and mathematics, and able to evaluate the work of humans.

The products postulated in this paper differ most strikingly from conventional computer applications in three ways. First, is the requirement for natural communication with humans via spoken or written natural language. This capability is prerequisite to virtually any application intended to provide mass personal access to computing power. Second is a problem-solving capability so that humans can specify a desired effect rather than having to supply a detailed procedure for obtaining that effect. Thirdly, many of the products require an ability to interact directly with their real-world environment through sensors that acquire information, and effectors that implement a desired change.

To establish the technical feasibility of the projected products, we compared the levels of competence that could be achieved with current AI capabilities (Table 5) with those required to achieve the stated performance goals. One might conclude from Table 5 that a reasonable start has been made for many of the products. It must be pointed out, however, that in the past it has proven very difficult to extrapolate the results of AI research. Most of the work cited in this table has involved limited domains of knowledge, small data bases, and ad hoc techniques. It is unknown at this time how many of the techniques are extendable to "real-life" data bases.

In predicting future capabilities in the artificial intelligence field, we were faced with the problem "flow do you know that a development in a subfield of AI represents significant progress?" Critics of the AI field such as Dreyfus (1972), and Fein (1964), have doubted the significance of much of the alleged progress in the field. In particular, Dreyfus has used the analogy of a man who, having succeeded in climbing a tree, claims that he has solved part of the problem of going to the moon.

There is no simple resolution to the problem, and indeed we cannot hope to settle the question here. (It may be that it is not possible to evaluate a development until some years have passed, and one sees the effects that the development has produced.) The following indicators of progress were useful in our evaluation:

Attainment. A project succeeds in doing something that was thought to be impossible or extremely difficult.

Theory-building. A project blends several apparently disjoint facts within a common framework that promises a productive line of research.

Unusual improvement. A significant improvement is made as far as speed, cost, memory size, etc. for an existing important accomplishment.

Impossibility proof. A demonstration that something can't work, or can't be performed in a certain manner closes a research path in which progress had been stymied.

Significant terminology or concept formation. The development of a language or representation upon which further work in the field can be founded.

Fundamental Problems in Achieving the Products

There are two types of difficulty in achieving the AI products, the first concerning present limitations of the AI field, and the other concerning the problem of formalizing areas of human endeavor which do not appear to be amenable to formalization. These two problem areas are discussed below.

Fundamental Problems in AI Technology— In examining current capabilities in the AI product areas, Table 5, three main problems appear to be:

(1) Acquiring, representing, structuring, and retrieving world knowledge,

(2) Attaining generality from ad hoc programs, and

(3) Problems in perception.

Acquiring, Representing, Structuring, and Retrieving World Knowledge—A crucial part of many of the AI products is the ability to acquire, represent. Structure, and retrieve large amounts of general knowledge about the world, ad hoe knowledge about specific problem domains, and self-knowledge regarding the system's own capabilities. The fundamental problem associated with large amounts of information is determining which subsot is relevant to the task at hand. This problem includes the choice of a represonlalion that presents the relevant data at an appropriate level of detail. The system must also be able to restructure the relevant knowledge so that, it can be efficiently applied to a specific problem. Knowledge is typically acquired in an incremental manner; as new facts are encountered in the performance of tasks, they are integrated-- whenever possible as generalisations of existing information, The system must bo able to determine when necessary knowledge isn't available, and how it can best be acquired.

Some other outstanding issues include the retrieval of knowledge storod in procedural form, the representation of common sense," and the utilizution of plausible inference.

Generality from Ad Hoc Programs—AI programs which aspire to great generality in a problem domain have most. often exhibited unaeceptably poor performance. Consequently, many AI workers have adopted the view that intelligence is best embodied in the collective etfect of many ad hoc procedures, each of which solves a particular problem in a limited context. Generality is then attained by accumulating a library of such procedures and selecting the most appropriate ones for any given situation. This approach is reasonable given the limited number of situations likely to arise in the restricted domains of present systems, When such systems fail, a human programmer can use his knowledge of the task requirements, the problem domain, and the machine's capabilities to add a routine appropriate for the new situation. Thus, one approach to generality is to provide a future system with similar knowledge so that it can emulate this problem-solving function of a human programmer. The user of such a system might be able to guide the machine in constructing its own programs,

Table 5

CAPABILITIES IN DERIVED PRODUCT AREAS
ACHIEVED (ACHIEVABLE) WITH CURRENT A.I. EXPERTISE

1. <u>Automated Inquiry System</u>

Automatic inquiry systems with inferential reasoning are currently practical for small data bases and for large well structured ones. Effective searches may require the human user to refine his query interactively.

2. <u>Automated Intelligence System</u>

Systems have been built that augment intelligence by clustering data, identifying critical factors in multivariant analysis (factor analysis), and by searching for inferential chains that relate two arbitrary facts in a symbolic data base. However, such systems cannot acquire their own knowledge. Inputs must be manually filtered for relevance and supplied in an appropriate internal representation.

3. <u>Talking Typewriter</u>

A voice typewriter could today provide reliable phonetic transcription for about 100 to 500 arbitrary words if spoken one at a time. Transcription of continuously flowing speech could be accomplished for the same number of words if the discourse were confined to a limited domain for which the machine had adequate semantic knowledge to resolve segmentation ambiguities. Present pattern classification techniques have proven unable to discriminate more than a few hundred categories. (This word limitation could be circumvented by the crude expedient of using spelled speech.)

4. <u>Automatic Language Translator</u>

Low quality automatic language translation is already commercially used to obtain crude but readable abstracts of foreign technical papers. High quality translation that captures all intended subtleties requires that the input utterance be first understood and then reconstituted in the second language. Such translation is thus feasible only in the limited semantic domains which have been successfully handled by existing language understanding programs.

5. <u>Automatic Identification Systems</u>

These systems now suffer from the same pattern classification problems that limit the voice typewriter. It is difficult to build systems that discriminate more than 50 to 100 categories (e.g., faces, fingerprints, voices). An interactive system in which humans provide descriptions of important features for a mechanized decision maker could perhaps be made practical in the near future.

6. <u>Mobile Robot</u>

Autonomous real-world robots are not yet feasible, although at least two laboratories are actively pursuing the following goals: (a) a mobile robot that plans and executes tasks requiring polyhedral solids to be pushed between otherwise empty rooms (future operation will include a real office environment); (b) a robot vehicle that can circumnavigate a closed roadway system.

7. <u>Machine/Animal Symbiont</u>

Several research projects are currently attempting to do pattern classifications using inputs obtained from the neurons of animals and even from human EEG signals. A small number of gross patterns can be discriminated, but the results at present are far too crude for practical use.

8. <u>Automatic Diagnostician</u>

Many diagnostic techniques have already been successfully automated, including chemical laboratory analyses, cell counting, x-ray and electrocardiogram interpretation, etc. Several interactive systems have been built which accept symptoms, test results, and patient history, and return diagnoses or specific requests for more data. Progress in this field is limited because a machine cannot yet observe a patient directly and because each application must be programmed from scratch. More fundamental problems exist due to the crude state of current medical knowledge.

9. <u>Personal Biological Model</u>

Many critical biological functions are already automatically monitored in hospital intensive care wards. The problem of making such capabilities available to everyone on a continuing basis is mostly one of economics. Further work by biologists in modeling the human system is needed to adequately simulate the effects of medication.

10. <u>Computer Psychiatrist</u>

Computerized psychiatry is not yet feasible. After the problem of natural language understanding is solved, attention can then focus on the equally challenging task of representing the semantics and pragmatics of human interpersonal relations.

Table 5 (Concluded)

11. Artificial Organs

Although the problem of controlling artificial organs with feedback from the body is currently more in the domain of control systems than A.I., future implementations may involve A.I, concepts. Present experimentation with feedback controlled pacemakers and, also, the life support systems in satellites can be viewed as first examples of this promising technique.

12. Computer Arbiter

A system has been built that simulates decisions of the supreme court with some degree of success. The computer bases its decisions on precedents of law and biases of the justices disclosed in past decisions. If those biases were replaced by a formalized statement of the principles of justice, one might have the basis for an automatic arbiter. However, such formalization has not yet been attempted.

13. Industrial Robot

Robot manipulators are being used in increasing numbers on automobile assembly lines, to do repetitive tasks like spot welding, which can be preprogrammed and which operate without feedback. The addition of simple, visual and tactile sensors would significantly broaden the range of application. For example, the General Motors Research Lab successfully demonstrated a system that could mount wheels on a hub, using visual techniques to align the wheel with the studs.

14. Voice Order-Taker

A voice order-taker, like the voice typewriter, can now be built to handle about 500 items. If the catalog is mare extensive, items would have tn be specified by giving code numbers. Limited voice response by such a system, using combinations of prestored words and phrases, is now available.

15. Insightful Weather; Insightful Economic Model
16.

Predicting weather or economic activity is still as much an art as a science. The main problem with automating such activities is that the experts themselves have yet to agree on the most important decision criteria or even on the most relevant input parameters. It is perhaps more feasible to simulate a particular expert, but even hero the expert probably would be unable to express formally the subjective criteria he uses to reach decisions. The machine is further handicapped in areas like weather analysis, where important inputs are only available in pictorial form (e.g., cloud cover photographs). In these domains, an interactive system could be used, where the machine's decision is based on manually interpreted input data.

17. Robot Chauffer

Sensors and systems are now available that can augment driving skills (e.g., collision avoidance radar, anti-skid braking computer, etc.). Moreover, laboratory vehicles have been made to follow a white line and programs have been written to detect the edge of a road and to discriminate planar shadows from real obstacles.

18. Robot Tutor

The best existing CAI systems allow mixed initiative interactions, wherein the student can alter the course of instruction by asking the machine unanticipated questions about this subject area to test his own under—standing. Researchers are currently planning systems which will ask probing questions to model a student's comprehension of a subject area, and then plan a customized tutorial strategy intended to transform the existing conceptual structure into a desired one. The success of such research is very dependent on progress in natural language understanding and psychological investigations of knowledge representations and learning.

19. Universal Game Player

Many game playing programs have been written in the course of studying heuristic programming and learning. Machines can play champion level checkers and dominoes; reasonable chess (B-level), Kalah, and Scrabble; and poor Go. However, each game was individually programmed using heuristics deduced from careful human introspection and representations painfully derived to allow effective utilization of those heuristics. It is still premature to contemplate a general game player that can be taught to play a new game as one would teach a human opponent, or significantly improve its internal representations, and thus performance, as a result of playing.

20. General Factotum

General-purpose humanoid-type robots remain, for the time being, in the realm of science fiction. However, until very recently, no A.I. research has been seriously directed toward this goal. The mobile robot effort described under Product 6 could be viewed as a beginning.

21. Creation and Evaluation System

The evaluation of human artistic endeavors is perhaps the most subjective of all human judgements. Consequently, the criteria of judgement are among the most difficult to formalize as computer algorithms, except for a few general notions like harmony or color balance. Perhaps the computer's inability to evaluate its own creative efforts in human terms explains the primitive state of computer-generated art and poetry. (See Gip and Stiny [l973J for a discussion of aesthetic systems.)

for example by suggesting promising heuristics, providing necessary knowledge, and evaluating performance on test cases. Such an interactive problem solver should have a motive-guessing capability with which to attempt solutions to incompletely formulated problem requirements. Indeed, a large part of the solution to many problems (like most of those arising in AI) lies in properly formulating the problem.

Perception Problems—A number of difficult problems arise in real-world perception (e.g., speech, vision) that for the most part have been systematically eliminated from the limited domains of present systems. These problems include sensory overload which conceals objects of interest in a torrent of irrelevant detail, the related problem of partitioning interesting "figures" from the general background, forming internal representations of complex or amorphous real-world objects, and finally matching such descriptions against a large data base of known objects. The foremost problem appears to be in partitioning the input into meaningful entities in the context of the current problem. In vision, objects of interest must be isolated from the background. In Speech, a continuous waveform must be divided into segments corresponding to words. The problem is well known to Gestalt psychologists who have observed that the interpretation of parts or features of an object is often determined by the interpretation of the whole. Thus machines, in proceeding from the parts to a whole can be overwhelmed by the combinatorics of examining every alternative interpretation for each part. Recently, one of the authors had some success in applying dynajnic-prograramiTig-tj'pe techniques to reduce the combinatorics of a brute-force search for the best interpretation of a scene, according to a model of the desired object. It may be practical to use such techniques in a general scene analysis system, to follow up hypotheses resulting from a preliminary organization (Fischler, 1973),

The Problem of Value Formalization—Several of the products postulated require that a formalization of aesthetics, mores, ethics, and other judgmental areas be available. For example, Product 12, the computer arbiter, requires not only the ability to understand the subtleties of human emotion,* but must also have available some codification of the moral, ethical, and legal rules of society ao as to be able to give advice, or render a verdict. The creation and valuation system. Product 21. requires a formalism which somehow expresses man's concepts of excellence in the arts.

There is little to report in the way of progress in these areas. Aside from some work in music by computer, von Foerster and Beachamp (1969), some elementary attempts at computer poetry and prose, and some rough analysis of court decisions, there is little hope (or danger, depending on one's point of view) that the problem of formalizing such areas of human endeavor will be solved in the near term.

Even if such formalization could be made, society would be faced with the embarrassing problem of revealing the inconsistencies inherent in any organization of of human beings. We know that there are many situations in which some people are considered "more equal" than others. For example, white collar fraud involving theft from a company is frequently dealt with less severely than blue collar theft from the same company. Appearances and attitudes often affect judgments, as do family connections or personal acquaintances. Although the adversely affected segment of the pbulic seems to be willing to tolerate such inequities, due to ignorance, lack of power, or indifference, one doubts whether the same passive behavior would occur if an attempt was made to incorporate explicitly the present defects of society into an automated arbitration system. A conflict would undoubtedly result with those having a vested interest in the status quo,

4. Societal Implications

Decisions concerning the rate of technological development, time of appearance of a product, and the societal implications are judgmental in nature, and for this portion of the study we used the Delphi technique, a method for systematically soliciting and collating informed judgments on a particular topic. Under this procedure, participants respond to a series of questionnaires interspersed with summaries of the responses by group members to earlier questionnaires. The method differs from simple polling procedures in that the feedback obtained from the respondents allows the questions to be modified (e.g., a question may be too vague), and tends to prevent "snap-judgments" since the respondent can see how his answer stands with respect to the other participants.

As indicated in Section 2, we first partitioned the products into the societal areas which they most affect: human information processing, security, health, law. production and manufacturing, commerce, education, social interaction, and aesthetics. We then derived applications and implications for each product, where an application is a statement of how the product would be used by society, and an implication is the effect on society of these applications. Thus, for Product 1.1, Industrial Robot, one of the applications listed was Robots for tedious assembly line tasks," while one of the implications was, "displacement of blue-collar workers."

We used a format similar to that of De Brigard and Helmer (1971), as shown in Fig. 2. For each product, we indicated the social impact sector, the applications and implications, and the AI categories. The information solicited was:

(1) What revisions should be made in the product list?

(2) What revisions should be made in the applications and implications for each product?

(3) What is the potential significance of each product?

Either via natural language or by means of sensors, such as a "lie detector" sort of device.

See Lipinaki et al (1972) for an interesting discusslot of expert interrogation.

(4) What is the respondent's best estimate of the prototype and commercial date of the product?*

(5) What is the likelihood of each implication and application?

(6) What is the desirability of each implication and application:

The questionnaire was sent to members of the International Joint Artificial Intelligence Council and to other experts in the AI field throughout the world. A total of sixty questionnaires was mailed out, and twenty-one responses were obtained. An additional mailing was sent *to* the San Francisco IEEE Systems, Man, and Cybernetics Society, and a total of twenty questionnaires were received. The analysis of results given below is based on two rounds of responses from the first group; the complete results will be presented in a forthcoming IEEE report. Where interesting differences between the two groups occurred, they are also briefly noted.

Results of the Delphi Study

The Delphi comments were reviewed manually, while the time estimates and likelihood and desirability of the products were tabulated by a simple computer program. These results were presented directly on the questionnaire, as shown in Fig. 2 tor review by the participants. The left side of the "house-like" figure represents the lower quartile, the roof peak is the median, and the right side is the upper quartile.

General Observations on the Results-~Due to lack of space, the complete 13 page Delphi Questionnaire will not be reproduced here. Instead an overall summary of the Delphi responses is given in Table 6, The results were obtained by weighing each response with the self-ranking of expertise supplied by each respondent for each product. ■ In the table, products were first, grouped according to the median rating they received for potential significance (high, medium, and low), and within these significance groups, they were listed in order of prototype date, The overall desirability is a rough summary of the responses concerning desirability for each product.

The following results are of interest:

(1) Product potential Significance

Half of the AI products were thought to be of high potential significance and most of the rest fell in the medium potential significance category.

(2) Rate of Progress

Respondents were generally very optimistic concerning technology development. As shown in Table 6, thirteen of the products have median dates for prototype product development of 1985 or earlier. In the detailed responses, only one product had an upper quartile in the "never" range (P7>. It is interesting to note, however, that in going from the first to the second round, there was a very slight, but noticeable, shift toward conservatism in product dates as well as the expected observation of a reduction in the width of the houses indicating a sharpening of concensus.

(3) Prototype vs. Commercial Dates

For the products having prototype dates of 1985 or earlier there is a fairly systematic translation of 5 1 2 years as one moves from median prototype to median commercial date.

(4) Likelihood and Desirability

There is also a strong mood of optimism among the respondents regarding the likelihood and desirability of the various products. Only five products had implications whose desirability was deemed "very detrimental" (P1, P2, P5, P6, and P18), and only one product (P7) had an nverall detrimental rating.

(5) Products Potentially Dangerous to Society

Singled out for special attention were products that were simultaneously probable and detrimental for one of their applications or implications, i.e., hud a median estimate of 1ikelihood in the range (possible, probable, highly prohable) and a median desirability in the range (detrimental, very detrimental). These products and their associated problem area(s) are indicated below;

- PI - Automated Inquiry System (opportunity for censorship)
- P2 - Automated Intelligence System (use by government to monitor actions of citizens)
- P5 - Automatic Identification System (use by government to monitor the civilian population)
- P6 - Mobile Robot (robot soldier; aggressive actions by some nations)
- P7 - Machine/Animal Symbiont (use for weapons systems sensors; alteration of human values)
- P10- Computer Psychiatrist (use of system by government to influence behavior)
- P18- Robot Tutor (possible indoctrination of Students by government)
- P20- General Factotum (possible emotional impact on children)

Note that in asking for judgments concerning time of commercial appearance, we are asking not only whether a product can be made, but also when the commercial forces of need plus profitability will cause a product to appear on the market.

A weight of 1 through 5 was applied to (I) unfamiliar, (2) casually acquainted, (3) familiar, (4) quite familiar, and (5) expert ratings.

Table 6

SUMMARY OF DELPHI RESULTS

Products	Median Prototype Date	Median Commercial Date	Desirability (*)
High potential significance			
P5 - Automatic identification system	1976	1980	+
P8 - Automatic diagnostician	1977	1982	++
P13- Industrial robot	1977	1980	++
P1 - Automated inquiry system	1978	1985	+
P9 - Personal biological model	1980	1985	++
P11- Computer-controlled artificial organs	1980	1990	+
P18- Robot tutor	1983	1988	+
P16- Insightful economic model	1984	1990	++
P2 - Automated intelligence system	1985	1991	0
P20- General factotum	2000	2010	+
Medium potential significance			
P14- Voice response order-taker	1978	1983	+
P15- Insightful weather analysis system	1980	1985	++
P3 - Talking typewriter	1985	1992	+
P6 - Mobile robot	1985	1995	0
P4 - Automatic language translator	1987	1995	+
P12- Computer arbiter	1988	1995	+
P10- Computer psychiatrist	1990	2000	0
P17- Robot chauffeur	1992	2000	+
P21- Creation and valuation system	1994	2003	+
Low potential significance			
P19- Universal game player	1980	1985	+
P7 - Animal/machine symbiont	2000	2010	-

(*)
++ Very favorable
+ Favorable
0 Balanced
- Detrimental

Table 7

TIME PREDICATIONS OF AI EXPERTS (DERIVED FROM MICHIE [1973])

	Estimated Years				
	5	10	20	50	> 50
Computing systems exhibiting intelligence at adult human level	0	1	16	20	26
Significant industrial spin-off from AI research	30	28	4	1	2

(6) Differences Between AI Council and SMC Results

Although not universally the case, AI Council respondents tended to be slightly more optimistic and have a somewhat stronger concensus than the SMC group, both as far as prototype and commercial development dates, and the likelihood and desirability of the applications of the products. Perhaps this reflects the fact that many members of the AI Council are actually engaged in the research themselves and have a vested interest in the results. Yet one cannot disregard the collective opinion of those experts who are closest to the field.

Relation to other Studies—There have been many studies concerned with the impact of technology on society, e.g., the effect of computers and data banks on privacy, Westin (1972) and the effects of automation on employment, Borodin and Gotlieb (1972), Our study was not designed to quantify the impact of AI products on society; we can only obtain informed opinion concerning the feasibility of a comprehensive set of AI products and the general nature and spectrum of such impact. Thus, it is similar in spirit to the Delphi study of DeBrigard and Helmer (1971) which examined the societal consequences of twenty physical and biological breakthroughs.

Another study of interest is reported by Michie (1973), in which 65 experts in artificial intelligence from England and the U.S. were queried. Two estimates of interest, concerning level of AI ability and the impact of AI on industry are reproduced in Table 7. The entries in the table represent the number of experts who responded to the items for each time estimate. It will be noted that almost all of the experts felt that computing systems exhibiting human intelligence are at least twenty years away, and that significant AI industrial spin-off will occur in less than ten years.

5. Conclusions

Two types of conclusions are given below;

(1) General societal implications of the AI products

(2) Recommendations for an AI review board.

Societal implications

In examining the societal implications, we find that the most important possible long-term effects of the increasing development and application of AI technology are as follows:

(1) A decreased need by most persons for direct contact or interaction with other human beings—That is, it may be possible for intelligent machines of the future to supply not only intellectual stimulation or instruction, but also domestic and health care, social conversation, entertainment, companionship, and even physical gratification.

(2) The need arises to formalize algorithmically some of the ethical and empirical rules and tradeoffs that society observes implicitly, but which are rarely made explicit—For example, it may be necessary for a robot to have available rules which dictate how to trade-off life for property, e.g., when is it allowable to wreck an automobile to avoid killing an animal?

(3) The erosion or elimination of uniquely human activities which typically distinguish man from other animals or machines—There will be a profound psychological impact as humans begin to wonder whether there are human capabilities which are truly unique.

(4) There will develop application areas which are not feasible without machine aid, such as accurate weather prediction or economic forecasting.

(5) There will be many mundane activities which will be cheaper or otherwise more desirable fur machine accomplishment, with the resultant displacement of human workers.

(6) The automation of the mechanics of government, education, law, and health care could imply a concentration of decision-making responsibility which could in turn lead to a powerful elite,

Societal Controls

In examining some of the implications of the AI-oriented products, we noted that some of these products will bring with them questions of safety, privacy, and ethics.

Safety—The safety factors are obvious: Is a designer to be allowed to pr€>duce autonomous devices which are free to navigate over city streets and in homes and offices without some review procedure? for devices which require an algorithmic formulation of Judgmental factors, who shall decide whether the formulation is anti-social? Furthermore, how can such formulations be validated, and by whom?

AI development has now reached a stage where such questions are no longer in the science-fiction realm. For example, some of the mobile robots already developed have enough speed and power to cause possible injury to humans encountered by the robot in the course of performing its tasks. Any device which includes the characteristics of mobility, speed, strength, and unpredictability can be extremely hazardous, especially in an uncontrolled environment or where children are present.

Privacy—Some of the products postulated include the ability to interrelate large and diverse data bases using advanced deductive techniques. Proper safeguards must be imposed on the collection and use of ouch data base systems to insure that the right of privacy of the individual is respected.

Ethics Problems of ethics arise in automated counseling systems which advise a student concerning career choices. Has anyone verified that the counseling program is not biased against certain groups, or unaware of the status of certain career fields? Another example of an ethical question arises in ELIZA-like systems, Welzenbaum (1966), for psychiatric counseling which the uninitiated individual assumes has depth and sophistication, not realizing that the program operates using very limited word patterns and phrases.

Somehow, these potential dangers to society must be communicated to decision makers in both the public and private sectors. Baram (1973) has discussed technology assessment and social control,

> "It is now time to...develop a coherent framework for the social control of technology and ensure that forthcoming processes of technology assessment and utilization will be systematic and humane."

Perhaps one vehicle for keeping the public and decision makers informed about AI devices might be an AI Ethical Review Board, similar to the review boards used in the medical profession, which deal with safety and ethical problems. Members of the board would be selected from the AI community, but would also include legal, political, and other advisors, as required. The board would recommend legislation, establish standards, suggest safety guidelines tor the use of robots in human environments, and advise investigators in the AI field by means of journal notices.

Acknowledgements

The overall Technological Forecasting and Assessment Project of the San Francisco Chapter of the IEEE SMC Society was headed by 0. Firschein; L, S. Coles directed the Delphi Study Project. The product list and applications and implications for the Delphi were derived jointly by the authors.

Roy Amara and Andy Lipinski of the Institute for the Future were invaluable in educating us to the ways of Delphi. Monroe Pastermack and Peter Schwartz of the World Future Society were also quite helpful in this regard. Members of the AI Center at SRI were consulted throughout the project.

The initial funds for the study were provided by the IEEE, and the analysis and computer processing were funded under National Science Foundation Grant GJ-37696-

References

A. P. Ambler, et al., "A Versatile Computer-Controlled Assembly System," Proceedings Third International Joint Conference on Artificial Intelligence, Stanford University, Stanford, California (1973).

N. A. Badre, "Computer Learning from English Text," Memo ERL-M372, Electronics Research Laboratory, College of Engineering, University of California, Berkeley, California, (December 20, 1972).

M. S. Baram, "Technology Assessment and Social Control," Science, Vol. 180, No. 4085, pp. 465-473 (Hay 4, 1973).

A. N. Broers and M. Hatzakls, "Microcircuits by Electron Beam," Scientific American, pp. 34-44 (November, 1972).

M. J. Cetron and B. Bartocha, Methodology of Technological Assessments, (Gordon & Breach, 1972).

E. Charnlak, "Toward a Model of Children's Story Comprehension," Artificial Intelligence Technical Report 266, Massachusetts Institute of Technology, Artificial Intelligence Laboratory (December 197Z).

L. S. Coles, "The Application of Theorem Proving to Information Retrieval," Proceedings Fifth Hawaii International Conference on Systems Sciences, Honolulu, Hawaii (January 1972).

R. DeBrigard and 0. Helmer, "Some Societal Consequences of Technological and Scientific Developments," Industrial Application of Technological Forecasting, M. J. Cetron and C. A. Ralph, eds., (Wiley 1971).

H. Dreyfus, What Computers Can't Do: A Criticism of Artificial Reason, (Harper & Row, New York, 1972).

E. Ejiri, et al., "A Prototype Intelligent Robot that Assembles Objects from Plan Drawings," IEEE Trans. Cornp., pp. 161-170 (February 1972).

G. W, Ernst and A. Newell, "GPS: A Case Study in Generating & Problem Solving," ACM Monograph, (Academic Press, New York, 1969).

G. Falk, "Scene Analysis Based on Imperfect Edge Data," A.I. Journal (Summer 1972).

E. Feigenbaum and J. Feldman, Computers and Thought, (McGraw Hill, New York, 1963).

L. Fein, "The Artificial Intelligencia," IEEE Spectrum, Vol. 1, No. 2, (February 1964).

J. A. Feldman, et al. , "The Use? of Vision and Manipulation to Solve the Instant Insanity Puzzle," Proceedings Second International Joint Conference on Artificial Intelligence, London (1971).

J. A. Feldman, et al., "Recent Developments in SAIL—An ALGOL-based Lanugage for Artificial Intelligence," Proceedings AFIPS Conference FJC, Vol. 41, Part 2 (1972).

R. E. Fikes, et al., "Learning and Executing Generalized Robot Plans," Artificial Intelligenco, Vol. 3, No. 4 (Winter 1972).

R. E. Fikes and N. J. Nilsson, "STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving, Artificial Intelligence, Vol. 2, No. 314 (Winter 1971).

M. A. Fischler and R. A. Elschlager, "The Representation and Matching of Pictorial Structures," IEEE Trans, on Computers (January 1973) .

C. C. Foster, "A View of Computer Architecture," Communications of the ACM, Vol. 15, No. 7, pp. 557-565 (July 1972).

J. M. Forrester, Urban Dynamics, (MIT Press 1969).

J. M, Forrester, World Dynamics, (Wright-Allen Press, Cambridge, Massachusetts 1971).

J. Cips and G. Stiny, "Aesthetic Systems," AIM-189, Computer Science Department, Stanford University, Stanford, California (January 1973).

C. C. Gotlieb and A. Borodin, Social Issues in Computing (Harper and Row 1973).

R. Greenblatt, "The Greenblatt Chess Program," Proceedings of the FJCC (1969).

P. E. Hart, et al., "Artificial Intelligence—Research and Applications," Annual Technical Report, Artificial Intelligence Center, Stanford Research Institute, Menlo Park, California (1972).

P. Hayes, "A Logic of Action," Machine Intelligence 6, B. Meltzer and D. Michie, eds., pp. 495-520 (American Elsevier, 1971).

C. Hewitt, "PLANNER: A Language for Proving Theorems in Robots," Proceedings of the First International Joint Conference on Artificial Intelligence (1969).

M. Kelly, "Visual Identification of People by Computer," Memo AI-130, Computer Science Department, Stanford University, Stanford, California (July 1970).

A. J. Lipinski, et al., Computer-Assisted Expert Interrogation on Current Methods Development," Proceedings First International Conference on Computer Communication pp. 147-154 (October 1972).

J, McCarthy and P. Hayes, "Some Philosophical Problems from the Standpoint of Artificial Intelligence," Machine Intelligence 4, B, Meltzer and D. Michie, eds., (Edinburgh University Press, 1969).

J, McCarthy, "Programs with Common Sense," (November 1958) and "Situations, Actions, and Causal Laws," [1963] reprinted together as Ch. 7 of Semantic Information Processing, M. L. Minsky, ed. (MIT Press 1968).

D. Michie, "Machines and the Theory of Intelligence," Nature, p. 241 (February 23, 1973).

M, L, Minsky, ed., Semantic Information Processing, (MIT Press 1968).

M. L. Minsky and S, Papert, "Artificial Intelligence Progress Report," Memo AI-252, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts (1972).

A. Newell, "Limitations of the Current Stock of Ideas about Problem Solving," Proceedings of Conference on Information Handling, A. Kent and 0. Taulbee, eds., pp. 195-208 (Spartan Books, 1965).

N. J. Nilsson, "A Mobile Automaton: An Application of Artificial Intelligence Techniques," Proceedings of First International Joint Conference on Artificial Intelligence, Washington, D.C., pp. 509-520 (1969).

R. L. Paul, "Modelling, Trajectory Calculation and Servoing of a Computer Controlled Arm," AIM-177, Artificial Intelligence Laboratory, Stanford University, Stanford, California (1972).

D. R. Reddy, et al., "The Hearsay Speech Understanding System," Proceedings Third International Joint Conference on Artificial Intelligence, Stanford, California, (1973).

J, L. Ryder, "Heuristic. Analysis of Large Trees as Generated In the Game of Go," A.I. Memo 155, Artificial Intelligence Project, Stanford University, Stanford, California (1971).

J. F. Rulifson, et al., "QA4: A Procedural Calculus for Intuitive Reasoning," Artificial Intelligence Center, Project 8721, Stanford Research Institute, Menlo Park, California (November 1972).

A. J. Samuel, "Some Studies in Machine Learning Using the Game of Checkers," IBM Journal of R and D, Vol. 3, No. 3, (1959).

R. C. Schank, "Finding the Conceptual Content and Intention in an Utterance in Natural Language Conversation, Proceedings Second International Joint Conference on Artificial Intelligence (British Computer Society, London 1971).

J. R. Slagle, "Artificial Intelligence: The Heuristic Programming Approach," Programs that Play Kalah, Section 2.5, pp. 26-28 (McGraw Hill, New York 1971),

G. Sussman, "Hacker, Some Preliminary Results," MIT Artificial Intelligence Laboratory (1972).

G, Sussman and D. V. McDermott, "From PLANNER to CONNIVER—A Genetic Approach," Proceedings AFIPS Conference Vol. 41, Part 2 (FJCC 1972).

F. B. Thompson, "English for the Computer," Proceedings FJCC, pp. 349-356 (1966).

F. B, Thompson, et al., "REL: A Rapidly Extensible Language System," Proceedings 24th National ACM Conference (1969).

M. Turoff, "Delphi and Its Potential Impact on Information Systems," FJCC, Vol. 39 (AFIPS Press, Montvale, New Jersey 1971).

H. von Foerster and J. M. Beauschamp, Music by Computer, (Wiley 1969).

J. Weizenbaum, "On the Impact of the Computer on Society," science, Vol. 176, No. 609 (1972).

J. Weizenbaum, "ELIZA, A Computer Program for the Study of Natural Language Communication between Man and Machine," Comm. ACM, Vol. 9, No. 1, pp. 36-45 (January 1966).

A. F. Westin, ed., Databanks in a Free Society, (Quadrangle Books, New York 1972).

T. Winograd, "Understanding Natural Language," Cognitive Psychology, Vol. 3, No. 1 (January 1972).

P. H. Winston, "Learning Structural Descriptions from Examples," MAC-TR76, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts (September 1970).

P. H. Winston, "The MIT Robot," Machine Intelligence, Vol. 7 (Edinburgh University Press 1973).

W. A. Woods, "Procedural Semantics for a Question-Answering Machine," Proceedings AFIPS Conference FJCC, Vol. 33, Pt. 1 (1968).

Appendix A

The Delphi Methodology. We found that the Delphi methodology has the following advantages and disadvantages.

The problem of participation. There are two basic problems involved in the solicitation of responses on a voluntary, nonpayment basis: (1) approximately five hours of expert time is required to supply answers for all three rounds of the Delphi, and most experts do not have much time to spare for such extra-curricular studies, (2) there is the problem of the expert who disagrees with parts or all of the Delphi structure or orientation, but does not feel that the effort required to revise the questionnaire is worthwhile. Such an expert may respond to the questionnaire, but in a half-hearted manner.

The problem of evaluating expertise. A basic problem that arises in a Delphi study is the individual's rating of his expertise. Because the true expert may rate himself modestly, while the novice inflates his capabilities, we find a tendency for the responses to cluster in the "moderately expert" range. We thus lose the greater importance of the true expert, while inflating the importance of the novice's response. Although there does not seem to be any valid way to overcome this difficulty, an attempt has been made to 'calibrate' the responses using a technique developed by Lipinski [Andrew Lipinski, Institute for the Future, Menlo Park, California, Personal Communication].

Delphi as a communications tool. As has been indicated by Turoff (1971), Delphi can be considered as a process which allows the establishment of a meaningful group-communication structure. The questionnaire then serves as an entree to the expert, and enables response to be obtained that would ordinarily not be available. We found that communication is maximized if a personal interview is available after the expert has finished a round of the questionnaire, because comments and analyses which would not appear in written form could be captured by the study team.

Delphi as an organizing tool. We found the categorization and organization required to derive the questionnaire a very useful exercise in helping to analyze the field of artificial intelligence. Because we are forced to derive a meaningful product list, we had to examine carefully the AI capabilities, postulate products, and iterate the capabilities-product process until a satisfactory list was obtained. After the questionnaire was prepared in draft form, it was tried on A.I experts at SRI, and on Delphi experts at the Institute for the Future, Menlo Park, California, and Pacific House Associates, Palo Alto, California. Their suggestions lead to several modifications in the questionnaire and the proposed product list. Thus we found the preparation of the Delphi study a unique organizational tool, one which forced us to confront the basic problem areas in AI in a more efficient manner.

SOCIAL IMPACT SECTOR: Human Information Processing — 2nd Round

(Note: Occurrences of underlining throughout the second-round product statements, applications, and implications has been used to denote revisions suggested by the first-round respondents.)

PRODUCT	MIGHT RESULT IN		Likelihood Of Application or Implication: Highly Probable	Probable	Possible	Doubtful	Almost Impossible	Desirability of Application or Implication: Very Favorable	Favorable	Balanced	Detrimental	Very Detrimental
1. AUTOMATED INQUIRY SYSTEM An automatic information retrieval system, using a man/machine dialog to determine user needs, which can search its data base to present the user with specific information or "facts" rather than references to other sources. Potential significance: ☒ high; ☐ medium; ☐ low. AI Categories: 1.2, 1.3, 1.4, 2. Expertise-level ____ Prototype Date ____ Commercial Date ____ 70 80 90 70 80 90	A1:	Library fact and reference retrieval										
	A2:	Facility which provides the user with editing, file manipulation and sharing, reminder, and current awareness capabilities.										
	I1:	Increased utility of data bases, since the data is better organized and retrieved.										
	I2:	Benefits to science, technology, and the arts due to better data dissemination.										
	I3:	Greater opportunity for censorship.										
	I4:	Might displace persons in secretarial service, technical aid, editorial assistance, etc.										
	I5:	Decreased need for conventional information sources, such as newspapers and magazines.										
DIFFICULTIES ANTICIPATED												
COMMENTS												
2. AUTOMATED INTELLIGENCE SYSTEM General augmenter of human intelligence, capable of automatically monitoring ongoing streams of input data, coordinating facts, and making logical inferences to obtain insights and alert the human as appropriate. Potential significance: ☒ high; ☐ medium; ☐ low. AI Categories: 1.-6. Expertise-level ____ Prototype Date ____ Commercial Date ____ 80 90 80 90 2000	A1:	Applications in social and natural sciences, medicine, law, publishing, government, etc. to research broad topics and policies.										
	A2:	Use by government and other organizations to determine the effects of policy changes on world events so that appropriate actions can be taken.										
	I1:	Use by government to monitor actions of citizens.										
DIFFICULTIES ANTICIPATED												
COMMENTS												

FIGURE 2 EXAMPLE OF THE DELPHI QUESTIONNAIRE