Stevens, Hallam (2022) *The business machine in biology: the commercialization of AI in the life science.* IEEE Annals of the History of Computing, 44 (1) pp. 8-19.

https://doi.org/10.1109/MAHC.2021.3104868

The business machine in biology – the commercialization of AI in the life sciences

Hallam Stevens

Published in: IEEE Annals of the History of Computing, 44 (Jan-March 2022): 8-19.
https://www.computer.org/csdl/magazine/an/2022/01/09514407/1w7acZfdVde
[Special issue: “Expert systems and the commercialization of artificial intelligence”]

Introduction

Although expert systems do find a place in the history of AI, the most celebrated discoveries in AI revolve around the success of machines in specialized fields such as mathematics and in “intellectual” games.[1] As Nathan Ensmenger has pointed out, chess has long been considered an exemplary problem of AI research.[2] Other histories have focused on programs such as the General Problem Solver and the Logic Theorist as central to AI’s history.[3] More recent attention to AI’s success in chess, Jeopardy, and Go reinforce the sense that the most important advances in AI belongs to rather abstract domains of pure reasoning. Against this background, the turn to machine learning approaches in the past decade suggests that AI has only recently been successfully applied to “practical” problems such as image recognition, machine translation, or advertising.

More detailed attention to the history of expert systems presents a significantly different picture of AI research. In particular, it describes a field thoroughly and

---

[1] See Pamela McCorduck. 2004. *Machines who think.* 2nd ed. Routledge, chapter 12; Nils J. Nilsson. 2009. *The quest for artificial intelligence: a history of ideas and achievements.* Cambridge University Press, chapter 18.

[2] Nathan Ensmenger. 2011. “Is chess the Drosophila of artificial intelligence? A social history of an algorithm.” *Social studies of science* 42, no. 1: 5-30.

[3] Dick, Stephanie. 2011. “AfterMath: The work of proof in the age of human-machine collaboration” *Isis* 102, no. 3: 494-505.

deeply engaged with solving practical and commercial problems. AI work in expert systems emerged from and was, in turn, influential for a range of applied fields. Significantly, the 1980s was a period of extensive hype over expert systems, with many inside and outside the field believing they would revolutionize business operations. Like most crazes, the craze for expert systems quickly burned out. Nevertheless, such systems left important legacies in various domains and especially in biology and business.

This paper traces one important trajectory in the history of expert systems. From the 1960s to the 1990s, an interdisciplinary group of scientists and engineers centered on Stanford University showed how computer programs could successfully reason their way through complex problems in a range of practical domains. Although this story begins with the one of the pioneers of expert systems – Edward Feigenbaum – it ultimately came to involve a range of actors from chemistry, medicine, molecular biology, and business.

This history is of particular importance because of the ways in which it coupled AI simultaneously to both biomedicine and industry. The geneticist Joshua Lederberg, Nobel Laureate in Medicine, collaborated with Feigenbaum and others to connect AI to the life sciences. Biology became a crucial test-bed for expert systems and, in the long term, these systems had a transformative effect on biology. Feigenbaum and his collaborators and students brought biology and computing together in especially powerful ways. We now take for granted that biology can be computerized – we have whole sub-disciplines such as bioinformatics, biocomputing, and computational biology devoted to the task of studying life as information. The computer systems and software that Feigenbaum and his colleagues developed played an important role in establishing the possibility of these kinds of work.

The team that emerged at Stanford also brought biology and business together in new configurations. This is something else we now take it for granted: living things such as tissues or cells can have economic value and a start-up can just as well begin with a novel protein as a novel computer chip. In 1980, Feigenbaum, Peter Friedland, Douglas Brutlag, and Laurence Kedes started IntelliGenetics Inc., one of the earliest biotechnology start-ups. Unlike other early biotech companies such as Genentech and

Biogen, IntelliGenetics was focused on computation. In particular, it aimed to bring the power of non-numeric computation and expert systems to bear on problems in molecular biology. This work laid the groundwork for further developing the application of expert systems in other business domains. The successes and failures of IntelliGenetics tell a story of the emerging connections between biology, business, and computing. It suggests that the computerization of biology and the commercialization of biology were closely intertwined.

The aim of this paper is not only to call greater attention to the history of expert systems within the history of AI, but also to suggest some reasons *why* this history has been relatively overlooked. My argument here is that expert and knowledge systems were in some sense the victim of their own success. The aim of expert systems was bring the knowledge of experts "inside" the system itself. This internalization of expertise led to a black-boxing not only of the knowledge itself, but also of the mechanisms and rules through which that knowledge was applied. Because expert systems became so successful (and successfully woven into commercial technologies) their very ubiquity has rendered them relatively invisible. In particular, the integration of expert systems into large, proprietary software systems (e.g. TurboTax) makes them especially invisible. By detailing the commercialization of one particular expert system, this paper attempts to sheds light on the ways in which artificial intelligence systems have already been "built in" to the worlds we inhabit. As Feigenbaum put it, they became "part of the woodwork." This no doubt has implications for our current renewed moment of AI hype.

This narrative draws on published reports and articles, oral histories, and a range of archival sources including manuscripts from the collections of Feigenbaum, Lederberg, and Douglas Brutlag (all housed at the Stanford Library Archives and Special Collections) as well as GenBank records from the National Institutes of Health.

The Stanford context

During the 1960s, Stanford University was on a mission to transform the way it conducted research. Led by Fred Terman, a pioneer of Silicon Valley and the

University's Provost, Stanford aggressively sought out new talent, new sources of money, and new areas of work.[4] Following MIT's example, Terman aimed to build connection with both industry and the defence establishment.[5] Biomedicine and computing were central to these plans. In biomedicine the University recruited the Nobel Laureates Joshua Lederberg and Arthur Kornberg.[6] In the new field of computer science, they hired John McCarthy from MIT's Artificial Intelligence Laboratory and Feigenbaum from the University of California Berkeley. Stanford's "laissez faire philosophy" promoted novels forms of interdisciplinary collaboration. The University, Lederberg recalled, was happy to supports such research "so long as we were able to secure funding."[7] Military and industrial associations also focussed attention on applied problems that could often not be solved within the confines of traditional disciplines. And indeed, Stanford proved to be a uniquely suitable setting for such boundary crossing for decades to come.

One project that emerged brought together Lederberg with the biochemist Carl Djerassi. In the midst of the 1960s space race, Lederberg obtained funds from NASA to study exobiology – life in outer space. If it existed, how could such life be detected? After all, it might not be possible to bring a sample of lunar or Martian soil back to earth. What was needed was an automatic, computerized way of analyzing a sample for amino acids, a telltale sign of life?[8] Djerassi was an expert in using mass spectrometry to analyze organic compounds. On earth, the output of the mass spectrometer needed to be interpreted by a chemist. Could artificial intelligence programs play the role of the chemist in space? Djerassi recalls that Lederberg wanted

---

[4] See C. Stewart Gillmor. 2004. *Fred Terman at Stanford: building a discipline, a university, and Silicon Valley.* Stanford University Press.

[5] Stuart W. Leslie. 1993. *The cold war and American science: The military-industrial-academic complex at MIT and Stanford*. Columbia University Press.

[6] Gillmor, *Fred Terman*, p. 354.

[7] Joshua A. Lederberg. 1990. "How DENDRAL was conceived and born" Proceedings of Association for Computing Machinery Conference on History of Medical Informatics, Bruce I. Blum and Karen Duncan, eds., ACM Press, pp. 14-44. See also: Simcha Jong. 2006. "How organizational structures in science shape spin-off firms: the biochemistry departments of Berkeley, Stanford, and UCSF and the birth of the biotech industry" *Industrial and corporate change* 15, no. 2: 251-283.

8 Lederberg, "How DENDRAL was conceived and born"; Joshua A. Lederberg. Georgia L. Sutherland, Bruce G. Buchanan, Edward A. Feigenbaum. 1969. "A heuristic program for solving a scientific inference problem: summary of motivation and implementation." Stanford Artificial Intelligence Project, Memo AIM-104, November. Box 57, Folder 60, Joshua A. Lederberg Papers, National Library of Medicine. [hereafter JAL papers]. For a more complete account of DENDRAL see Joseph November. 2012. *Biomedical computing: digitizing life in the* United *States*. Johns Hopkins University Press.

to, “use [me] as an example and say, let the computer learn how Carl Djerassi thinks, and see whether they could put that into a programmed language.”[9] If the knowledge (and reasoning process) of the biochemist could somehow be encoded into a computer program, perhaps it could automatically interpret the output of the spectrometer.

This is where Feigenbaum, the AI expert, came into play. Feigenbaum had come to AI through thinking about human learning and psychology. He had completed his PhD under Herbert Simon at Carnegie Tech, creating a model called the Elementary Perceiver and Memorizer (EPAM) that became an important theory of memory and learning. After teaching organization theory at the University of California at Berkeley for five years, Feigenbaum’s was recruited to Stanford in 1965.[10] Stanford would allow Feigenbaum to pursue his love of computers. The mass-spectrometry problem was an ideal place for the psychologist-turned-computer scientist to test out some of his new ideas. Feigenbaum became part of the team that included Lederberg, Djerassi, and Bruce Buchanan to create DENDRAL: a computer program that stored part of the knowledge of mass spectrometry of a chemist and attempted to emulate his reasoning.

The design of DENDRAL was based on the notion that experts usually *do not* reason deductively to reach conclusions. Rather, according to DENDRAL’s approach, they utilize “heuristics” or rules of thumb to whittle down the number of plausible solutions. Feigenbaum believed that the kinds of tasks that humans routinely perform every day rely on a huge amount of knowledge about the world. This included not only facts about the world, but also rules of good practice, judgment, and guessing. A computer’s memory might be too small to store all of this knowledge; but, if the problem was confined to a limited domain, such as organic chemistry, and the computer was given as many facts and heuristics as possible in this field, then we might expect it to be able to reason effectively within this domain. Feigenbaum

---

[9] Carl Djerassi. Undated. “Using artificial intelligence in chemistry” Video interview. http://www.webofstories.com/play/carl.djerassi/88;jsessionid=53989C7B3511EA429C7E42DAC769D18C

[10] Len Shustek. 2010. “An interview with Edward Feigenbaum” *Communications of the ACM* 43, no. 6: 41-45.

called these programs “expert systems” since they became equivalent to an expert in a particular area.[11]

To determine the identity of a molecule from a mass spectrograph spectrum, for example, DENDRAL would apply heuristics (derived from expert knowledge) that would discount large numbers possibilities. By applying a sufficient number of these general rules, the program could arrive at one (or a small number) of solutions. As Harry Collins has pointed out, such heuristics are still articulable and codifiable rules and necessarily still omit much of the “tacit” or “cultural” knowledge that a human expert draws upon in solving a problem.[12] Nevertheless, the notion of reasoning via heuristics became an important model for thinking about thinking over the next several decades.[13]

The DENDRAL program did not succeed in discovering life on other planets. It did, however, give Lederberg a taste for computing. In 1966, he applied for and received funds from the National Institutes of Health to fund a dedicated computer for the Stanford Medical Center (prior to this Lederberg and Feigenbaum had mostly been using computers at the Stanford Artificial Intelligence Laboratory).[14] The Advanced Computer for Medical Experimentation (ACME) provided a shared resource for any of Stanford’s researchers who wished to pursue biomedical research using a computer. By 1972, ACME hosted over two hundred biological and medical projects including investigations of drug interactions, sensory information processing, and EEG data.[15] By then too, Lederberg and Feigenbaum had even more ambitious plans: a newer and faster computer that would connect medical researchers not just at Stanford, but around the US. The Stanford University Medical Experimental

---

[11] For a broader history of expert systems see Nilsson, *The quest for artificial intelligence*, chapter 18.

[12] Harry M. Collins. 1990. *Artificial experts: social knowledge and intelligent machines.* MIT Press. Chapter 7.

[13] Not only were heuristics deployed in expert systems, but also, like the LHASA system described by Evan Hepler-Smith, they caused chemists to reflect on their own thought processes (according to Lederberg, it motivated a “fresh study of the conceptual structure of organic chemistry”). Lederberg, “How DENDRAL was conceived and born” 14. See also, Evan Hepler-Smith. 2018. “A way of thinking backwards”: computing and method in synthetic organic chemistry. *Historical Studies in the Natural Sciences* 48, no. 3: 300-337.

[14] Joshua A. Lederberg. 1965. “Advanced Computer for Medical Research,” Application for Research Grant, NIH from Stanford School of Medicine, 9/30/1965. Box 60, Folder 25, JAL papers.

[15] Joshua A. Lederberg. 1972. “ACME Computing Facility, Annual Report FY 1972.” Box 60, Folder 35, JAL papers.

Computer – Artificial Intelligence in Medicine (SUMEX-AIM) would be connected to new computer networks such as the ARPANET and be available via modem over telephone lines.[16] Like ACME, it was funded by the NIH and, like ACME, the aim was to provide a widely shared computing resource for biomedical research.

SUMEX-AIM supported a variety of computation-based biomedical projects including the diagnosis and treatment of diseases such as glaucoma, blood infections, and pulmonary dysfunction, as well as the determination of three-dimensional protein structures from x-ray crystallographic data.[17] Among these was a project focused on molecular biology. The invention of recombinant DNA technique in the early 1970s and the development of reliable DNA sequencing methods had opened up a variety of new experimental possibilities to the molecular biologist. DNA could now be cut, cloned, copied, and rejoined. "Molecular biology was undergoing an enormous revolution," Peter Friedland recalled, "turning it from a purely qualitative science to one where massive amounts of data were now becoming available."[18]

Just as DENDRAL had attempted to reason like a biochemist, Feigenbaum and Lederberg thought computer programs might be helpful in reasoning through these new DNA-based problems. They teamed up with molecular biologists, Douglas Brutlag (from Stanford's biology department) and Laurence ("Larry") H. Kedes (from Stanford's Medical School), who developed MOLGEN (for 'molecular genetics') to run on the SUMEX-AIM computer.[19] MOLGEN incorporated a 'knowledge base' – a set of information and rules about biological molecules that would help a molecular biologist to plan an experiment. Say, for example, you wanted to clone a particular gene – inputting the appropriate sequences into the MOLGEN software, it could tell

---

[16] 1972. "Proposal for Stanford University Medical Experimental Computing Facility (SUMEX)" Submitted to Biotechnology Resources Branch of the National Institutes of Health, 1 June. Box 63, Folder 43, JAL Papers.

[17] Feigenbaum, Edward. 1981. SUMEX: Stanford University Medical Experimental Computer Resource. RR-00785. Annual Report – year 08. 1 June. Box 64, Folder 26, JAL Papers.

[18] Peter Friedland and Laurence Kedes, "Molgen-IG statement," personal correspondence, 7 May 2014.

[19] Peter E. Friedland, Laurence Kedes, and Douglas Brutlag. Undated. "MOLGEN – Application of symbolic computation and artificial intelligence to molecular biology." Hard binder, Box 63, EAF papers. The MOLGEN project is also described in detail in Timothy Lenoir. 2002 "Science and the Academy in the 21st Century: Does Their Past Have a Future in an Era of Computer-Mediated Networks?" in Wilhelm Vokamp, ed., *Ideale Akademie: Vergangene Zukunft oder konkrete Utopie?* Berlin: Akademie Verlag, 2002, pp. 113-129.

you which vectors and restriction enzymes would be most suitable. Molecular biologists across the US could log into SUMEX-AIM to use the MOLGEN software to assist with their research.

Like DENDRAL, MOLGEN sought to encode the knowledge of an expert (in this case a molecular biologist) into software. This required both "rules," derived from expert knowledge, and "heuristics" designed to interpret and apply that knowledge to specific problems. These heuristics existed as separate data structures that, in principle, could be applied to many different types of problems. It was these data structures that would be passed down, in modified forms, to future expert systems in a variety of domains.

IntelliGenetics

At Stanford in the late 1970s – just as MOLGEN was being initiated – the mood amongst biomedical scientists was one of excitement. The university's engineers and physical scientists had been spinning off their inventions into industries for years, contributing much to the development of Silicon Valley in the process. But now, suddenly, there seemed to be opportunities for life scientists to tap into these Bay Area networks of capital and power and commercialize their ideas too. As Doogab Yi has argued, the middle of the 1970s saw a profound transformation in how scientists, public officials, and university administrators perceived the potential for biological knowledge to be transformed into profit and serve the public interest.[20] Such changes were closely centered on Stanford and other Bay Area universities.[21] "It was like water off a duck's back to start a company," Feigenbaum recalled later.[22]

Brutlag was deeply immersed in this emerging culture. Studying DNA replication under the supervision of Arthur Kornberg, Brutlag received his PhD from Stanford in 1972. After spending two years in Australia, he returned to Stanford as junior

---

[20] Doogab Yi. 2015. *The recombinant university: genetic engineering and the emergence of Stanford biotechnology*. University of Chicago Press.

[21] Genentech, for example, spun out of work on recombinant DNA at Stanford. See Sally Smith Hughes. 2013. *Genentech: the beginnings of biotech*. University of Chicago Press.

[22] Interview with Edward Feigenbaum, 9th January 2008, Palo Alto California.

faculty.[23] Brutlag's steps into the new experimental worlds of DNA and RNA made him acutely aware of the potential value of software for molecular biologists. In the years after World War II, molecular biology had emerged as an information science. Closely tied to the disciplines of cybernetics and information theory, molecular biologists reimagined living systems as information systems.[24] Institutional shifts also played an important role. Growing investment in computers by universities and governments in the 1960s made biologists like Brutlag and Lederberg more aware of the potential rewards of deploying them.[25]

Brutlag's first involvement with computers came when he was introduced to a member of Lederberg's research group who was attempting to apply computers to the problem of genome mapping. They agreed that this was a problem ripe for computer methods. Through this collaboration, Brutlag also met two of Feigenbaum's students: Mark Stefik and Peter Friedland.[26] Stefik and Friedland were working on the MOLGEN project. "I must admit I was skeptical at first," Brutlag remembers, "but the more I learned about artificial intelligence methods and the successes of DENDRAL…, the more I was convinced of the importance of these methods to molecular biology and to the fields now known as genomics and bioinformatics."[27]

Part of Brutlag's attraction to the MOLGEN project was the importance of the SUMEX-AIM computer. His collaborations with computer scientists had allowed him to experience first hand the power of networked computers: scientific progress could be radically sped up when information was just a few keystrokes or an email away. Brutlag wanted to bring this sort of connectivity to the biological community.[28] Working with Kedes, they added tools for analyzing DNA, RNA, and protein sequences to MOLGEN. Perhaps most importantly, the system hosted remotely-accessible copies of DNA and protein sequence databases from Los Alamos National Laboratory and the National Biomedical Research Foundation.[29] MOLGEN came

---

[23] Douglas Brutlag. CV. https://cap.stanford.edu/profiles/viewCV?facultyId=4624&name=Doug_Brutlag
[24] Lily E. Kay. 2000. *Who wrote the book of life? A history of the genetic code.* Stanford University Press.
[25] November, *Biomedical computing.*
[26] Douglas Brutlag, personal correspondence, 26 March 2014.
[27] Ibid.
[28] Ibid.
[29] Ibid.

online at a time not only when more and more biological information was becoming available, but also as new forms of computer-driven and data-driven research were beginning to become possible. Bruno Strasser's account of the origins of GenBank, for example, suggests how databases presented new moral economies that challenged older ways of doing and knowing in biology.[30] MOLGEN was part of broader shifts towards informatization of biology and it quickly became popular amongst biologists, attracting a growing number of users through SUMEX-AIM.[31]

The demand was so high, in fact, that biologists were soon using up over half of all SUMEX-AIM's computer time. But SUMEX-AIM's mission was AI research in biomedicine, not acting as a database for biologists. When it became apparent that pharmaceutical companies were also using MOLGEN via SUMEX-AIM, the AIM committee began to object. By early 1980, it was becoming clear that the 'service' parts of the MOLGEN project would have to find another home.[32] For most of the MOLGEN group, the obvious answer was to house it in a new company: revenues would come from charging both pharmaceutical companies and academics for use of the service.

The founding of IntelliGenetics brought together computerization and commercialization – these two powerful trends within the life sciences at Stanford. Both had, in different ways, emerged from the revolutionary new recombinant methods for manipulating DNA.[33] In this way, the coupling of business to computers via DNA made sense. In September 1980, Brutlag, Feigenbaum, Kedes, and Friedland incorporated IntelliGenetics.[34] The four founders quickly hired a graduate student from the Stanford Business School, Marty Hollander, to write their business plan. The immediate aim was to sell computer resources to molecular biologists. In the short

---

[30] Bruno Strasser. 2011. "The experimenter's museum: GenBank, natural history, and the moral economies of biomedicine" *Isis* 102, no. 1: 69-96. See also, Hallam Stevens. 2013. *Life out of sequence: a data-driven history of bioinformatics*. University of Chicago Press.

[31] Edward A. Feigenbaum. 1980. "Molgen Report," Report to Elke Jordan, NIH. 8 September. Box 18, Folder 17, Edward A. Feigenbaum Papers, Special Collections and University Archives, Stanford [hereafter EAF papers].

[32] Glyn Moody. 2004. *Digital code of Life*: *how bioinformatics is revolutionizing science, medicine, and business.* Wiley. p. 18.

[33] On the climate at Stanford and its relationship to biology see: Yi, Doogab. 2015. *The recombinant university: genetic engineering and the emergence of Stanford biotechnology*. University of Chicago Press.

[34] Peter Friedland and Laurence Kedes, "Molgen-IG statement," personal correspondence, 7 May 2014.

term, this meant adapting the MOLGEN software for commercial use and making it available to biotechnology and pharmaceutical companies and whoever else would pay.[35] The first hurdle was to raise capital. Feigenbaum had professional connections to a venture capital fund connected to the Schlumbergers, a French family that had made their fortune in oil. Feigenbaum made contact with Eugene Delutio, the fund's representative in New York, enquiring about the possibility of an investment in IntelliGenetics. "I asked for only one million," Feigenbaum recalled. "He said it wasn't worth his while to do it for one million, but he would do it for two million." IntelliGenetics was up and running.[36] The ease with which this funding was acquired suggests much about the dense networks of influence and capital that permeated Stanford in the early 1980s.

In the longer term, IntelliGenetics aimed to offer a total computing package: software, support, training, and even hardware for biologists. "These firms" Hollander wrote in the business plan, "require computerized assistance for the storage and analysis of very large amounts of DNA sequence information which is growing at an exponential rate and will continue to do so for the foreseeable future."[37] Expanding information meant an expanding need for computers and a growing market for IntelliGenetics.

BIONET

The market for computerized tools for genetic engineering was not as bullish as the company expected or hoped. Although IntelliGenetics was able to gain the custom of big names, including the pharmaceutical company Smith-Kline-French and the chemical and biotechnology giant Monsanto, this did not immediately translate into large sales.[38] IntelliGenetics anticipated that most of its revenue would come in the form of payment for 'connect time' to their computer systems via dial-up connections and commercial networks such as TELENET. "What we didn't realize," Brutlag recalled, "is that a lot of pharmaceutical firms … didn't want to license logistics from a third party, but instead wanted to develop the programs in-house. They thought they

---

[35] Douglas Brutlag, personal correspondence, 26 March 2014.
[36] Interview with Edward Feigenbaum, 9th January 2008, Palo Alto California.
[37] IntelliGenetics. "Business plan," 8 May 1981. p. 5. Box 13, Folder 2, Douglas Brutlag Papers, Special Collections and University Archives, Stanford [hereafter DLB papers].
[38] Interview with Edward Feigenbaum, 9th January 2008, Palo Alto California.

could do better themselves than licensing from other places."[39] As more and more biologists realized the importance of bioinformatics – and as more university labs gained access to cheap personal computers – the demand for outside computing help diminished.

By 1984, IntelliGenetics was writing to Stanford's Office of Technology Licensing, admitting they had been over-optimistic and pleading for a renegotiation of their license for the MOLGEN software. The company had expected sales of $16.9 million in 1983, with earnings of $3.4 million; in fact, they had revenues barely over half a million and had suffered a loss of $730 000. Connect time to the IntelliGenetics computers had only brought in $50 000 per month.[40] This shortfall encouraged IntelliGenetics to explore several other directions for its business. The first, and least successful, was the sale of "genetic engineering workstations," branded as BION.[41] These were computers purchased by IntelliGenetics, loaded with MOLGEN and other software, and re-sold to molecular biology labs. As Feigenbaum recalls, they were "forced into a business we weren't good at" and sales were poor – the company managed to sell just ten between mid-1983 and mid-1984.[42]

More successful were IntelliGenetics' efforts to secure government sponsorship for their work. In December 1981, the NIH extended a "Request for Proposals" for the construction and operation of a national nucleotide sequence database.[43] IntelliGenetics work in developing computer resources for biology made them strong candidates and they submitted a joint proposal with the Los Alamos National Laboratories (LANL). IntelliGenetics bid was not successful and the contract for GenBank was awarded in 1982 to Bolt, Beranek, and Newman, a computer company based in Cambridge, MA (also working in collaboration with LANL). But the NIH's original plan had called for two parts: a database for sequences and an accompanying set of centralized computing resources for biologists (including network access to

---

[39] Glyn Moody, *Digital code of life*, p. 18.

[40] IntelliGenetics, "Form S-1," p. 4. Box 13, Folder 3, DLB papers; IntelliGenetics, "A survey of computer use by molecular biologists" p. 7. Box 63, Folder 8, EAF papers.

[41] "Research accelerated through the use of molecular biology workstation" IntelliGenetics brochure. Box 63, Folder 9, EAF Papers.

[42] Interview with Edward Feigenbaum, 9th January 2008, Palo Alto California.

[43] On the history of GenBank see Bruno J. Strasser. 2011. "The experimenter's museum: GenBank, natural history, and the moral economies of biomedicine." *Isis* 102, no. 1.: 60-96.

software and hardware). The NIH had abandoned the second part for financial reasons, but IntelliGenetics thought it might be possible to convince them to revive it.[44]

In October 1982, IntelliGenetics submitted a second proposal to the NIH suggesting that they pay IntelliGenetics to provide their software to the academic community of molecular biologists.[45] On the one hand, the NIH was worried about placing an academic resource in the hands of a for-profit business – would information and know-how be able to flow freely? On the other hand, they perceived that allowing IntelliGenetics to run the service comprised "a unique and potentially powerful coupling of intellectual and commercial forces." Providing a "commercial quality service" to academic users could create greater efficiencies and even new modes of technology transfer.[46] In the end, the arguments in favor of IntelliGenetics prevailed. In March 1984, the NIH awarded them $5 million to operate BIONET, providing access to the MOLGEN tools, online copies of databases (such as GenBank), and software libraries.[47]

With government sponsorship, IntelliGenetics managed to successfully deploy commercial computer operations in the molecular biology community. This brought together public funding, academic biology, computing, and "start-up" business together into a unique configuration. Through BIONET, IntelliGenetics also made an expert system available to the biomedical community, demonstrating some of the possibilities of the advanced forms of data manipulation and reasoning that this form of AI offered.

IntelliCorp

---

[44] Douglas Brutlag and David Kristofferson. (undated). "BIONET: an NIH computer resource for molecular biology." Box 15, DLB Papers.

[45] Ibid.

[46] Privileged communication, NIH special study section, March 17-19 1983. Box 15, Folder: Bioinformatics: Brutlag: BIONET (1), DLB papers.

[47] Dennis H. Smith, Douglas Brutlag, Peter Friedland, and Laurence H. Kedes. 1986. "BIONET(TM): national computer resource for molecular biology" *Nucleic Acids Research* 14, no. 1: 17-20.

Meanwhile, the company was having success in another line of business too. This revenue stream had nothing to do with biology at all. IntelliGenetics lack of success in attracting biologists and biotech companies to use their software forced them to look for work in other places. MOLGEN was supposed to be an artificial intelligence system for molecular biology. But it had emerged as a part of a larger project to build computer-reasoning systems based on large quantities of domain-specific knowledge.

MOLGEN was a test of the expert systems idea in molecular biology. But, in theory at least, this approach could be applied to any field. This meant that it might be possible to build a software tool that could be repurposed for different kinds of tasks. Such software would make it possible to drop in a set of domain-specific knowledge and you would have a ready-made artificial intelligence system that could reason with that knowledge. This idea emerged from several of the Stanford projects. Work had begun work on just such a general system, called EMYCIN, as an offshoot from the medical diagnosis system MYCIN ("E" stood for "essential" but the system was also known as "empty MYCIN" since it had the domain knowledge removed).[48] Likewise, Stefik and Friedland had developed something called the 'Units Package' as part of their MOLGEN work.[49] With business struggling, Feigenbaum suggested that IntelliGenetics shift its business focus from computational biology to software engineering expert systems. In a sense, it was MOLGEN with the molecular biology taken out. What was left was a set of data structures corresponding to rules of inference that could logically connected elements of knowledge. At IntelliGenetics, they called this a Knowledge Engineering Environment (KEE).[50]

Since the knowledge that could be inputted into KEE could be *anything*, this presented a multitude of new business opportunities. IntelliGenetics could help companies to develop specialized, AI-based expert systems for their own particular fields. This is exactly what they began to do as their work in molecular biology

---

[48] Edward A. Feigenbaum. 1980. "Knowledge engineering: the applied side of artificial intelligence" Stanford Heuristics Programming Project Report No. HPP-80-21, Department of Computer Science Report No. STAN-CS-80-812. Available at: https://stacks.stanford.edu/file/druid:cn981xh0967/cn981xh0967.pdf

[49] 1980. "MOLGEN funding renewal application." Folder 44, Box 2, EAF Papers.

[50] IntelliCorp. 1984. Annual Report to Securities and Exchange Commission. Commission File No. 2-87037. Folder 11, Box 61A, EAF Papers.

stagnated. In 1984, in recognition of the new direction of their business, the company was renamed IntelliCorp. IntelliGenetics persisted as a wholly-owned subsidiary.[51]

One of the first customers of KEE was Northrop. The aviation manufacturer used IntelliCorp's KEE to improve their manufacturing process. The Northrop system drew on a vast database of information about tools, materials, and techniques to generate plans for manufacturing plane parts. Through the expert system, "the corporation is gradually gathering a living library of planning expertise. Knowledge resident in the system represents Northrop's accumulated, collective, manufacturing wisdom: It represents how well we can make aircraft..." The KEE became a way of understanding of how to manufacture planes.[52]

The KEE seemed to be applicable to a range of business problems. For instance, a New York bank (Manufacturers Hanover Trust Co.) sponsored the development of an expert system for currency trading. Using IntelliCorp's KEE, the bank spent several months interviewing its currency traders to extract relevant knowledge. This information was then distilled into 350 "rules." "Building an expert system for trading is extremely difficult," *Computerworld* magazine reported, "because most traders rely on a mixture of hunch, instinct, and fundamental indicators rather than a set of fixed rules."[53] This is exactly where the KEE was supposed to come into its own – it could combine experience and rule-of-thumb knowledge in the form of the "rules" with market data and mathematical models to quickly make trading recommendations. Just as DENDRAL and MOLGEN had reasoned with chemical and molecular biological knowledge respectively, these expert systems deployed heuristics derived from business domain experts.

Opportunities came from a range of other industries too. Arthur Anderson & Co. contracted IntelliCorp to build a software tool that would allow Securities and Exchange Commission filings to be delivered electronically and another to automate

---

[51] IntelliGenetics. 1985. "IntelliCorp common stock offering." 17 December. Folder: "Board of Director, materials, annual reports," Box 64, EAF papers.

[52] Edward Feigenbaum, Pamela McCorduck, H. Penny Nii. 1988. *The rise of the expert company: how visionary companies are using artificial intelligence to achieve higher productivity and profits.* Macmillan, pp. 26-29.

[53] Alan Alper. 1988. "Brokerage seeks to trade on AI" *Computerworld* 23 May.

the analysis of financial statements.[54] NASA employed the company to develop a computer system for diagnosing faults in life support systems. AT&T used the KEE to assist with microchip circuit design. And GTE used IntelliCorp's software to create a system for improving the maintenance of central office telephone switches.[55] Applications emerged in diagnostics, planning, scheduling, configuration and process control.

The demand for expert systems grew rapidly. Along with a group of nineteen computer scientists, in 1981 Feigenbaum spun out another commercial venture: Teknowledge.[56] Like IntelliCorp, Teknowledge also developed expert systems for a range of customers including Rockwell, General Motors, Elf-Aquitaine Oil, Motorola, and Phillips Petroleum. By the mid-1980s, expert systems had become the newest and hottest fad, expected to revolutionize the business world. "We are about to see the next explosion," Feigenbaum told *Newsweek* in 1983, "which is the application of computers to reasoning." *Fortune* magazine tried to downplay the hype, but admitted, "Expert systems are suddenly the biggest technology craze since genetic engineering."[57] Artificial intelligence, here in the form of knowledge systems, was the next big thing.

As with many such fads, expert systems failed to live up to these big promises. One problem was that systems (including IntelliCorp's KEE) required dedicated, powerful, and expensive computers to run. Such computers were a significant investment even for a large company. Although hundreds of expert systems were built, IntelliCorp and Teknowledge did not become the next IBM or Microsoft.

What did happen to expert systems, however, is more interesting. In order to reason about things in the real world (airplane parts, molecules, or currency prices), systems such as Stefik and Friedland's Units Package and the KEE had to create representations of these things inside the computer. That is, the software had to have

---

[54] See Box 47, Folder 28, EAF Papers.

[55] IntelliGenetics. 1985. "What is KEE?" Folder: "Q&A for IJCAI conference," Box 64, EAF papers.

[56] Teknowledge Inc., Business Plan, September 1, 1982, p. 2. Box 65, Folder: Teknowledge Business Plan, 1982, EAF papers.

[57] Alexander. 1984. "Why computers can't outthink the experts" *Fortune*, 20 August, pp. 105-118. Quotation p. 105.

lines of code corresponding to particular objects and telling the computer how to treat such objects: what is it possible to do with such an object and how does it relate to other objects? Such ‘frame languages’, as they were called, were the basis of knowledge engineering and expert systems. But, by the mid-1980s, this kind of thinking had also become the basis of a powerful new way of developing software: object-oriented programming languages. Stefik was one of the first to see this connection and collaborated on an AI-based object-oriented programming language (called LOOPS).[58] Expertise on expert systems was partly absorbed into this field: by the late 1980s, more and more of IntelliCorp’s business was in object-oriented programming.

IntelliCorp took IntelliGenetics’ intellectual property and took the “knowledge” out of this. This shell became a flexible framework for developing other expert systems that could be applied to a range of domains, becoming incorporated into business software, business operations, databases, programming languages. What remained was a specific model – represented in data structures -- for taking expert knowledge and encoding into rules that could be used to solve practical problems. Once they had been set up, these systems allowed experts and their knowledge to “disappear” into the black box of the software itself.

GenBank

But the biotech side of IntelliCorp’s business had not disappeared. After all, IntelliGenetics had a contract with the NIH. In May 1986, IntelliCorp sold a controlling stake in IntelliGenetics to the Amoco Technology Corporation for $4 million.[59] However, there continued to be crossover between IntelliCorp’s corporate business and its work on computer systems for biology. IntelliCorp’s expertise in knowledge engineering had emerged out of its work on biology, and now this

---

[58] On LOOPS see M. Stefik, D.G. Bobrow, S. Mittal and L. Conway. 1983. “Knowledge programming in Loops: report on an experimental course” *The AI Magazine* 4, no. 3, pp. 3-13. Mark Priestley has written about how functional programming developed from an early AI paradigm. Here we see another form of programming developing from another AI paradigm. See Mark Priestley. 2017. “AI and the origins of the functional programming language style” *Minds and machines*. DOI: 10.1007/s11023-017-9432-7.

[59] 1995. “Oxford Molecular Acquired IntelliGenetics, Inc.” *Probe* 4(3-4).

biological work began to be influenced by the industrial orientation of its main business.

The company developed MOLGEN into the 'IntelliGenetics Suite' of software.[60] As one of the few commercially available software packages for molecular biology, it became important in supporting the growing community of biologists using computers. The software itself still bore the marks of its origins as a knowledge engineering tool – the search, record-keeping, and data-analysis functions were designed to allow biologists to manage and manipulate the large amounts of molecular data they were producing in their laboratories.[61] These functions were guided by the heuristic rule-following approaches that had been passed down from DENDRAL and MYCIN.

BIONET – the national computer network now run by the company – represented even grander ambitions. IntelliGenetics imagined BIONET as a complete knowledge resource for molecular biologists: it included not only the means to access and share software, but also a portal to various databases of biological information. IntelliGenetics used knowledge engineering methods to attempt to integrate this data and make it widely accessible.[62] The IRX language, developed as part of BIONET, provided access to bibliographic as well as biological data. It allowed users on the network to pose natural language queries to rapidly search databases such as GenBank (for DNA sequences) and Medline (for biological literature).[63] BIONET became an online KEE for biological work.

The GenBank database itself soon became part of this effort. The initial contract between the NIH and BBN/LANL ran for five years from 1982. When the contract came up for renewal, IntelliGenetics was in a strong position. BBN had fallen well behind in entering DNA sequences into the database from the published literature. Both the NIH and the biological community were unhappy. IntelliGenetics' proposal for the new GenBank contract highlighted their expertise in developing practical,

---

[60] Joel Huberman. 1988. "Computer power for the rest of us." Box 12, Folder 7, DLB Papers.

[61] Ibid.

[62] Ibid.

[63] "Set BIONET in context" (undated) Folder: Bioinformatics: Brutlag: BIONET (3), Box 15, DLB papers.

sophisticated, user-friendly systems for synthesizing and managing large amounts of data. Many of the IntelliGenetics/ IntelliCorp staff associated with the GenBank project (including Michael Kelly, Edward Brayman, David Roode, David Kristofferson, Brutlag, Kedes, and Friedland) had had significant experience working on commercial projects and brought this experience and sensibility with them. Indeed, the IntelliGenetics proposal cited several "related" commercial projects that would "provide a broad base of experience, personnel, and testing grounds that will greatly enhance our implementation of GenBank."[64] In other words, IntelliGenetics showed how GenBank could be improved by becoming an expert system.

This strategy paid off: the five-year contract for GenBank (running from 1987 to 1992) was awarded to IntelliGenetics (collaborating with LANL). Crucial in the NIH's decision were IntelliGenetics' plans for developing an automated means for biologists to enter their DNA sequences into the database (BBN/LANL had keyed sequences by hand or used magnetic tapes sent by mail).[65] IntelliCorp had developed similar expert systems for entering and parsing SEC filings and financial documents. As with those systems, automated entry was a complex matter. Entries had to be in standard format, correctly annotated, and free of errors. A expert system-based approach could deploy rules for checking and correctly entries as they were uploaded. IntelliGenetics system reduced the backlog in database entry from two years to twenty-four hours.[66]

IntelliGenetics also had ambitious plans for restructuring the GenBank database to accommodate the much larger volumes of data that it now needed to handle. In 1987, IntelliGenetics began to transform GenBank into a relational database. Relational databases organize their information into a series of tables that are linked or 'related' by sets of logical connectors. This reorganization allowed GenBank to represent much more complex relationships between elements of biological information than had been possible with GenBank's old 'flat file' format. It also allowed users to perform

---

[64] IntelliGenetics. 1987. "A proposal for the next five years of the Genbank nucleic acid sequence database" Response to RFP#NIH-GM-87-04,Technical proposal, submitted by IntelliGenetics, March 1987, p. 14. NIH archival documents in possession of the author [hereafter NIH documents].
[65] Ibid.
[66] IntelliGenetics. 1988. "GenBank author entry software (Authorin project): Requirements specification" Version 1.17 (March). NIH Documents.

more sophisticated queries on information in the database.[67] The updates to GenBank were remarkably similar to the transformations that took place at Northrop and other companies: both used expert systems know-how to restructure and reorganize data to make it available to a wider range of users and usable for a wider range of purposes. Both in business and in biology, IntelliCorp and IntelliGenetics restructured and linked information in ways that made it more accessible and productive.

Business rules

By the time IntelliGenetics had started to implement its changes to GenBank, the company was already moving out of the spotlight. In 1987, Congress funded the new National Center for Biotechnology Information (NCBI) at the NIH. The NCBI would assume control of GenBank in 1992, when IntelliGenetics' contract expired.[68] In 1990, IntelliCorp completed its sale of IntelliGenetics to Amoco and in 1995 the company was acquired and absorbed into the Oxford Molecular Group.[69]

But IntelliGenetics' transformation of GenBank turned out to be crucial for biology. Even before Human Genome Project officially got underway in 1990, its organizers realized that one the main challenges it would face would be the management of information.[70] The knowledge engineering of IntelliGenetics, both at BIONET and at GenBank, contributed much to solving the informatic problems involved with sequencing the three billion base pairs of the human genome. The speed and efficiency with which information could be processed, managed, searched, and analyzed was critical to the project's success.[71] Many of the concepts that became central to the field of bioinformatics – user-interactive data entry and retrieval, standardized languages and protocols, database integration and cross-referencing, interfacing with the biomedical literature – were developed in and through BIONET and GenBank.

---

[67] IntelliGenetics, "A Proposal for the next five years" p. 14.

[68] On this transition see Stevens, *Life out of sequence*.

[69] 1995. "Oxford Molecular Acquired IntelliGenetics, Inc." *Probe* 4 (3-4).

[70] Stevens, *Life out of sequence*.

[71] On the importance of speed in the human genome project see: Fortun, Michael. 1999. "Projecting Speed Genomics." Michael Fortun, Everett Mendelsohn (eds.), Practices of Human Genetics: International and Interdisciplinary Perspectives, Sociology of the Sciences Yearbook 19. Dordrecht: Kluwer, pp. 25–48.

The company's work in biology got absorbed into the mainstream of bioinformatic practice. This was part of a broader transformation of molecular biology (and particularly genomics) into a highly data-driven discipline. Increasingly volumes of data drove the needs for new tools to store and manipulate that data. This, in turn, allowed for further increases in data volumes. As I have argued elsewhere, the growth of data transformed the kinds of questions that biologists asked and the forms of knowledge it produced.[72]

Something similar happened in the business world too. Feigenbaum recalled the impact of his work on the business world: "When it propagated into industry it got called 'Business Rules'… it became part of the woodwork."[73] Business school courses these days rarely mention 'expert systems' and executives are unlikely to be impressed by talk of 'knowledge bases' transforming their business. But the legacy the 'knowledge engineering' persists. In the 1990s, IntelliCorp was hired by the German firm SAP to help them introduce artificial intelligence and object-oriented technologies into their software.[74] SAP was interested in applying IntelliCorp's expertise in developing advanced versions of distribution, accounting, human resources, and manufacturing software. In 1997, SAP decided to embed IntelliCorp systems into its business resource automation software. IntelliCorp continued to develop and market SAP software until 2019 when its core products were sold to Tricentis.

SAP software is ubiquitous across a range of industries. Expert systems were designed to "absorb" expert knowledge and black box in within the system. This is exactly what happened as KEEs were integrated into larger commercial software systems. Behind the scenes, then, IntelliCorp's technology – expert systems and object-oriented software – continue to exert a strong influence on how businesses do their work.[75]These systems are based on a particular model of AI, tracing its roots to

---

[72] Stevens, *Life out of sequence*.
[73] Interview with Edward Feigenbaum, 9th January 2008, Palo Alto California.
[74] "IntelliCorp Inc. history" Available at: http://www.fundinguniverse.com/company-histories/intellicorp-inc-history/
[75] IntelliCorp's product "LiveCompare" is still utilized for migration and upgrades of SAP systems. This product, as well as IntelliCorp's other remaining assets, were sold to Tricentis in May 2019

DENDRAL, in which expert knowledge could be encoded into heuristic rules that could then be applied to solve specific domain problems.

Conclusions

The notion of 'knowledge engineering' laid some the intellectual groundwork for many kinds of software systems and approaches that we now take for granted, especially in the business world. As we have seen, this includes object-oriented programming, and relational databases. This is not to claim that the origins of object-oriented programming and relational databases can be traced entirely to expert systems; of course, these powerful tools have complicated trajectories that have been documented elsewhere. However, the role of expert systems in these stories shows how AI has become an integral part of our software environment. It has become a ubiquitous tool.

These expert systems also laid important foundations for developments in biology and biomedicine. They made critical contributions to the ways in which biological databases were built, structured, and used by biologists. The tools that emerged from IntelliGenetics became the core of the emerging fields of bioinformatics and computational biology. By permitting the large-scale submission and organization of DNA sequence data, they also became an important enabler of the Human Genome Project (and hence foundational for the field of genomics more generally). MOLGEN, BIONET and GenBank were important elements of wider reorientation towards data and computing that reshaped large parts of biology from the 1980s onwards.

What is perhaps most important about this story is the way in which expert systems *simultaneously* impacted both business and biology. These were not separate or distinct activities; the early 1980s was a moment at which biology was *becoming* business. Expert systems helped to broker this development, shuttling concepts, aspirations, funding, and personnel back and forth between the two. Expert systems linked together humans and machines in new configurations to solve practical

---

("Tricentis adds new capability for accelerating digital transformation with SAP applications" 8 May, available at: https://www.tricentis.com/news/livecompare-acquisition-sap/).

problems, making modes of computerization in the life sciences into modes of commercialization. As such, the work of Feigenbaum and his colleagues used expert systems to reorganize relations between disciplines , experts, knowledge, and machines. In particular, it drew new forms of expertise, money, and institution into biology and biomedicine. This suggests how the computerization of biological work involved deploying new practices that were closely linked to the corporate uses of computers for speeding-up, making efficient, and centralizing control.

This invisibility of these systems should be understood as part of their success. Expert systems attempted to "black box" expertise, making it possible to make certain kinds of technical decisions without human input. In this sense, the very aim of an expert system was to bring the expert inside the system make the expert "disappear" into the system itself. The invisibility of "expert reasoning" and "business rules" is the system working as it should.

At the outset of this essay, I suggested that the story of expert systems could provide lessons about our current moment of AI hype. Certainly, the 1980s fervor for expert and knowledge systems suggests that it is hardly the first time that AI has been perceived as transformational for business. AI seemed to provide a kind of "magic" that could upend the way businesses made decisions and organized their work. This magic faded quickly. But this was not because expert systems failed. Rather, knowledge and expert systems dispersed into many processes, products and systems that we no longer think of as AI. Knowledge bases and relational databases and object-oriented programming became mundane parts of "what computers can do."

Our present-day enthusiasm for machine learning and deep-learning is likely to dissipate in the same way: not so much because they stop working, but rather because they become so successful that they become ubiquitous. At this point, it becomes harder to brand it as "artificial intelligence"; like an expert system, it becomes mundane. In this sense, the history of expert systems is an example of the strange elusiveness of AI: the engineering of the never-quite-there. But this does not mean that AI systems do not have important historical effects – indeed, as these systems are rebranded or fade into the woodwork, their impact becomes harder to see. This has implications for our present debates about the transparency of AI systems. Expert systems were designed to hide themselves, to make the “expert” invisible. The same is likely to be true of many machine learning systems: to constantly examine their inner workings, to keep the black box open, is to undermine the very purpose of their operation. Expert systems demonstrate how the long-lasting effects of AI may be in changing our expectations about what computers do, where they belong, how we relate to machines, and what kinds of problems they can be applied to.