The Director of Central Intelligence

Washington, D.C. 20505

Intelligence Research & Development Council

File

Executive Registry

84-4001

12 July 1984

STAT STAT

MEMORANDUM FOR: [redacted] Executive Director of Central Intelligence

FROM : Philip K. Eckman
Chairman, AI Steering Group

SUBJECT : 1983 AI Symposium Summary Report

STAT

1. Enclosed is a summary report of our 1983 Symposium on Intelligence Applications for Artificial Intelligence. The unclassified report is intended to be a synopsis of the major themes and issues which emerged during our Symposium last December.

2. In some cases, detailed viewgraphs from individual presentations were provided in the Proceedings which were distributed at the Symposium. Video tapes of selected key speakers are available through the CIA Self Study Center, [redacted] A complete set of audio tapes was also recorded. STAT

3. Over 600 people attended last year's Symposium. Plans are underway to continue this annual event with a third Symposium in the Spring of 1985 at a place and date to be determined.

STAT

[redacted]

Philip K. Eckman

B239

Cys to DCI
DDCI

DCI EXEC REG

ARTIFICIAL INTELLIGENCE SYMPOSIUM

"INTELLIGENCE APPLICATIONS OF AI"

December 6, 7, & 8, 1983

A REPORT

Symposium Co-ordinator: SMART SYSTEMS TECHNOLOGY, Inc.
Specializing in Artificial Intelligence Education and System Implementation

- CONTENTS -

1. EXECUTIVE SUMMARY

1.1 SYMPOSIUM OBJECTIVES

There is a strong current trend in both DoD and the Intelligence Community, toward exploratory development of computer-based intelligent systems for processing data. This trend is a deliberate response to an urgent requirement: senior defense and intelligence managers have long foreseen that the continually increasing rates of acquisition of all-source data can not be matched by any comparable increase in the number of intelligence analysts available to process and interpret the data. **A central objective of Artificial Intelligence is to enable computers to emulate the intellectual functions that humans employ in analyzing and interpreting data. For these reasons the Intelligence Community is interested in assessing the potential of Artificial Intelligence as a technological key to computer-based intelligent systems for intelligence applications.**

In December 1983 the Artificial Intelligence Steering Group (AISG) of the Intelligence Research and Development Council sponsored the second Annual Symposium on Intelligence Applications of AI. The Symposium had several objectives:

- To provide a forum for the exchange of ideas and experiences concerning the technology of artificial intelligence (AI) and how it might be applied to information collection, processing, and analysis within the Intelligence Community.
- To provide the audience of intelligence professionals with a better understanding of the current state of the art in AI, a more focused appreciation for where and how AI can be applied to the intelligence business, and an increased awareness of who the relevant AI players are -- both inside and outside the Government.

The Symposia also provide a series of snapshots in time against which the community can measure the pace of progress in AI research and technology, and the rate of AI technology transfer into intelligence applications.

The presentations in these Symposia were selected to provide a balanced view of industry trends, academic research, and intelligence applications. In order to allow in-depth treatment of a few topics, the first two Symposia have concentrated primarily on Computer-Assisted Image, Signal and Speech Interpretation, Expert Systems, Intelligent Data Bases and Overviews of Intelligence Requirements. Deliberate omissions to date include robotics and AI applications to resource allocation and planning.

Approved For Release 2008/09/04 : CIA-RDP86M00886R000500040010-5

As intelligence applications of such technologies mature they will be covered in future Symposia.

The presentations are organized in terms of six major themes. These themes are concerned with the Intelligence Community's requirements for computer-based intelligent systems, the significance of AI research and technology for realizing these systems, and the R&D investment strategies and technology transfer tactics required to make it all happen. **These themes will significantly affect the Community's future courses of action concerning AI research and technology.** In this Executive Summary we will present the highlights of the speakers' remarks as they relate to these major themes, indicating where the speakers provided novel points-of-view or interesting technical detail. More detailed summaries of the individual presentations are presented in Sections 2.1 through 2.6 of this report.

1.2 MAJOR THEMES

Theme #1: Requirements Analysis: Some points-of-view concerning the Intelligence Community's requirements for computer-based intelligent systems.

The speakers agreed that computer assistance is critical in fusing, analyzing and interpreting vast quantities of all-source data. The use of computer-based job performance aids to increase analyst productivity was emphasized by several speakers. **Interesting detail** was presented by Mr. David McManis, NIO/W, concerning the urgency of the requirements for natural language processing, for analyst aids in accessing data bases, for better warning indicator methodologies, and for unified concepts of information handling. **Emphatic expressions** of the urgency of improved information handling systems and analyst aids were made by high-level managers in Defense and Intelligence.

Theme #2: AI Research and Applications: The history and scope of AI, and summaries of some current research trends.

AI research and development spans a very wide spectrum of activities with multiple objectives; primarily, the practical objective of making machines smart in order to make them more useful, and the scientific objective of understanding the computational nature of human intelligence.

Today's Expert Systems are "idiot savants", demonstrably useful in strictly circumscribed applications; DEC's XCON is a landmark. Some Natural Language Understanding systems are demonstrably useful; however, existing systems are far from being able to recognize and respond to the user's goals as distinguished from responding literally to queries. AIC's INTELLECT is the most successful commercial Natural Language product to date. **Interesting detail** provided by Dr. Patrick Winston on MIT vision research from which new principles of signal-to-symbol transformation are emerging. These new "scale-space transform" techniques may eventually facilitate building expert systems for signal understanding. **Interesting detail** by Dr. Winston on Dr. Michalski's research on "symbolic clustering" and on Dr. Winston's own research on learning and reasoning by analogy. The long-range payoff of such research may be realized in expert systems which learn from experience, and which recognize analogies between current and past situations and respond accordingly. Additional **interesting detail** was provided by Dr. Rodney Brooks concerning Stanford's approach to model-based automated image interpretation using ACRONYM. Partially successful initial results were reported; the technology is still experimental.

Promising developments in the area of computer hardware and programming environments for AI. Several Lisp machine vendors: LMI, Symbolics, and Xerox, demonstrated very high-quality AI workstations and Lisp programming environments. DEC announced that Common Lisp and OPS5 will soon be available as supported products for the VAX. Apollo and IBM personal workstations also were shown as interesting low-cost alternatives for limited AI applications. The cost of AI workstations is still high reliability is improving, and performance is generally quite good.

Theme #3: Intelligence Applications of AI: Some examples of current AI R&D projects in Defense and Intelligence.

Seven AI research and development projects with clearly defined defense or intelligence objectives were presented. Three were directed towards automated signal processing and interpretation. Each of these systems seeks to represent human experts' knowledge about the engineering and operational backgrounds of signals and to reason with this knowledge in combination with traditional signal processing algorithms.

Three of the projects were directed towards automated image interpretation. All of these systems seek to represent collateral information known to experts about objects and events in the scene and to reason with this information in combination with traditional image processing techniques.

The seventh system seeks to reason with military experts' knowledge about targets and battlefield operations in combination with mathematical optimization algorithms to achieve an objective of optimal resource allocation.

Only one of the seven systems described (TRW's signal sorter tuner) is claimed to be functioning productively in an operational environment. The projects are fairly recent starts, and none as yet is close to the level of accomplishment claimed for DEC's XCON. For the image and signal-related systems the intrinsic difficulty of signal-to-symbol transformation and the very broad background knowledge required for these domains, are obstacles to operational success.

Theme #4: Investment Strategies by Industry: Some examples illustrating how U.S. and Japanese industry are structuring their investments in AI.

Speakers from seven major U.S. corporations agreed that industry should and will invest in AI because it is probably important, useful, and profitable. Industry is making a major thrust in expert systems for planning, diagnostics and data interpretation. **Interesting detail** was provided on DEC's corporate thrust towards building expert systems modules for in-house applications in sales, manufacturing, diagnostics, and system configuration. These modules will eventually be integrated into a corporate-wide composite system. **Interesting detail** on Westinghouse's corporate strategy of "learning AI technology by doing it", training its own AI engineers in-house, facilitating technology transfer by establishing close ties with universities (CMU) and by establishing in-house Centers-of-Excellence. **Interesting detail** on treating situations where complex combinatorics are required and where "humans find infinitely many ways to screw things up." **Emphatic** testimonials by DEC on expert systems' proven value to the corporation, "comparable to the effects of the assembly line on Ford Motor Company."

As to the Japanese ICOT Project: **Key observation** made by Dr. John Alan Robinson that the complex and detailed project plan published by ICOT obscures the simple central themes of the project. The themes are: a commitment to make Logic Programming the central language for symbolic computing and a commitment to make the world's fastest (10^9 logical inferences per second) logic computer by the 1990's. In Dr. Robinson's view, this is rather like a dramatically enhanced version of the U.S. Lisp Machine projects which also seek to provide very high performance, low-cost symbolic processing tools.

Significant prediction: the Japanese may achieve their engineering goals with the super-Logic Programming Machine, but will find the AI sub-goals (Vision, Speech, Expert Systems) just as difficult to accomplish as we have.

Theme #5: Investment Strategies by Defense: A summary of DARPA's Strategic Computing Program.

DARPA disclaimed that the Strategic Computing program is a U.S. response to the Japanese ICOT program. **Interesting detail** was provided by Dr. Robert Kahn, Director of DARPA's Information Processing Techniques Office, on the program's general goals. Like ICOT, these goals include symbolic processing capabilities of 10^8 logical inferences per second with generic AI software for vision, speech, and data base applications. Unlike ICOT, there was no expressed commitment to Logic Programming as a central theme.

Also, unlike ICOT, there is a specific focus on three major military application areas: Navy battle management, autonomous battlefield vehicles, and automated aircraft cockpits. **Contrast** with ICOT: while the ICOT project is dominated by a central theme, and implemented by a principal team of forty AI researchers, the DARPA program is largely directed toward facilitating the efforts of unspecified and independent research groups from all over the U.S.. These independent groups will have access to computing hardware and facilities for designing and manufacturing innovative microprocessors and multi-processors.

Theme #6: The Tactics of AI Technology Transfer: The pragmatics of starting and maintaining AI R&D facilities.

The observations of six experienced R&D managers concerning in-house AI Centers as a mechanism for facilitating AI technology transfer may be summarized as follows:

- In-house AI facilities are a very good mechanism for developing AI capabilities within an Agency,
- The facilities' research programs should be closely coupled to the Agency's mission,
- The scarcity of qualified AI personnel is a critical problem . . . you have to train some of your own,
- Long-term investment commitment by high-level management is essential for stability and success of the AI facility,
- Computing equipment and software should be selected to facilitate sharing software with universities and other research groups.

1.3 CONCLUSIONS

The national technology base for building intelligent systems is substantial and rapidly growing. Practically every major high technology company in the U.S. is engaged in exploratory development projects related to computer-based intelligent systems, both for internal use and as products. Growth in the national AI technology base is spurred by continuing price/performance improvements in computer hardware, by DoD spending on intelligent systems R&D, by consumer expectations for "smart" products, by several AI technology successes in knowledge-based systems and database interfaces, and by the perceived need to compete with Japanese technology thrusts. It is virtually certain that AI technology growth will continue and will accelerate. Generally, private industry should not expect immediate major payoffs from its investments in AI.

The Intelligence Community has urgent and specific requirements for computer-based intelligent systems. The massive data-acquisition capabilities of our collection systems exceed, by at least a factor of ten, the intelligence processing capabilities of the severely limited number of experienced intelligence analysts available. Intelligence Community working groups have identified specific problem areas where computer-based job performance aids for intelligence analysts would have high payoff. These include: data base access and automated data-base construction; foreign language translation; image and signal analysis; intelligence resource allocation and planning; capturing unique knowledge and skills in knowledge-based systems; and aids to data fusion and interpretation. AI R&D is addressing each of these areas, but no generic solution has been found for any of these applications. The most significant progress to date is in data base access and in knowledge-based systems, where several commercially viable systems have been developed by industrial firms focusing on very specific problems.

The Intelligence Community's current program in AI technology applications has not reached critical mass. Several leading corporations including Westinghouse and DEC are exerting, on their own behalf, AI efforts substantially larger than the Intelligence Community's. Factors that senior corporate managers judge to be essential to successful major AI programs include: firm, long-term management commitment and project goals that are both technically achievable and significant to the corporation's business.

Approved For Release 2008/09/04 : CIA-RDP86M00886R000500040010-5

Some corporations address the current scarcity of experienced AI technologists through staff training programs, in-house centers of AI expertise, and links with universities. By contrast, the Intelligence Community is currently pursuing perhaps a dozen relevant but relatively small exploratory development efforts in AI systems. These are primarily conducted by outside contractors. An AI orientation program for Intelligence Community managers has recently been initiated. While these steps are in the right direction, the overall effort still lags far behind critical mass when measured against the magnitude and urgency of the Community's requirement to reduce its information-processing overload. **We should be preparing system specifications, initiating substantial development programs, and training technical personnel now if we expect to see computer-based intelligence in our information processing systems by the mid-to-late 1990s.**

* *

These remarks conclude our Executive Summary of the December 1983 Symposium on Intelligence Applications of AI. In Sections 2.1 through 2.6 following, more detailed summaries of the major presentations are provided.

STAT

For access to the original data the reader is referred to the Symposium Proceedings and to video and audio tapes made of the Symposium, available through the Symposium Chairman [redacted] Office of Research and Development, Central Intelligence Agency, Washington, D.C. 20505 [redacted]

STAT

2. DETAILED SUMMARIES OF SYMPOSIUM PRESENTATIONS

Artificial Intelligence is the most rapidly expanding research and development activity in applied computer science today. Specialized AI computers and AI software systems have migrated from university research laboratories into exploratory development projects in several government laboratories and in practically every major high-technology industry in the U.S. The presentations in this Symposium were selected to present a balanced overview of current industry trends, academic research, and government programs related to intelligence applications. In the following section of the report, the presentations are organized in terms of six general themes. In the spirit of presenting objective data for Intelligence Community managers and decision makers, we have attempted to make this report a reasonably neutral account of what the speakers actually said.

2.1 REQUIREMENTS ANALYSIS: Some points-of-view concerning the Intelligence Community's requirements for computer-based intelligent systems.

2.1.1 Speaker: Dr. Philip K. Eckman
Chairman, Intelligence Community AI Steering Group

Topic: Symposium Objectives

Summary: Dr. Eckman foresees that in the coming decade, the amount of all-source information will increase tenfold and the number of analysts only by 20%. To paraphrase Von Neumann: it's an insult to a man to have him do a job that a machine can do as well. The issue, in Dr. Eckman's view, is how to make people more productive and reserve for them tasks that are uniquely human.

2.1.2 Speaker: Dr. Richard D. DeLauer,
Under Secretary of Defense for Research and Engineering

Topic: AI and Defense

Summary: Dr. DeLauer perceives AI and supercomputers as vital to the defense and intelligence communities. Supercomputers and high speed processing can help DoD and the Intelligence Community overcome our adversaries' numerical advantage by advancing our national capabilities in the areas of training, information handling, heads-up displays for cockpits, fusion of sensor data, and battle management. In Dr. DeLauer's view, we need to rapidly fuse, analyze and distribute large amounts of intelligence data; AI will be important to us by helping make this less labor intensive. As a direct result of this Symposium, Dr. DeLauer and his staff intend to assess and collate ideas received from the Symposium audience and will use these ideas as a source of guidance on possible new ways of using AI and directing further research in AI.

2.1.3 Speaker: David Y. McManis
National Intelligence Officer for Warning

Subject: Intelligence Requirements

Summary: What is missing to date, in Mr. McManis' view, is a unified concept of information handling, and that's really what AI should be about.

The NIO/W is concerned mainly with warning, in the extended sense . . . strategic and tactical warning, warning of "suprise attack" of the Pearl Harbor variety, and with other forms of threat, such as technology breakthroughs or political and economic instabilities. Mr. McManis observed that we do a generally good job in current and long-term intelligence; the missing area is the six-month time frame for warning and forecasting. Several areas that need to be addressed include:

- Automated data bases, such as bibliographic, biographic, order-of-battle, and geographic data bases.

  There is a need to get the analyst out of the loop, and to automate the process of building data bases. ELINT data is well suited to this approach. With respect to open source text, there is an urgent need for natural language processing, even if only 50% of the job gets automated, as with machine translation.

- Analytic access: in Mr. McManis' view, it is incredible to watch senior analysts having to weave their way through the labyrinths of our data base systems. Better terminals and better tools for manipulating data should be developed which remove the barriers between the analyst and the data.

- Indicator methodologies for recognizing swing events that flag harmful situations are urgently required.

  Mr. McManis suggested that ORD might provide some assistance in applying expert systems against the Indicator Methodology problem. Some "warning models" are being built to help the NIO/W understand what the key indicators are, and how to collect against them. NIO/W is interested in imagery targets with low activity levels but high indicator value; comparing yesterday's pictures with today's may prove vital.

- Presenting information to decision makers so as to answer their questions quickly: Here a tremendous amount of help is needed in providing multiple alternative views of data in response to different types of questions and different levels of security classification.

2.1.4 Speaker: Mr. John N. McMahon
Deputy Director of Central Intelligence

Topic: AI and Intelligence

Summary: Mr. McMahon recalled that twenty-five years ago the Intelligence Community talked about the information explosion and automated information processing, but history has shown the Community is smarter at building big collection systems than at processing information. AI can make significant contributions in machine translation and photo interpretation. The Community's experiences with AI to date have been somewhat "spotty"; the results have been marginal. Nevertheless, in Mr. McMahon's view the capability is out there, but we haven't been smart enough to harness it.

Threats to U.S. national security include such diverse phenomena as Soviet technological advances in lasers and particle beams; Soviet military buildups; the cascading effect of the Third World's $625 billion dollar debt; the growth of economic competition from Europe and Japan. Intelligence is required to keep track of all this, and in Mr. McMahon's view, the AI Community can contribute to U.S. National Security by helping the Intelligence Community handle such problems. Mr. McMahon observed that: "the Intelligence Community can't use poverty as an excuse, since we've had a 75% budget increase in the last three years. Let's do something with it!"

2.2 OVERVIEW OF AI RESEARCH AND APPLICATIONS: History and scope of AI and summaries of some current research trends.

2.2.1 Speaker: Dr. Patrick H. Winston
Director, Artificial Intelligence
Laboratories, MIT

Topic: Overview of Artificial Intelligence
Research and Applications

Summary: In Dr. Winston's view, AI is defined by its objectives, rather than by its techniques. The objectives are: making machines smart in order to make them more useful, and understanding the nature of intelligence . . . which Dr. Winston referred to ironically as the "Nobel Laureate motive".

When we understand something thoroughly it no longer seems mysterious; when we understand how AI programs work, they may cease to appear "intelligent". Dr. Winston related an anecdote concerning Dr. Slagle's first symbolic integration program, a precursor to the MACSYMA system developed at MIT in the 1960s. Dr. Winston taught the program in his AI course at MIT; after the class one student exclaimed: " . . that program isn't really intelligent . . . it integrates the same way I do."

Dr. Winston summarized current research and applications in five AI subdomains: Expert Systems, Natural Language Processing, Vision, Signal Understanding, and Learning.

Expert Systems are "applied logic". Three major categories include:

- Expert systems for identifying things; the MYCIN diagnostic expert system is an example,
- Expert systems for configuring things; DEC's XCON system is an example,
- Expert systems for interpreting signals; the Schlumberger dipmeter advisor is an example.

Dr. Winston expressed suprise that so much has been accomplished with such simple tools as rule-based systems. Today's expert systems are "Idiot Savants", brilliant in narrow domains, otherwise helpless. The next generation of systems must be able to shift among alternative points-of-view; must degrade gracefully as the knowledge base is reduced; must be able to break their own rules, build models, and most significantly learn from experience.

A good domain for Expert Systems is one where: specialized knowledge is required as distinct from generalized "common sense"; where humans require an hour or so to solve the problems; where an expert is committed to the project; and where the subject matter is systematized to the extent that there are books or manuals on it.

Natural Language Understanding requires a broad spectrum of varieties of intelligence and knowledge. We now have operational systems that do useful things, providing access to data bases in an easy, natural way. We are a long way from having systems that understand the speaker's intent and frame-of-reference sufficiently to respond to queries in a helpful rather than a literal fashion. As with Expert Systems, Dr. Winston expressed surprise that so much has been accomplished with such simple tools as semantic grammars and frames. INTELLECT was singled out as the most successful commercial natural language product. Frame-based systems have been applied to such tasks as skimming newspapers . . . an application that might be of great value to the Intelligence Community. Because the narrative may misfit an inappropriate frame, frame-based systems can make hilarious mistakes and must be used with great care.

Vision and Signal Understanding: Dr. Winston observed that work in vision has led to a principle for signal-to-symbol transformation that may enhance expert systems for signal understanding. The spatial second dervatives of deliberately blurred images cross zero at points in the image where humans tend to perceive edges. By parameterizing the degree of blur and plotting the loci of such zero crossings the image is transformed into "blur versus zero crossing" space or scale space in Dr. Winston's terminology. This leads to a natural segmentation of images and signals. The transformation is invertible, and with this transformation, signals can be divided into pieces, much as the human eye would divide them up.

Learning: Dr. Winston illustrated Dr. Michalski's program for separating diseased from non-diseased soy-bean plants, "discovering" the combination of properties that discriminates among classes. The program searches a vast tree of possible combinations of properties, using induction heuristics . . . a sort of "symbolic cluster-analysis algorithm", and in this sense it accomplishes a symbolic learning task. Dr. Winston proceeded to describe his research in the area of "learning by analogy". One of his programs takes as input very simple stories . . . a precis of Macbeth was used as an example . . . and outputs a semantic network relating actions by the play's principal characters. The arcs in this net represent causal relations. The nodes can then be generalized, and the relationships treated as general principles.

When other new stories are mapped onto the resulting structure the principles may become hypotheses about the new situations, arrived at "by analogy" with the initial structure. For example, similar motives might be imputed to protagonists in sufficiently similar narratives.

Dr. Winston views this research in analogical reasoning as a very tentative but promising first step into a new area . . "like being at Kitty Hawk in 1903, planning the program that will lead to the 747". He has extended the approach from stories to "object form and function," and is writing programs that may enable robots to infer useful properties of an object . . . such as "being liftable" . . . from analogies and formal resemblances to other known objects. Computation is a major impediment to extending such toy learning systems to real applications. The sorts of five-order-of-magnitude advances which are the goals of the DARPA Supercomputing program are required here.

In summary, Dr. Winston observed, progress in AI is a locus of points in "Applications versus Knowledge space," as shown below.

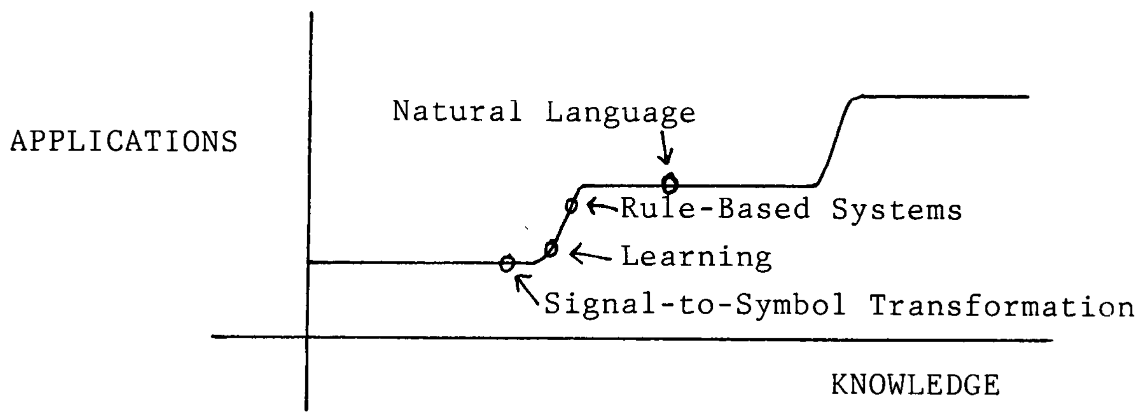

Most of AI is still at or near the first riser of this step-like curve. Our attitude towards AI, in Dr. Winston's view, should be one of "restrained exuberance".

2.2.2 Speaker: Dr. Rodney Brooks
Assistant Professor
Computer Science Department
Stanford University

Topic: Overview of AI Applications in Computer Vision and Image Understanding

Summary: Dr. Brooks addressed the problem of automated identification of objects in an image.

There are two traditional approaches: image based, i.e. processing and segmenting the image, working in the direction of inferring the objects that might have given rise to the observed image data; or model based, inferring from an a-priori object model the image that it might have produced, and comparing the inferred image with the observed image data.

The ACRONYM system developed by Dr. Brooks and Dr. Binford at Stanford University, represents a model-based approach to image understanding. The primitive components from which complex object models are built are generalized cylinders. These parametrically defined tubular structures are combined to form more complex objects . . . aircraft models, in the example under discussion.

The parameters of a camera model, together with the parameters defining the length, cross-section and curvature of the generalized cylinders, suffice to compute classes of predicted images.

Dr. Brooks experimented with digitized aerial photographs of aircraft in commercial airports. Generalized cylinders were fit against the edge-enhanced imagery, and the parameters of the cylinders were checked for consistency against the camera model and against object models for known aircraft types. The experiment met with limited success. Failures to recognize some objects resulted in part from failures in edge detection and in part from the computational impracticability of applying rule-based consistency checks against the entire image, rather than against pieces of it. In Dr. Brooks' very forthright summary he asserted that this is still an experimental technology; it's hard to make the computations parallel, and the performance currently is poor.

2.2.3 Speaker: Dr. Raj Reddy
Director, The Robotics Institute
Carnegie-Mellon University

Topic: Overview of AI Applications in Robotics and Speech Understanding

Summary: No matter what area of AI we're working in, Dr. Reddy observed, knowledge representation is the key ingredient. For example, Herb Simon once undertook as an experiment the task of encoding the knowledge in one chapter of a physics text (on statics). After six months he had developed twelve condition-action rules which sufficed to solve nineteen out of twenty-two of the problems at the end of the chapter. What we still don't know how to do in AI is to "give the computer this textbook and tell the machine to distill the knowledge and solve the problems."

Consider the variety of types of human knowledge. A small fraction is algorithmic. Perhaps ninety-five percent is informal, derived from examples. In robotics applications the extreme difficulty we find in constructing a machine like the hexapod walker, for example, underscores the need for "understanding" how we walk, how we see, or how we process speech.

In general, Dr. Reddy observed, we still don't fully understand what knowledge is involved in these actions, how it is represented, or how the knowledge is used by the organism to solve problems.

Dr. Reddy proceeded to describe the CMU Robotics Institute. The staff numbers approximately one hundred and fifty, with approximately fifty PhDs and sixty university students. Westinghouse is a major industrial contributor to the Institute. The major research themes are:

- Manufacturing Facilities of the Future,
- Operations in Hazardous Environments.

In the first area, knowledge-based simulation looks promising. AI-based simulators for training Westinghouse production-plant operators cost approximately one million dollars to build and are saving the company approximately five million dollars per year. Other general problem areas include: how to build flexibility into a production facility; resource allocation . . . what resources are required by a facility for a product; allocation of control . . . what operations should be autonomous, and where should human supervision be employed. In general, Professor John McDermott of CMU observes that "white-collar robotics" is easier than "blue-collar robotics".

In the Hazardous Operations area, Sutherland and Moravec are designing a six-legged walker and a three-wheel autonomous rover. The control problem for the six-legged walker is still not completely solved. Moravec's stereo vision comparator for robotic vehicles highlights requirements for faster computation: as Dr. Reddy remarks, it "takes one step, then thirty minutes to think, takes one step, then thirty minutes to think . . . ."

2.2.4 Speaker: Dr. Gian-Carlo Rota
Professor of Applied Mathematics
and Philosophy, MIT

Topic: Observations and Speculations on
the Effect of the Computer on Scientists

Summary: Dr. Rota observed that in computer chess, the winning programs have proved to be those that employ brute force search rather than relying on ingenious heuristics. "It's a sad commentary on the human condition that reliability wins over genius!" Dr. Rota proceeded to make three principal observations regarding AI:

- The simple-minded machine is most likely to succeed, e.g. computer chess,
- There is no reliable way of telling beforehand whether an AI program is easy or hard . . . there's no better test than trying to build things; that's when we find out whether our "common-sense" descriptions match reality,
- Advances in AI have come from the "hard" rather than the soft sciences.

As to international competition, Dr. Rota observed that in the U.S., the free exchange of information among scientists, and our superior university system, give our research establishment an advantage over the Russians and the Japanese. A serious national problem, however, is that we allow our best technology to "drain" overseas. Dr. Rota expects Hitachi to announce early in 1984 a machine with five times the capability of a CRAY1. Fujitsu will announce a machine with ten times CRAY capability. These rapid advances by the Japanese computer industry are due to U.S. failure to protect this technology.

In Dr. Rota's view, two new frontiers for AI technology are:

- transition from digital to holographic computation,
- the unexpected relevance of the theory of random nets and non-deterministic programming to AI computation.

2.2.5 Speakers: Dr. John Vittal - Xerox
Mr. William Kaiser - Apollo
Mr. Steve Lazerowich - Symbolics
Mr. Robert Abramson - DEC
Dr. David Yun - E-Systems
Dr. Frank Spitznogle - LMI

Topic: Industry Panel: Description and Schedule of AI Demonstrations

Several computer manufacturers made AI processors available during the three days of the Symposium for demonstrations of AI tutorial programs and working AI systems. These were available during the mid-day breaks and following the afternoon sessions in the Tunnel adjacent to the main Auditorium at CIA Headquarters.

2.2.5.1 Xerox Summary: Dr. John Vittal described Xerox AI program development tools . . . primarily Interlisp-D and Smalltalk, and the Xerox personal workstations with high-resolution graphics. The emphasis is on programmer productivity enhancement through interactive programming. The Xerox demonstrations include:

- Interlisp D, the Lisp programming environment,
- LOOPS (object-oriented programming system with constraints facilities),
- RABBIT (a mechanism for accessing databases),
- TRILIUM (a knowledge-based system for designing copier interfaces),
- FORMS (a system to facilitate designing forms),
- PAPERWORKS (a system to facilitate annotating reports),
- MAPS (a system for annotating maps),
- SMALLTALK (an object-oriented program development environment),
- ANALYST (an analyst aid, facilitating interaction with maps, text, and mail)

2.2.5.2 Apollo Summary: Mr. William Kaiser described the workstation architecture: 32 bit VLSI processor and bit-map displays, with individual workstations interfaced through a local area network. Their major market is computer aided design; the Apollo supports FORTRAN, PASCAL, and C, color graphics, and Bell and Berkeley Unix.

AI-related demonstrations include:

- T, (a dialect of Lisp developed at Yale University)
- Portable Standard Lisp
- DUCK (the Smart Systems Technology logic programming system for developing expert systems)
- KES (the Software A&E expert system application generator)
- SIL (the SILMA programming language, for high-speed graphics applications)

2.2.5.3 Symbolics Summary: Mr. Steven Lazerowich described the Symbolics demonstration programs:

- Thoughtsticker (the Pangaro, Inc. tool for representing alternative views of concepts)
- SARSYS (A TASC, Inc. system for radar image interpretation)
- KNOBS (the MITRE experimental system for air mission planning)
- INCA (VERAC's DARPA-sponsored data fusion project)
- SAGE (Symbolics' document support system)

2.2.5.4 DEC Summary: Mr. Robert Abramson described the DEC AI Technology Center in Hudson, Mass. DEC is a user and developer of AI systems, such as XCON, and intends to be a vendor of systems such as COMMON LISP and OPS under VMS. DEC demonstrations include:

- o XSEL (an expert system for guiding DEC sales representatives in selecting components; XSEL calls XCON for configuration assistance)
- o OPS (Production System)
- o COMMON LISP (DEC's implementation)

2.2.5.5 E-Systems Summary: Dr. David Yun of Southern Methodist University, consultant to E-Systems, described a low-cost, high-availability knowledge-base system development environment on the IBM personal computer. The hardware/software combination cost is approximately three thousand dollars.

2.2.5.6 LMI Summary: Dr. Frank Spitznogle described several features of the Lambda machine and the systems to be demonstrated. The systems include:

- o Lisp Machine programming environment on LMI's Lambda machine,
- o PROLOG (compiled into Lisp Machine lisp),
- o TI Natural Language System,

The system architecture is based on a high-speed bus (the NU-bus). The Lisp processor is interfaced to this bus; a 68010-based UNIX processor is also interfaced to the bus, allowing either UNIX or Lisp programming.

The LMI implementation of PROLOG is expected to run extremely fast and will be integrated into the Lisp programming environment. A natural language parser and a "smart arithmetic" demonstration were used to present PROLOG concepts.

2.3 INTELLIGENCE APPLICATIONS OF AI: Some examples of current AI R&D projects in Defense and Intelligence.

These presentations were made during the second day of the Symposium, in sessions classified at the Secret level. In order to facilitate distribution, this report is to be kept unclassified. This section summarizes the unclassified aspects of the presentations.

2.3.1 Speaker: Dr. James Slagle
Senior Scientist, Navy Center for Applied Research in Artificial Intelligence

Topic: BATTLE, an Expert Advisor for Weapons Allocation

Summary: At the Navy Center for Applied Research in Artificial Intelligence (NCARAI), several research and exploratory development AI projects are underway:

- o Expert Systems in Combat Management,
- o Expert Systems for equipment troubleshooting,
- o Target Classification, using ISAR data on vessels,
- o Natural Language Processing, applied to Navy message automation,
- o Multi-sensor fusion,
- o Adaptive control.

Dr. Slagle described the BATTLE system in the context of combat management. BATTLE has been implemented at NRL by Jim Slagle and several colleagues at NCARAI. The objective of the system is to provide improved weapons allocation plans for Marine Corps Artillery and Air Support. Data supplied by forward observers allows real-time updating of the plans in response to target damage reports. The system performs two sorts of computation: analysis of the effectiveness of weapon-target combination and complete multi-weapon, multi-target allocation plans.

The weapons effectiveness calculation is implemented by a computation network, which is a generalization of PROSPECTOR. Rules are specified by military experts and stored in a data base. The allocation plan is computed by an algorithm which searches efficiently through the space of possible allocation plans for optimal, or near-optimal, allocations which maximize damage to targets. BATTLE has been tested by a Marine Corps artillery expert and was judged to produce valid allocation plans for real problems in reasonable time. Other problems, such as assigning individuals to jobs in an organization or parceling out computational tasks in a multi-processing environment, may be amenable to similar approaches.

2.3.2 Speaker: Dr. Gerald A. Wilson and
Dr. Robert Drazovich
Advanced Information and Decision
Systems, Senior Computer Scientist

Topic: Computer-Based Assistants for Science
and Technology Analysis

Summary: Dr. Wilson described the prototype Expert Assistant for Science and Technology Analysis (ASTA). The objective of the ASTA project is to develop an interactive computer-based "intelligent assistant" to facilitate analysis of shipboard radars. The ASTA computer terminal presents the user with a sequence of frames requesting parameters on the radar under analysis. The user extracts these parameters from photographic, radar intercept or other available data, and enters the required information into the frame. ASTA proceeds to "reason" with this data, using facts and hypotheses about radars stored in its knowledge base. The system may prompt the user for additional data as required for its analysis.

ASTA has four major components: the dialog manager which controls the user/system interface; the activity support manager, which maintains dynamic data bases of user hypotheses and general radar knowledge; the information manager which controls access to local and external static databases; and the support tool manager which accesses radar analysis algorithms. The rule-based sytem architecture underlying ASTA is MRS, developed by Dr. Genesereth at Stanford University.

ASTA is currently under development. The system handles only a fraction of the information which could be provided and only "knows about" a small number of radars. However, Dr. Wilson observed that the results are quite promising and that experienced radar system analysts are pleased with the way in which ASTA helps them to record information in a structured way.

Dr. Drazovich described the Advanced Digital Radar Image Exploitation System (ADRIES). The goal of the ADRIES project is to develop an interactive computer-based image exploitation workstation to facilitate analysis of synthetic aperture radar (SAR) imagery. This development is in the context of a DARPA program involving eight companies. A final demonstration is scheduled for the Summer of 1986. The workstation is intended to assist the user through:

- o "smart" techniques for automating some SAR image analysis tasks,
- o detecting and classifying some tactical targets (missile sites, regiments, etc.),
- o supporting the use of collateral information on the SAR image to narrow the area search.

The ADRIES workstation will enable the analyst to display SAR imagery at several resolutions. Collateral information will be used by the system to narrow the area search for specific targets to "most likely" local regions which will then be magnified and enhanced. The collateral information includes current battlefield intelligence, enemy order-of-battle, and operating characteristics of the targets. The development system architecture consists of a VAX 750 mainframe, VICOM image processor, and a Symbolics 3600 Lisp Machine, all interfaced through Ethernet.

2.3.3 Speaker: Mr. Mark Williams
Senior Staff Engineer,
ESL, Inc.

Topic: AI in Signal Processing

Summary: Mr. Williams described a research project underway at ESL, exploring the application of expert systems concepts and techniques to signal collection and analysis. In traditional signal processing, Mr. Williams observed, the "front-end processor" extracts signal features and makes statistically based decisions using pre-determined algorithms. The objective of the expert systems approach is to increase the flexibility of the signal processor in dealing with noisy or non-standard situations by enabling symbolic reasoning about the signal, the environment, the available collection of alternative processing techniques . . . much as a human expert might. Mr. Williams represented the idealized "intelligent signal processing paradigm" in the schematic shown below.

INTELLIGENT SIGNAL PROCESSING PARADIGM

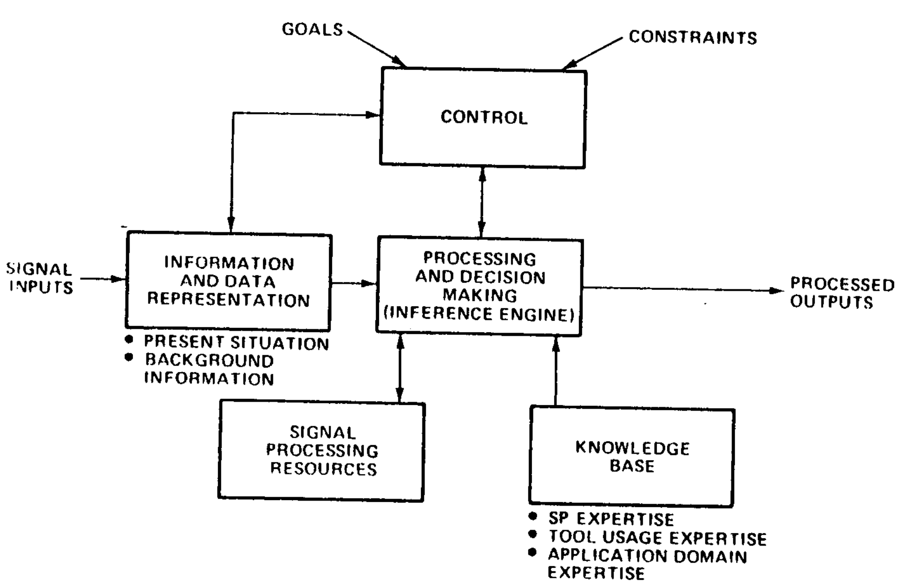

ESL is currently experimenting with these AI approaches to signal classification. The development system architecture consists of a signal digitizer, VAX mainframe, and Xerox 1100 (Dolphin) Lisp Machine connected via Ethernet.

2.3.4 Speaker: Dr. Donald Close
Program Manager
Hughes Aircraft Company

Topic: Image Understanding

Summary: Dr. Close described the overall objective of the Hughes program as a demonstration of the feasibility of automated port monitoring using operational imagery. Eighteen months into the project, a Phase I demonstration was conducted on a general purpose computer, using the best available software developed up to that time by the image understanding research community.

The ACRONYM system was selected for the project. ACRONYM is described in greater detail in Dr. Rodney Brooks' presentation on model-based image interpretation (Section 2.2 of this report). The initial experiments were valuable and instructive. On the positive side: the prototype systems successfully aligned different images of the same scene, successfully detected objects of probable interest, and correctly identified the object classes, thereby validating the model-based approach to image interpretation. On the negative side: the ACRONYM system was designed as an image-understanding research tool and was found to have a number of specific deficiencies for this particular application; the available object models were too limited, and the system's handling of rules and data structures was less flexible than required.

For the Phase 2 technology development follow-on task (an on-going two-year effort), Dr. Close and his colleagues intend to incorporate a detailed camera model into the system; to add knowledge-based planning and control; to incorporate three-dimensional computer graphics object modeling and image processing algorithms to facilitate the image intensity and texture processing judged to be necessary for object detection and classification in the port monitoring application.

2.3.5 Speaker: Dr. Dick Kruger
Senior Scientist
Science Applications, Inc.

Topic: AI Applications in Synthetic Aperture Radar Image Interpretation

Summary: Dr. Kruger presented an outline of SAI's recent research and development initiative in rule-based synthetic aperture radar (SAR) image exploitation. The objective of the effort is to use terrain knowledge and collateral information to narrow the area search for targets in SAR imagery. The underlying data base consists of DMA terrain maps of the Fulda Gap area, combined with SAR imagery of the same area taken during ongoing military exercises (REFORGER, 1981). Mobility data is derived from the terrain maps. Rules relating vehicle characteristics to terrain mobility are used to characterize certain geographic regions as "denied", thereby directing the search for targets to more likely areas. Initial results appeared promising and the rule set is currently being extended. The OPS-5 production system was used for the underlying rule-based-systems architecture.

2.3.6 Speaker: Mr. John C. Kelly
Senior Staff Engineer,
TRW Defense Systems Group

Topic: An Operational Artificial Field Engineer for Tuning a Signal Sorter

Summary: Mr. Kelly's organization collects and processes signals which are sorted into "bins" based on the features of their time-versus-frequency plots. In the traditional approach to signal sorting, very large numbers of signal feature extraction parameters are manually set on the signal sorter and subsequently fine-tuned for specific applications. This labor-intensive tuning process requires expert field engineers and is expensive.

TRW is developing an experimental rule-based system to facilitate the manual tuning process. Mr. Kelly observes that the rule-based tuner has been in operation at a field site since 1983; it is successful, and is estimated to have saved two man-years of engineer labor to date. Mr. Kelly described his attitude as . . . euphoric! The rule-based tuner continues to evolve; the rule set is growing from seventy-seven rules in 1983, to an expected two hundred rules by February, 1984. The system is implemented in PASCAL.

2.3.7 Speaker: [redacted] STAT
National Photographic Interpretation Ctr.

Topic: AI Applications in Image Analysis

Summary: [redacted] presented a preview of a report by the Exploitation Research and Development Committee (EXRAND) on AI applications to classified imagery. The report is scheduled for publication by the end of December, 1983. The objective of the report will be to provide: STAT

- o guidance for focused activities by the Intelligence Community,
- o information for senior managers,
- o system implications of bringing AI into image data systems.

The details of [redacted] presentation will not be included in the present unclassified report. STAT

Approved For Release 2008/09/04 : CIA-RDP86M00886R000500040010-5

2.4 INVESTMENT STRATEGIES BY INDUSTRY: Some examples illustrating how U.S. and Japanese industry are structuring their investments in AI.

2.4.1 U.S. Industry Panel:

Dr. Charles Herzfeld, Vice President and Director of Research and Technology, ITT Corporation
Dr. Edward C. Taylor, Director of Requirements Analysis, TRW Defense Systems Group
Ms. Norma Abel, Digital Equipment Corporation
Dr. Carl Smith, Staff Computer Scientist, Shell Development Company
Dr. Floyd Hollister, Senior Member, Technical Staff, Computer Science Laboratory, Central Research Laboratories, Texas Instruments, Inc.
Dr. Carl Love, Senior Consultant, Corporate Planning, Westinghouse Electric Corporation
Dr. Dan Schutzer, Vice President, Citibank

Topic: Why is Private Industry Investing in AI?

2.4.1.1 ITT Summary: Dr. Herzfeld commented that twenty years ago, he may have been the first person to mention the new science of AI in Congress . . . today, that child begins to walk. Industry invests in AI for three reasons: it is probably important, useful, and profitable. Expert systems have enormous potential to enable high quality output from less skilled individuals. ITT is developing rule-based systems to pinpoint flaws in the switching systems the company manufactures. Pattern processing is an important area that's not quite ready for development yet. AI will be enormously important to industry in system design, and in the programming environment. Operational planning in business . . . first simulating a corporation, then running it as simulated . . . may only be ten years away. And finally, Dr. Herzfeld observed, "indicators and warnings" tools should be translated from the national security application into industrial management tools for running big companies.

2.4.1.2 TRW Summary: In the view of Dr. Taylor, the most significant factors limiting the use of advanced weapon systems are the cognitive limitations of the human brain. "Augustine's Law", so called, describes the exponential distribution of skills and aptitudes. This law is illustrated in the following graph:

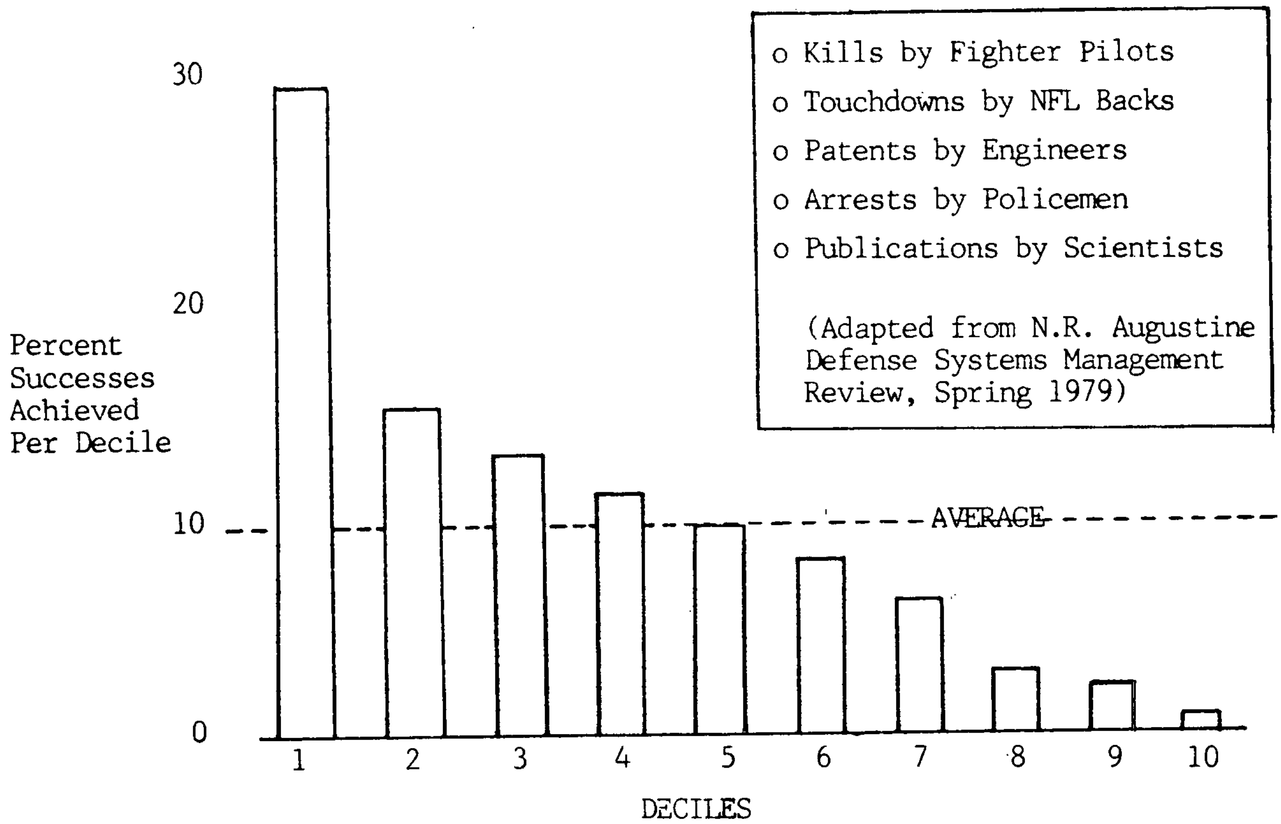

Statistics show that thirty percent of the total number of successes in many competitive events are scored by the top ten percent of competitors; the lowest three deciles together account for only ten percent of total successes. In combat, for example, this indicates that only the top ten percent of operators will use advanced information processing and display tools successfully.

It is therefore important to develop expert systems and computer-based intelligent advisors that will emulate those qualities that make the top decile of experts unique.

2.4.1.3 DEC Summary: In Ms. Abel's view, the positive effects of expert systems on DEC have been enormous . . . in her words, comparable to the effects of the assembly line on Ford Motor Company.

XCON, DEC's expert system for computer configuration, is only one member of a large family of expert systems currently under development at DEC . . . sales systems, manufacturing systems, service systems are also under development. In less than 5 years, DEC hopes to have these systems interconnected, as illustrated below.

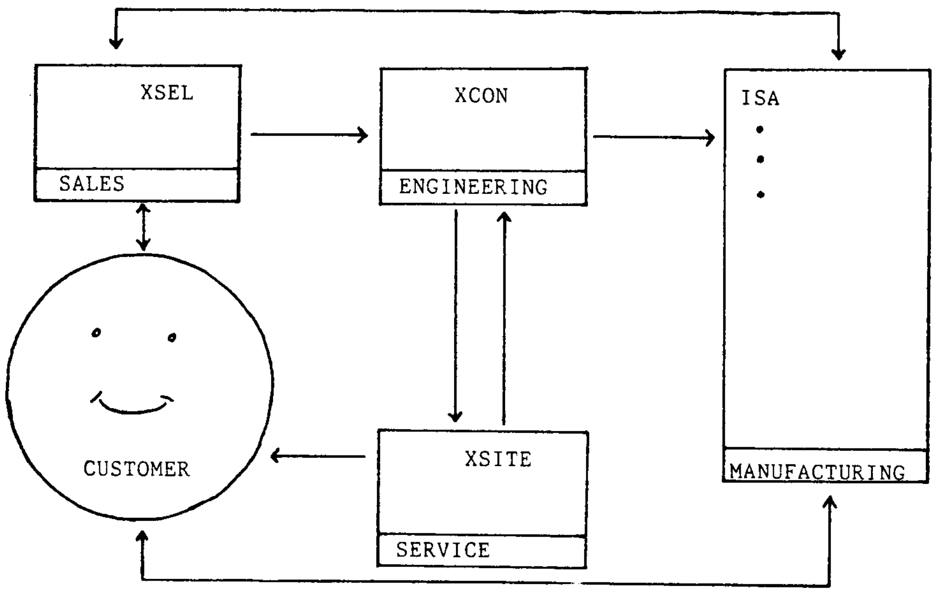

The Knowledge Network DEC Envisions

Before XCON, all DIGITAL's systems went through a process of final assembly and test during which trained technical editors configured VAX systems according to customer order forms, out of component parts. XCON has enabled the company to save money over the last two years by reducing the need for this expertise. This has changed the way the corporation does business, and DEC is introducing other rule-based expert systems and planning systems in manufacturing and sales. The rule-based systems are generally written in OPS-5; the planning systems, generally the farthest from completion, are developed in CMU's FRL, or frame representation language, used to simulate distribution networks.

What has DEC learned about industral applications of expert systems? Ms. Abel observed that:

- They're state-of-the-art, and therefore risky.
- They're very large software systems and need all the support that goes with maintaining gigantic systems.
- Technology transfer is a constant, continuing process.
- Human and motivational issues of getting systems accepted in the working environment are as difficult as the technical issues.
- It takes a lot of corporate investment because the systems change the way a corporation functions. With XCON, DEC was very lucky!
- The process needs a variety of skills . . . a few AI researchers, and a much larger number of the applications experts already available to the corporation.

2.4.1.4 Shell Summary: Dr. Smith from Shell Development Company observed that Shell evaluates emerging technologies in terms of three basic issues:

- The Scope of the technology . . . what does it do that is not better done with other technologies?
- The Prospectus . . . what do we expect to gain from the technology?
- How much will it Cost to achieve the technology?

Shell considered three major components of AI for internal applications: Robotics, Natural Language, and Expert Systems. Dr. Smith indicated that there are three things a company can do with respect to emerging technologies:

- nothing!
- low-level surveillance, acquiring useful systems when someone else has developed them,
- do some development itself.

In Natural Language and Robotics, Shell is taking the second option. In Expert Systems, Shell is taking some initiative, and has developed two expert systems . . . one of which is called NDS, a communication network fault diagnostic system developed for Shell by a contractor and still in the prototype stage.

Dr. Smith asserts that he is enthusiastic about AI, particularly about the tools AI has produced: the graphics environments of the Lisp machines and the Methodology of Expert Systems as a way of organizing heuristic procedures.

2.4.1.5 TI Summary: Texas Instruments, Dr. Hollister observed, is probably the world's third largest computer company. TI manufactures its own chips; collects, processes and interprets oil exploration data; manufactures metallized and electronic products as well as consumer products like Speak-and-Spell. TI views AI as a technology as revolutionary as the transistor and the company invests several million dollars in AI research each year. Since 1978, Natural Language has been a major focus of TI's AI research. TI's overall objective here is to enable people to interface with computers, data bases, and other information-processing products without requiring specialists in data base programming to be on hand. Other AI activity areas at TI are Expert Systems, Speech Understanding Systems, Planning Systems and Symbolic Computing.

TI and Western Digital are probably the major providers of data processing services to the oil industry. TI has experts who interpret "G-LOG data" to infer geological structure, and the company would like to capture that expertise in computer programs. As another potential application area, TI manufactures its own silicon for computer chips. There are a few TI employees in "Augustine's upper decile" when it comes to adjusting the process to produce pure silicon; that valuable expertise should also be captured. Automated chip design for testability, and automated chip inspection and flaw detection are also expert systems targets at TI.

TI relies on robotics for production cost reduction in the process of assembling and testing calculators.

In summary, Dr. Hollister observed, TI justifies its investment in AI in terms of expected improvements in existing products as well as innovative new products. The impact from AI is expected to be felt in the late 1980's and 1990's.

2.4.1.6 Westinghouse Summary: Dr. Carl Love commented that Westinghouse's investment in AI is almost exclusively in developing knowledge-based systems technology. For scientific research, Westinghouse is relying on the university community to do the job.

Today, Dr. Love observed, the US is a knowledge-based economy. The successful manufacturing operations aren't the largest, but the smartest . . . those which use knowledge. Westinghouse has a strategy with respect to AI. Its general principles are:

- The company can't sit back and wait while AI is developed by others, then go out, buy the technology, and drop it like a bomb on the company's problems. Realistically, the company has to learn the technology by doing it.

- Westinghouse should develop links with leading universities; the company has derived true benefits from its relationship with CMU.

- A company has no alternative but to provide training for its own people. Westinghouse has a substantial training program in place and is looking at a component for training program managers, in addition to training the people who will actually do the work.

- Westinghouse is establishing two or three internal AI Centers of Excellence, as well as a distributed base of technical competence in AI.

Generic Westinghouse expert systems application projects include:

- order entry systems - Westinghouse is learning a lot by talking with DIGITAL, about treating situations where complex combinatorics are required and where "humans find infinitely many ways to screw things up."

- the areas of equipment testing, equipment diagnostics, and documentation of software.

- expert knowledge capture, as for example in the case of experienced power systems engineers nearing retirement age.

- computer perception, as for example in automated inspection of welds.

- manufacturing and scheduling, a joint project with CMU.

As to marketable AI products, Westinghouse sees enormous opportunities for embodying technical service products in expert systems which, unlike consultants, are readily portable; also, in systems which support human operators in controlling nuclear facilities. In addition, Dr. Love believes industrial knowledge-based systems may be instrumental in achieving substantial reduction in production costs.

2.4.1.7 Citibank Summary: Dr. Dan Schutzer described Citibank as a business characterized by:

- rapidly changing environments,
- need for timely information,
- requirements for data capture and decision support,
- large volumes of data to analyze,
- scarcity of experts.

In bond transactions, for example, the Citibank analyst searches rapidly through large volumes of data for profitable opportunities. In the Citibank working environment, Dr. Schutzer believes AI technology will make contributions in the areas of:

- capturing expertise in making profitable bond trades,
- capturing expertise in making long-term predictions,
- man-machine interfaces,
- development environments for rapid prototyping of experimental systems.

2.4.2 The Japanese Fifth Generation Computing Project

Speaker: Dr. John Alan Robinson
Professor of Computer and Information Science, Syracuse University

Subject: Japan's Fifth-Generation Computing Project: Objectives, Status, and Prospects

In describing the Japanese Fifth Generation project, Dr. Robinson drew the following analogy: Isaiah Berlin has contrasted the Hedgehog and the Fox . . . the Hedgehog knows one big thing, the Fox knows lots and lots of little things. The central thesis of the Japanese Fifth Generation Project asserts that lots and lots of the little pieces that constitute modern computer programming and software engineering can be unified by the one big idea of Logic Programming.

In Dr. Robinson's view, the Japanese project is a "symbolic Sputnik"; the U.S. could be doing it too, but we haven't accepted the idea that logic ought really to be at the center of computer science.

The Institute for New Generation Computer Technologies (ICOT) is the central facility for the Japanese project which was initiated in 1982 and directed by Dr. K. Fuchi and Dr. Moto-Oka. The staff consists of between forty and fifty very capable and highly motivated computer scientists. The project has borrowed heavily from the European developments in logic programming and PROLOG, notably from the work of Robert Kowalski and David Warren.

Ambitious and very detailed project plans have been publicized by ICOT. The goal of the project is to develop by the 1990's a super-computer with 10^{12} bytes of main memory and capable of 10^{9} logical inferences per second (lips). The machine language may combine both logic and Lisp. The project has seven major themes and more than two dozen sub-themes, shown in the following table.

Basic application systems 1	1–1) Machine translation system 1–2) Question answering system 1–3) Applied speech understanding system 1–4) Applied picture and image understanding system 1–5) Applied problem solving system
Basic software systems 2	2–1) Knowledge base management system 2–2) Problem solving and inference system 2–3) Intelligent interface system
New advanced architecture 3	3–1) Logic programming machine 3–2) Functional machine 3–3) Relational algebra machine 3–4) Abstract data type support machine 3–5) Data flow machine 3–6) Innovative von Neumann machine
Distributed function architecture 4	4–1) Distributed function architecture 4–2) Network architecture 4–3) Data base machine 4–4) High-speed numerical computation machine 4–5) High-level man-machine communication system
VLSI technology 5	5–1) VLSI architecture 5–2) Intelligent VLSI CAD system
Systematization technology 6	6–1) Intelligent programming system 6–2) Knowledge base design system 6–3) Systematization technology for computer architecture 6–4) Data base and distributed data base system
Development supporting technology 7	7–1) Development support system

Many of the sub-themes are major AI projects in their own right, such as speech and picture understanding systems. ICOT has developed each of these project plans in considerable detail.

In Dr. Robinson's view, this complex roadmap camouflages the true simplicity of the overall plan. During the first phase (1982-1985), Dr. Fuchi intends to develop a PROLOG workstation. This will be ICOT's version of the U.S. Lisp Machine, and will be an essential tool for subsequent phases of the project. During phase two (1985-1989), they intend to build a parallel inference machine. During the third phase (1989-1992) they intend to develop the "super-parallel inference machine," described above.

In summary, Dr. Fuchi has placed his bet on the central idea of Logic Programming. If they continue on this path, Dr. Robinson believes they will get where they intend, at least in terms of their engineering goals. With the ambitiously stated AI goals such as speech and picture understanding, success is less predictable . . . here they'll encounter the same sorts of difficulties that we do. Dr. Fuchi's adoption of Logic Programming is to be commended, in Dr. Robinson's view. Many of the problems of proving correctness of programs simply go away, with this approach. If the specifications are formulated in logic and the programs are deduced from the specifications, then the programs are bound to be as correct . . . or as faulty . . . as the original specifications. Logic Programming is also ideal for extending the concepts of relational data bases, for knowledge representation in expert systems, and for natural language parsing. In Dr. Robinson's opinion, the Europeans, and now the Japanese, have been quicker to grasp the advantages of logic programming for such applications than the U.S. has been.

2.5 INVESTMENT STRATEGIES BY DEFENSE: A summary of DARPA's Strategic Computing Program.

Speaker: Dr. Robert Kahn
Director, Information Processing
Techniques Office, DARPA

Topic: DARPA Strategic Computing Program

Summary: Dr. Kahn observed that the DARPA strategic computing project is aimed at Defense needs for intelligent machines. The major program goals are:

- o to speed up current logic machines by four orders of magnitude, from the present capability of 10^4 to a target capacity of 10^8 logical inferences per second.
- o to develop generic software packages for computer vision, speech, intelligent data bases and other functional areas.
- o to focus technology on three military demonstration systems . . . naval battle management, autonomous land vehicles, and automated aircraft cockpits.
- o to increase the number of qualified faculty and students in computer science.

The U.S. Government is concerned with supercomputers for symbolic computing and for numeric computing as well. There are supercomputing programs in each of these areas.

The DARPA program is concerned mainly with advancing the state-of-the art in AI and symbolic computing rather than with numerical computing. A Defense Science Board task force, headed by Professor Joshua Lederberg, is preparing reports on:

- o high impact areas of symbolic computing within Defense,
- o requirements on the technology base, imposed by the goals of the supercomputing project,
- o how to introduce AI and supercomputing into Defense,
- o what the Defense investment strategy should be, in regard to supercomputing.

These reports should be ready in 1985.

On the issue of numerical supercomputing, the White House, represented by Dr. Keyworth of OSTP, has set up a Federal Co-ordinating Committee on Science and Engineering Technology Policy, known as the Picksett Committee, which is addressing and reporting to OSTP on these issues.

In planning the strategic computing program, DARPA developed a requirements analysis by matrixing six major military applications against twelve generic AI/Symbolic Computing functionalities. This matrix is shown below.

MATRIX OF SYSTEM CAPABILITIES vs. MILITARY APPLICATIONS

	AUTONOMOUS VEHICLE	BATTLE MANAGEMENT & ASSESSMENT	PILOT'S ASSISTANT	TERMINAL HOMING	AUTOMATED DESIGN & ANALYSIS	WAR GAMING
VISION	R	O		R		
SPEECH		O	R			
NATURAL LANGUAGE	R	R			O	R
INFORMATION FUSION	R	R	R			
PLANNING & REASONING	R	R	R		R	R
SIGNAL INTREPRETATION	R	R	R			
NAVIGATION	R			R		
SIMULATION/MODELING	R	R			R	R
GRAPHICS/DISPLAY		R	R		R	R
DB/IM/KB	R	R	R	R	R	R
DIST. COMMUNICATION	R	R	R			R
SYSTEM CONTROL	R	O	R	R		

R REQUIRED CAPABILITY
O OPTIONAL

Kahn 19

Battle Management emerged as the "richest" problem area, requiring essentially all the AI functionalities; Intelligent Data Handling referred to as Data Bases and Information Management and Knowledge Systems (DB/IM/KS), emerged as a functionality required by all the military applications.

DARPA's strategy will be to stimulate technology growth by making computing tools readily available to the academic and industrial research communities, and by allowing these communities to procure the tools they need. In the microelectronics area, DARPA intends to emphasize a "fast turnaround program," within which chip design and multiprocessor designs can be submitted to fabrication facilities over the ARPANET, (the Mead-Conway approach), with production and delivery to the designer targeted for four to six weeks.

Modularity of AI systems is another goal; DARPA would like to make it fairly simple to "plug together" the vision sub-systems, speech sub-systems, and knowledge based sub-systems emerging from AI research into integrated AI systems.

2.6 THE TACTICS OF AI TECHNOLOGY TRANSFER: The pragmatics of starting and maintaining AI R&D facilities.

Panelists:

Mr. George Lukes, Physical Scientist, Research Institute
U.S. Army Engineer Topographic Laboratories
Dr. James S. Albus, Chief, Industrial Systems Division National Bureau of Standards
[redacted] Image Research Scientist STAT
Office of Research and Development, Central Intelligence Agency
Dr. Jude E. Franklin, Manager, Navy Center for Applied Research in Artificial Intelligence
Dr. Northrup Fowler, III, Computer Scientist, Rome Air Development Center
Dr. David Brown, Assistant Director, Advanced Computer Systems Department, SRI International

Topic: What are the Problems and Requirements for Starting and Maintaining a Government AI Facility?

2.6.1 USAETL Summary: Mr. Lukes and his associates at ETL are tasked with deriving mapping and terrain data from aerial imagery. This problem domain is "data rich and information poor." ETL has been exploring various techniques of automation for two decades. Mr. Lukes believes that the in-house AI facility is an effective means to transfer AI technology to ETL. Mr. Lukes observed that the single most difficult problem facing the AI facility is the scarcity of qualified people. The second most important factor is maintaining a realistic, long-term management commitment to the AI facility; results won't come easily, quickly, nor cheaply. Management must commit both to capital investment . . . estimated at \$100K to \$150K per research workstation; and to developing human resources . . re-orienting and re-training some of ETL's best people in new AI technology.

In Mr. Lukes' view, ETL has a critical near-term need to develop a cadre of talented people who are sufficiently knowledgeable about AI that they can interact effectively with the AI research community and the AI technology base. This necessitates hands-on experience, and developing this cadre is the current focus of ETL's in-house AI laboratory.

The emphasis of ETL's engineering development efforts is on interactive, computer-assisted workstations. These workstations are both hard-copy (stereoscopic aerial photography) and soft-copy (digital image processing). ETL is concentrating on interfacing the workstations with the laboratory's AI research capabilities.

2.6.2 NBS Summary: In Dr. Albus' view, there are several widespread misconceptions concerning the current realities and future possibilities of AI and robotics that may deter a Government agency's management from investing in in-house AI and robotics development:

- o Misconception #1: The futuristic world of intelligent machines depicted in "Star Wars" is here today, and further investment is not required.
- o Misconception #2: AI and robotics are science fiction, and investments made today in these ideas are a waste of effort and money.
- o Misconception #3: AI and robotics are just commonplace engineering. There's no science to it, and it's not a proper investment of R&D funds.
- o Misconception #4: AI and robotics are appropriate areas for industrial product development. Since the private sector will pour in all the money that's needed, government doesn't need to get involved in R&D.

Dr. Albus followed with specific advice on starting in-house laboratories in Government agencies:

- o The program should be tightly coupled to the agency's mission. This may take some ingenuity! At the NBS Industrial Systems Division, the robotics program is tied to measurement standards and computer interface standards in manufacturing.
- o The laboratory manager should pick the right problem, making sure it is important to the agency's management, and that it is do-able. In Dr. Albus' ironic words, "Be lucky!"
- o The manager should be able to make the case that he knows what he's doing, that the laboratory's activities have scientific merit and are critical to the agency's mission.

In the long term, this technology may bring about the "factory of the future," a hybrid system merging industrial robotics with AI principles. The intelligent factory may significantly impact world industry and the world economy.

STAT

2.6.3 ORD Summary: [redacted] observed that ORD has traditionally encouraged its intelligence and technology experts to contract out research and development on Agency problems in their areas of expertise. ORD's in-house AI lab is a venture in a new direction, assembling within the Agency the right mix of people and technology to explore and develop symbolic processing for imagery and signals. The ORD AI Lab is faced with industry competition for scarce human resources; security complicates this problem.

In [redacted] view, ORD is following directions similar to those outlined by Mr. Lukes in training and developing its own knowledge engineers. The ORD AI Lab will concentrate on providing a resource environment where individuals can come on rotational assignments from operational elements to work with R&D teams of ORD personnel and contractors using in-house computers and software. STAT

The Laboratory is now completing installation of a VAX 11/780 and Symbolics 3600 and Xerox 1100 Lisp Machines. AI software and hardware selection is largely driven by current activities in the AI research community and the major AI funding agencies, like DARPA. Wherever there is a significant development of a system or tool that fits an Agency problem, that tool will be acquired and adopted. In Dr. Gifford's view, building the AI laboratory will be an evolutionary effort . . . the best way to start is to plunge in and begin working problems.

2.6.4 NCARAI Summary: The Navy Center for Applied Research in Artificial Intelligence (NCARAI) was initiated in 1980, largely due to the Chief of Navy Research, Admiral Baciocco, who conceived the Center as a means for coupling basic AI research to practical Navy applications. Dr. Jude Franklin agreed with Mr. George Lukes' assertion that the scarcity of qualified people is a primary obstacle to starting an AI facility. NCARAI has a dynamic program of recruiting visiting scientists from Universities, granting part-time appointments, and hosting visiting researchers from Army, Navy, and Air Force laboratories.

In selecting its portfolio of research problems, NCARAI has not emphasized the computer vision and robotics areas on the grounds that other government facilities already have successful ongoing programs in these areas. NCARAI has an active research program in Expert Systems, including the Marine Corps "BATTLE" weapons allocation system, an electronics troubleshooting system, an automated code generator for automatic test equipment, an amphibious assault planning system, and a system to assist radar image interpreters in classifying targets imaged with inverse synthetic aperture radar.

In Natural Language, there is an active project addressing automated Navy message summarization and prioritization. In Distributed Problem Solving, a project is underway to generate composite pictures of battle scenarios from multi-source data (including radar, esm, and acoustics). NCARAI is also working with NISC/Suitland to develop an expert system for analyzing and assessing submarine activity.

Now, and in the future, Dr. Franklin intends to emphasize computer environment growth and to encourage technical publications by NCARAI staff. A computing environment which allows for a machine dedicated to each project is one of the most important factors in attracting researchers to the center.

2.6.5 RADC Summary: Dr. Fowler began by emphasizing "scaling" as a critical problem for military applications of AI. Using secondary storage memory for the sake of example, Dr. Fowler observed that for all the current successful AI projects the knowledge bases reside in virtual memory at run time. In the real world of Air Force problems, the knowledge bases will be extremely volatile and very large. AI research does not adequately address the issue of volatile data and very large memory requirements.

With regard to the pros and cons of in-house AI facilities, Dr. Fowler cited the report of the Packard committee. In May, 1983, this committee visited and subsequently critiqued a number of Government laboratories. The Committee observed that university laboratories tend to be superior to Government laboratories in basic research, while industry laboratories are often superior to Government laboratories in developing practical applications. Before starting yet another government laboratory, Dr. Fowler advised that a careful examination be made of the tradeoffs between doing the application in-house, at a university, in industry, or a combination of all three.

2.6.6 SRI Summary: Mr. Brown described SRI's AI R&D program planning projects for NASA and the Army. Based on a review of NASA programs (circa 1980), SRI recommended that NASA/Goddard concentrate on AI applications in the area of Space Tracking and Data Systems. Space Mission Planning and Monitoring is another prime candidate; in particular, the use of planning systems on the NASA space station. SRI's report included a proposed management and support structure for a NASA research group to be situated at the AMES Research Center.

A second SRI study performed for the U.S. Army in 1981 focused on battlefield applications of AI and robotics and identified the required research to enable these applications.

Appendix 1

AI Symposium Program, 1983

ARTIFICIAL INTELLIGENCE SYMPOSIUM

"INTELLIGENCE APPLICATIONS OF AI"

Sponsored By

U.S. INTELLIGENCE COMMUNITY

December 6, 7, & 8, 1983

MAIN AUDITORIUM
CIA HEADQUARTERS
LANGLEY, VA.

By Invitation Only

Symposium Co-ordinator: SMART SYSTEMS TECHNOLOGY, Inc.
Specializing in Artificial Intelligence Education and System Implementation

Program - First Day, Morning Session, Tuesday, December 6, 1983

STAT

Symposium Chairman:

Executive Secretary, AI Steering Group
Office of Research and Development
Central Intelligence Agency

8:30	Dr. Philip K. Eckman Chairman, AI Steering Group Office of Research and Development Central Intelligence Agency	Symposium Objectives
8:45	Dr. Richard D. DeLauer Under Secretary of Defense for Research and Engineering	AI and Defense
9:00	Dr. Patrick H. Winston Professor of Computer Sciences Director, Artificial Intelligence Laboratories MIT	Overview of Artificial Intelligence Research
10:00	Coffee Break	
10:15	Dr. Patrick H. Winston, MIT	Overview of AI Applications
11:00	Industry Panel: Xerox, Apollo, Symbolics, DEC, E-Systems, LMI	Descriptions and Schedule of AI Demonstrations

Program Note: Several computer manufacturers have made AI processors available during the three days of this conference for demonstrations of AI tutorial programs and working AI systems. These will be available during the mid-day breaks and following the afternoon sessions in the Tunnel adjacent to the Auditorium. Schedules of these demos are posted in the Auditorium Lobby. Additional demos may be scheduled by arrangement with the manufacturers' representatives.

12:00 Lunch Break - System Demos in Auditorium Tunnel

Program - First Day, Afternoon Session, Tuesday, December 6, 1983

Session Moderator: Lt. Col. James T. Jones
Vice Dean
Defense Intelligence College

1:30	Dr. Rodney Brooks Assistant Professor Computer Science Department Stanford University	Overview of AI Applications in Computer Vision and Image Understanding
2:30	Coffee Break	
2:45	Dr. Raj Reddy Professor of Computer Science Carnegie-Mellon University Director, The Robotics Institute	Overview of AI Applications in Robotics and Speech Understanding
3:45	Panelists: Dr. Patrick H. Winston, MIT Dr. Raj Reddy, CMU Dr. Rodney Brooks, Stanford	Current Trends in AI Research
4:15	End First Day Program - System Demos in Auditorium Tunnel	

Program - Second Day, Morning Session, Wednesday, December 7, 1983

Session Moderator: Commander Paul E. Tilson, Jr.
U. S. Navy

THIS SESSION WILL BE CLASSIFIED SECRET

8:30	Mr. David Y. McManis National Intelligence Officer/ Warning Central Intelligence Agency	Intelligence Requirements
9:15	Dr. Gian-Carlo Rota Director's Office Fellow Los Alamos National Laboratories Professor of Applied Mathematics and Philosophy MIT	Observations and Speculations on the Effect of the Computer on Scientists
10:00	Coffee Break	
10:15	Dr. James Slagle Naval Research Laboratory	"BATTLE" Resource Allocation Project: Expert Advisor for Weapons Allocation.
10:45	Dr. Gerald A. Wilson Senior Computer Scientist Advanced Information and Decision Systems	Computer-Based Assistant for Science and Technology Analysis
11:30	Mr. Mark Williams Senior Staff Engineer ESL Inc.	AI in Signal Processing
12:00	Lunch Break - System Demos in Auditorium Tunnel	

Program - Second Day, Afternoon Session, Wednesday, December 7, 1983

Session Moderator: [redacted] STAT
Image Research Scientist
Office of Research and Development
Central Intelligence Agency

THIS SESSION WILL BE CLASSIFIED SECRET

1:30	Dr. Donald Close Program Manager Hughes Aircraft Company	Image Understanding Project
2:00	Dr. Bobby Hunt Chief Scientist Science Applications, Inc.	AI Applications in Synthetic Aperture Radar Image Inter-pretation
2:30	Coffee Break	
2:45	Mr. John C. Kelly Senior Staff Engineer TRW Defense Systems Group	An Operational Artificial Field Engineer for Tuning A Signal Sorter--The User's Viewpoint
3:30	[redacted] STAT EH/R&AD/DPG/NPIC Central Intelligence Agency	AI Applications in Image Analysis
4:00	End Second Day Program - System Demos in Auditorium Tunnel	

Program - Third Day, Morning Session, Thursday, December 8, 1983

Session Moderator: Mr. Norman Glick
Chief Techniques Staff
National Security Agency

8:30	Dr. Robert Kahn Director, Information Processing Techniques Office - DARPA	DARPA Strategic Computing Program
9:15	Dr. John A. Robinson Professor of Computer and Information Science Syracuse University	Japan's Fifth-Generation Computing Project: Objectives, Status and Prospects
10:00	Mr. John N. McMahon Deputy Director of Central Intelligence Central Intelligence Agency	AI and Intelligence
10:15	Coffee Break	
10:30	Panel: What are the Problems and Requirements for Starting and Maintaining a Government AI Facility?	

Panelists:

Dr. George Lukes, Physical Scientist, Research Institute
US Army Engineer Topographic Laboratories
Dr. James S. Albus, Chief, Industrial Systems Division
National Bureau of Standards
[redacted] Image Research Scientist STAT
Office of Research and Development
Dr. Jude E. Franklin, Manager, Navy Center for Applied
Research in Artificial Intelligence
Dr. Northrup Fowler, III, Computer Scientist,
Griffiss Air Force Base
Mr. David Brown, Assistant Director, Advanced Computer Systems
Department, SRI International

12:00 Lunch Break - System Demos in Auditorium Tunnel

Program - Third Day Afternoon Session, Thursday, December 8, 1983

Session Moderator: [redacted] STAT
Chief - Development & Implementation Division
DIA

1:30 Industry Panel : Why is Private Industry Investing in AI?

Panelists:

Dr. Charles Herzfeld, Vice President and Director of Research and Technology, ITT Corporation

Dr. Edward C. Taylor,, Director of Requirement Analysis TRW, Defense Systems Group

Mr. Dennis E. O'Connor, Senior Group Manager for Intelligent Systems Technologies Group, Digital Equipment Corporation

Dr. Carl Smith, Staff Computer Scientist, Shell Development Company

Dr. Floyd Hollister, Senior Member, Technical Staff, Computer Science Laboratory, Central Research Laboratories, Texas Instruments, Inc.

Dr. Carl Love, Senior Consultant, Corporate Planning, Westinghouse Electric Corporation

Dr. Dan Schutzer, Vice President, Citibank

2:30 Coffee Break

2:45 Industry Panel Continues

STAT

3:30 [redacted] Future Directions for the AI Steering Group
Executive Secretary, AI Steering Group
Office of Research and Development
Central Intelligence Agency

4:00 End Third Day Program

Appendix 2

AI Symposium Program, 1982

NOVEMBER 1982
AI SYMPOSIUM
INTELLIGENCE APPLICATIONS OF ADVANCED COMPUTER AND INFORMATION TECHNOLOGY: FOCUS ON ARTIFICIAL INTELLIGENCE

Sponsored by:
OFFICE OF SCIENTIFIC WEAPONS RESEARCH
AND
OFFICE OF RESEARCH AND DEVELOPMENT

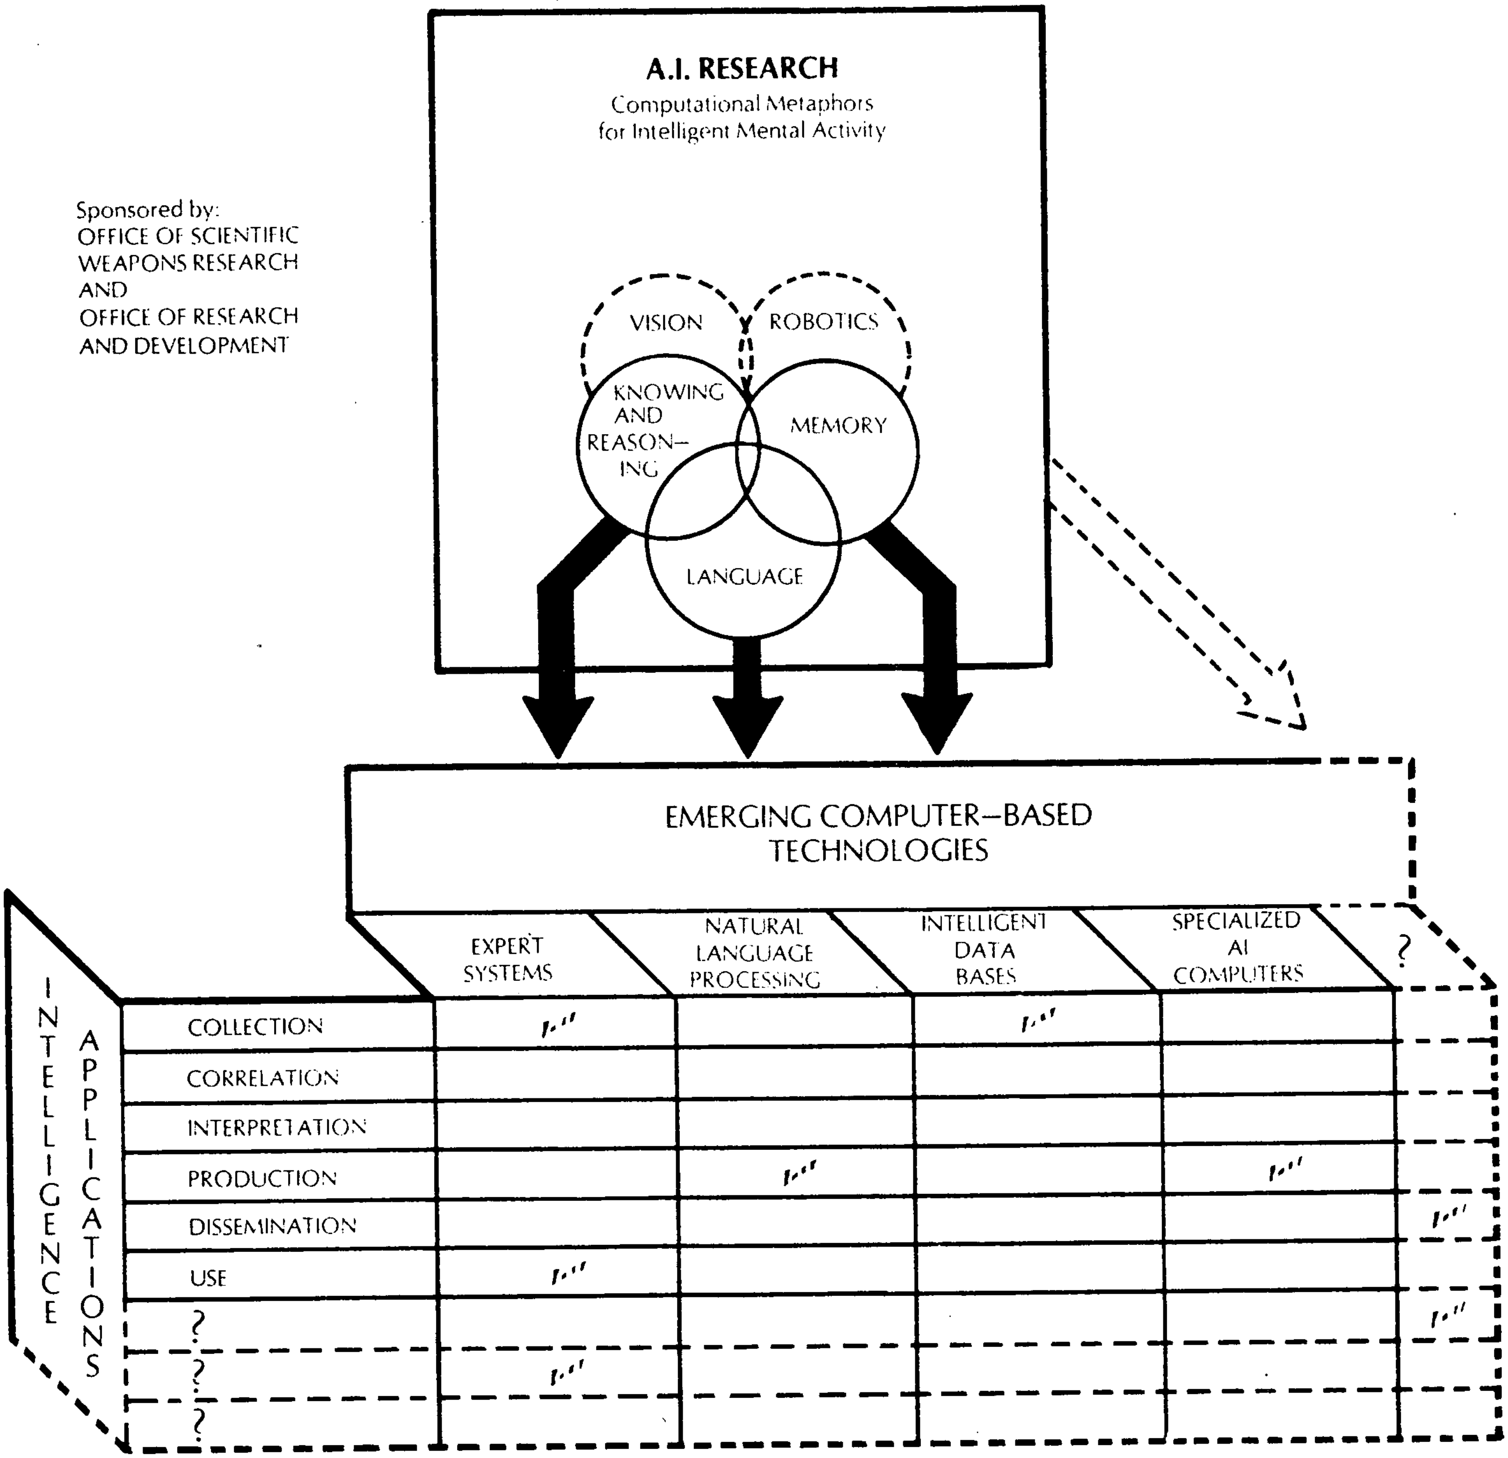

Symposium Co-ordinator: SMART SYSTEMS TECHNOLOGY, Inc.
Specializing in Artificial Intelligence Education and System Implementation

Tuesday, Nov. 30, 1982

	8:30	Opening Remarks by Evan Hineman, Deputy Director, Science and Technology
9	8:45	History and Overview of AI Research and Technology Dr. Drew McDermott, Yale University
10	10:10	Break
	10:25	The Xerox DOLPHIN Computer XEROX
11	10:50	The SYMBOLICS Lisp Machine SYMBOLICS
	11:20	To be announced
12	11:50	Lunch
1	1:15	History and Overview of Expert Systems Research and Technology Dr. Harry Pople, University of Pittsburgh
2	2:40	Break
3	2:55	AI Research and Technology at DEC Digital Equipment Corporation
4	3:45	A Case History: Building an Expert System for Ocean Surveillance and Target Tracking VERAC Corporation
	4:30	

Lunch: During this period the Lisp Machines will be available for additional demonstrations

Wednesday, Dec. 1, 1982

	History and Overview of Machine Inference and Intelligent Data Bases Dr. J.A. Robinson, Syracuse University
9:55	Break
10:10	DADM: A Knowledge-Based Data Retrieval System System Development Corporation
10:55	Break
11:05	INTELLECT: A Natural Language Data Base Query System Artificial Intelligence Corporation
11:50	Lunch
1:15	AI Research and Development at TI (Attendance Restricted to Gov't) Texas Instruments Corporation
2:05	Break
2:15	Rule-Based Information Retrieval AI&DS
2:45	Script-Based Situation Assessment Hughes Research Laboratory
3:30	Break
3:40	Intelligence Applications of AI Panel Moderator: STAT Office of Research and Development
4:30	

- CONTENTS -

STAT

CONTENTS (continued)