THE SRI ARTIFICIAL INTELLIGENCE CENTER

A Brief History

January 24, 1984

Technical Note 317

By: Nils J. Nilsson, Director

Artificial Intelligence Center
Computer Science and Technology Division

This paper was prepared as an introduction to the COMTEX/Scientific DataLink Microfiche Edition of the SRI Artificial Intelligence Center Technical Notes. Also to appear in the AI Magazine.

SRI International

333 Ravenswood Ave. • Menlo Park, CA 94025
(415) 326-6200 • TWX: 910-373-2046 • Telex: 334-486

THE SRI ARTIFICIAL INTELLIGENCE CENTER

A BRIEF HISTORY

Charles A. Rosen came to SRI in 1957. I arrived in 1961. Between these dates, Charlie organized an Applied Physics Laboratory and became interested in "learning machines" and "self-organizing systems." That interest launched a group that ultimately grew into a major world center of artificial intelligence research—a center that has endured twenty-five years of boom and bust in fashion, has "graduated" over a hundred AI research professionals, and has generated ideas and programs resulting in new products and companies as well as scientific articles, books, and this particular collection itself.

The SRI Artificial Intelligence Center has always been an extremely cohesive group, even though it is associated with many contrasting themes. Perhaps these very contrasts are responsible for its vitality. It is a group of professional researchers, but visiting Ph.D. candidates (mainly from Stanford University) have figured prominently in its intellectual achievements. It is not part of a university, yet its approach to AI has often been more academic and basic than those used in some of the prominent university laboratories. For many years a vocal group among its professionals has strongly emphasized the role of logic and the centrality of reasoning and declarative representation in AI, but it is also home to many researchers who pursue other aspects of the discipline. Far more people have left it (to pursue careers in industry) than are now part of it, yet it is still about as large as it has ever been and retains a more or less consistent character. It is an American research group, supported largely by the Defense Department, but, from the beginning, it has been a melting pot of nationalities.

Our early history actually had little to do with computer science. By 1959, Charlie Rosen's Applied Physics Lab at SRI was working on the problem of using electron-beam machining techniques to manufacture circuits consisting of millions of tiny field-emission triodes. The details of how these triodes were to be manufactured are not important to this story, but Charlie knew that not every one of them would work and that, furthermore, the task of wiring up so many components to do something useful would be an extremely difficult undertaking. He had been following research in "self-organizing" systems—systems that organized themselves on the basis of local affinities and in response to global requirements. Perhaps the organizing principles of these systems could be applied to microcircuit design.

One of the people doing work in self-organizing systems was Frank Rosenblatt, a Cornell psychologist with an interest in constructing models of neural circuits. Frank had invented a model called a "perceptron" that consisted of linear threshold elements connected together in diverse ways. In one of his instantiations of a perceptron, some threshold elements received inputs from others through resistances that could be varied by motor-driven potentiometers. These variable resistances simulated variable synaptic strengths between neurons; changing the resistances was a way of having the machine learn. Frank was dissatisfied with the cost and slowness of motor-driven potentiometers and sought Charlie's advice about more elegant electrically variable "weights." Alfred E. (Ted) Brain, a physicist working for Charlie at SRI, suggested the use of multiaperture magnetic cores whose flux density could be varied continuously. Intrigued by these ideas, Charlie, Ted and George Forsen, an electrical engineer in Charlie's lab, decided to build a "perceptron" of their own, MINOS I.

MINOS I was nearly finished when I arrived on the scene in the summer of 1961. It was configured as a "trainable" pattern recognition device. It could be presented with a pattern of binary inputs (arranged in an 8 x 8 array), and its human "teacher" could then insist on a desired binary response. If it responded incorrectly, it was "trained" by a process that automatically adjusted appropriate units among its set of 72 magnetic weights. After just a few iterations through a set of training patterns, the machine's weights would usually converge on values that would not only recognize training patterns accurately, but would also give the correct answers to training patterns that were perturbed by minor random errors. Well, that was all pretty fascinating to me—a person with some background in statistical detection theory and information theory. I decided to throw my lot in with these pioneers, and we would all try to understand better what was going on with these kinds of machines.

It became apparent to all of us that, quite apart from their interest as neuropsychological models, these machines could best be comprehended by using the tools of statistical decision theory. They implemented specific kinds of decision rules, and their "training" procedures could be viewed as techniques for collecting sample statistics of the input patterns. We continued our collaboration with Frank Rosenblatt and with David Block, an applied mathematician at Cornell who was interested in the mathematical properties of the dynamics of these systems. We also formed close ties with Bernard Widrow and Norman Abramson of the Stanford Electrical Engineering Department and with their students. Bernie Widrow was experimenting with similar kinds of devices that he called "Adalines" and "Madalines," using them as adaptive controllers. Norm Abramson was interested in machine learning as a statistical problem. One of Norm's students, Tom Cover, and another Stanford student, Bradley Efron, (both of whom later became Stanford professors) worked closely with us. I taught courses on these subjects at both Stanford and UC Berkeley. Notes for these courses ultimately led to my little book, *Learning Machines*, published in 1965.

By 1963 several new people, including Richard Duda, had joined the Applied Physics Laboratory, and I had become the leader of a component called the "Learning Machines Group." We had government contracts from the Office of Naval Research, the Rome Air Development Center, and from what was then called the Army Signal Corps to explore the theory underlying these machines and their potential applicability to problems of pattern and signal recognition. We constructed a larger learning machine, called MINOS II, with 6600 variable magnetic weights.

None of us had any background in the new field of computer science, although we were well aware that decision techniques similar to those implemented in our machines were also being programmed on computers. Bernie Widrow's group at Stanford had a small IBM computer (an IBM 1620 with a storage access channel) on which they ran "simulations" of their Adalines and Madalines. Since we were all under the impression that computer simulations were too slow for practical applications of these techniques, we preferred the analog computation being performed in our circuitry. We found ourselves siding with Frank Rosenblatt in countering Marvin Minsky's complaint that money was being wasted on building these special-purpose machines when computers could so easily handle not only these kinds of calculations, but more powerful "heuristic" ones as well. Nevertheless, some of us learned to program; I remember writing learning-machine programs for the SRI Burroughs 220 machine in a strange language called "Balgol." Most of what we knew about "artificial intelligence" in those days was what we gleaned from Minsky's article "Steps Toward Artificial Intelligence."

In late 1963, satisfied that MINOS II was nearing completion, Charlie became absorbed by a more grandiose vision. What would it be like, he wondered, to build a large learning machine whose inputs would come from television cameras and other sensors, and whose outputs would drive effector motors to carry the machine purposefully through its environment. He wrote a memo describing plans for such a robot or "automaton," as he called it. During 1964 we spent a lot of time planning a robot research project and discussing the idea with possible sponsors. As interest in computer science and artificial intelligence grew, we were ready to concede that our robot ought be equipped with heuristic computer programs as well as pattern-recognizing learning machines. Marvin Minsky spent a couple of weeks with us in August 1964 helping us prepare our proposal.

At about this time, Charlie attended a summer course on LISP taught by Bertram Raphael, one of Minsky's doctoral students. After receiving his degree, Bert went to the University of California at Berkeley to work with Ed Feigenbaum, who had recently completed his Ph.D. studies at Carnegie. When Ed shortly left Berkeley for Stanford, Bert decided to join us at SRI. He taught us all LISP and helped with our robot proposal.

Charlie, Bert, and I were working hard trying to make prospective sponsors enthusiastic about robots. We had some success with Ruth Davis, who was in charge of an office in the Defense Department responsible for research and development, and with Ivan Sutherland, director of the new Information Processing Techniques Office (IPTO) of the Advanced Research Projects Agency (ARPA). Finally, in November 1964, we received an informal request from ARPA to bid on a research program to develop "automatons capable of gathering, processing, and transmitting information in a hostile environment." Several research organizations submitted proposals. Ours was successful, and work began in 1966 on what was to become "Shakey" the robot. The AI labs at MIT and Stanford began ARPA-supported work on computer-controlled, vision-guided mechanical arms at approximately the same time.

While awaiting funding for our robot work, we arranged telephone access to the Q-32 computer at the System Development Corporation in Santa Monica. Bert supervised my initial attempts at LISP programming on this computer, and together we developed a rather elaborate robot simulation program called REX. One of Tom Cover's Ph.D. candidates at Stanford, Cordell Green, became interested in Bert's thesis work on semantic information retrieval; Cordell came to SRI one summer to work under Bert. Cordell built a question-answering system (QA1) for the Q-32 based on an ad-hoc theorem-proving methods.

Our continuing work on pattern recognition was also taking on a more computer-oriented tone. We had acquired an SDS 910 computer that we used for programming several learning-machine algorithms. John Munson, a physicist who joined us after completing his doctorate under Luis Alvarez at Lawrence Laboratory, quickly adopted the 910 and helped bring us into the computer age. The 910 was also used to control the MINOS II hardware in a configuration we called MINOS III. The latter included an optical "preprocessor," designed by Ted Brain, that replicated the input image 1000 times to create 1000 differently sampled inputs to the learning machine. The entire configuration was delivered to the Signal Corps in 1968. Probably the most impressive achievement of our pattern recognition work was a system that combined trainable pattern-recognition techniques with a FORTRAN compiler to recognize hand-printed FORTRAN code characters on coding sheets.

Peter Hart joined our group in 1966. Peter had done his Ph.D. studies under Tom Cover at Stanford on "nearest-neighbor" pattern classification schemes. He soon became interested in computer vision [20, 36]*, and he and Dick Duda took over responsibility for providing the vision subsystem for Shakey. At this time we were also improving our computational capabilities, first with the SDS 930 computer and then with the time-shared version—the SDS 940. The first complete prototype of Shakey, a mobile cart with a television camera and an optical range finder, was controlled from the SDS 940 over radio and TV links [19, 39, 40]. Since our work and interests were expanding well beyond pattern recognition, we renamed ourselves the "Artificial Intelligence Group" in 1966.

---

*Numbers in brackets refer to the SRI AI Center Technical Notes listed at the end of this article.

Much of our early work with Shakey was directed at studying "navigation algorithms" for calculating routes across floor space cluttered with various obstacles. We explored several techniques for representing key points in space as nodes in a graph, so that graph-searching methods could be used for route planning. Out of these studies, Peter, Bert, and I developed the A* algorithm, with a heuristic component in its evaluation function. (For robot navigation problems, the heuristic component was set equal to the straight-line distance to the goal location.)

Cordell had read about John Alan Robinson's "resolution principle" for automatic theorem proving and decided to apply it in the new version of his question-answering system, QA2. As this project was progressing at a fast pace, we decided that the "high-level" problem solving Shakey had to perform in determining how to achieve its goals could be handled by QA2. Cordell had shown how the state transformations needed for robot problem solving could be found with a theorem prover. One of Cordell's dissertation advisers at Stanford was John McCarthy. John had earlier proposed a system he called the "advice-taker"; Cordell's work could be regarded as an attempt to implement such a system. Our lab has had a close collaborative relationship with John McCarthy ever since. Cordell's Ph.D. thesis on problem solving through resolution theorem proving had a significant impact on the field and, in a sense, anticipated the subsequent development of logic programming [4, 8].

The Artificial Intelligence Group had by now grown to over a dozen people, and I was gradually becoming a manager instead of a researcher. I decided that I preferred being the latter, so I stepped down as head of the group. (Anyway, since Charlie had become more interested in AI than in the rest of his Applied Physics Laboratory, he found himself devoting more and more time to the AI Group. Eventually the latter was promoted to laboratory status, with Charlie as its director.) For my part, I wanted to spend more time thinking about problems involving heuristic search and about theorem proving, and I also wanted to write a book about the use of these methods in AI. The best way to get started on the book was to volunteer to teach an AI course at Stanford, which Art Samuel and I did together in 1969. (The book finally appeared in 1971.) While at Stanford, I met two more students, J. M. (Marty) Tenenbaum and Tom Garvey. Tom soon joined us to help Cordell program his theorem prover and, subsequently, to do his doctoral research with us on machine vision [117]. After receiving his Ph.D., Marty also came aboard to continue his vision research at SRI [84, 87, 95, 121, 123].

Motivated by the requirement that Shakey be able to see something of its surroundings, our pattern recognition work was evolving into more general algorithms for machine vision. A French visitor, Claude Brice, and Claude Fennema, of our group, developed some interesting techniques for finding homogeneous regions in scenes [17]. These regions then formed the basis for associating parts of the scene with models of objects, such as blocks, floors, walls, and doorways.

Richard Fikes joined us in 1969 after finishing his Ph.D work at Carnegie-Mellon University [14]. Rich and I, reacting to obvious efficiency problems in applying QA2 to robot problem solving, began thinking about other ways of synthesizing robot plans. We came up with STRIPS (for STanford Research Institute Problem Solver), a system that employed theorem proving to decide whether or not an action could be applied in a given state, but that used more direct and nonlogical methods for predicting what the world would be like after an action was applied [43]. By this time, Cordell's theorem-proving system had progressed to QA3 [15], which we made a subroutine of STRIPS. The design of STRIPS was also heavily influenced by Newell, Shaw and Simon's General Problem Solver (GPS). In the meantime, we had acquired a Digital Equipment Corporation PDP-10 computer. A new and final Shakey system was developed on the PDP-10. Shakey's software consisted of several new vision algorithms, lower-level action routines, and an enhanced STRIPS system that was also able to learn and use generalized versions of previously computed plans [64, 70]. We could also give our robot certain commands in English, using a program written by Steve Coles that translated English sentences into logic [41]. This new system was the subject of a 1972 film entitled "Shakey: Experiments in Robot Planning and Learning," produced by Peter Hart, Rich Fikes, and me.

After a stint as the Shakey project leader, Bert Raphael decided that he didn't mind being a manager as much as I did, so he was a natural candidate to take over the AI Lab when Charlie Rosen stepped down in 1971 to develop a project that would apply robotics to industrial assembly and inspection problems. Bert ran the group until 1976 when Peter Hart took over.

Bert's AI Lab was rocked by a series of major challenges that affected it in several ways. ARPA discontinued its support of our work on robots and asked us to seek more practical applications of our AI research. (Charlie, meanwhile, was able to secure NSF and commercial funding for a different kind of robot project—industrial robots.) We participated in an ARPA program in "speech understanding" and built up an excellent group in natural-language processing. We began work in expert systems and developed the PROSPECTOR program. We also began work in "image understanding," a term used to describe an ARPA program for applying scene analysis and vision techniques to photo-interpretation and cartographic aids. (Even with all of this activity, Bert did manage to find time to write a book about AI entitled *The Thinking Computer.*)

Throughout this period, AI research everywhere was being subjected to intense pressure to prove its relevance to practical problems. Luckily, we had several talented people who helped us survive and prosper. Earl Sacerdoti, a Stanford doctoral student in computer science, did his dissertation work at SRI on robot problem solving. First he made some major improvements in STRIPS that enabled it to solve planning problems hierarchically [78]. Then he conceived of an entirely different hierarchical planner that he called NOAH [101, 109]. We had intended to use NOAH in a sequel to the Shakey system, but ARPA said "no more robots." Charlie Rosen and Earl cooked up an idea for a project that wouldn't actually be concerned directly with a robot, but would still allow us to pursue robot problem solving and robot vision. They proposed a "computer-based consultant" (or CBC) that could serve as an expert on how to assemble, dismantle, and repair complex electromechanical equipment [94, 99]. The problem solving and planning for these actions would be the same as if we had built a robot to perform them—except that, in the CBC, the system told a human apprentice what to do instead of controlling motors. We thought that the system would be of interest to the military because of its massive training needs and the increasingly serious problem of poor performance by lower-level technicians. ARPA bought the idea—perhaps because they couldn't immediately figure out a reasonable way to say no and because, after all, the project sounded plausible. That contract allowed us to continue work on NOAH and on machine vision until 1975, when ARPA finally did say no.

One of the main points of Cordell Green's thesis was that theorem-proving systems could be used for much more than simple "question-answering." Or, to put it another way, one could ask such a system some rather complicated questions—such as what sequence of actions a robot should perform to accomplish a given goal, or what procedures would solve a given programming problem. One of the uses of QA3 was to "synthesize" programs. At the same time that Cordell was using QA3 for this purpose (about 1968), Richard Waldinger was finishing his Ph.D. dissertation at Carnegie-Mellon University on a similar technique. Since joining our lab in 1969, Richard has authored or coauthored a number of important papers on robot problem solving [26, 107], program verification [86, 132], and program synthesis [34, 52, 98, 156, 177, 246, 260] using logical methods. He has collaborated on much of this work with Zohar Manna of Stanford University. They are now finishing a two-volume textbook on the logical foundations of programming. (Richard is also famous around SRI for his generous hospitality in serving freshly brewed coffee to as many of us as can squeeze into his office about four o'clock every afternoon.)

After receiving his Ph.D. degree from Stanford in 1969, Cordell and two other Stanford doctoral students working at SRI, Jeff Rulifson and Bob Yates, began developing a successor system to QA3 called, naturally, QA4 [42, 48, 50, 60, 65, 73]. Cordell had become inspired with the idea of "programming in logic," and Jeff had kept in close touch with Carl Hewitt's work on PLANNER at MIT. QA4 was something like a version of PLANNER, called MICRO-PLANNER, but had several additional features. Richard Waldinger also worked on QA4 and used it to build a program verification system [86]. Earl Sacerdoti suggested that our INTERLISP system could be augmented to endow it with most of the QA4 features, and QLISP was the successful result [81, 118, 120]. Earl finally wrote his NOAH system in QLISP. The QA4, QLISP, and robot-problem-solving people kept in close touch with similar work at other laboratories. "Problem-solving workshops" in the Sierra Nevada, at Pajaro Dunes, California, and at Loch Tay, Scotland, gave SRI researchers opportunities to talk with people like Carl Hewitt, Jerry Sussman, and Terry Winograd of MIT, Harry Barrow, Donald Michie, Rod Burstall, and Bob Boyer of Edinburgh, John Alan Robinson of Syracuse, and Alan Kay of Utah.

Around 1970, Cordell reported for active duty as an Army lieutenant as part of his ROTC commitment. After a short stint learning to drive tanks, he was reassigned to ARPA as a program manager. In that capacity, Cordell began a program of research in "speech understanding" in 1972. Several contractors were chosen to participate in this project—among them the SRI AI Lab. Don Walker, a linguist from Mitre, joined us to become the Speech-Understanding System project leader. Much of the technical work on this project during its early years was done by Bill Paxton and Ann Robinson. Bill designed several speech-understanding systems and developed a new parsing system called a "best-first parser" [92]. His Ph.D. dissertation at Stanford grew out of his work at SRI [134].

During the mid-'70s, several other natural-language people joined us. Barbara Grosz, a Ph.D. candidate at UC Berkeley, finished her dissertation at SRI [150, 151]. Her work on focus in discourse and its effect on the interpretation of referring expressions launched one of the leading themes of the SRI natural-language group [150, 151, 185, 188, 292]. Jane Robinson, a linguist from IBM, became our chief "grammarian" [97, 112]. Gary Hendrix, a doctoral student of Bob Simmons at the University of Texas, initiated our interest in semantic networks [105, 164].

The Speech-Understanding Project continued for five years until 1977. SRI worked jointly with the System Development Corporation to build an integrated system—SRI on "high-level" syntax, semantics, and pragmatics, SDC on acoustic processing and phoneme recognition. Although our collaboration never resulted in a complete, integrated system, the SRI portion was completed and formed the basis for subsequent natural-language research at SRI [91, 110, 111, 141, 142]. A description of much of this work is contained in a book edited by Don Walker, entitled *Speech Understanding.*

After ARPA discontinued its support of speech understanding, we were able to obtain a small amount of funding from NSF for basic research in what we called Task-oriented Dialogue Understanding Systems (TDUS). Ann Robinson played a key role in organizing our efforts to write that proposal and in running the subsequent project. These were lean years for our natural-language people. Increasingly called upon to help with ARPA-sponsored applications-oriented work, they received little support to pursue what we all knew were important basic research goals. It was difficult to acquire and keep topflight researchers. We did manage to hire Jerry Hobbs—but lost Bill Paxton. Gary and Barbara attempted to secure major funding from NSF but were turned down. Nevertheless, the TDUS work and the process of writing the large NSF proposal did reinforce our picture of "language as purposeful action" [210]. And we completed a large computer grammar of English (DIAGRAM [205]), along with a companion parser called DIAMOND.

During 1974, we were asked by ARPA to develop a rather detailed five-year plan for our CBC project. We were also encouraged to have the system give advice about a more obviously military piece of equipment than the air compressor we had been using as a focus for our work. ARPA was simultaneously becoming more insistent about applying AI techniques to so-called "command and control" problems. Our detailed CBC plan notwithstanding, ARPA decided in 1975 to terminate SRI's work on the CBC project and to "redirect" us toward more applied activities. (I was not enthusiastic about this new course, so I stepped down as leader of the ARPA project.)

Our ARPA work split into two components at this time. Marty Tenenbaum and Tom Garvey were successful in conceiving a project to apply vision techniques to photo-interpretation [127]. This was funded under ARPA sponsorship as part of its "image-understanding" program. Although its rationale was rooted in specific applications of military importance [137, 158, 196, 200, 244], the program nevertheless encouraged and supported a wide range of excellent basic research in machine vision at several universities and at SRI. Marty was soon joined by Harry Barrow, who decided to leave Edinburgh University to work with us. Marty and Harry began a long period of collaboration in vision research that led to many important papers [108, 137, 157, 221, 222]. (Incidentally, Harry Barrow obtained an early copy of DEC-10 Prolog from Edinburgh that some of us experimented with. It also attracted the attention of a Japanese visitor, Koichi Furukawa, who later played an important role in the adoption of Prolog as the kernel language of the Japanese Fifth-Generation Project.)

Earl Sacerdoti and Gary Hendrix were instrumental in organizing a second ARPA project to develop a natural-language interface to distributed data bases. This system, called LADDER [140, 145, 152, 163, 224], was based on Gary's semantic grammar LIFER [135, 138]. (Many of these acronyms were invented by Earl—our champion nomenclator.) The LADDER project benefited from our expertise in natural-language processing, but, since it had an applications orientation, it did not underwrite the basic research we wanted to conduct. Daniel Sagalowicz (another of Tom Cover's Ph.D. candidates) came over to us upon obtaining his degree and, after a stint doing systems programming for our computers, joined the LADDER project and became our expert on data bases.

I decided to write another book. Again it was time to give a seminar at Stanford, which I did in the fall of 1976. This time I met Bob Moore, who was finishing his MIT Ph.D. work at Stanford. I also met a Stanford student named Doug Appelt who was looking for a dissertation topic. Bob taught Doug and me about problems in reasoning about knowledge and about his new techniques for doing such reasoning. These methods were developed in Bob's thesis [191] and later, in Doug's Ph.D thesis [259, 312], applied to the problem of generating English sentences. Bob joined SRI to work on the LADDER project and, ultimately, to pursue his interests in reasoning and representation. Doug came over to SRI to do his Ph.D. work under Gary Hendrix. (My book finally appeared in 1980.)

While worrying about ARPA's apparent lack of interest in our CBC project, we became intrigued by what Ted Shortliffe was doing at Stanford in applying AI ideas to the problem of medical diagnosis. We had let him use the SRI PDP-10 computer to develop a program, MYCIN, that gave advice regarding certain bacterial infections and their treatment. In the summer of 1974, Peter Hart and Dick Duda wondered what other application areas might be appropriate for this kind of program. Peter and Dick had just finished a book entitled *Pattern Classification and Scene Analysis*, neither of them was inextricably involved in the ARPA projects, and we were all interested in finding some way of extending our sponsorship base that would make us more independent of the shifting currents at ARPA.

Peter hit upon the idea of building an expert system (as such programs came to be called) to give advice to geologists about ore deposits. Contact was made with several economic geologists (those who are experts in ore deposits), notably Charles F. Park, Jr. at Stanford and a Stanford doctoral student in geology, Alan Campbell. We had worked out a scheme [136] for encoding geologic knowledge in semantic networks (similar to those being studied by Gary Hendrix) and in production rules (similar to those used in MYCIN). An early version of an expert system for "Mississippi Valley"-type lead-zinc deposits was developed on the basis of some ideas developed by Peter, Dick, and me for dealing with uncertain information [124]. Peter and Dick were able to secure funding from NSF and the U.S. Geological Survey for expanding this system and adding knowledge about other types of ore deposits. This work led to the PROSPECTOR system [155]. Several additional people soon joined the project, including Rene Reboh, John Gaschnig, and Kurt Konolige.

After I had disentangled myself from ARPA work, Dick Duda and I also attempted to develop a similar system to give advice about agricultural pest management. We developed a simple prototype system but were unable to obtain funding for further work on what we thought was an important problem area. The U. S. Department of Agriculture did not seem anxious to support AI research.

Charlie Rosen, ever energetic, had by the mid-1970s developed a strong project in industrial robotics within the AI Lab. This project served as a model for many similar robotics research efforts initiated later in university and industrial research laboratories. Our work in industrial robotics also provided us with a focus for the research we were pursuing in machine vision [175, 193, 216, 223, 234, 262, 274]. Obtaining a Unimate robot arm, some small computers, and vision devices, Charlie assembled a crew of about a dozen people. Whereas most of the rest of us were quite content to make intellectual contributions to AI, Charlie really wanted to exert a tangible influence on industrial practice [66, 133, 174]. He received funds, advice, and problem descriptions from several industrial "affiliates" (such as General Motors and Westinghouse) that contributed financial support to the project. Charlie was convinced that some of the trainable pattern-recognition techniques of our early history could be employed in a simple but useful robot vision system for recognizing industrial parts. Dick Duda worked out a trainable decision-tree procedure, while Gerry Agin developed the "binary" machine-vision algorithms that were embodied in what came to be called the "SRI Vision Module" [103, 193]. In 1978, Charlie left SRI to found the Machine Intelligence Corporation (MIC), which would manufacture and market vision modules (and, later, robot systems). David Nitzan thereafter assumed the directorship of SRI's industrial-automation program.

The AI Lab had by now grown to over 40 people and needed more managerial help. Peter Hart became an associate director early in 1976, and we subsequently organized ourselves into several major program areas. Charlie Rosen (and then David Nitzan) headed the "Industrial-Automation Program," Gary Hendrix became the leader of the "Natural-Language Program," and Marty Tenenbaum took over the "Machine Vision Program." In December 1976, Bert left the lab to take on another position at SRI, and Peter Hart took over as director. Not too long after that, Earl Sacerdoti became associate director to help with managerial matters.

Peter's years as director of the AI Center (as it was renamed) were not jostled by quite as many severe external crises as were Bert's. We either had in hand or would obtain substantial support for all of our research programs. New ARPA projects in natural language, KLAUS, and TEAM, were launched. The TEAM (Transportable English Access Medium) project had an applications-oriented goal—the development of easily transportable systems for natural-language access to data bases [254, 263, 279, 293]. The KLAUS (Knowledge-Learning and -Using System) project allowed the natural-language people (finally) to pursue more basic research [230, 257, 268, 270, 286]. It had as its goal the development of reasoning techniques for (1) interacting with users in English about programs and data, and (2) accepting instructions and absorbing new knowledge stated in English. New projects supported by ONR and by AFOSR allowed us to begin work on distributed AI systems [232, 294], as well as on new planning systems called SPOT and SIPE [245, 258, 266]. PROSPECTOR, which became one of the most extensive expert systems ever developed, proved its efficacy by predicting the existence of a hitherto unknown molybdenum deposit in the state of Washington. SRI was chosen as the site of the ARPA image-understanding testbed—a facility that integrated much of the vision work being supported by ARPA [277, 298, 299, 300, 311]. The industrial-automation program attracted over 30 industrial affiliates and played a major role in helping industry enter the robotics age. The AI Center grew to over 50 researchers.

By 1979, there were hints amidst this success of big changes in the offing. Gary and Earl were becoming restless. Keeping a close eye on developments at MIC, they were wondering whether it might not be exciting to join a small company that seemed to be poised for explosive growth. In 1980, Earl did leave to join MIC and, in 1981, Gary left to found Symantec, a company that would market natural-language data base systems. MIC had by then attracted some other members of the Center's industrial-automation program, and Ann Robinson and Norm Haas later joined Gary at Symantec.

Peter himself began to muse about what it would be like to have an AI laboratory in which researchers could concentrate more on actual research and less on the art of interacting with government sponsors. The idea of some form of corporate support for AI research seemed tempting. In 1979 Schlumberger began talking to Peter about establishing a well-funded, well-equipped AI research laboratory. Might there be some people in the SRI AI Center who would move to Schlumberger to form the nucleus of such a new lab? Finally, Peter agreed to leave SRI and set up a new AI laboratory within Fairchild (a Schlumberger subsidiary in Palo Alto). Dick Duda, Marty Tenenbaum, and Harry Barrow decided to accompany him. Bad news! Most of the key people who had seen us through the tumultuous '70s—Raphael, Hart, Duda, Hendrix, Sacerdoti, Tenenbaum, Barrow, Rosen—were now either gone or on their way out.

While Peter was thinking about leaving SRI, he and I talked at length about who would run the AI Center and about the leadership problem in general. I had decided not to leave and was therefore a candidate—the only candidate—to take over the lab. Trapped! In a job I didn't really want. Oh well, it might be fun—and it would give me a chance to rebuild the lab in the direction in which I thought it should go.

Peter left in the summer of 1980. The biggest immediate impact was on our Machine Vision Program because Marty Tenenbaum and Harry Barrow left with him; they were soon followed in fact, by three other "visionaries," Andy Witkin, Lynn Quam, and Dave Kashtan. (Quam and Kashtan subsequently returned.) Marty Fischler, who came to the AI Center in 1977, took over our Machine Vision Program and has rebuilt it around a fine group of over a dozen researchers who are investigating a wide range of topics in machine vision [213, 223, 261, 271, 276, 280, 306].

I found that I enjoyed my role of "King Arthur" to the finest group of AI researchers anywhere. In some ways, we have been able to revive some of the spirit of John McCarthy's old Stanford AI Lab. Paul Martin and Dave Wilkins brought along some of its traditions when they joined us. I have insisted, more strongly than my predecessors had wanted to, that we pursue mainly basic research on fundamental AI problems; others can do (and are doing) applications. We moved into the personal-workstation age by obtaining several Symbolics 3600 Lisp Machines. We spun off the group that was concentrating on industrial robot applications; it could exist better on its own as a separate laboratory.

I have been aided in running the lab by some very capable people who quickly learned a great deal about research management. Daniel Sagalowicz became my assistant director and made everything function smoothly and efficiently. About a year after Gary Hendrix left, Barbara Grosz took over the Natural-Language Program. We formed a new program area under Tom Garvey, called AI Technology, to pursue applications of AI to "situation assessment" [307] and other military problems. Several new people joined us, including Stan Rosenschein, John Lowrance, Mark Stickel, C. Ray Perrault, David Warren, Fernando Pereira, Michael Georgeff, and Lauri Karttunen. But new companies continued to attract some of our staff away: David Warren, Rene Reboh, and finally (alas!) Daniel Sagalowicz.

Besides the Technical Notes in the present collection, the SRI AI Center continues to generate a high output of publications—in the AI Journal, in books, and in AAAI, ACL, and IJCAI proceedings. (Fifteen papers were presented at the 1983 IJCAI meeting in Karlsruhe.) The very strong support for basic research in AI provided by ARPA and other government agencies resulted in a number of new achievements in commonsense and evidential reasoning; in reasoning about the beliefs and knowledge of other agents; in new insights about the semantic content of natural-language utterances; in more powerful grammars, parsing strategies, and semantic translation schemes for natural language; in speech-act-based strategies for English-sentence generation; in vision; in planning and in program synthesis.

A major development of 1983 was the establishment at Stanford University of the Center for the Study of Language and Information (CSLI) funded by the System Development Foundation. CSLI will provide a focus for basic research on computer and human languages, and on their interrelationship. Participants include people from the Stanford departments of philosophy, linguistics, and computer science, as well as from local research organizations. About fifteen people from the SRI AI Center will work half-time or more on CSLI activities. Barbara Grosz played a major role in organizing CSLI and in writing the proposal that led to its funding. (After CSLI was established, however, Barbara decided to give up her role as director of the Natural-Language Program to devote full time to her research interests.) In my opinion, work at CSLI will touch the essence of what AI is all about and will thus have an extremely important effect on the SRI AI Center. It almost certainly will strengthen our orientation toward basic research and will put us in contact with the resource we need most—bright, innovative students.

So, here we are, still trying to keep to an internally conceived course while the sponsoring winds now and then compel a tack. By and large, we have been successful at pursuing what we've thought best to pursue; we're confident that eventually we'll even get back to work on intelligent robots. Another unfinished task for us is to demonstrate convincingly to the rest of the AI world the significance of logic as a basis for knowledge representation and reasoning.

Anyway, all things considered, I'm rather happy with the way things have turned out and am thinking about writing another book.

ACKNOWLEDGMENTS

Many of the persons named in this retrospective narrative helped me recall our history more accurately than I could possibly have done alone. I also want to thank all the other people at SRI, both researchers and managers, who have provided such a fertile and supportive environment for our research. We would not be telling this story had we not had farsighted and understanding sponsors. Bill Huber of the Army Signal Corps and Marshall Yovits of ONR helped launch us. Fred Dion and John Sammon of the Rome Air Development Center, and Richard Wilcox of ONR sponsored much of our theoretical work on pattern recognition. There would not have been an AI Center at all (in fact, not even AI itself!) without ARPA. A succession of ARPA-IPTO directors—Ivan Sutherland, Bob Taylor, Larry Roberts, J. C. R. Licklider, Dave Russell, and Bob Kahn—has furnished most of our support. Bernie Chern at NSF and several large corporations funded our industrial-robotics work. NSF and the U. S. Geological Survey sponsored our development of PROSPECTOR. Several projects have been funded by Kent Curtis and his colleagues at NSF. Charles Smith of the System Development Foundation deserves special mention for his role in providing support for CSLI. Bill Price of AFOSR was another important sponsor. The "grand old man" of all AI sponsors, of course, is, Marvin Denicoff of ONR. Marvin has been the patron of some of the most exciting AI research and of many books.

References

The following SRI Artificial Intelligence Center technical notes are cited in the text:

[4] Green, C., " Application of Theorem Proving to Problem Solving." (March 1969).

[8] Green, C., "The Application of Theorem Proving to Question-Answering Systems." (May 1969).

[14] Fikes, R., "REF:ARF: A System for Solving Problems Stated as Procedures." (September 1969).

[15] Garvey, T. D., and R. Kling, "User's Guide to the QA3.5 Question-Answering System." (October 1969).

[17] Fennema, C., and C. R. Brice, "Scene Analysis of Pictures Using Regions." (November 1969).

[19] Munson, J., "The SRI Intelligent Automaton Program." (January 1970).

[20] Duda, R. O., and P. E. Hart, "Experiments in Scene Analysis." (January 1970).

[26] Waldinger, R. J., "Robot and State Variable." (April 1970).

[34] Manna, Z., and R. J. Waldinger, "Towards Automatic Program Synthesis." (July 1970).

[36] Duda, R. O., and P. E. Hart, "A Generalized Hough Transformation for Detecting Lines in Pictures." (August 1970).

[39] Rosen, C. A., "An Experimental Mobile Automaton." (July 1970).

[40] Nilsson, N. J., "A Mobile Automaton: An Application of Artificial Intelligence Techniques." (March 1969).

[41] Coles, L. S., "An Experiment in Robot Tool Making." (October 1970).

[42] Rulifson, J. A., J. F. Derksen, and R. J. Waldinger, "QA4 Working Paper." (October 1970).

[43] Nilsson, N. J., and R. Fikes, "STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving." (May 1971).

[48] Rulifson, J. F., J. A. Derksen, and R. J. Waldinger, "A Problem Solving Language." (November 1970).

[50] Rulifson, J. A., "Preliminary Specification of the QA4 Language." (April 1970).

[52] Manna, Z., and R. J. Waldinger, "On Program Synthesis and Program Verification." (November 1970).

[60] Rulifson, J. F., "QA4 Programming Concepts." (August 1971).

[64] Raphael, B., "Robot Research at Stanford Research." (February 1972).

[65] Derksen, J. A., J. F. Rulifson, and R. J. Waldinger, "QA4: A Language for Robot Planning, Among Other Things." (May 1972).

[66] Rosen, C. A., "Robots, Productivity, and Quality." (May 1972).

[70] Fikes, R., P. E. Hart, and N. J. Nilsson , "Learning and Executing Generalized Robot Plans." (July 1972).

[73] Rulifson, J. F., J. A. Derksen, and R. J. Waldinger, "QA4: A Procedural Calculus for Intuitive Reasoning." (October 1972).

[78] Sacerdoti, E. D., "Planning in a Hierarchy of Abstraction." (June 1973).

[81] Reboh, R., and E. D. Sacerdoti, "Preliminary QLISP Manual." (August 1973).

[84] Tenenbaum, J. M., "On Locating Objects by their Distinguishing Features in Multisensory Images." (September 1973).

[86] Waldinger, R. J., and K. N. Levitt, "Reasoning About Programs." (October 1973).

[87] Tenenbaum, J. M., et. al., "An Interactive Facility for Scene Analysis Research." (January 1974).

[91] Walker, D. E., "The SRI Speech Understanding System." (April 1974).

[92] Paxton, W., "A Best-First Parser." (April 1974).

[94] Nilsson, N. J., "Plan for a Computer-Based Consultant System." (May 1974).

[95] Tenenbaum, J. M., et. al., "ISIS: An Interactive Facility for Scene Analysis Research." (June 1974).

[97] Robinson, J. J., "Performance Grammars." (April 1974).

[98] Manna, Z., and R. J. Waldinger, "Knowledge and Reasoning in Program Synthesis." (November 1974).

[99] Hart, P. E., "Progress on Computer-Based Consultant." (January 1975).

[101] Sacerdoti, E. D., "The Nonlinear Nature of Plans." (January 1975).

[103] Agin, G. J., "An Experimental Vision System for Industrial Application." (June 1975).

[105] Hendrix, G. G., "Expanding the Utility of Semantic Networks through Partitioning." (June 1975).

[107] Waldinger, R. J., "Achieving Several Goals Simultaneously." (July 1975).

[108] Barrow, H. G., and J. M. Tenenbaum, "Representation and Use of Knowledge in Vision." (July 1975).

[109] Sacerdoti, E. D., "A Structure for Plans and Behavior." (August 1975).

[110] Walker, D. E., "Progress in Speech Understanding Research at SRI." (August 1975).

[111] Paxton, W., and J. J. Robinson, "System Integration and Control in a Speech Understanding System." (September 1975).

[112] Robinson, J. J., "A Tuneable Performance Grammar." (September 1975).

[117] Garvey, T. D., "Perceptual Strategies for Purposive Vision." (September 1976).

[118] Wilber, M., "A QLISP Reference Manual." (March 1976).

[120] Sacerdoti, E. D., "QLISP: A Language for the Interactive Development of Complex Systems." (March 1976).

[121] Barrow, H. G., and J. M. Tenenbaum, "MSYS: A System for Reasoning About Scenes." (April 1976).

[123] Tenenbaum, J. M., and H. G. Barrow, "Experiments in Interpretation-Guided Segmentation." (March 1976).

[124] Duda, R. O., P. E. Hart, and N. J. Nilsson, "Subjective Bayesian Methods for Rule-Based Inference Systems." (January 1976).

[127] Garvey, T. D., and J. M. Tenenbaum, "Application of Interactive Scene Analysis Techniques to Cartography." (September 1976).

[132] Manna, Z., and R. J. Waldinger, "Is 'Sometime' Sometimes Better Than 'Always?' Intermittent Assertions in Proving Program Correctness." (June 1976).

[133] Nitzan, D., and C. A. Rosen, "Programmable Industrial Automation." (July 1976).

[134] Paxton, W. H., "Experiments in Speech Understanding System Control." (August 1976).

[135] Hendrix, G. G., "LIFER: A Natural Language Interface Facility." (December 1976).

[136] Duda, R. O., "Semantic Network Representation in Rule Based Inference System." (January 1977).

[137] Barrow, H. G., T. D. Garvey, and J. Kremers, "Interactive Aid for Cartography and Photo Interpretation." (January 1977).

[138] Hendrix, G. G., "LIFER Manual: A Guide to Building Practical Natural Language Interfaces." (February 1977).

[140] Sacerdoti, E. D., "Language Access to Distributed Data with Error Recovery." (April 1977).

[141] Walker, D. E., "An Overview of Speech Understanding Research at SRI." (April 1977).

[142] Paxton, W., "A framework For Speech Understanding." (June 1977).

[145] Sagalowicz, D., "IDA." (June 1977).

[150] Grosz, B. J., "The Representation and Use of Focus in a System for Understanding Dialogs." (August 1977).

[151] Grosz, B. J., "The Representation and Use of Focus in Dialogue." (July 1977).

[152] Hendrix, G. G., E. D. Sacerdoti, D. Sagalowicz, and J. Slocum, "Developing a Natural Language Interface to Complex Data." (August 1977).

[155] Hart, P. E., and R. O. Duda, "PROSPECTOR-A Computer Based Consultation System for Mineral Exploration ." (October 1977).

[156] Waldinger, R., and Z. Manna, "Synthesis: Dreams --> Programs." (November 1977).

[157] Barrow, H. G., and J. M. Tenenbaum, "Recovering Intrinsic Scene Characteristics from Images." (April 1978).

[158] Quam, L. H., "Road Tracking and Anomaly Detection in Aerial Imagery." (March 1979).

[163] Sacerdoti, E. D., "A Ladder User's Guide." (May 1978).

[164] Hendrix, G. G., "Encoding Knowledge in Partitioned Networks." (June 1978).

[174] Rosen, C. A., "Machine Vision and Robotics: Industrial Requirements." (November 1978).

[175] Tenenbaum, J. M., H. G. Barrow, and R. C. Bolles, "Prospects for Industrial Vision." (November 1978).

[177] Manna, Z., and R. J. Waldinger, "A Deductive Approach to Program Synthesis." (December 1978).

[185] Grosz, B. J., "Focusing and Description in Natural Language Dialogues." (April 1979).

[188] Grosz, B. J., "Utterance and Objective: Issues in Natural Language Communication ." (June 1979).

[191] Moore, R. C., "Reasoning About Knowledge and Action." (September 1979).

[193] Bolles, R. C., "Part Acquisition Using the SRI Vision Module." (September 1979).

[196] Tenenbaum, J. M., H. G. Barrow, R. C. Bolles, M. A. Fischler, and H. C. Wolf, "Map-Guided Interpretation of Remotely-Sensed Imagery." (September 1979).

[200] Fischler, M. A., J. A., Tenenbaum, and H. C. Wolf, "Detection of Roads and Linear Structures in Low-Resolution Aerial Imagery Using A Multisource Knowledge Integration Technique." (October 1979).

[205] Robinson, J. J., "DIAGRAM: A Grammar for Dialogues." (February 1980).

[210] Robinson, A. E., D. E. Appelt, B. J. Grosz, G. G. Hendrix, and J. J. Robinson, "Interpreting Natural-Language Utterances in Dialogs About Tasks." (March 1980).

[213] Fischler, M. A., and R. C. Bolles, "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography." (March 1980).

[216] Barnard, S. T., "Automated Inspection Using Gray-Scale Statistics." (March 1980).

[221] Tenenbaum, J. M., M. A. Fischler, and H. G. Barrow, "Scene Modeling: A Structural Basis For Image Description." (July 1980).

[222] Barrow, H. G., and J. M. Tenenbaum, "Reconstructing Smooth Surfaces from Partial, Noisy Information." (July 1980).

[223] Bolles, R. C., "Locating Partially Visible Objects: The Local Feature Focus Method." (August 1980).

[224] Sagalowicz, D., "A D-LADDER User's Guide." (September 1980).

[230] Hendrix, G. G., "KLAUS: A System for Managing Information and Computational Resources." (October 1980).

[232] Konolige, K. G., "A First-Order Formalization of Knowledge and Action for a Multiagent Planning System." (December 1980).

[234] Bolles, R. C., "Three-Dimensional Locating of Industrial Parts." (January 1981).

[244] Smith, G. B., "Detection of Rivers in Low-Resolution Aerial Imagery." (June 1981).

[245] Robinson, A. E., and D. E. Wilkins, "An Interactive Planning System." (June 1981).

[246] Manna, Z., and R. Waldinger, "Deductive Synthesis of the Unification Algorithm." (July 1981).

[254] Armar, A., B. J. Grosz, and D. Sagalowicz, "A TEAM User's Guide." (December 1981).

[257] Grosz, B. J., "Research on Natural-Language Processing at SRI." (November 1981).

[258] Wilkins, D. E., "Parallelism in Planning and Problem Solving." (January 1982).

[259] Appelt, D. E., "Planning Natural-Language Utterances to Satisfy Multiple Goals." (March 1982).

[260] Manna, Z., and R. J. Waldinger, "Special Relations in Program-Synthetic Deduction." (March 1982).

[261] Barnard, S. T., and M. A. Fischler, "Computational Stereo." (March 1982).

[262] Bolles, R. C., "Recognizing and Locating Partially Visible Objects: The Local Feature Focus Methods." (March 1982).

[263] Grosz, B. J., et. al., "TEAM: A Transportable Natural-Language System." (April 1982).

[266] Warren, D. H., "A View of the Fifth Generation and Its Impact." (August 1982).

[268] Stickel, M. E., "A Nonclausal Connection-Graph Resolution Theorem-Proving Program." (October 1982).

[270] Grosz, B. J., et. al., "DIALOGIC: A Core Natural-Language Processing System." (November 1982).

[271] Barnard, S. T., "Interpreting Perspective Images." (November 1982).

[274] Barnard, S. T., "Inspection of Printed Circuit Boards for Part Integrity." (January 1983).

[276] Fischler, M. A., and H. Wolf, "A General Approach to Machine Perception of Linear Structure in Image Data." (February 1983).

[277] Hanson, A., "The DARPA/DMA Image Understanding Testbed User's Manual." (February 1983).

[279] Grosz, B. J., "Team: A Transportable Natural-Language Interface System." (March 1983).

[280] Pentland, A., "Fractal Textures." (March 1983).

[286] Stickel, M. E., "Theory Resolution: Building in Nonequational Theories." (May 1983).

[292] Grosz, B. J., "Providing a Unified Account of Definite Noun Phrases in Discourse." (July 1983).

[293] Martin, P., et. al., "Transportability and Generality in a Natural Language Interface System." (August 1983).

[294] Konolige, K. G., "A Deduction Model of Belief." (August 1983).

[298] Laws, K., "The DARPA/DMA Image Understanding Testbed Programmar's Manual." (February 1983).

[299] Hanson, A., "The DARPA/DMA Image Understanding Testbed System Manager's Manual." (August 1983).

[300] Hanson, A., "Managing the IU Testbed Under Eunice VMS." (August 1983).

[306] Smith, G. B., "The Relationship between Image Irradiance and Surface Orientation." (September 1983).

[307] Lowrance, J., and T. D. Garvey, "Evidential Reasoning: An Implementation for Multisensor Integration." (September 1983).

[311] Hanson, A., "Overview of the Image Understanding Testbed." (October 1983).

[312] Appelt, D., "Planning English Referring Expressions." (October 1983).