Chapter 4
The Life and Times of a Successful SRI Laboratory: Artificial Intelligence and Robotics

Background

Throughout most of its life SRI has been, for better or worse, a collection of individual laboratories or centers, with each following its own path. Doing so has usually meant advancing the state of the laboratory's particular art or pursuing the opportunities offered by the R&D marketplace. On many occasions, SRI has risen to the challenges of cross-laboratory and cross-discipline efforts, thereby realizing one of its unique attributes. But, for the most part, SRI has been technically segregated, closely resembling the university research environment from whence it came.

In this section, we depart from our general theme of selected SRI project impacts on industry and society to describe how one laboratory, centered on a specific discipline, has fared at SRI. We take a longitudinal look at one of SRI's most successful and noteworthy labs, and how one project often led to another over time, sometimes deliberately, but often incidentally. This nearly 40-year perspective provides a comprehensive history of one lab's major projects. The discipline is one that is both modern and mercurial, artificial intelligence (AI), and the SRI laboratory is the Artificial Intelligence Center (AIC).

AI is one of those terms that has gone in and out of favor, mainly depending on whom you are talking to and when. But whether it is called artificial or machine intelligence or something else, giving a computer system the ability to perform acts that at least seem intelligent will continue to be a worthwhile and compelling area of research. For over 40 years, SRI has been home to one of the world's premier AI research centers, a center that has consistently helped define and invent the changing spectrum of science and technologies that make up this evolving field.

While machines are infinitely better than humans at some tasks, giving a computer even a few of the more modest capabilities humans take for granted, such as facial recognition, common sense, the flexible interpretation of a task, and language, has proved extremely difficult. Progress in achieving those ends will almost certainly continue to be piecemeal and evolutionary. Here, we give a brief account of how SRI approached a number of facets of AI. A more detailed account of the genesis of the AIC is given in a 1984 history by one of its former directors, Nils Nilsson.[A]

Leadership Responsibilities

As with most laboratories at SRI, the AIC divides responsibilities for winning and for carrying out projects. The Center is divided into program areas staffed by research professionals whose fields of work are aligned with the particular program. The Program Manager is responsible for making sure:

- That the field of research being pursued is not only current but advances the state of the art.
- That the people in the program area are capable of advancing the technology.
- That the members of the program are collectively billing projects at about 80% of their available time.
- That, when asked, will collaborate with other programs or laboratories on projects where program expertise is applicable.

The Director of the Center is responsible for its technical and financial well being and, most of all, for its continued existence. That means starting new programs when the opportunity presents itself, stopping programs before they are moribund, assuring the quality of the staff and their work, and otherwise defining the culture and morale of the group. This pattern pretty well represents all the labs and centers across SRI, but the AIC enjoys a higher than average technical reputation and an enviable longevity. At a staff level of approximately 80, the AIC is a healthy size, big enough to smooth out the vagaries of individual contract transitions and small enough that just two levels of management are needed and that a single area of interest can be pursued. In the discussion that follows, we relate how a group of this size has made its impressive way in the helter-skelter world of contract research.

How AI Began at SRI—Learning Machines

The possible utility of this nontraditional engineering approach to problem solving came to Charlie Rosen around 1959 when he was head of SRI's Applied Physics Laboratory. For an early electron-beam machining technique, SRI had conceived of an approach to manufacturing electronic circuits by using huge arrays of field-emission triodes.[1] Because each triode was about a micron (10^{-6} m) in size and because there were thousands of them, it became clear that, given the fabrication methods of the day, not all of the triodes would work. Given the high numbers of these emitters, was there a way to make them self-organize to work around the imperfect ones that were sure to occur?

A staff member from Cornell Aeronautical Research, Frank Rosenblatt, a psychologist by training, visited SRI about that time and described to Rosen and Ted Brain the principles of a new concept he called *perceptrons*. These were elementary "learning machines" whose architecture and logic units grossly imitated the brain's neurons and their functions. He claimed that his perceptron systems could learn, by being presented with many samples over and over again, to recognize many different patterns automatically; the system's internal connections were thus being adapted or trained until the system "learned" to identify each pattern. Perceptrons (and, at Stanford, Bernard Widrow's analogous Madeline systems) were pioneering systems that formed the basis of an important class of "intelligent" machines that came to be called *neural networks*. Rosenblatt visited SRI seeking help in developing inexpensive logical elements (threshold logic units) needed for the proposed construction of a large perceptron. It was conjectured that, to do anything worthwhile, a very large parallel-operating machine would be necessary. (At that time, digital computers with the power and speed required to simulate such a large system were not available.)

These new ideas were immensely stimulating. Supported by Division head Jerre Noe, Rosen and Brain began promoting this project and within months had secured initial funding from the Office of Naval Research (ONR). Active staff recruiting ensued, resulting in one of the first groups anywhere working in AI. The group included Nils Nilsson, Dick Duda, John Munson, George Forsen, David Hall, and Dick Singleton—some of whom went on to become well-known in this field. Soon, with additional funding from the U.S. Army Signal Corps, the U.S. Air Force, the U.S. Navy, and others, SRI became the largest research group in the world working on perceptron-like systems. A highlight of this research was the delivery to the Signal Corps in the mid-1960s of a system called MINOS II,[2] a large self-contained learning machine developed and built at SRI (see Figure 4-1). The logic units, which were based on ferrite multi-aperture cores, invented by SRI's Hew Crane, served as variable analog weights that effectively changed the connectivity between logic units during training.

MINOS II also included a novel optical preprocessor, consisting of 1024 lenses that could replicate as many optical images from a TV screen or projector. The 1024 images were then sampled in parallel using masks to extract one important feature from each image. These features were then combined into 100 elements that delivered a state of +1, -1, or 0 depending on whether the signal exceeded a threshold or not. The preprocessor output, then, made up the input to the trainable part of the perceptron system. That part of the system used a matrix of 6600 magnetic weights to identify objects about which it had been trained.

These early "learning machines" were essentially used as pattern-recognition systems. They were applied to the recognition of military targets shown, for example, in aerial photographs, and to the classification and recognition of the salient features in time-varying signals such as radar, sonar, spoken words and phrases, hand-printed isolated characters, faces, and the like.

[1] This is the same technology of field-emitting devices described in Chapter 7.

[2] In this particular setting we have chosen not to define the acronyms associated with individual software/hardware systems. They are simply too numerous and in many cases would be meaningful only to those in the AI community.

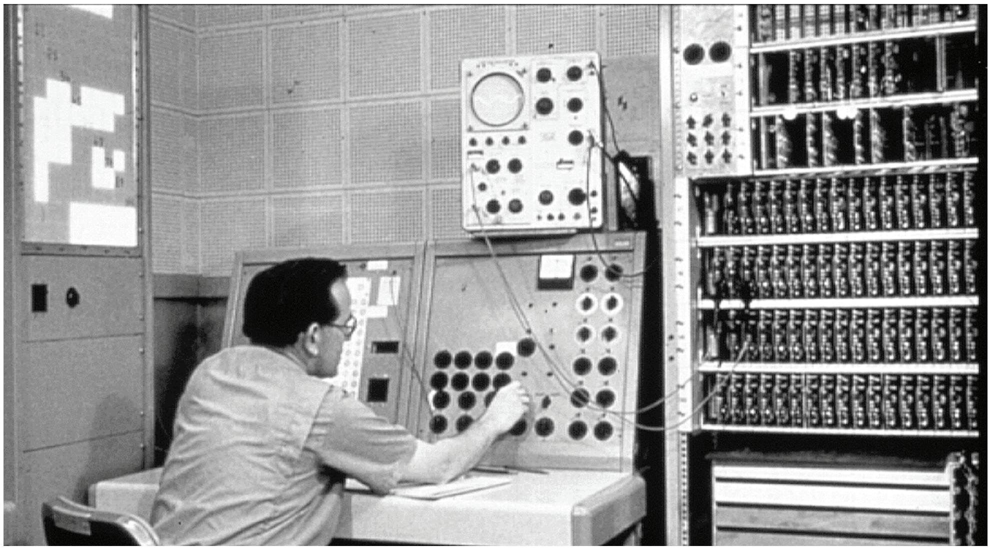

Figure 4-1. MINOS II preprocessor and Ted Brain (circa 1964).

in November 1968. Papers describing its capabilities in context-based, hand-written code recognition and other areas were published.[3] However, perceptron-like work took a backseat to emerging digital processors, and sponsorship simply dried up. Nonetheless, continuity for an operating group such as the AIC is essential, and a contract research lab demands no small amount of anticipation of such changes in funding.

The funding for learning machine research began to slow drastically in the late 1960s, partially because influential MIT researchers claimed they could prove that these neural nets were "a blind alley," and that mainstream AI work should be based on digital computing. More recently, their "proof" has been shown to be valid only for the earliest of the perceptron architectures, and neural net research has regained some interest. At that time, however, SRI researchers and others in the field were unable to overcome an important barrier to progress for such machines: the need to automatically and adaptively train the elaborate architectures needed for more complex tasks, particularly those with more than two layers of logic units.

During the development of the MINOS family of analog-based learning machines, a curious but meaningful transition occurred. To better predict what MINOS might do under a given set of teaching patterns, simulations of the analog system were written on emerging general-purpose digital computers. This was done in part because of the inordinate time it took to set up and operate the MINOS machines. But early on it was recognized that the simulation was running faster than the machine it was simulating. That was just one of the many indicators of the growing power of digital machines. The last of the MINOS series, MINOS III, thus became partly digital. The overall MINOS III package, consisting of an SDS 910 digital computer, the MINOS II analog learning machine, and the 1024-lens preprocessor, was delivered to the Signal Corps

Entering into Robotics

So, in the mid-1960s, Rosen and Nilsson initiated an in-depth, internal study at SRI to define a new AI program to replace the fading perceptron work. After 3 months, the concept of an intelligent mobile robot system was developed and formed the basis for a new SRI program. This "intelligent" machine, later named *Shakey,* was to serve as an R&D test-bed for key AI subsystems; namely, machine vision and scene analysis, natural language, theorem proving, planning and problem solving, and—because the machine was a mobile automaton—navigation and obstacle avoidance (see Figure 4-2). SRI promoted the program at the U.S. Department of Defense (DoD) for 18 months and finally succeeded in obtaining sponsorship from the Advanced Research Projects Agency (ARPA).

[3] John Munson, Richard Duda, and Peter Hart, "Experiments in the Recognition of Hand-Printed Text, Parts I and II," ACM Fall Joint Computer Conference, December 1968, San Francisco. The MINOS III system was "trained" to read handwritten FORTRAN code and, with the introduction of context-based reasoning, it did surprisingly well. *SRI Journal*, No. 19, March 1968, stated that when averaged over many writers, it could recognize about 85% percent of instructions correctly. If trained on a single writer, it had accuracy as high as 97%, even though the printing was untutored and constrained only by a standard code sheet.

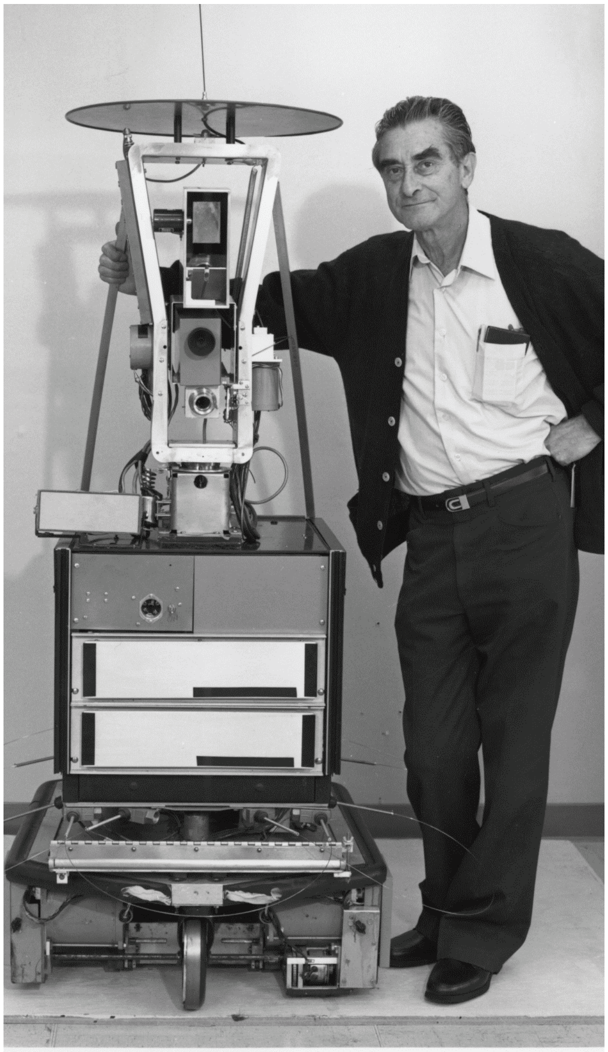

Figure 4-2. Shakey the Robot and Charlie Rosen.

Shakey and its environment proved a fertile ground for new concepts in AI. It had obvious needs for things such as sensors and navigation approaches and for not-so-obvious elements that would give it a measure of autonomy. Sophisticated planning systems, capable of entertaining hierarchical goals (performing a mission expressed through hierarchical sets of subtasks and in the process not destroying itself) were also required.

The Center took a sophisticated approach to the capabilities needed by a wandering robot. For example, advanced problem solvers were built that employed some of the world's first automatic theorem provers. Fortuitously, these had been developed in 1963-64 using a new technology in logical systems brought to the AIC from Stanford by Cordell Green. The technology influenced the design of STRIPS, the SRI problem solver used in Shakey's navigational system, and in hierarchical planning systems like NOAH[4] that were intended to help Shakey consider and successfully satisfy multiple goals. This work continued to evolve, becoming a logical basis for both reasoning problems and planning systems and with more than 25 professionals participating in the program for more than 10 years. Perhaps the most significant contribution from the Shakey era was a search algorithm that became the basis for the huge literature on optimal search, including present route finding systems such as MapQuest. All this helped establish SRI as a world-class center in AI research.

In the early 1970s, Rosen resigned as head of the AIC and turned over supervision to Bert Raphael, who was later followed by Peter Hart, Nils Nilsson (see Figure 4-3), Stan Rosenschein and Ray Perrault. Rosen wished to see some pragmatic results from AI and initiated a new SRI program aimed at transferring some of the technology learned in the Shakey work to industrial automation. He believed that recently introduced industrial robots could be greatly improved by incorporating sensors, computer controls, and new training methods, and that these machines could then act as "smart" material-handling systems and perform simple assembly tasks and inspection. After a year's promotional activity, a joint National Science Foundation and SRI Industrial Affiliates program was started to develop industrial robot systems that would increase productivity and quality in manufacturing. This program grew to include over 25 major industrial firms, including General Electric, Westinghouse, General Motors, Ford Motor Company, Unimate, 3M, Digital Equipment, Lockheed, and others. For a decade SRI's program in robotics was a model for cooperative research by government, universities, and industry. Furthermore, SRI's program was enormously influential as a source of technology in this field and as a convenient center where key people in industry could meet three to four times a year and share experiences to their mutual advantage. Many industrial research groups had their start by participating in these exchanges and by viewing the periodic laboratory demonstrations at SRI and at the various participating companies.

[4] NOAH was an early hierarchical planning system that derived from work on a Stanford thesis on robot problem solving that Earl Sacerdoti had done while at SRI. Unfortunately, NOAH was developed about the time ARPA was losing its interests in robotics.

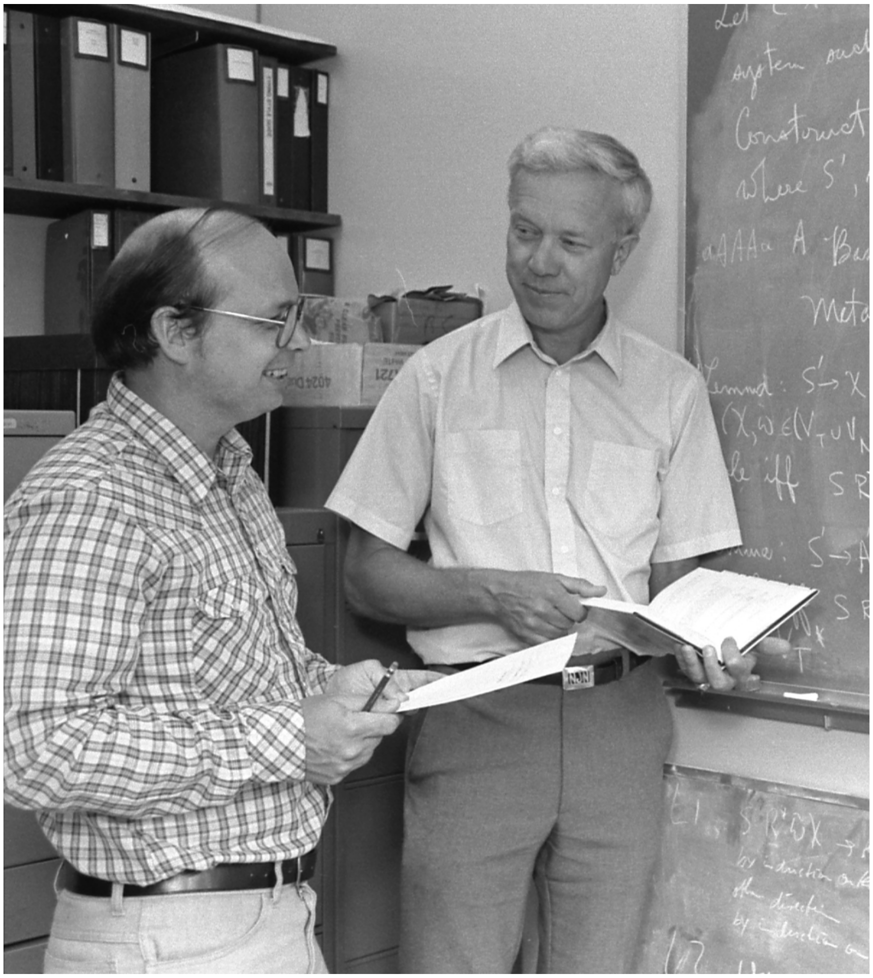

Figure 4-3. AIC Director Nils Nilsson (right) and Daniel Sagalowicz.

Because much of the SRI work done in the early days was not patented but, in fact, widely reported, subsequent research around the country on unmanned vehicles still incorporates ideas and devices developed in the work on Shakey. In its November 20, 1970 issue, *Life Magazine* carried an article by Brad Darrach whose headline, "Meet Shakey, the first electronic person," provoked much teasing. The article in the July 1971 issue of *Fortune* was more circumspect, saying "The Stanford machine, a similar one at M.I.T., and the mobile robot Shakey…are actually all experiments in solving problems through the techniques of scene analysis."[5]

Rosen's group developed the vision system first used with Unimate industrial robots and developed several end-effectors or "hands" equipped with tactile, force, and torque sensors; one hand was equipped with a proximity sensor. As part of a robot for package handling, a two-sweep method for reading any-orientation bar code was developed by John Munson and patented by SRI, but SRI never pursued infringements on that patent. Later, package-handling machines led to the long series of mail-handling research projects pursued by the then separate Robotics and Mechanical Engineering Laboratories.

Another reason for adopting a more commercial bent in robots at SRI was that in the early 1970s ARPA, the prime sponsor of much of AIC's robotics work, went through one of its periodic "relevance" transitions; desiring more practical solutions from its programs, it discontinued robotics research. Consequently, it was time to redefine the Center's future. Searching out new technology and who can build it, understanding the changing intent of a range of sponsors, and relying a bit on serendipity are common approaches SRI research labs take to stay in business. The best research labs make sure that new, trial research areas not only have open and challenging vistas, but also lend themselves to solving important client problems. Only then will they obtain significant levels of funding.

Further Marketplace Adaptation—Moving Toward Language- and Knowledge-Based Systems

Thus, the resourceful AIC leadership and staff adapted by entering two new areas of work: one concerned the creation of programs that captured, then supplemented human knowledge; and the other sought to get

[5] Shakey now enjoys a permanent home at the Computer History Museum in Mountain View CA and was recently "elected" into the Carnegie Mellon University's Robot Hall of Fame.

machines to understand natural human language. These initiatives resulted in SRI's first participation in:

- The machine embodiment of knowledge or expert systems.
- Automatic speech recognition and understanding.
- Text-based natural language understanding.
- Understanding the content of photographic images.

The early to mid-1970s saw the blossoming of these new fields, rich with unknowns and research opportunity. The work on expert systems grew out of an AIC attempt to finesse ARPA's departure from robotic systems by shifting to the design of software that might ultimately control robots. Thus was born the notion of a computer-based consultant (CBC) that would serve as an "expert" on assembly or disassembly of electromechanical equipment. The CBC could be used either as a training aid or real-time helpmate for military personnel or in actually programming a robot. ARPA bought into the idea...at least for a while.

The AIC's work in natural language understanding grew out of a new program that had been initiated at DARPA in 1972 by a former AIC researcher, Cordell Green, then on active duty there. ("Defense" was added to ARPA in 1972 but its mission didn't change appreciably and the two terms are interchangeable here.) The program was called Speech Understanding Research (SUR). Don Walker led the SRI work, which attracted a host of outstanding people in language, including Barbara Grosz, Jane Robinson, Bill Paxton, Gary Hendrix, and later Ray Perrault. Though DARPA would terminate the speech program by 1976, it did continue work in areas such as natural-language interfaces with databases. The limits of natural language understanding continued to be extended with the arrival at AIC of Bob Moore and Doug Appelt with their notions of *reasoning* about knowledge.

Concurrent with these two new programs came another initiative from the Center's Marty Tenenbaum and Tom Garvey to help automate photographic analysis. They saw the job of photo interpretation as one ripe for automation and were anxious to apply new ideas in vision research. They began to do so in 1976, but the problems were sufficiently complicated that they are still being addressed today. DARPA was interested in this work and later created a long-term program in "image understanding," We return to this program later.

Expert Systems

DARPA did not renew the CBC project, and Government-sponsored robotics work at SRI was truly dead. But the leaning of CBC toward the encoding of human expertise and a new attempt at Stanford to create an expert system for advising on bacterial infections (MYCIN) led Peter Hart and Dick Duda to wonder about other applications. A secondary goal was to try to avoid the funding vacillations the Center had experienced with DARPA. This new area of knowledge-based systems produced one of the world's first examples of an expert system, with a system that provided consultation services on mineral deposits. Naturally, it was called PROSPECTOR. Using production rules and semantic networks from the natural language efforts, it was completed in 1977. PROSPECTOR represented a significant advance in the state of the art of expert systems because of the data structures used, the techniques for updating probabilities, and the extensive geological application models developed. The program housed information on more than 20 types of ore deposits. A similar number of experts were interviewed for approximately 50 hours each. One of the types of deposits was molybdenum and, using this segment of the program, a large extension to a previously sampled but unmined bed of molybdenum was predicted and found on Mt. Tolman in the state of Washington. Unfortunately, the bed lies under an area used for tailings from a smelting process and thus cannot be exploited economically. Major participants in PROSPECTOR were Hart, Duda, John Gaschnig, Rene Reboh, Nilsson, and Kurt Konolige.

Though not many complete software systems have emerged as purely expert systems, the evidence/rules concept they are based on has become an important tool of computer science and is embedded in many programs. One, a heavily used expert system that emerged in another applied AI group at SRI, was the Automated Air-Load Planning System (AALPS).[6] From their extensive work with the

[6] AALPS was created in SRI's Information Telecommunications and Automation Division around 1984, and the prototype version was completed 2 years later. The project leader there was Debra Anderson.

Army's 82nd Airborne Corps at Ft. Bragg, North Carolina, SRI staff recognized the critical need for moving matériel rapidly. Using expert system techniques, SRI was able to capture in AALPS much of the knowledge the military airlift loadmasters used to load air cargo, including correct weight distribution and order of egress. Because AALPS also contained the size and weight characteristics for all airborne matériel and cargo aircraft configurations, the program was able to compute loading plans in seconds to minutes, whereas previously used manual approaches required days or weeks. Manifests created using AALPS eventually were accepted by the Army and the Air Force Military Airlift Command. AALPS saw extensive use in the Gulf War and is still used throughout DoD.

Not all expert systems are as successful as AALPS in correctly modeling known criteria or events, however. Many situations can occur where the input information is incomplete, inexact, and uncertain. To cope with this uncertainty, a new reasoning method, developed by John Lowrance and Tom Garvey, was invented to deal with more realistic evidence that doesn't fit nicely into the framework of the rule-based algorithms on which all expert systems were then built. The pioneering method, which was called evidential reasoning, drew on the various representational forms of uncertainty.[7] These were embodied in a program called GISTER that has been used by many other SRI expert system applications as well.

Another outgrowth of the AIC's expert system work was a slightly different problem orientation called "procedural reasoning." This method of imparting supplemental expertise in problem solving addressed problems in which a strong set of procedures or sequences had to be followed, often in real time (e.g., to meet safety concerns). To tackle this type of problem, Mike Georgeff created the Procedural Reasoning System, which had a role in the Space Shuttle and the paper describing it was recently honored by the AAAI, two decades after it was published in 1987.

Natural Language Understanding

The other technical initiatives begun in the early 1970s were research into both natural language and image understanding, research that proved to be enduring at SRI. The natural language work was first vectored toward database query and how such systems could be more easily and quickly built. Part of the desired capability was enabling domain-independent transport; that is, the ability to move the processing engine from one subject area to another with a different associated vocabulary. First, Hendrix developed LIFER, a system for English access to databases based on semantic grammar. This was followed by Hendrix's, Sacerdoti's, and Daniel Sagalowicz's LADDER, an English interface to a distributed database, under which began the development of DIALOGIC, a large grammar of English implemented in an augmented context-free grammar framework. DIALOGIC then became the basis for TEAM, another English interface with databases that provided a new and easier approach to tailoring interfaces to new domains.[8] The Knowledge Learning and Using System (KLAUS) and Parsing and Translation system (PATR) were used for parsing and generating natural language based on constraints using a method called unification. Developed by Stuart Shieber, Fernando Pereira, and Lauri Karttunen, they were the predecessors of the Core Language System (built at SRI's Cambridge Centre in England), of Gemini, the Air Traffic Information System (ATIS), and of several other systems worldwide.

In 1983, the AIC's natural language program and SRI joined a research consortium dedicated to exploring the fundamentals of language and the information it and machines shared. Joining such outside consortia has not been common at SRI, but since the Institute had an excellent reputation in this small and growing community, the collaboration with other local campuses was natural. The consortium included SRI, Stanford University, and the Xerox Palo Alto Research Center; the new organization was called the Center for the Study of Language and Information (CSLI). Its seed funding came from the required passing of

[7] Although the distinction between representational forms or models of uncertainty may be open to debate, the ones handled were normal probabilities (often Bayesian), the Dempster-Shafer model, and fuzzy logic.

[8] The main authors of DIALOGIC and TEAM were Barbara Grosz, Jane Robinson, Jerry Hobbs, Bob Moore, Paul Martin, and Fernando Pereira (see Figure 4-3 from the *SRI Journal*, 2(6), August 1982).

Figure 4-4. Jane Robinson, Barbara Grosz, and Bob Moore of the Natural Language Program.

funds to the public interest by the nonprofit to for-profit conversion of the System Development Corporation in 1968. The Center is located at Stanford, and with the three members contributing directly and sharing other, non-CSLI results, a serious research agenda began that continues over two decades later. The CSLI not only offered new opportunities for fundamental work but also attracted new talent to the AIC.

The AIC contributed important research to the CSLI areas of unification grammars and situated automata. Unification grammars are a type of formal specification of a language wherein grammatical units (words, phrases, sentences, etc.) are given complex formal descriptions such as sets of attribute-value pairs. They are called "unified" because they permit the representation of grammatical knowledge independent of the language-processing algorithm used. Situated automata do not constitute a natural language concept, but involve a method for defining the behavior of automated objects such as a robot. In situated automata, an agent is specified in declarative terms (e.g., performance norms or limits or protective conditions). This specification is then compiled to machine operations that exactly support the specification. The digital machine can then, for example, operate in a provably time-bounded way; an ability important in situations where it may endanger itself.

Another branch of natural language understanding work at SRI was text extraction, which searches normal newspapers, books, or compendia for information about a specific topic. In this case the opportunity was presented by DARPA in 1984 and the AIC was prepared to respond. The collegial approach that DARPA often took—that is, letting several different universities or research centers work collaboratively on different parts of a difficult problem—was modified in this case. The contracts in this program required a periodic "bake-off" or competitive demonstration of the performance of each contractor's systems. These were the so-called Message Understanding Conferences or MUCs. SRI participated in this program from its inception and created FASTUS, one of the most capable systems to date. FASTUS performs a semantic search through textual forms such as the *Wall Street Journal* and automatically completes a user-defined information template. Because of FASTUS's relevance to information search systems fostered by the Internet, SRI is now exploring the commercialization of the system.

In a related vein, the 1980s saw a return by DARPA to the problem of automatic speech recognition.[9] Computers were becoming increasingly capable of meeting the intrinsic real-time requirements of speech recognition, but it was also an opportunity to explore new algorithms. This second entry into speech would become more important to SRI and is discussed elsewhere (see Chapter 2). SRI's participation in this new program began in another laboratory, with the AIC participating because of its continuing interest in natural

[9] The first speech program at DARPA, SUR, had run its intended 5-year course by 1976 and had been terminated. This new program at DARPA was not a continuation of the earlier work.

language understanding. Not surprisingly, the structure and vocabulary of natural language systems can also play a potentially important part in speech recognition. For this reason, Gemini was built at SRI as a parsing and semantic interpretation system to supplement the already existing speech recognizer. In this way and others SRI has continued to advance the state of the art in natural language understanding for more than two decades.

SRI-DEVELOPED HUMAN-COMPUTER INTERACTION TOOLS

Available Capability	When
Pointing devices, predominantly the mouse	From about 1965
Automatic handwriting recognition	From about 1978
Automatic speech recognition	From about 1990
Natural language understanding for text and speech	From about 1977

Human-Computer Interaction

But clearly not all convergences of technology pan out, even when they appear made to order to do so. In early 1988, another research opportunity seemed at hand and it stemmed in part from unique capabilities (see table) that SRI had already developed—a set of computer interaction tools. Given these capabilities, it seemed a propitious time to create an environment in which many of these advances in information technology could be integrated to build better and more complete approaches for humans' interactions with their computers. Computers were gaining so much power that a goodly part of their processing resources could be used to improve the ability of the machines to understand what their users were trying to do. Arguably, the dominant mode for future human-machine interaction is speech, and SRI had one of the best laboratories in the world for automatic speech recognition.

Accordingly, an informal computer dialog laboratory was formed to bring some of these modalities together and to explore experiments in machine interaction. While the lab never met its promise, a few things did emerge. A semiconductor fabrication process control system with natural language input, called Shoptalk, became one of the first systems to automatically determine context as part of its natural language component. Another outcome was a closer relationship between speech recognition and natural language understanding, then housed in two separate laboratories. From this foray into human-machine interaction came a simple but powerful notion, arrived at experimentally: if a computer user is seeking an object such as a piece of information and sees a means of securing it on-screen, the easiest thing for humans to do is to point to it, whereas if the means isn't displayed, the easiest thing to do is to ask for it. This simple finding explains the importance of speech in human-machine interaction, particularly if such interaction is to be natural for the human and not simply convenient for the machine.

Most humans are quite flexible in how they approach a task, particularly one with which they are not familiar. Part of that flexibility is the way we can easily adapt to different communications modalities. That trait, then, should be reflected in the design of human-computer interactions. To match machine capability with different human modalities, the notion of software agents was suggested. Nevertheless, although a useful construct for such a functional communications inventory, it didn't catch on in the field for another 4 years or so.

Software Systems Employing Agents

Given the multiple modalities involved in human-machine interaction, some of which are complex for a machine to recognize, a specialized piece of software was created to handle each mode. To convey that modularity and independence, each module was called an agent. An agent could represent almost any capability or functionality such as speech recognition, email, handwriting recognition, or natural language understanding. In 1993, the AIC's Adam Cheyer created what is known as an Open Agent Architecture (OAA), a framework within which different functional agents can collaboratively or competitively vie for completing a task. Figure 4-5 shows an early family of agents that supported a computer user. The facilitator is but another agent that knows agent capabilities and can adjudicate which agent gets assigned which task. That assignment is flexible, however, in that if the preferred agent is unavailable, another may be eligible and assigned. This approach represents one of the most flexible and sophisticated approaches to the use of software agents, and it has formed the basis for dozens of applications at SRI in which human-computer interaction are important. These applications have ranged from controlling and processing video streams, spoken-language interfaces with simulators, to the control of semiautonomous robots using speech and gestures.

Nor are these agents confined to operation within a single machine. In our increasingly connected world, agents in one location need to negotiate with those in another location to carry out users' requests. One example is the smart refrigerator that was built at SRI. The "fridge" was programmed to have a certain complement of food inside. For tracking purposes, the food items were bar-coded so that by noting their comings and goings, the fridge was aware of what it did and didn't have. The agent in charge of that accounting, then, was free to communicate with an agent in the family car to notify the occupant that something was needed from the store. It was the agents, not the processors, that had a functional relationship that was intended to satisfy some high-level user need.

This flexible concept came to be used widely across SRI, both in the AIC and in other labs. As many as a dozen different programs with important human-machine interaction used OAA. Whether the term "agents" will survive in the evolution of software systems is open to question, but there is no doubt that the concept of a software entity, programmed to be aware of a context reflective of its owner's interests, will grow in some fashion under some rubric.

Image Understanding

One evolutionary line in the AIC's map of technology migration (shown toward the end of this chapter) is computer vision. Progress along that line requires a computer to be able to identify the objects that make up a given scene. "Image understanding" is the term used to define the ability of a machine not just to identify objects or conditions in two-dimensional (2D) or three-dimensional (3D) scenes, but even to deal with their roles, relationships, and activities. It means applying *a priori* and scene-derived knowledge to

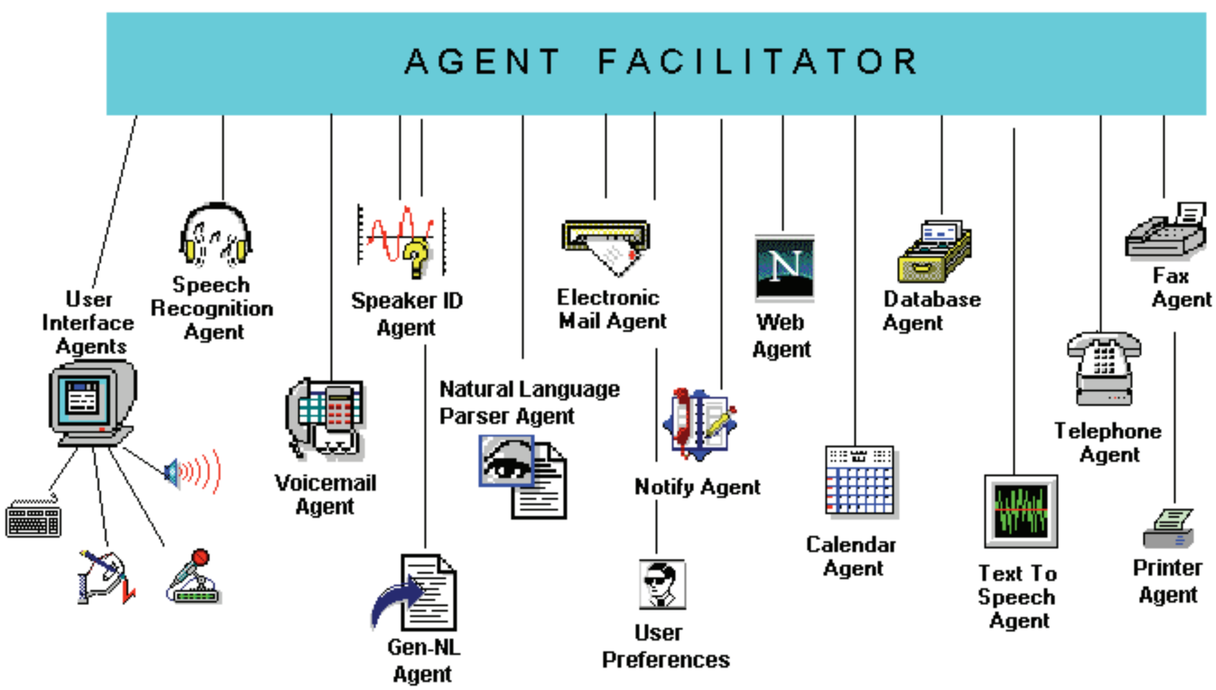

Figure 4-5. Organizational schematic of Open Agent Architecture.

automatically recognize objects of interest and what their presence implies. These are often very difficult tasks to perform and, like language, which humans take for granted, present enormous challenges to machines. A good example of image understanding is automatically finding and following a road in a photograph and then computing its exact location to update or upgrade a map; that is, making an approximate position on a map into an exact one. Automatically determining the location and dimensions of a building in a photograph is another. One of the earliest investigations into the field of image understanding was Garvey's Stanford Ph.D. thesis where he used the contextual relationships between objects to help identify and locate them. Other Center members involved early on were Marty Tenenbaum and Harry Barrow.

Somewhat akin to image understanding is the 3D construction of a scene from extracted information. The art of scene synthesis using a variety of available inputs such as digital maps, photographs, and objects extracted from photos or movies had begun in earnest at SRI by 1980. The chief architects of this work were Marty Fischler and Lynn Quam. The work began with programs that first employed the stereo pairs that were being derived from analog photographs and then the emerging digital terrain data sets. Programs were written that could form visual representations of terrain and its landmarks, and then view that scene from an arbitrary point. These were given names like IMAGECALC and TERRAINCALC, and they were photogrammically rigorous; that is, they derived all aspects of the computed scene from a precise knowledge of the location and perspective of the original source. After years of work and under the DoD-sponsored RADIUS program, Quam, Tom Strat, Aaron Heller, and others created one of the most sophisticated and accurate image manipulation programs anywhere. It is called 3DIUS.

Thus, over the 20 or so years of this program SRI has contributed:

- Methods for the top-down, goal-driven automatic exploitation of image content
- The generalized 3D derivation and representation of a scene from stereo pairs
- The extraction, modeling, and presentation of 3D objects within a scene using stereo algorithms, and a host of intrinsic characteristics such as shape, surface orientation, range or depth, and color or shading
- The construction of 3D topography from digital maps and photographs, with the ability to view the composite scene from any point
- Representation of natural scenes using fractals and a pliable, 3D, equation-based form called superquadrics
- Aids to autonomous navigation using image matching, plus scene analysis through the use of perspective changes due to lateral movement.

One of today's most sophisticated capabilities in this field is the ability to sense a 2D or 3D environment with enough precision for an autonomous vehicle to move safely within it. In 1987, the Center created a method that uses motion to discriminate the location of important objects. The processing has the arcane name of epipolar analysis. Other more conventional sensors such as those employing acoustics, laser rangefinders, and infrared detectors have also been used.

Another important product of image processing is that of visualization. With roots in the terrain representation programs discussed above comes a recently developed capability called TerraVision. TerraVision, created by Yvan LeClerc and others, is an interactive terrain visualization system that allows users to navigate, in real time, through a 3D graphical representation of a real landscape created from digital terrain elevation data and aerial images of that same landscape. The program is unusual in that it can deal with huge datasets (terabytes) that can be distributed over a wide-area network. From such a collection of sources, it can potentially produce a high-resolution model of the entire earth showing various types of imagery and cultural features. To enable TerraVision's wide use, it employs Virtual Reality Modeling Language (VRML) to store all of its terrain data, thus enabling users with a standard VRML plug-in to view TerraVision datasets over the Web.

THE PRICE FOR 10 MINUTES OF TELEVISION "FAME"

Scientific American sponsors a public television program called *Frontiers*. To help interpret the many scientific and technical wonders it explores, it has as its host the actor Alan Alda. In 1994, the producers wanted to cover robots. They called well-known research universities such as MIT and Carnegie-Mellon, but were informed that the field was not far enough along to demonstrate anything impressive. Eventually, they called Kurt Konolige at SRI who had become interested in the residual robotics work here. (Stan Rosenschein, a former Center director, had left some time before with the robotics talent to form a new company called Teleos.) Kurt told the show's producers to come, and then set about to add capabilities to the SRI robot, Flakey, that would permit an untutored person to interact with it. Flakey already knew how to map its environment—the hallways of SRI—and to reason about doing certain tasks. But interacting with it took some technical skill. Speech recognition capability was under development on Flakey, but, except for its sonar sensors, it was blind.

With perhaps elevated confidence, SRI proceeded to finish the speech command modules in Flakey and add a stereoscopic vision capability—all within 6 weeks (see Figure 4-6a,b,c). The vision software drew on a program from Interval Research Corporation across town, and the remaining work was done at SRI. The result was a stunning success. Alda could speak to Flakey with not only the commands directly relevant to navigation such as stop, turn right, and follow me, but could even inform and task it using high-level notions such as "get the budget file from Karen." At one point Alda wanted to see if Flakey would obey a command spoken by its own internal speech synthesizer. The command words were input via keyboard, spoken by the robot, and on listening to the command, Flakey executed it faithfully. These moments of "fame" are not cheap. The film crew was here for over 8 hours to get enough footage for less than 10 minutes of edited video. *Frontiers* returned to SRI for two other episodes, one on telepresence surgery (see Chapter 5) and one on artificial muscle.

Robots Revisited

Over the nearly two decades between the early days of Shakey and the mid-1980s, many advances were made that were relevant to mobile robotics. First, the power and affordability of computer workstations grew enormously, opening the way for a truly stand-alone and autonomous robot. Second was new software that addressed important needs such as real-time sensing, mapping, and, in particular, planning systems that enabled rapid decision-making. SRI was responsible for many of these advances, as mentioned above. The first, non-real-time planners emerged at SRI, including one of the first hierarchical and nonlinear planning systems. Folding in experience in rule-based and procedural expert systems like PRS, SRI developed a real-time, reactive planning system called SAPHIRA. Also, as mentioned, came innovations such as reasoning about uncertain knowledge or evidence. The Center has also done fundamental work in more esoteric concepts such as nonmonotonic (you-can-change-your-mind) reasoning, fuzzy logic, and reductions in the search space of certain theorem provers used in reasoning and planning. Then there is the flexible and resilient method of aggregating software modules already mentioned under OAA. All these helped lay the groundwork for the emergence of SRI's second-generation mobile robots.

In fact, several robots define the present state of mobile robotics at SRI. The most capable one was created first in the early to mid-1980s and became affectionately known as *Flakey*. Its name was a casual reflection to its being the sequel to the much earlier Shakey and, quite independent of its evolving reliability, the name stuck. In addition to its normal navigation sensors, Flakey has a number of human interface properties such as speech recognition and synthesis and stereovision (see the above box for an interesting account of their genesis). This suite of tools, assembled under OAA, lets it interact intelligently with its surroundings, including humans, and gives it perhaps the most comprehensive aggregation of features of any robot ever.

As mobile robots become increasingly sophisticated, robot competitions have been established to test both a robot's sensing and reasoning capabilities. Along with universities, some government labs, and a few commercial companies, SRI has entered a number of these friendly competitions. The first of these was the AI professional society's first Robot Exhibition and Competition held in San Jose, California in July 1992.[10] Here SRI's Flakey finished second to

[10] Sponsored by the American Association for Artificial Intelligence.

Figure 4-6. Some of SRI's more recent robot equipment (clockwise from upper left): Flakey, a vision module called "An Extra Pair of Eyes," Pioneer, and a gaggle of Centibots.

the University of Michigan's robot in a competition that consisted of navigating a cluttered environment, identifying objects in that environment, and following instructions to visit several sites in a specific order. Other competitions have been entered and some won. A 1996 event in particular revealed SRI's innovative spirit. The competition involved learning a complex space with many rooms of different purpose and performing a set of tasks, including reporting an appointment for a meeting to a "professor." All competitors used one robot, except SRI, which used two. The two diminutive Pioneers (see Figure 4-6(d)) were able to share the tasks, and won the competition, halving the completion time taken by the second-place robot. Obviously, the ability of the two robots to coordinate the tasks entailed both risk and, as it turned out, reward.

The latest notion in robot research in the AIC is teamwork. Figure 4-6(e) shows a working family of small mobile robots with two capabilities. One type has laser mapping tools, and the other, simpler ones move in the region the first robots have mapped to accomplish specific tasks collaboratively. The robots, which are called Centibots, perform a kind of surveillance and tasking role in otherwise hazardous environments. As of early 2004, the SRI Centibots, 100 strong, have successfully mapped and performed tasks under a DARPA competition and in a building environment never seen before. Charles Ortiz and Regis Vincent have led this new effort.

Under more recent DARPA field tests, Curt Konolige, Bob Bowles, and their SRI colleagues have won competitions where their SRI robot found its way fastest through unmapped outdoor terrain.

Bioinformatics and a New Concept for Databases

Although AIC projects generally entail substantial derivative innovation, here is an example of how an absolutely new initiative can take wing. Born of necessity and of the creative talents of computer scientist Peter Karp, a different kind of database has been evolving over the past decade or so. As is often the case in fields touched by AI, ways have been found to build representational or symbolic forms that in effect encode human knowledge and make it more amenable to computer processing. Such representation can be applied to imprecisely known or evolving fields of work or, as we have seen earlier, even to language itself. To be successful, however, it is critical for the chosen representation to model the target system with high fidelity and also be able to be read and understood by the specialists who deal with it, if not by laymen. In the accelerating biology research field, where new functions of cells at the molecular level are constantly being revealed, a means to record such progress is needed. Traditionally, such records take the form of published papers, but Karp and his colleagues are building another, more efficient and revealing way to define progress, by depicting ongoing discoveries in the way a specific cell functions. The first cell chosen is *Escherichia coli*, or *E. coli*. Scientists use this widely studied cell to understand cell workings in general, and now that the observed metabolic processes in the cell have been supplemented by its genomic sequencing, there is much to depict.

Karp's and colleagues' representational mechanism is called a pathway database, and the resulting functional picture it constructs describes the numerous biochemical reactions and enzymes that constitute *E. coli's* life. These reactions involve molecular transport, cell metabolism, and the complex networks that regulate cell function. Which cellular proteins accelerate a given chemical reaction? Which chemical compounds inhibit or activate those enzymes, and by what physical mechanisms? If done well, the representation not only accurately depicts what is observed but also helps explore possible but unconfirmed pathways in the cell's metabolism.

But even the simple, one-celled *E. coli* is complex. It has a metabolic network that involves 791 chemical compounds involved in 744 enzyme-catalyzed reactions. Small wonder that its complexity needs some type of model and an understandable representation of it. But this complexity also underscores other reasons for a consistent symbolic representation: first, because no one person can assimilate all that is going on, such a dynamic database helps convey less precisely known or related aspects of an evolving model; second, such a representation permits the expression of qualitative theories about how and why certain functions exist; and, third, it enables the creation and use of a formal, precise ontology.[11] Given the development of symbolic processing in the AI world, this model can also be computationally exercised to test the different theories advanced. New genomic sequences for *E. coli* can be used to predict new pathways and their relative importance. Last, to help biologists interact with the database, an online interface brings the model and its representation into more standard English. This new pathway database program is called EcoCyc and it has been followed by MetaCyc and BioCyc, two collections of computationally-derived metabolic pathways and enzymes for hundreds of organisms. These may be the harbinger of what is needed to understand infinitely more complex biological systems as they become defined.[B]

Attempts at Commercialization

Given AIC's leading edge developments, history should be marked by examples of commercial impacts from its work. Since the early 1980s that has indeed been the case. Commercialization forms have ranged from the outright sale of individual software packages to licensing. At the same time, some individuals chose to follow their entrepreneurial instincts elsewhere. In light of SRI's current emphasis on commercialization, we briefly review the AIC's history in this area.

Though not always with concerted effort, the AIC has been exploring the commercialization of its software for the past 20 years, with about 30 products involved. Figure 4-7 shows the magnitude of the return on these efforts. Though the curves don't reveal the details of how successful individual products were, as normally the case most of the income has been attributable to just a few packages. In fact, 90% of the total of the nearly $11.5 million income shown in the curve came from just four software "products."

The software placed for commercialization had several important characteristics. Three of the four successful products were written and proffered by just one person, Dave Kashtan, whose software design efforts were *not* derived from the AIC's research efforts *per se*. Kashtan was an excellent systems programmer in the AIC who understood the operating systems of

[11] An ontology is a catalog of the types of things that are assumed to exist in a domain of interest from the perspective of a person using a language to describe and talk about that same domain.

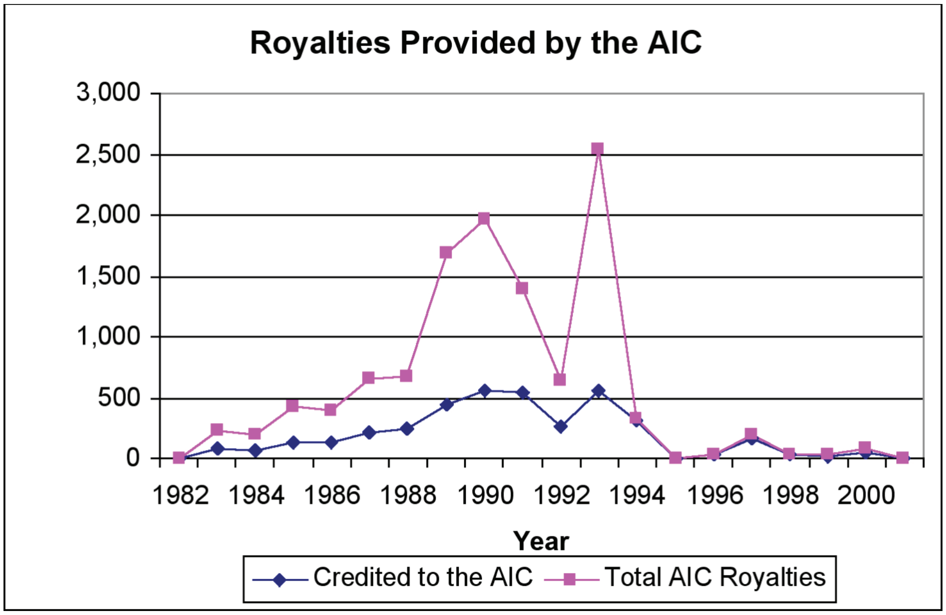

Figure 4-7. The AIC's royalty history since 1982.

the day. The widely used VAX series of computers from Digital Equipment Corporation served essentially two communities. The commercial sector used VMS as its operating system, and the other, dominated by the research community, favored UNIX. Because a lot of software was being written for the UNIX system that wouldn't run under VMS, Kashtan wrote a version of UNIX that would and called it EUNICE. EUNICE and another utility program called MULTINET, which took several of the communications protocols emerging in the government research world and made them operate under VMS, became popular. But they had nothing to do with research in AI. The next most popular software, the IMAGECALC/3DIUS/CME series mentioned earlier, accounted for perhaps 6% of the total return and was based on reasoning about imagery.

The third characteristic of the most profitable software was that it was eventually taken outside SRI to become the centerpiece of a new company called TGV,[12] where normal maintenance and other product services could be offered with fewer distractions. Thus, the most lucrative software from the AIC was utilitarian, both in what it did and how it was supported.

The fourth characteristic of the commercialization income stream in the AIC is that it did not correlate well with the SRI administration's press for commercialization as conveyed under the added impetus from Chairman Cook and President Sommers. The ability to increase commercialization opportunities naturally depends on just what research or other offerings are extant *within SRI*. That occurrence is somewhat problematical and the ability to forecast subsequent success is even more elusive.

Having said that, one commercialization effort has emerged since the new SRI emphasis that has a good chance of success. The new start-up is called Discern Communications, and its product is an outgrowth of the natural language understanding program and the years of DARPA sponsorship. Discern is addressing the need people have for product or service support from sources ever more devoid of human interaction. Whether through the telephone or online, these information sources are increasingly large, complex, and, in the case of technical subjects, arcane.[13] The frustration level can be excruciating, and people become desperate to talk with someone.

But economics have favored the use of automation to human inquiry. So, without human response, interaction, and convergence, what can be done? Suppose that the automated responder could let people ask for what they need in more natural terms and, at the same time, unambiguously understand what has been said. Both natural language understanding

[12] TGV stood for "Two Guys and a Vax" and was eventually bought by Cisco.

[13] According to Discern, by 2006 corporations will be creating 200 terabytes of information per day, of which 80% will be unstructured data.

algorithms and the speed of processors now offer that potential. Discern has built the best system to date for interpreting a phoned or online question and finding direct answers from multiple data sources. Gone is the overbearing cascade of menus, both verbal and visual. Discern dynamically understands the context and syntax of the question in terms of the enterprise's information sources and generates sentences "on the fly" containing the answer.

The new company gained seed round funding from SRI and first round investments from Spanlink Communications. Its browser-based product was chosen as "Product of 2002" by Technology Marketing Corporation's *Interaction Solutions Magazine*. As of May 2004, Discern has been integrated into Spanlink in return for stock considerations to SRI.

As a final note on the commercialization area, considerable effort must be expended to market each SRI "product." Although some of these efforts are accounted for in tracking development costs, most are not. Even less accountable are the extra hours principals devote to making research prototypes presentable and reliable. In most cases, the real economic and opportunity costs of commercialization will never be known.

AIC Alumni

What of those that left the AIC to seek their equity fortunes elsewhere? In all, nine companies have been founded at some point by leaders from the AIC. One of the earliest, Symantec, is still going, albeit not with the AI orientation with which it began. Kestrel Institute is also in business, as is the more recent conversion of Discern into Spanlink. TGV, Teleos, and Interop were bought out by larger companies, and the remaining three failed—not a bad batting average in this particular sport. While they did not start companies, many other AIC alumni nevertheless went to the commercial world at levels of high responsibility in AI-related or other companies. Others entered academia either here or in Europe. As evidence of the academic quality of work in the AIC, some became full professorships at institutions such as Harvard and Stanford.

A Recap—Maintaining Continuity in the Uncertain World of Contract Research

The preceding account shows in several ways how a vital and innovative lab, through deliberate initiatives and adaptation, can continue productive work for over three decades. Figure 4-8 summarizes the technical pathway with a "genealogy chart" of the various AI technologies the AIC has created over its history. The listing along the left-hand side shows in sequence the AI categories in which the AIC engaged from its beginning to the present. Note the first box to the right of the category represents a *new* initiative created both to expand the science and maintain the business aspects of the Center. A measure of the quality of the new initiatives has been their staying power, supported, of course, by adaptation. As the flow of what is technically possible changes with time and as new, and old, problems become solvable, new opportunities present themselves for those who are prepared.

Probably the most important requisite for such continuity is that the staff has what might be called "competent curiosity." It is not enough just to dream of a next direction for exploration; researchers must also have the skills both to understand the state of the art and to know how to extend it. Associated with all such advances are the degree of foresight involved and the quality of the foundation laid. Is the work incremental and predictable, or is it fundamentally different from what has gone before? The proof of how important a next step can be is often evidenced by the lab's technical reputation, which, in turn, is often defined by its publishing record. The AIC, through its adventurous advances in its field, has acquired an international reputation for good, fundamental work.

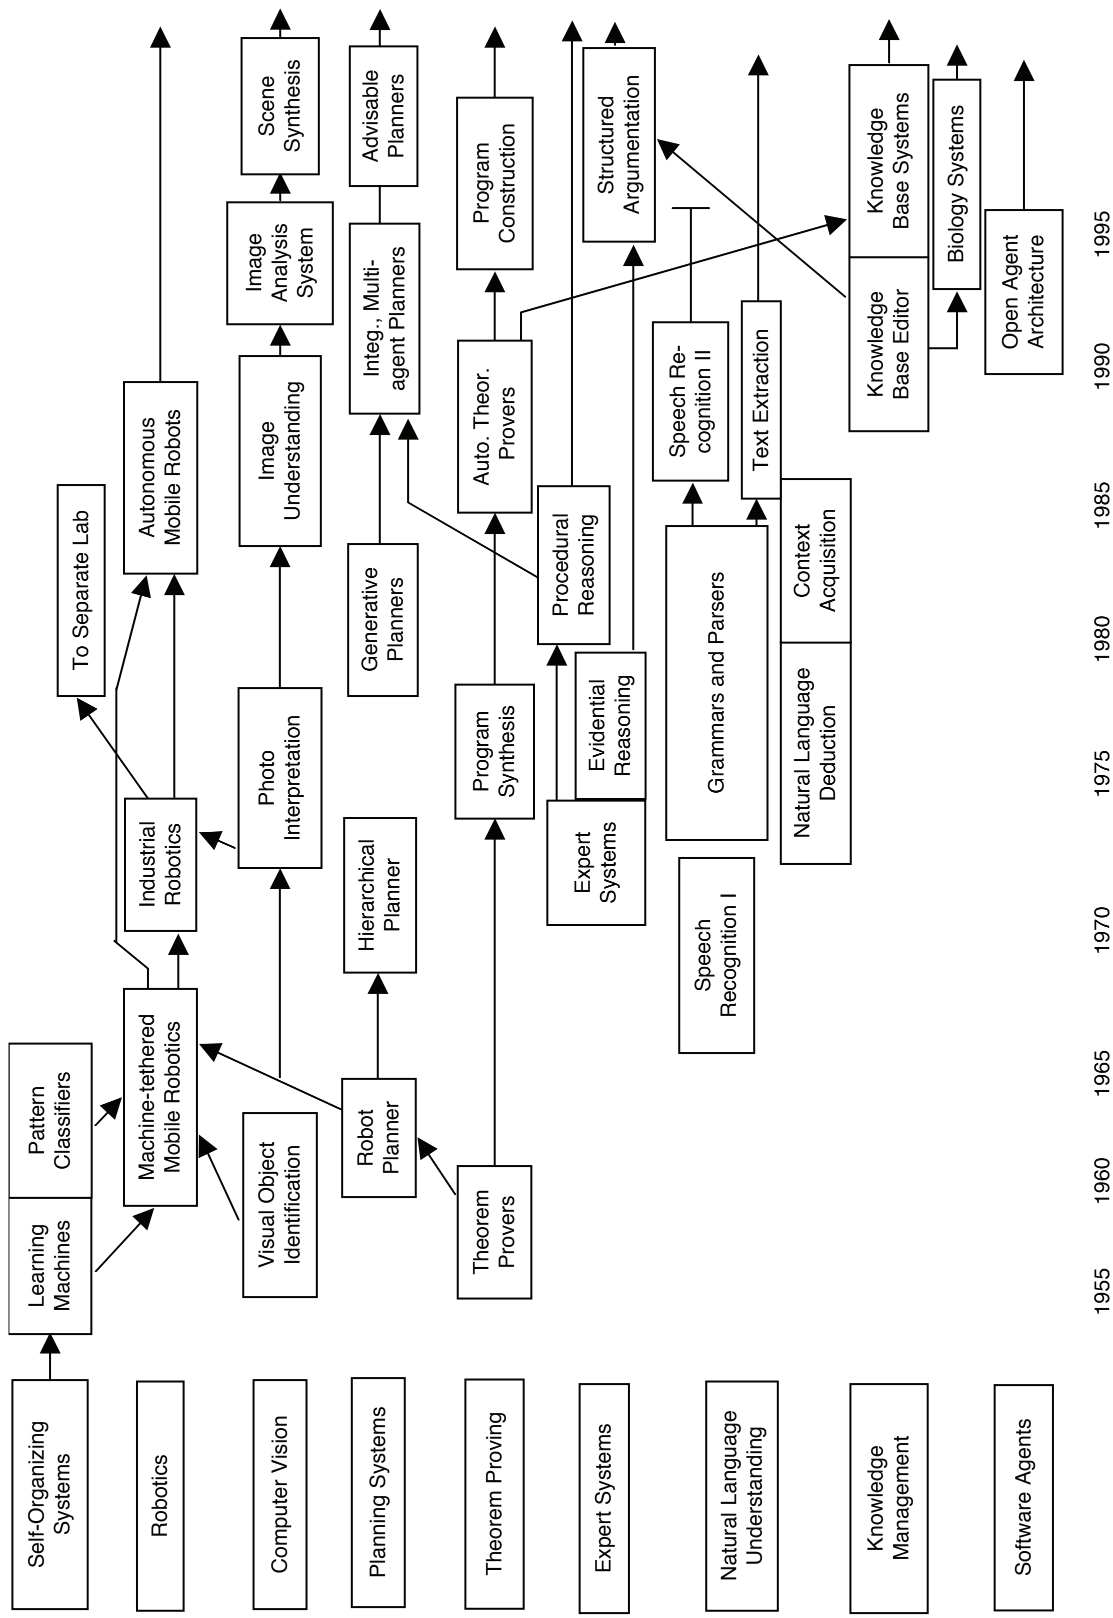

Figure 4-8. Technology Migration in the AIC from 1957 to 2000.

The next requisite for success at SRI is adaptation. We have related how the lab reacted to changes in the priorities of an important funding client like DARPA. When a program was to continue beyond one funding cycle, the lab came up with ideas about just what needed to be extended, reworked, or stopped. As always, a good, visionary researcher needs to understand what is important to the client in addition to trying to advance his or her own specialty. A good, long-horizon client permits such exploration. But many of SRI's clients have specific near-term goals instead. What does work is a deft adaptability that still leaves an interesting research problem on the table.

When a sponsored program first shows signs of winding down, it is time to lay new research agendas in front of a sponsoring agent, and in so doing try to determine or to influence the next likely course of work. Such adaptation is modulated, however, by the general progress of technology. What is now or predictably will be possible? With the dizzying advances in computing technology, what power is likely to become affordable soon? The advent of affordable automatic speech recognition was in part a result of a technical approach, but mostly attributable to the declining cost of powerful computing hardware.

SRI's adaptation may mirror a client's adaptation that comes partly through SRI's informal counsel and guidance. A good example is the approach to building the tools of automatic planning. Up to the mid-1990s, the approach sought by Government clients was an automatic planning system that, after considering all the inputs, produced an executable plan—sort of a black box approach. For realistic plans, this approach would often have required huge and complex descriptions of possible resources and conditions. Under SRI guidance, the approach was modified to make plan generation more like human planning; namely, repeated human interaction with machine-generated options whenever the plan became uncertain or complex. This approach both produces better plans and keeps the ownership of the plan more defined. Creating this kind of trusted and involved role with a client is a big boost to project continuity.

Ray Perrault, the present AIC director and one with the longest tenure, added the perspective noted above (see Figure 4-9). He also indicated that the SRI staff often assumes a rather unexpected role in the *client's* community, that of bringing its organizational elements together where otherwise they might not be. Such isolation of client subunits applies mostly to the Government, but not exclusively. Fulfilling this amalgamator's role requires the respect of the separate client elements and the objectivity that is SRI's hallmark. The role is, of course, informal, but it can bring closure faster than would happen through internal measures, and it helps position SRI as a useful player acting in the overall client's best interest. This

Figure 4-9. Current and long-standing AIC Director, Ray Perrault. As with many group leaders at SRI, Ray also has a worldwide reputation, in this case in natural language representation.

integration happened often in the RADIUS program.

Finally, one additional point needs to be made regarding the lab's enviable continuity—a point that is not exactly self-evident: How can able researchers, driven by their individual interests to advance their art, provide the continuity and longevity the AIC has seen in a world where applied, rather than basic, science is increasingly demanded? One way has been to subordinate personal scientific interests for a time until they can be resurrected in later work. Another is to maintain those interests, but employ them flexibly as new projects and contexts unfold. Of course, it helps when the number of projects is large and overlapping and, perhaps most importantly, when the researcher is in a widely used, essentially inevitable area of technology.[14]

In July 2003, the Center won perhaps its largest single contract ever because of its willingness to assemble the requisite internal and external talent to meet evolving research opportunities. The new project, the Cognitive Agent that Learns and Observes (CALO), is rich in research potential and yet framed in a way that targets practical needs. CALO intends to draw broadly from existing knowledge about agent technology and about critical supporting components in machine learning, natural language processing, knowledge representation, behavioral studies, planning, and human-computer interaction. SRI will integrate knowledge in these areas conveyed by the best U.S. universities and companies in this field.[15] The sponsor, DARPA, is looking for revolutionary ways that computers can support decision-makers. This award is a solid affirmation of the value of research talent, the ability to draw a diverse team together to meet a client need, and, perhaps most importantly, the ability to provide a combination of both attributes. That combination has kept the AIC at the top of its field for nearly half a century!

By any SRI measure, therefore, the AIC is an outstanding research organization. It has built and maintained a reputation that attracts good people. It has also shown its ability to apply high levels of technical skill to problems of client interest, thus advancing science and meeting clients' needs simultaneously—a profoundly good prescription for success in high-quality contract research.

[14] In this discussion of volition in what a researcher works on, one aspect of being a lab or center at SRI needs mention: its relationship to other labs and their ongoing research operations. Since over this time period, artificial intelligence often seemed to pose a new and fruitful solution to chronic engineering problems, it was natural for other researchers to occasionally approach the AIC leadership seeking collaboration and commitments. Depending on the AIC leadership at the time and the specifics of the case, these requests were either rebuffed or accommodated. The reasons for denying such collaboration often lay in not wanting to defocus those who were trying to advance their art, to avoid classified work, or less defendable reasons. Such reactions are strongly influenced by the center director but also include the preferences of the staff directly affected by the request. One of the traits of a research atmosphere is to protect an engaged researcher from top-down direction on how his or her time should be spent. One such request for collaboration came in an exploration of new work and when denied led the requesting researcher to leave SRI in frustration. This struggle of when and how to conduct interdisciplinary work is endemic to a contract research organization and so is treated in a bit more depth in Appendix E.

[15] Among the 16 universities are MIT, Stanford, the University of Texas, and Carnegie Mellon; companies such as Boeing are also represented.

Endnotes

[A] Nils J. Nilsson, *The SRI Artificial Intelligence Center—A Brief History*, AIC Technical Note 317, SRI International, January 24, 1984.

[B] Peter D. Karp, “Pathway Databases: A Case Study in Computational Symbolic Theories,” *Science*, Vol. 293, September 14, 2001, pp. 2040–44. As of April 2002, a dozen or so pathway/genome databases have been derived, the last being the sequencing of *Agrobacterium tumefaciens*.