Early AI Applications at Schlumberger

Reid G. Smith and Eric J. Schoen, *i2k Connect Inc., Houston, TX, 77079, USA*
Jay M. Tenenbaum, *Cancer Commons, Mountain View, CA, 94041, USA*

Schlumberger was one of the early companies to invest in artificial intelligence (AI). In this article, we summarize the work done on AI applications in the first twenty years of that investment. We highlight contributions to the company's well-known oilfield services business, as well as the technology businesses it has since divested. We also include contributions beyond Schlumberger, a bibliography of publications, and reflections from our personal experience.

Schlumberger has been at the forefront of technology in its oilfield services business since the late 1920s. Throughout its history, the company has invested in many leading-edge technologies in the ongoing quest to create and leverage differentiating technology to better serve its customers. Examples include nuclear, electromagnetic, and sonic measurements, chemical treatments, telemetry, and networking.

In the late 1970s, the first successful practical applications of artificial intelligence (AI) caught the attention of the senior management of the company; for example, Prospector, an expert system developed at SRI International for mineral exploration and resource evaluation [19], and MYCIN, an expert system developed at Stanford University to assist physicians with clinical decision-making in the selection of therapy for patients with infections [17].

In response, Schlumberger started two research efforts: a small group in the company's Schlumberger-Doll Research laboratory in Ridgefield Connecticut, in 1978; and a larger center in Palo Alto California—the Fairchild Laboratory for Artificial Intelligence Research, in 1980.

Jean Riboud, then Chairman and President of Schlumberger, summarized the company's interest [44].

"This technical revolution—the artificial intelligence—is as important for our future as the surge in oil exploration. It will force us to design new tools, it will change the capabilities of our services, it will multiply the effectiveness of our instruments. It will change the order of magnitude in the size of our business."

In this same presentation, Riboud laid out the rationale for the purchase of Fairchild Semiconductor, arguing that semiconductors would be the key to Schlumberger's measurement technology.

Thus, began a journey that continues to this day.

It will be clear from our story that while AI pervaded the thinking of the day, it provided more of a context than a single focal point. The work was much broader, extending to a wide range of computational technologies. This should come as no surprise: even for systems where the focus *is* AI, the AI functionality is always delivered as a part of an integrated system.

It will also be clear that the impact is much broader than a simple accounting of the products that followed from Schlumberger's research investment.

We have included a Timeline that describes the major steps that Schlumberger took over the years to pursue a variety of projects in the AI Field.

TIMELINE—ORGANIZATIONS AND RESEARCH THEMES

To better understand the details of Schlumberger's early AI projects and the context in which the work was done, this section presents a brief Schlumberger AI timeline, including a synopsis of the organizations in which the research was carried out and their research themes. See [56], for a more comprehensive list of people and publications.

In 1977, Schlumberger Board Member Jerome Wiesner, then President of MIT, suggested to Jean Riboud, Chairman and President of Schlumberger, that the company look into AI. Conversations followed with Marvin Minsky and Patrick Winston at MIT, who

1058-6180
Digital Object Identifier 10.1109/MAHC.2022.3149469
Date of publication 8 February 2022; date of current version 18 March 2022.

suggested discussions with Edward Feigenbaum and Bruce Buchanan of the Heuristic Programming Project at Stanford.

1977: SDR (Schlumberger-Doll Research, Ridgefield, CT). SDR is the original Schlumberger research lab, founded in 1948. AI work started in 1977, with a small team led by Bud Frawley.

1977: Bud Frawley visited Stanford. This was the start of a collaboration that extended beyond the period covered in this article. Several former Stanford students joined Schlumberger over time.

1978: Knowledge Engineering Conference, hosted at SDR, by Michel Gouilloud, director of the laboratory [30].

This meeting brought together several Schlumberger research and engineering managers with AI experts, including Raj Reddy (CMU), Edward Feigenbaum, Penny Nii, and Bruce Buchanan (Stanford), Patrick Winston and Randall Davis (MIT), Peter Hart (SRI), and David Barstow (then at Yale, later at SDR).

This was the meeting where dipmeter interpretation was selected as the most suitable problem for the application of the most practical AI techniques of the day (i.e., expert systems). Machine learning was considered, but not pursued, due to a lack of sufficient data, the relative immaturity of the field, and the constraints of computers in 1978.

These meeting notes are from conversations with participants, including Bruce Buchanan, Edward Feigenbaum, Al Gilreath, Peter Hart, and Chip Letton.

1978–1989: SDR AI research was carried out in the Systems Science department. Over the years, the department heads were Jim Baker, followed by Peter Will, and Stan Vestal. Themes included AI (automatic programming, expert systems, knowledge representation and reasoning, machine learning), graphical modeling, databases, and distributed systems [8].

1979: Ed Feigenbaum and Penny Nii spent part of a sabbatical year in Clamart and Montrouge, France, where they worked with geologist Jacques Harry and others on a lithological expert system application based on the EMYCIN framework [14].

1980–1985: FLAIR (Fairchild Laboratory for Artificial Intelligence Research, Palo Alto, CA, USA). Peter Hart was the director. The initial team included Marty Tenenbaum, Harry Barrow and Dick Duda (from SRI International), and Allan Schiffman (from Fairchild).

Themes included perception and graphics (including vision and speech), cognition (knowledge representation and reasoning), AI architectures, automated test equipment (ATE), engineering/manufacturing (e.g., integration frameworks [early service-oriented architectures], concurrent engineering, automated process planning for machine tools, dynamic constraint optimization CAD systems) [15].

1983: Hart and Duda departed to form Syntelligence. Marty Tenenbaum took over as laboratory head. Tony Ley took over as head of Fairchild Research.

1984: First commercial use of Dipmeter Advisor

1985–1989: SPAR (Schlumberger Palo Alto Research). SPAR was founded in 1985 with the staff from FLAIR. The laboratory was conceived as an advanced development center for Schlumberger's technology businesses. Themes included those from FLAIR with increased emphasis on computer-aided design (ECAD and MCAD) and ATE. The directors were Charles Smith, followed by Philippe Lacour-Gayet, and Reid Smith.

1989–1997: SLCS (Schlumberger Laboratory for Computer Science, Austin, TX and Montrouge, France, both cities that already hosted Schlumberger product development centers.).

SLCS was conceived as an advanced software center for all business units of the company. It was formed from the SDR Systems Science department staff and the staff from SPAR. Reid Smith was the director. Jean-Claude Bernard led the team in Montrouge.

Themes included high-performance computing, AI, and three-dimensional (3-D) modeling/visualization. SLCS also led the effort to improve software quality and software development productivity across all units of the company.

In 1994, the center was renamed to Schlumberger Austin Research (SAR), with Bill Preeg as director.

In 2003, Schlumberger closed SAR and the neighboring Austin Systems Center. The company transferred the staff from both to its main U.S. engineering center in Houston to better integrate advanced computational technologies into the company's product development programs.

PRODUCTS / SYSTEMS / TECHNOLOGIES

One way to measure impact is in terms of systems and technologies that were integrated into commercial products and services. This section contains a few representative examples, each associated with one of the major divisions of Schlumberger in the day: Schlumberger Oilfield Services and Schlumberger Technologies.

The contributions described in this section are structured as shown below.

SCHLUMBERGER OILFIELD SERVICES

The Oilfield Services division was the company's core business dating to its founding and is the portion of the company that exists today. Oilfield Services comprises products and services, from exploration through production, to help exploration companies and operators to "deliver reservoir performance sustainably" [49].

Dipmeter Advisor

The first commercial application of AI in Schlumberger was the Dipmeter Advisor [23], [53], [54], also the first AI project undertaken by the company, primarily at SDR. The Dipmeter Advisor system attempted to emulate human expert performance in the interpretation of the dipmeter well-logging tool that measured the conductivity of rock in a number of directions around the borehole. That tool was central to geological interpretation in the day. The task for a dipmeter interpreter is to translate these signals into conclusions about the subsurface geological features (e.g., faults, deltas, fans) that are indicators of the location of oil and gas-bearing formations. The Dipmeter Advisor was eventually sold as a service at 0.50 $US/foot.

Pattern detection algorithms proved to be as central to high performance as were the interpretation rules of the human expert [53]. Good patterns were found to be essential input for the expert reasoning system. Put another way, losses in the transition from signals to symbols are hard to overcome later in the processing pipeline with expert interpretation knowledge. Of course, the human expert, legendary oil finder J. A. Gilreath, performed both tasks at an expert level. The initial expert system implementation did not.

The development of the pattern detection algorithms was guided by the earlier experience of one of the authors in sonar signal processing.

The pattern detection algorithms were later built into the Log Quality Management System [40], [42], elements of which were incorporated into Schlumberger's logging truck systems (CSU-F and Maxis).

In a theme that permeates all of the systems highlighted in this article, the amount of effort that had to be devoted to the non-AI components was dominant. Here are four examples:

1) The Dipmeter Advisor demonstrated the challenge of infrastructure: integrating field-acquired data and the data center computing (VAX-VMS systems) with nontraditional hardware (Xerox "Dolphin" D-Machines) and software (Interlisp-D). In 1982, SDR was the first nonacademic research laboratory outside of Xerox PARC to purchase and deploy D-machines. Before these could be productive in Schlumberger's environment, SDR had to implement a full Ethernet network stack within the VAX-VMS operating system for Xerox's PUP protocol family (a predecessor to TCP-IP), ranging from packet transport and reliable streams to virtual terminal support and the Leaf page-level file access protocol.
2) Approximately half of the coding effort in the Dipmeter Advisor went into the interactive graphics system (see Figure 1). For some clients, the ability to scroll and interact with log graphics in real time (as opposed to paper charts and highlighter pens) was the most important and useful element. And long after the dipmeter tool itself was eclipsed by more advanced logging tools, the interactive graphics ideas continued to live on in subsequent Schlumberger systems.
3) To assist in commercialization, the Dipmeter Advisor was translated from Interlisp-D to Common Lisp, so that it could run on more widely available Sun workstations.
4) As another example of the kind of "infrastructural" work what was essential in the day, Eric Schoen and his colleagues wrote an entire window system facade in Lisp for two different backends. The team abstracted the Lucid Common Lisp window system on top of Lucid and on top of the NeWS toolkit (from James Gosling, who later went on to develop Java). For the NeWS toolkit, the Lisp-defined window system

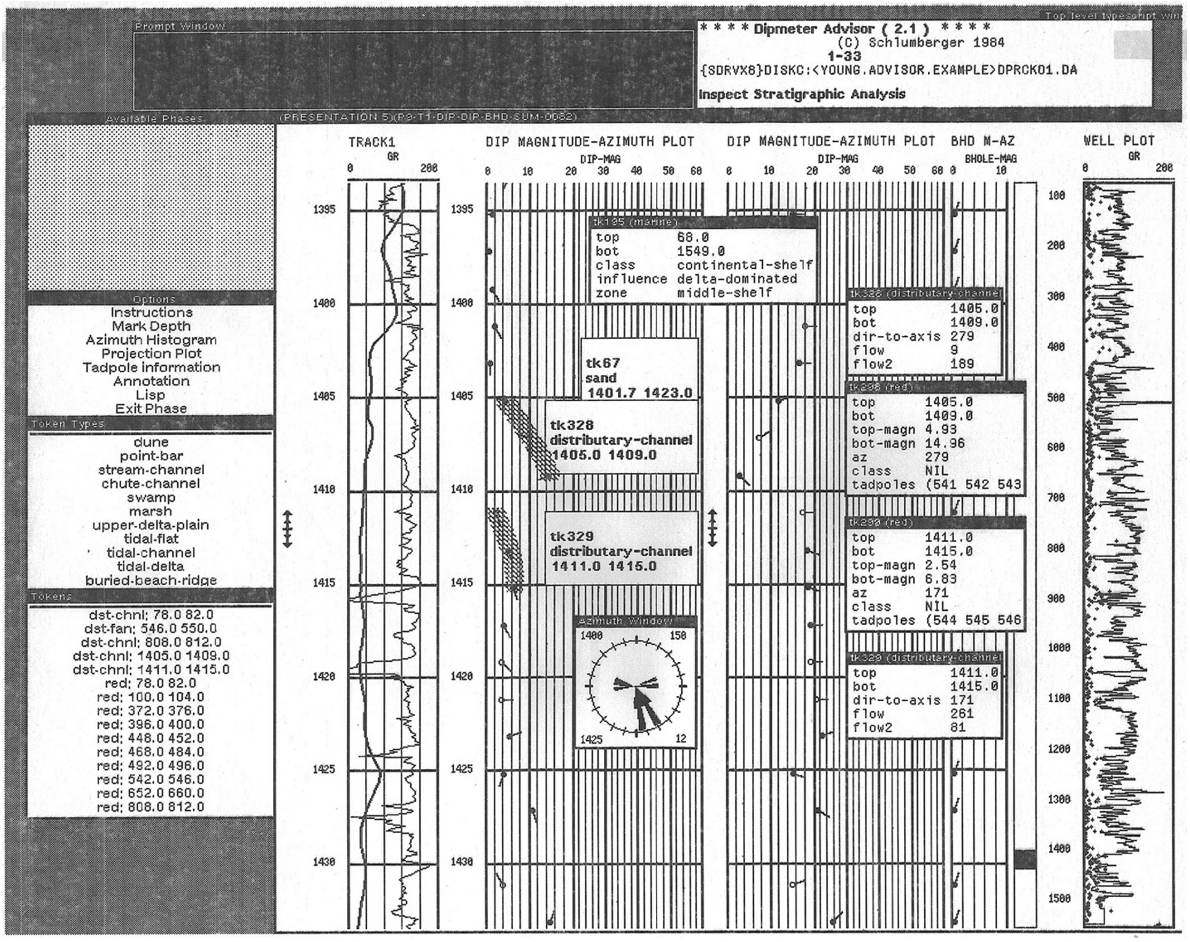

FIGURE 1. Dipmeter advisor stratigraphic analysis display.

was transpiled to PostScript and loaded onto the NeWS window system.

Knowledge Representation Technology

An essential component of the Dipmeter Advisor was a knowledge representation and programming toolkit called "HyperClass," containing: STROBE [57], a knowledge representation language; and Impulse [50], a knowledge base editor for STROBE. STROBE inherited both a legacy of frame-structured knowledge representation languages such as Units [59] and Loops [13], and concepts such as message passing and reflection from the evolving field of object-oriented programming. Impulse, written in STROBE, was an extensible knowledge acquisition framework, ultimately for both ontologies and rules.

Within Schlumberger, ideas from STROBE—translated to C—informed the configuration system embedded into Schlumberger's Maxis well logging system. Translated to Common Lisp, HyperClass was also used in early e-Commerce applications spun out of Schlumberger Technologies division (see below) and licensed by Sun Microsystems for its Symbolic Programming Environment [29].

GeoFrame Reservoir Characterization System

The Dipmeter Advisor architecture was later extended and incorporated into the Crystal system for reservoir characterization (modeling that characterizes underground reservoirs based on their ability to store and produce hydrocarbons). The goal was to develop an environment that integrated existing large, traditional interpretation programs (often involving the application of numerical models) with expert advisory systems that employed symbolic methods. The new architecture utilized declarative task description as a way of structuring the user interface. Explicit representation of control and data flow and hierarchical task/subtask relationships were used to guide an end user through a complex computation and to provide a uniform framework for application developers [55], with computations split between the interactive Lisp workstation UI and a backend VMS system to run standard Schlumberger log interpretation software.

Crystal led to Pleiades for early field deployment and was eventually redeveloped (in C and C++, with an Oracle database) and extended to become the GeoFrame Reservoir Characterization system [32]. After two decades, GeoFrame was ultimately supplanted by Petrel, which itself informs new cloud native architectures such as Schlumberger's DELFI, which Schlumberger characterizes as a "cognitive E&P (Exploration & Production) environment" [25]. Again, the key learnings back in the mid-1980s were not related to AI *per se*. Sun workstations had just appeared and were more practical for field deployment than Lisp machines. Replacing object-based knowledge bases with much more mature databases resulted in scaling, integrity, and ease of maintenance.

Lightweight Processes

The Crystal system was distributed between Lisp workstations for end-user interface and VAX/VMS systems for backend processing. The architectures of both portions relied upon lightweight multithreading. Schlumberger researchers were exposed to lightweight processes in Interlisp-D, which had come about from what the Xerox PARC people learned from Niklaus Wirth and Modula 2. Wirth spent a sabbatical year at PARC and worked on Mesa. The ideas made their way into Interlisp-D.

At Schlumberger, Eric Schoen built lightweight multithreading into the Vax side of Crystal—written in BLISS and integrated with MAINSAIL (a machine-independent version of the Stanford AI Lab SAIL language) used for application-level programming in the VAX portion of Crystal. Today, the model would be considered a "fiber" programming model, based upon the cooperative multithreading concepts inherited from Modula 2. In Crystal, there were ongoing reliability issues with this approach, because MAINSAIL had a garbage collector that was not ready for a lightweight multithreading model.

This threading model, however, survived as OVM: the OmniPak Virtual Machine. That was Scott Guthery's port of Schoen's BLISS implementation into the C language. OVM made it into Austin's Common Software Library, used by both GeoFrame and the acquisition software used in the CSU logging trucks. Schoen was later involved in reimplementing OVM's cooperative API atop native Posix threads while safely allowing preemptive scheduling; through the 1990s, this was used to implement a distributed cache-coherent memory architecture among simultaneous GeoFrame users sharing a common project database.

Massively Parallel 3-D One-Pass Seismic Migration

At SLCS, two Thinking Machines Corporation Connection Machines were purchased (initially a CM-200, later a CM-5). These machines, originally intended for AI applications, grew out of the doctoral research of Danny Hillis at the Massachusetts Institute of Technology (MIT) AI Laboratory in the early 1980s. They were purchased to support 3-D modeling at SLCS (see the "Physics-Based Computer Graphics" section). However, one of the most significant results was a groundbreaking approach to seismic processing. Massive parallelism made 3-D one-pass depth migration practical [33]. Prior to that time, for a typical seismic survey—a terabyte of input data on tapes—the typical 3-D processing sequence might take about thirty weeks on a large-scale supercomputer, with four to five weeks of processing time devoted to migration alone. The new system achieved almost a 10-fold increase in efficiency. Modern seismic processing clusters are not entirely dissimilar: They leverage GPUs equivalent to the digital signal processors of the CM-5 but with a few orders of magnitude more performance and memory, and high bandwidth low-latency industry-standard interconnects between cluster machines.

Another notable step that ushered in a new approach to seismic acquisition took place in 1996. A joint industry-government team demonstrated real-time collaboration between Houston and Schlumberger's Geco-Diamond vessel, 130 miles offshore in the Gulf of Mexico. The team included major oil companies, network providers, NASA, Livermore and Sandia National Laboratories, and the Texas Medical Center. Advances in satellite communication, networking, digital collaboration, and supercomputing enabled onshore geophysicists to quickly visualize processed data and direct acquisition of new data, based on the initial results [10], [51].

SeisClass

Starting in the early 1990s, a team based in Stavanger, Norway, developed a seismic classifier toolbox to predict reservoir parameters from seismic data attributes. It incorporated a variety of machine learning and multivariate statistical analysis techniques, including supervised Bayesian classification, and unsupervised classification via neural networks and k-means clustering [58]. The package was commercialized as SeisClass in Schlumberger's Charisma GeoFrame in 1997–98.

The technologies used in SeisClass formed the foundation for work in several areas of seismic analysis, including 4-D, automated fault extraction, and depositional environment analysis. They are today recognized as important elements in moving seismic attribute analysis from exploration into exploitation (Per Salomonsen, Personal Communication, 2021).

Stream Machine

In the early days of AI at Schlumberger, automatic programming was a central research theme. At SDR, it was pursued by David Barstow, Elaine Kant (both of whom studied under Cordell Green at Stanford University), and other team members [8].

Three application threads were pursued: The first was quantitative log interpretation, where the key issue is the interaction among diverse types of knowledge (log interpretation, formal mathematics, programming techniques), and the target language and architecture. The second thread was data acquisition software used during logging, where the key issue is programming techniques for handling real-time situations with feedback loops. This work became the Stream Machine, described here. The third thread, SciNapse, is discussed in the next section.

The Stream Machine is a software architecture designed for data acquisition and process control under real-time constraints. It supports depth-based processing that can be near real time, can deal with multiple measurements at multiple scales of resolution, from sensors offset vertically from one another, and aligned for the application of models to calculate petrophysical quantities that are not measured directly. The work stemmed from a collaboration between the SDR automatic programming team and colleagues working on the next generation of logging trucks at Schlumberger's engineering center in Austin, TX [9].

The architecture is still present in Schlumberger logging applications today.

And, in keeping with the recurring theme of this article, we started out to apply AI to Schlumberger's businesses, but much of the effort was expended on infrastructural and other tasks. In the case of the Dipmeter Advisor, Crystal, and lightweight processes, the first step was to build all the networking infrastructure (in particular, VMS kernel extensions to support layers 4–7 of the OSI X.200 model) to put Digital Equipment Vax computers on Ethernet.

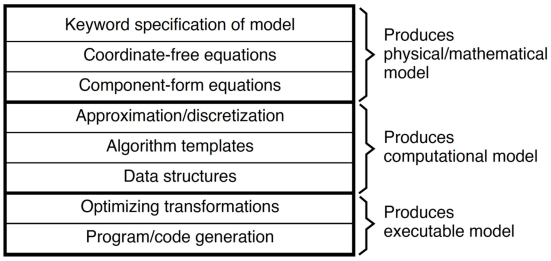

FIGURE 2. Levels of specification and abstraction [2].

SciNapse

The arrival of massively parallel computers suggested a third thread for Schlumberger's work in automatic programming. Elaine Kant and her colleagues at SLCS constructed SciNapse (originally Sinapse), a knowledge-based, domain-specific system, to synthesize mathematical-modeling software. From specifications of partial differential equation and numerical methods, SciNapse generated finite-difference programs comprising kernel stencils and boundary conditions in Fortran-77, Connection Machine Fortran, and C, which were used by modelers working on the development of Schlumberger sonic logging tools [35].

Figure 2 shows the multiple levels at which specifications can be entered in the SciNapse code generation system.

Elaine later went on to start her own company, SciComp, and the system, licensed to major investment banks, automates the implementation of software programs for financial risk-management activities—from algorithms research to production pricing to risk control [1], [2].

SlurryMinder

A critical phase of oil well completion involves positioning cement between the outer surface of a metal casing and the sides of the well. This task is done by injecting a specially formulated cement slurry down the center of the casing and up the sides of the bore hole. The SlurryMinder knowledge-based system is a design tool to aid field engineers in creating globally consistent cement slurry formulations and to rapidly disseminate current well-cementing techniques. Five separate subsystems are connected by a set of system kernel routines: the human interface, the inferencing mechanism, the concentration routines, the explanation and warning subsystem, and some utility routines for database browsing and administration. Within the inferencing mechanism, there are 14 knowledge bases containing over 700 separate rules and

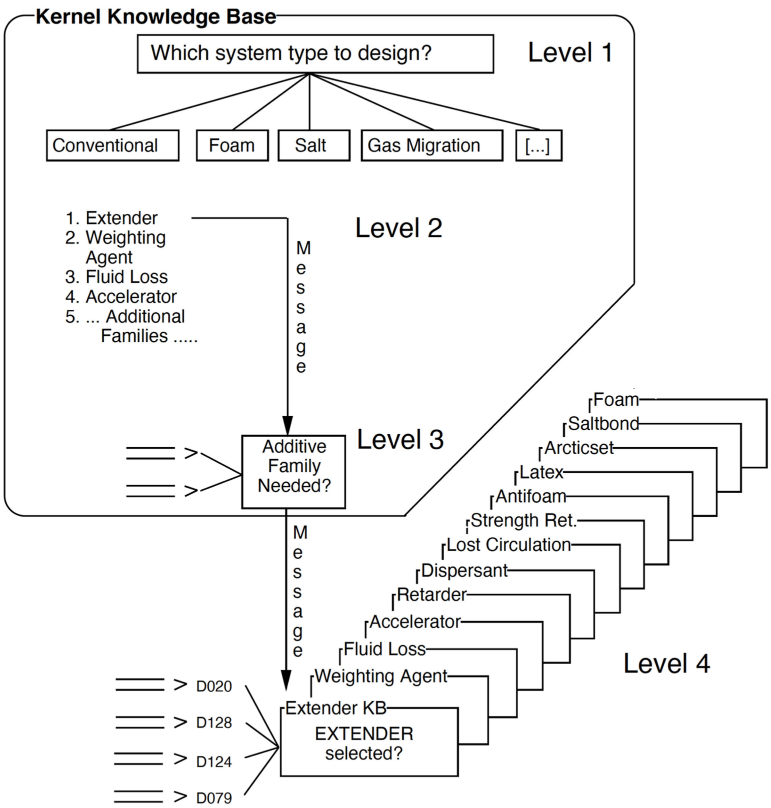

FIGURE 3. Simple schematic of the SlurryMinder architecture [37].

solution strategies. SlurryMinder was developed at the Dowell Schlumberger engineering center in Saint-Étienne, France. It won a deployed application award at the 1992 Innovative Applications of Artificial Intelligence Conference [37].

CemQuest

Another system focused on cementing is CemQuest (cement quality estimation), developed at Schlumberger's research center in Cambridge, England. This system enabled identification, characterization, and prediction of the variability of oil-field cements. CemQuest used neural networks to predict, directly from the Fourier-transform infrared spectroscopy spectrum, composition, particle-size distribution, and thickening times for certain cement-slurry formulations. It was applied in Schlumberger's regional field laboratories to detect and avoid cementing problems normally associated with cement quality and variability.

The authors explained the decision to use neural networks as follows: "The full relationships between the measurable properties of a cement powder and its slurry performance are not known and are expected to be complex, that is, highly nonlinear[. . .]. Therefore, it is best to choose a technique for finding such correlations that make as few assumptions as possible regarding their nature. ANNs offer the possibility of finding input-output correlations of essentially arbitrary complexity and consequently formed the basis for the AI methods we used in this work."

CemQuest won a deployed application award at the 1996 Innovative Applications of Artificial Intelligence Conference [22].

WaterCASE

WaterCASE is a case-based reasoning system for analyzing produced water from oil and gas wells [3], [5]. It helps engineers solve intricate water problems by linking identified problems with historically successful solutions. The system examines information from all sources including production history, reservoir descriptions, and logging results. An important aspect is that it makes allowances for missing data. This enables engineers to perform water-system analysis with only existing and sometimes incomplete datasets. A case-based reasoning architecture was chosen to take

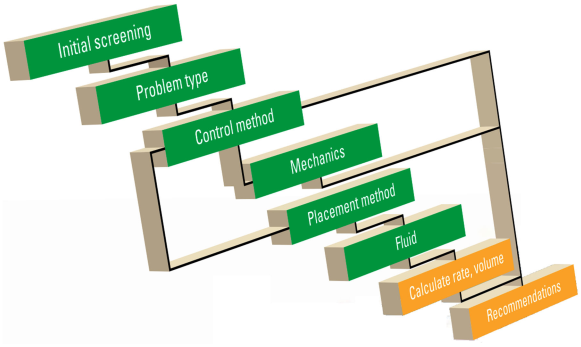

FIGURE 4. WaterCASE logical structure [3].

maximum advantage of Schlumberger's extensive documented experience with the treatment of produced water problems.

The user interface asks specific questions about symptoms and diagnostic test results that help process analysis of the water-control problem types. Once a sufficient set of answers is completed, problem types are identified and ranked according to their likelihood of incidence. The WaterCASE logical structure is shown in Figure 4.

Other Systems

AI was applied to drilling in a number of other oilfield systems in the early years of AI at Schlumberger. These include navigation for autonomous down-hole vehicles, an expert system for lost circulation, a quasi-Bayesian expert system for stuck pipe diagnosis, fault monitoring for downhole drilling motors using neural networks, stochastic grammars for interpreting what was happening on a rig from the surface data, and hidden Markov models for well control and artificial lift applications. See [12] and [26].

Schlumberger Technologies

The Schlumberger Technologies division represented a collection of technology-focused businesses, ranging from semiconductors (Fairchild) and semiconductor test systems, CAD/CAM software (Applicon), retail petroleum systems (gasoline pumps and point-of-sales systems), smart cards, and electronic metering for electricity, water, and gas. This section describes AI applications in these businesses before Schlumberger sold or spun them out in the early-mid 2000s to refocus on Oilfield Services.

IDS 5000 Electron Beam VLSI Diagnostic System

The IDS 5000 is an Integrated Diagnostic System that marries a noncontacting electron beam probe with a CAD model of a semiconductor chip to facilitate testing. A computer "maps" design features in the model into their physical locations on the chip under test, enabling the beam to measure internal signals in a nonloading, nondestructive way. This system compressed the troubleshooting or debug phase of chip development and reduced the number of design iterations necessary to produce a fully functional chip.

"The real key to success of the IDS 5000 is the integrated software which lets the user navigate effortlessly over the chip's terrain. A chip design engineer can sit down and, with little training, make measurements and diagnose problems rapidly. [. . .] Simply by selecting a particular wire in the layout window, with a computer mouse, the engineer causes the system to move automatically to the correct area of the chip and highlight the corresponding area of the layout diagram. This translates into huge efficiency gains for the engineer who can concentrate on diagnosis, and not on the tedious task of tracing wires through a maze of interconnections. Since engineers are able to correlate design data directly with chip real estate, troubleshooting time is cut dramatically for new chip designs." [47].

The IDS 5000 combined leading-edge physics, electronics, CAD, and computer vision to deliver an integrated solution [20], [21], [41]. The AI contribution was the computer vision system that enabled alignment between the CAD drawing and the circuits on the actual chip.

Electric VLSI Design System

Steve Rubin's unique VLSI ECAD system, sold commercially as "Bravo3VLSI," and later open-sourced, has been used by thousands of individuals at hundreds of institutions worldwide [27], [28]. Most ECAD systems use a geometry approach to integrated circuit layout; rectangles of "paint" are laid down on different layers to form the masks for chip fabrication. By contrast, Electric uses a connectivity approach; designers draw schematics by placing components (MOS transistors, contacts, etc.) and drawing wires (metal-2, polysilicon, etc.) to connect them [46].

In addition to the two systems described above, the FLAIR laboratory also pursued work on proving the correctness of digital hardware designs. See [6] and [7].

BravoMOST Mechanism Optimal Synthesis Tool

Schlumberger CAD/CAM's BravoMOST (Mechanism Optimal Synthesis Tool) was an interactive graphics package that aids the design of 2- and 3-D linkages.

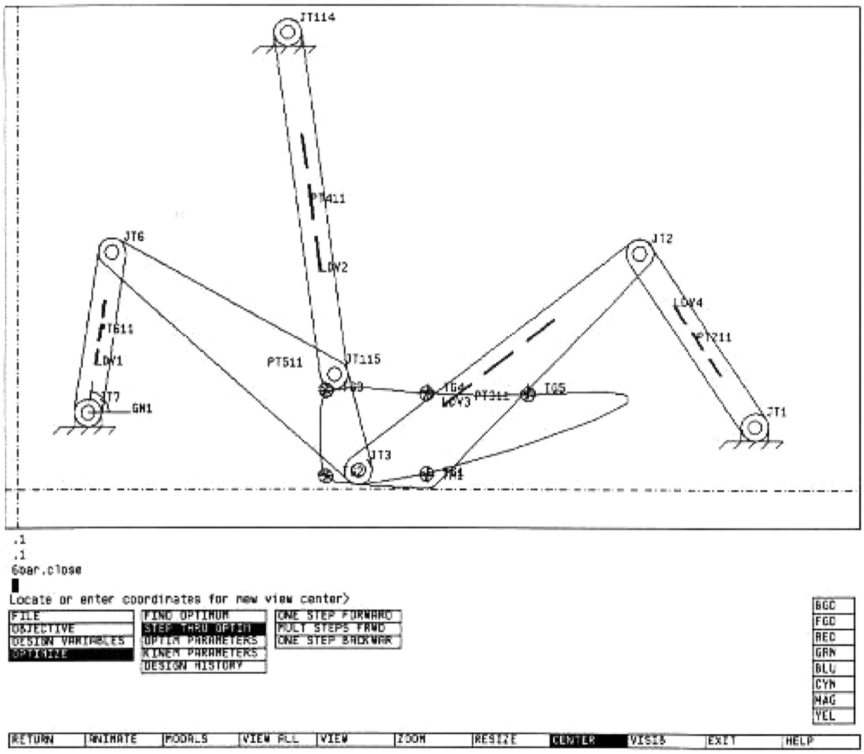

FIGURE 5. BravoMOST linkage design optimization, from [62].

It was the first product on the market that could solve complex 2- and 3-D, n-bar linkage problems [38], [48].

Geometric reasoning has been a central issue in AI from the early days of robotics. BravoMOST addressed a key geometric reasoning task, the geometric constraint satisfaction problem (GCSP): "Given a collection of geometric entities, or geoms, and constraints that describe how the geoms interact with each other, find their positions, orientations, and dimensions so as to satisfy all constraints simultaneously. Solving GCSPs is central to several related domains: describing mechanical assemblies, constraint-based sketching and design, geometric modeling for computer-aided design, and kinematic analysis of robots and other mechanisms" [38].

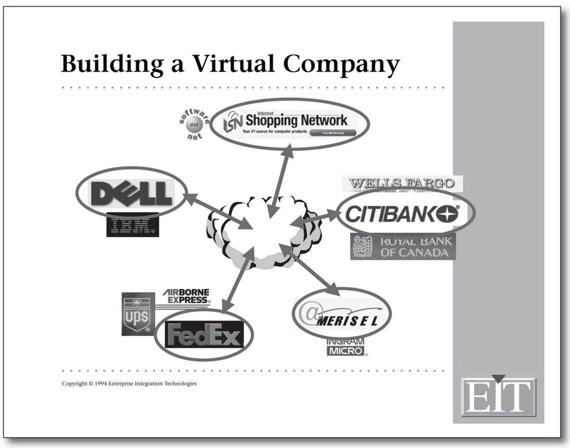

FIGURE 6. Virtual Company circa 1994 [60].

Knowledge Representation

Continuing work on KL-ONE, the well-known knowledge representation system, at FLAIR Ron Brachman and colleagues explored hybrid representation and reasoning systems combining logic with frame-based terminological components [16], [39], [43], now referred to as Description Logics. The principal themes of this work explored tractability and decidability of logic for inference in knowledge representation systems.

Impact Beyond Schlumberger

Thus far, we have focused on impact in terms of systems that went commercial inside Schlumberger, either on their own, or as components inside other platforms. But that is not the whole story.

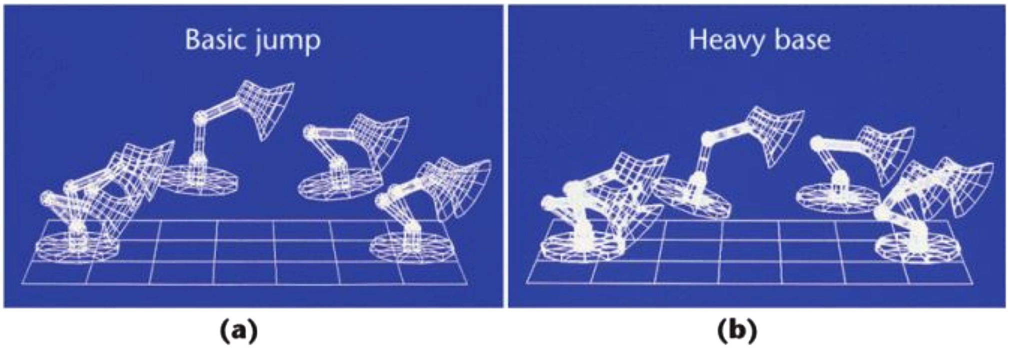

FIGURE 7. Luxo Lamp Animation, from [24]. Spacetime constraints. (a) The method generates the jumping lamp by minimizing deviations from Newtonian physics while trying to minimize the average muscle power that the lamp exerts. (b) Increasing the lamp base's mass makes it look heavier, without additional animator intervention.

Companies founded by people who worked in Schlumberger's research centers include the industry giants Cisco and Netflix, together with several others, such as Cancer Commons, CommerceNet, Enterprise Integration Technologies, i2k Connect, Pure Software, SciComp, Syntelligence, Veo Systems, and xCures. Industry recognition for ex-Schlumberger AI researchers includes: Computers and Thought Award (Hector Levesque—also a fellow of the Royal Society of Canada), 3 AAAI (Association for the Advancement of Artificial Intelligence) Presidents, 13 AAAI Fellows, and several AAAS, ACM, and IEEE Fellows.

In other words, the biggest impact has been made by talented people who built on their Schlumberger experience at other labs and in their own companies, to develop revolutionary products and services. In this section, we present three examples.

e-Commerce

The e-Commerce story begins with research conducted at SPAR in support of Schlumberger's CAD/CAM business.

Beginning around 1985, Marty Tenenbaum became intrigued with the idea of using the Internet to support large-scale engineering projects, such as the design of a 747 airplane, involving thousands of people at hundreds of companies [60].

Inspired by the distributed blackboard in the Hearsay speech understanding system, he and colleagues at Schlumberger and Stanford modeled this distributed design process as human and computer agents, collaborating through a shared model of the artifact. When an agent modified the design, affected agents were notified so they could critique the change or respond with changes of their own.

Over the Thanksgiving holidays in 1988, while running in the desert outside of Phoenix, he had an epiphany: the solution he had crafted was not limited to facilitating engineering and manufacturing. He had created a de facto Internet marketplace where individuals and companies could exchange products and services with unprecedented efficiency.

Armed with this new understanding, he left Schlumberger to form Enterprise Integration Technologies (EIT), with Jeff Pan, Jay Glicksman, and Allan Schiffman. Steve Harari, a marketing executive in Schlumberger's ATE division, joined to provide some business acumen and adult supervision. EIT became the first company to conduct a commercial Internet transaction (1992), a secure web transaction (1993), and an Internet auction (1993). In 1994, Marty founded the CommerceNet Consortium, an industry alliance to jumpstart e-Commerce. CommerceNet brought together fledgling Internet pioneers like Netscape, Amazon, and Yahoo, none of whom had more than a few dozen employees and connected them with industry stalwarts like Visa, FedEx, and IBM. These corporations provided the infrastructure needed to complete transactions, for example, to make payments and get things delivered. At its peak, CommerceNet had over 800 corporate members in more than 20 countries.

In 1997, Tenenbaum cofounded Veo Systems, which pioneered the use of XML documents for automating business-to-business transactions. Commerce One acquired Veo in 1998, and used its technology, with roots back to SPAR, to build industrial-scale marketplaces for the automobile, aerospace, and electronics industries and a dozen others.

Physics-Based Computer Graphics

At Schlumberger, the physics-based modeling work of Andy Witkin and team was thought to have several applications in the oilfield services business. It could be used to reconstruct geologic structures by correlating well logs across multiple wells in a field. It provided a way to use the logs as high-resolution boundary conditions to reconstruct detailed 3-D geology from low-resolution seismic data. And it could be used to simulate fluid flows within these structures.

But the ultimate impact was far broader, reflecting the foundational role of vision and graphics in AI and computer science.

The team's seminal papers on physics-based modeling ushered in a new generation of photo-realistic computer graphics and animated movies. Their contributions have been widely recognized in multiple best-paper awards, and a lifetime achievement award for Witkin at ACM SIGGRAPH. The entire team eventually wound up at Pixar, where they won two technical Academy Awards for modeling deformable objects (e.g., cloth, hair blowing in the wind). The team was also responsible for Pixar's now iconic bouncing Luxo Lamp animation, which was originally done for fun in the mid-80s on a Lisp Machine at SPAR.

"ACM SIGGRAPH recognizes Andrew Witkin with the 2001 Computer Graphics Achievement Award for his pioneering work in bringing a physics-based approach to computer graphics. His papers on active contours (snakes) and deformable models, variational modeling, scale-space filtering, space time constraints, and dynamic simulation are considered landmarks that have been inspirational to others and have shaped the field in such different areas as image analysis, surface modeling, and animation. Andy did his PhD. at MIT in the Psychology department. His thesis was about the perception of surface orientation from texture statistics. At Schlumberger Palo Alto Research, Andy developed the technique of "Scale-Space Filtering," which is a method for analyzing signals based on the changes in their inflection points under smoothing. The work has become a classic in the multi-resolution signal analysis literature. . . ."

For "Scale Space Filtering," see [63] and [64]. For "Snakes: Active Contour Models," see [36]. According to Wikipedia, it is the 11th most cited paper ever in Computer Science. Active contour models found their way into applications as diverse as seismic interpretation software for seismic horizon tracking in Schlumberger software, and in the mass market, in photo editing software such as Photoshop's "magnetic lasso."

And from the Academy of Motion Picture Arts and Sciences 78th Scientific & Technical Awards:

"To John Platt and Demetri Terzopoulos for their pioneering work in physically-based computer-generated techniques used to simulate realistic cloth in motion pictures.

Their 1987 article, "Elastically Deformable Models," was a milestone in computer graphics, introducing the concept of physically-based techniques to simulate moving, deforming objects." See [61].

OWL Web Ontology Language and the Semantic Web

OWL, the Web Ontology Language [65], is one of the key standards of the Semantic Web [11]. OWL's sublanguage variants including OWL DL descend directly from work on description logics; these trace to work by Brachman, Levesque, and Patel-Schneider during their time at FLAIR and SPAR.

REFLECTION

The early part of Schlumberger's AI journey has been documented in a number of publications, e.g., [4], [31]. The AI lessons learned have also been documented, e.g., [18], [52].

With the benefit of hindsight, we offer a few reflections on the challenges and learnings from the early days of corporate AI, based on our personal experiences at Schlumberger.

AI is always delivered as part of larger system.

Often it is a relatively small part, both from a functionality perspective and from a development/maintenance effort perspective. Indeed, the ratio can be as high as 90-10, when one includes systems architecture, infrastructure, networking, deployment, graphics/visualization, hardware, and many other essentials for developing, testing, delivering, and maintaining products and services to customers. This ratio was not well understood in the day, but it has, if anything, only increased between the standalone applications of the 1980s and present-day AI, embedded in every conceivable product and service.[1] As a result, there is a very fuzzy boundary between applied research and engineering.

[1]Highnam [34] underlines these points in an AI Magazine article on DARPA's vision. "Some of the work is deeply technical, such as handling pathologic failure modes, mitigating bias, and defense against adversarial attack. As much work, if not more, remains to establish the tradecraft and tools of system engineering when emerging AI technologies are inserted into enterprise systems."

Successful technology-based products and services arise from a deep understanding of both the issues faced by customers and the applicable science and technology.

But in those days at Schlumberger, and we suspect most other companies, there were few connections between corporate AI and computer science laboratories, and corporate marketing organizations or even better, customers. This was in stark contrast to the strong connections between other parts of Schlumberger research and marketing, engineering, and customers.

The boundaries between applied research and engineering product development in computationally based systems, like AI systems, are no fuzzier and harder to define than they are in the physics/mechanical engineering/electrical engineering/chemistry technologies that were engrained in the Schlumberger culture and experience of the day.[2] However, in those early days, software was new to the Schlumberger management culture. As one senior executive once said to one of us, "I am more comfortable with mechanical tools, where I can see how they work by looking at them, whereas I can't see software at all, let alone how it works." There was no experience to guide understanding on how to coordinate collaboration among software research & engineering functions. And AI in particular was viewed with skepticism by company management and by the other research and engineering groups—not dissimilar to the way in which it was viewed by the scientific and engineering communities outside the company. Continued support for software research and development in those days is a testament to the visionary leadership of the company.

Geography mattered in those days.

Schlumberger was even then a transnational enterprise, operating all over the world and made up of people from everywhere. So, there were relatively few cultural barriers to collaboration. However, the standard way to collaborate on and to transfer technologies between centers was to transfer people. This is always hard to do over extended periods, due to the consequent upheavals to family life, and so on. Those were also the early days of networking and of the local computers that preceded workstations and personal computers. Hence connectivity and compatibility were hard to come by. Different groups were on different networks until the mid-eighties and different research and engineering centers used different computers—so the code was not portable. It had to be recreated from the source.

Contrast with today, where communications technologies and the Internet make it normal to operate companies in a fully distributed manner. People living in widely disparate parts of the world come together virtually every day to collaborate on projects. Some companies, like our own i2k Connect and Cancer Commons, operate in a fully virtual manner, where the staff and the customers only occasionally meet face-to-face.

Looking back, we are grateful that we had the privilege of working with a such a diverse collection of talented people.

Schlumberger provided encouragement and a never-ending supply of interesting real-world problems to which to apply AI and other computational technologies. The company expected to apply AI technology in products and services, as it does with every other kind of technology. Indeed, that was what attracted us in the first place. We got our "10,000 hours" of hard knocks experience in what it takes to field practical AI . . . and more generally, computationally based systems and what it means to be responsive to customer needs.

OVER TIME, MEMBERS OF THE RESEARCH STAFF WENT ON TO GREAT ACHIEVEMENT AT UNIVERSITIES, OTHER RESEARCH LABS, ENTREPRENEURIAL VENTURES, AND GOVERNMENT.

The real legacy is the excellence of the people that were recruited into these labs. Over time, members of the research staff went on to great achievement at universities, other research labs, entrepreneurial ventures, and government. They carried with them projects and ideas that got their start at Schlumberger and went on to change the world.

Jean Riboud understood how early Schlumberger was to invest in AI, at a time when no one could fully grasp what the eventual outcomes and influences might be [45].

"As to artificial intelligence, we have hardly begun to understand what this abundant and cheap intellectual power will do to our lives. It has already

[2]"Industrial leaders [. . .] have realized that the traditional idea, that the sole goal of research is to create new knowledge, must be augmented by goals aimed at problem solving and effective processes for the commercial use of research results." (Schlumberger Research Visiting Committee, Personal Communication, 1997).

started to change physically the research laboratories and the manufacturing plants. It is difficult for the mind to grasp the ultimate consequences for man and society…"

ACKNOWLEDGMENTS

Thanks to the following people for generously sharing their first-hand experiences and insights: Harry Barrow, Jim Bennett, JC Bernard, Bruce Buchanan, Tony Crossley, Vu Duong, Edward Feigenbaum, Peter Hart, Chad Huston, Glenn Kramer, Chip Letton, Dennis O'Neill, Irwin Pfister, Trygve Randen, Steve Rubin, Per Salomonsen, and Allan Schiffman. We also thank the many people—staff members, students, consultants, and collaborators, with whom we shared this journey. We stand in awe of what they achieved at Schlumberger and beyond.

BIBLIOGRAPHY

[1] R. L. Akers, I. Bica, E. Kant, C. J. Randall, and R. L. Young, "SciFinance: A program synthesis tool for financial modeling," *AI Mag.*, vol. 22, no. 2, pp. 27–41, 2001, doi: 10.1609/aimag.v22i2.1559.

[2] R. L. Akers, E. Kant, C. J. Randall, S. Steinberg, and R. L. Young, "SciNapse: A problem-solving environment for partial differential equations," *IEEE Comput. Sci. Eng.*, vol. 4, no. 3, pp. 32–42, Jul. 1997, doi: 10.1109/99.615429.

[3] R. Arnold *et al.*, "WaterCASE, a case-based reasoning system for analyzing produced water, in managing water—From waste to resource," *Oilfield Rev.*, vol. 16, pp. 26–41, 2004.

[4] K. Auletta, *The Art of Corporate Success: The Story of Schlumberger*. New York, NY, USA: G. P. Pulham's Sons, 1984.

[5] B. Bailey, J. Tyrie, J. Elphick, F. Kuchuk, C. Romano, and L. Roodhart, "Water control," *Oilfield Rev.*, vol. 12, pp. 30–51, 2000.

[6] H. G. Barrow, "Proving the correctness of digital hardware designs," in *Proc. Nat. Conf. Artif. Intell.*, Aug. 1983, pp. 17–21.

[7] H. G. Barrow, "VERIFY: A program for proving correctness of digital hardware designs," *Artif. Intell.*, vol. 24, no. 2, pp. 437–491, 1984.

[8] D. Barstow, "Artificial intelligence at Schlumberger," *AI Mag.*, vol. 6, no. 4, pp. 80–82, 1985.

[9] P. Barth, S. Guthery, and D. Barstow, "The stream machine: A data flow architecture for real-time applications," in *Proc. 8th Int. Conf. Softw. Eng.*, Aug. 1985, pp. 103–110.

[10] G. R. Beering, "ARIES slashes survey time, distance," Amer. Oil Gas Rep., Jul. 1996, pp. 92–96. [Online]. Available: https://www.aogr.com/

[11] T. Berners-Lee, J. Hendler, and O. Lassila, "The semantic web," *Sci. Amer.*, vol. 284, no. 5, pp. 34–43, May 2001.

[12] T. Blanckaert and Y. M. Subervie, "Real-time data and expert system combine to develop artificial lift predictive diagnostics," in *Proc. SPE Prod. Oper. Symp.*, Mar. 2013, Paper SPE-164497-MS.

[13] D. G. Bobrow and M. Stefik, *The LOOPS Manual*, 1983. [Online]. Available: https://www.markstefik.com/wp-content/uploads/2011/04/1983-loops-manual-Bobrow-Stefik-part-1.pdf

[14] A. Bonnet and C. Dahan, "Oil-well data interpretation using expert system and pattern recognition technique," in *Proc. 8th Int. Joint Conf. Artif. Intell.*, 1983, pp. 185–189.

[15] R. J. Brachman, "Research at Fairchild," *AI Mag.*, vol. 4, no. 1, pp. 45–56, 1983, doi: 10.1609/aimag.v4i1.388.

[16] R. Brachman, "What IS-A is and isn't: An analysis of taxonomic links in semantic networks," *Computer*, vol. 16, no. 10, pp. 30–36, 1983, doi: 10.1109/MC.1983.1654194.

[17] B. G. Buchanan and E. H. Shortliffe, *Rule Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Reading, MA, USA: Addison-Wesley, 1984. ISBN 978-0-201-10172-0.

[18] B. G. Buchanan, R. Davis, R. G. Smith, and E. A. Feigenbaum, "Expert systems: A view from computer science," in *The Cambridge Handbook of Expertise and Expert Performance (Cambridge Handbooks in Psychology)*, K. Ericsson, R. Hoffman, A. Kozbelt, and A. Williams, Eds., Cambridge, U.K.: Cambridge Univ. Press, 2018, pp. 84–104.

[19] A. N. Campbell, V. F. Hollister, R. O. Duda, and P. E. Hart, "Recognition of a hidden mineral deposit by an artificial intelligence program," *Science*, vol. 217, no. 4563, pp. 927–929, 1982, doi: 10.1126/science.217.4563.927.

[20] S. Concina and N. Richardson, "IDS 5000: An integrated diagnostic system for VLSI," *Microelectron. Eng.*, vol. 7, no. 2–4, pp. 339–342, 1987.

[21] S. E. Concina and G. S. Liu, "Integrating design information for IC diagnosis," in *Proc. 24th ACM/IEEE Des. Automat. Conf.*, 1987, pp. 251–257, doi: 10.1145/37888.37925.

[22] P. V. Coveney, P. Fletcher, and T. L. Hughes, "Using artificial neural networks to predict the quality and performance of oil-field cements," *AI Mag.*, vol. 17, no. 4, pp. 41–53, Winter 1996, doi: 10.1609/aimag.v17i4.1239.

[23] R. Davis *et al.*, "The dipmeter advisor: Interpretation of geologic signals," in *Proc. 7th Int. Joint Conf. Artif. Intell.*, 1981, pp. 846–849.

[24] T. De Rose, “Andy Witkin: From computer vision to computer graphics,” *IEEE Comput. Graph. Appl.*, vol. 31, no. 1, pp. 102–104, Jan./Feb. 2011.
[25] DELFI Cognitive E&P Environment, in *Schlumberger Software*, 2021. [Online]. Available: https://www.software.slb.com/delfi
[26] B. du Castel, “Human-Centered oilfield automation,” *Schlumberger J. Model., Des., Simul.*, vol. 2, no. 43, 2011. [Online]. Available: https://www.researchgate.net/publication/216771630_Human-Centered_Oilfield_Automation, https://sites.google.com/site/bducastel/, https://sites.google.com/site/humanoilfieldautomation/
[27] Electric (software), in *Wikipedia*, 2021. [Online]. Available: https://en.wikipedia.org/wiki/Electric_(software)
[28] Electric Users Around the World, *Static Free Software (Home of the Electric VLSI Design System)*, 2021. [Online]. Available: https://www.staticfreesoft.com/electricLocations.html
[29] A. Endelman and S. Gadol, “The symbolic programming environment (SPE$^{\mathrm{TM}}$): A common lisp development environment for sun workstations,” *ACM SIGPLAN Lisp Pointers*, vol. 2, pp. 5–14, Oct. 1988, doi: 10.1145/1317250.1317251.
[30] E. A. Feigenbaum, P. McCorduck, and H. P. Nii, *The Rise of the Expert Company*. New York, NY, USA: Times Books, 1988.
[31] E. A. Feigenbaum., *Conf. Knowl. Eng. Appl. to Schlumberger Interpretation - Agenda Notes for talk*. The Edward A. Feigenbaum Papers: Work in artificial intelligence and computer science at Stanford University. Stanford University Libraries, Department of Special Collections and University Archives, 1978. [Online]. Available: https://exhibits.stanford.edu/feigenbaum/catalog/qn936rw8953
[32] “GeoFrame,” in *Ocean Drilling Program Logging Services*. New York, NY, USA: Lamont-Doherty Earth Observatory, Sep. 2002. [Online]. Available: http://www.odplegacy.org/PDF/Operations/Engineering/Logging_Tools/GeoFrame_IESX.pdf
[33] P. T. Highnam and A. Pieprzak, “Method of seismic migration using a massively parallel computer,” U.S. Patent 5,586,026, Dec. 17, 1996.
[34] P. Highnam, “The defense advanced research projects agency’s artificial intelligence vision,” *AI Mag.*, vol. 41, no. 2, pp. 83–85, 2020, doi: 10.1609/aimag.v41i2.5301.
[35] E. Kant, “Synthesis of mathematical modeling software,” *IEEE Softw.*, vol. 10, no. 3, pp. 30–41, May 1993.
[36] M. Kass, A. Witkin, and D. Terzopoulos, “Snakes: Active contour models,” *Int. J. Comput. Vis.*, vol. 1, no. 4, pp. 321–331, 1987.
[37] E. B. Kelly, P. Caillot, R. Roemer, and T. Simien, “SlurryMinder: A rational oil well completion design module,” in *Proc. 4th Conf. Innov. Appl. Artif. Intell.*, Jul. 1992. pp. 193–215.
[38] G. A. Kramer, “A geometric constraint engine,” *Artif. Intell.*, vol. 58, pp. 327–360, 1992.
[39] H. J. Levesque and R. J. Brachman, “Expressiveness and tractability in knowledge representation and reasoning,” *Comput. Intell.*, vol. 3, pp. 78–93, 1987, doi: 10.1111/j.1467-8640.1987.tb00176.x.
[40] P. Mullarkey, “An experiment in direct knowledge acquisition,” in *Proc. 8th Nat. Conf. Artif. Intell.*, Jul./Aug. 1990, pp. 498–504.
[41] H. K. Nishihara and P. A. Crossley, “Measuring photolithographic overlay accuracy and critical dimensions by correlating binarized Laplacian of Gaussian convolutions,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 10, no. 1, pp. 17–30, Jan. 1988.
[42] D. O’Neill, P. Mullarkey, and P. Gingrich, “The LQMS research project: A case study in knowledge-based systems development,” *Proc. IEEE Conf. Manag. Expert Syst. Programs Projects*, 1990, pp. 61–69.
[43] P. F. Patel-Schneider, “A hybrid, decidable, logic-based knowledge representation system,” *Comput. Intell.*, vol. 3, pp. 64–77, 1987, doi: 10.1111/j.1467-8640.1987.tb00175.x.
[44] J. Riboud, “Presentation to the New York Society of security analysts,” Mar. 20, 1980. [Online]. Available: https://reidgsmith.com/schlumberger/1980-03-20_Jean_Riboud_Presentation.pdf
[45] J. Riboud, “Book dust jacket quotation,” in *The Handbook of Artificial Intelligence Volume 1*, A. Barr and E. A. Feigenbaum, Eds., Stanford, CA, USA: William Kaufmann, Inc., 1981.
[46] S. Rubin, *Computer Aids for VLSI Design*. Reading, MA, USA: Addison-Wesley, 1987.
[47] Schlumberger Annual Report, 1989, pp. 34–37.
[48] Schlumberger CAD/CAM Division (Applicon), *BravoMOST: An Industry First in Linkage Optimization Software. Marketing Brochure*, 1990.
[49] Schlumberger, 2021. Accessed: Jan. 2021. [Online]. Available: www.slb.com
[50] E. J. Schoen and R. G. Smith, “IMPULSE: A display oriented editor for STROBE,” in *Proc. 3rd AAAI Conf. Artif. Intell.*, 1983, pp. 356–358. [Online]. Available: https://dl.acm.org/doi/abs/10.5555/2886844.2886919
[51] B. Schultz, “Oil team hits gusher,” in *Network World Collaboration*, Mar./Apr. 1996, pp. 14–18.

[52] R. G. Smith and J. Eckroth, "Building AI applications: Yesterday, Today, and Tomorrow," *AI Mag.*, vol. 38, no. 1, pp. 6–22, 2017.

[53] R. G. Smith and R. L. Young, "The design of the dipmeter advisor system," in *Proc. Annu. Conf. Assoc. Comput. Mach., 5th Gener. Challenge*, Jan. 1984, pp. 15–23, doi: 10.1145/800171.809598.

[54] R. G. Smith, "On the development of commercial expert systems," *AI Mag.*, vol. 5, no. 3, pp. 61–73, 1984.

[55] R. G. Smith, G. Lafue, E. J. Schoen, and S. Vestal, "Declarative task description as a user interface structuring mechanism," *IEEE Comput.*, vol. 17, no. 9, pp. 29–38, Sep. 1984.

[56] R. G. Smith, E. J. Schoen, and J. M. Tenenbaum, "Schlumberger computer science research bibliography," (Comprehensive list of publications from 1981-1996.), 1994. [Online]. Available: https://reidgsmith.com/schlumberger/Schlumberger_Computer_Science_Research_Bibliography.pdf

[57] R. G. Smith, "STROBE: Support for structured object knowledge representation," in *Proc. Int. Joint Conf. Artif. Intell.*, 1983, pp. 855–858. [Online]. Available: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.76.8108&rep=rep1&type=pdf

[58] L. Sonneland, P. Tennebo, T. Gehrmann, and O. Yrke, "3D Model-Based bayesian classification," in *Proc. 56th Eur. Assoc. Geosci. Eng. Meeting*, Jun. 1994, pp. cp-47-00331.

[59] M. Stefik, "An examination of a frame-structured representation system," in *Proc. Int. Joint Conf. Artif. Intell.*, 1979, pp. 845–852.

[60] J. M. Tenenbaum, "AI meets web 2.0: Building the web of tomorrow, today," *AI Mag.*, vol. 27, no. 4, pp. 47–68, 2006.

[61] D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer, "Elastically deformable models," *Comput. Graph.*, vol. 21, no. 4, pp. 205–214, Jul. 1987.

[62] P. R. Turner, "BravoMOST: Advanced technology for mechanism design," Applicon user group "Update" Newsletter, Florida, Feb. 1989.

[63] A. Witkin and J. M. Tenenbaum, "What is perceptual organization for?," in *Proc. 8th Int. Joint Conf. Artif. Intell.*, Aug. 1983, pp. 1022–1026.

[64] A. Witkin, "Scale-space filtering," in *Proc. 8th Int. Joint Conf. Artif. Intell.*, Aug. 1983, pp. 1019–1022.

[65] World Wide Web Consortium. OWL, 2007. [Online]. Available: https://www.w3.org/OWL/AI

REID G. SMITH is the Co-Founder and Chief Executive Officer of i2k Connect Inc. At Schlumberger, he was Vice President and Director for research laboratories in Cambridge, U.K., Palo Alto, CA, USA, and Austin, TX, USA. Later, he was a Vice President of Knowledge Management. He was the Senior Vice President, Information Solutions with Medstory (sold to Microsoft in 2007), and Enterprise Content Management Director with Marathon Oil Corporation. He is a Fellow of the Association for the Advancement of Artificial Intelligence and of the American Association for the Advancement of Science. He is also an IEEE Computer Society Distinguished Contributor and a Life Member of IEEE. Contact him at rgsmith@i2kconnect.com.

ERIC J. SCHOEN is the Chief Technology Officer of i2k Connect Inc. At Schlumberger, he retired as a Chief Software Architect. He contributed to the company's pioneering efforts to leverage knowledge-based systems, to the GeoFrame and Ocean platforms for reservoir characterization, to Schlumberger's software quality processes, and to its strategies for enterprise-scale architecture for data acquisition, transmission, processing, and delivery. Contact him at ejschoen@i2kconnect.com.

JAY M. TENENBAUM is the Founder and Chairman of nonprofit Cancer Commons and its for-profit spinoff, xCures Inc. At Schlumberger, he led the Fairchild Laboratory for Artificial Intelligence Research, and later the AI Laboratory at Schlumberger Palo Alto Research, where he was named the first Schlumberger Fellow. As an Internet commerce pioneer, he founded EIT (1990) and CommerceNet (1994), and later was an Officer and Director of Commerce One, Webify Solutions (sold to IBM in 2006), and Medstory (sold to Microsoft in 2007). He is a Fellow and former board member of the American Association for Artificial Intelligence, an Emeritus Consulting Professor of Computer Science at Stanford, and a former Director of Patients Like Me and the Public Library of Science. He has received numerous awards for his work in artificial intelligence and as a cancer patient advocate. Contact him at jmt@commerce.net.