Artificial Intelligence 59 (1993) 233–240
Elsevier

ARTINT 983

DENDRAL and Meta-DENDRAL: roots of knowledge systems and expert system applications

Edward A. Feigenbaum

Knowledge Systems Laboratory, Department of Computer Science, Stanford University, Stanford, CA 94305, USA

Bruce G. Buchanan

Computer Science Department, University of Pittsburgh, Pittsburgh, PA 15260, USA

During AI's first decade (1956–1966), the task environments in which AI scientists investigated their basic science issues were generally idealized "clean" task environments, such as propositional calculus theorem proving and puzzle solving. After the mid-1960s, a bolder and more applied inclination to choose complex real-world problems as task environments became evident. These efforts were both successful and exciting, in two ways. First, the AI programs were achieving high levels of competence at solving certain problems that human specialists found challenging (the excitement was that our AI techniques were indeed powerful and that we were taking the first steps toward the dream of the very smart machine). Second, these complex real-world task environments were proving to be excellent at stimulating basic science questions for the AI science, in knowledge representation, problem solving, and machine learning. To recognize and illuminate this trend, the *Artificial Intelligence Journal* in 1978 sponsored a special issue on applications of artificial intelligence.

Correspondence to: E.A. Feigenbaum, Knowledge Systems Laboratory, Department of Computer Science, Stanford University, Stanford, CA 94305, USA. E-mail: eaf@sumex-aim.stanford.edu.

Salient among the early applications of AI was the work of the DENDRAL group, applying AI to problems of the analysis of the mass spectra of organic molecules and the induction of new rules of mass spectral fragmentation. The paper "Dendral and Meta-Dendral: their applications dimension" [3] was solicited by the editor of the special issue. The article appeared in 1978, thirteen years after the start of the DENDRAL Project. In those thirteen years, many DENDRAL papers had been published in the literature of both AI and chemistry. Some of the results of DENDRAL and Meta-DENDRAL as applications to chemistry had been reported to chemists. The time was ripe for reporting these results not merely as chemistry but as applied AI, and the special issue provided us with the appropriate vehicle.

In this note, we will look both backward and forward from the 1978 publication date of the special issue. Though this note is necessarily short, two other and longer works have done this job thoroughly [8,9], the more recent paper having the advantages of perspective of time and experience with the technology transfer of DENDRAL to an industrial setting.

DENDRAL in the context of its time: more history

The DENDRAL Project began in 1965. Feigenbaum had been searching for a task environment in which to investigate processes of empirical induction (of models and theories from data) and had oriented his thinking toward finding such a task environment among the activities that scientists do. Lederberg, a geneticist whose work in 1965 on exobiology involved the mass spectra of amino acids, suggested the task of analyzing mass spectra—the formation of hypotheses of organic molecular structure from mass spectral data. Buchanan joined the effort shortly thereafter; his orientation was philosophy of science blended with AI, a concern for the nature of scientific discovery and the information processes underlying it.

DENDRAL work was largely experimental work. One of the earliest of the experimental results was also perhaps the most important. That was the *knowledge-is-power hypothesis*, which has become the slogan by which many in AI remember the DENDRAL Project. As we extended the limits of DENDRAL's abilities, what we found was that we needed, more than anything else, *more domain-specific knowledge* of chemistry and mass spectrometry (having more powerful AI problem solving methods was useful but not crucial to our success; more knowledge was crucial). Toward this end, we recruited the collaboration of Djerassi, a world-class specialist in mass spectrometry, and with Djerassi his team of researchers, visitors, and post-docs.

One of our important early motivations, investigating the processes of theory formation (in scientific work, and elsewhere) was postponed for several years. We made a decision (in retrospect correct) to achieve experimental results and gain system-building experience on a more concrete problem first: the hypothesis formation problem of inferring from one set of spectral data one (or a few) candidate molecular structure(s). Success with DENDRAL led us back to the original problem of theory formation, which now appeared in a quite specific and concrete form that was "meta" to DENDRAL (hence the project's name, Meta-DENDRAL). If the knowledge of mass spectrometry was crucial to the progress of DENDRAL, then we must codify it, and we knew of only two ways. Either work in the painstaking one-on-one fashion of our interaction with Djerassi's chemists (which has since become known as the knowledge acquisition part of knowledge engineering). Or, infer the knowledge directly from electronic libraries of mass spectral data and the known structures that gave rise to the data. The latter was the task of the Meta-DENDRAL learning program.

Did DENDRAL make a difference?

How did DENDRAL affect AI? The AI science of 1965 was largely dominated by the theme of the "*generality* dimension" of problem solving. At center stage was the program GPS, but the GPS model was being challenged by problem solvers based on theorem proving using the newly discovered and programmed resolution method. Hardly mentioned in the discussion was the role of knowledge in problem solving. For example, in a very important early paper by Newell, Simon, and Shaw [11] describing their chess playing program, the word "knowledge" is used only once, and that use is incidental to the paper's main line.

Because generality dominated the field's concerns, the issue of *levels of competence in performance* did not motivate the field very much in the early 1960s. That was odd, because the dream of human-level and beyond-human-level performance was strongly present at the birth of AI (e.g. predictions of world-class chess play; predictions of new theorems to be proved by AI programs; performance of an AI program on the New York State Regents geometry exam).

The DENDRAL group was strongly focused on the performance dimension of AI and played a key role in reinstating the view (the goal, the dream) that AI programs can perform at the level of the most competent humans performing the task (and perhaps beyond). It did this with enough specificity that its results could be extended by the group itself, and by others. For the DENDRAL group, the extension was to MYCIN, then to EMYCIN (the software generalization), then to applications in medicine, engineering, molecular biology, x-ray crystallography, submarine detection, etc. This system-building experimental AI effort, sustained over a period of fifteen years, engendered concepts, methods, techniques, software, and (not least) credibility for the approach. A new sector of AI was born, as well as a sector of the software industry.

How important was this to AI? In 1985, when the Editors of the *Artificial Intelligence Journal* asked key AI scientists for their views of the most important happenings in AI over the previous decade, Allen Newell responded:

> There is no doubt, as far as I am concerned, that the development of expert systems is the major advance in the field during the past decade.... The emergence of expert systems has transformed the enterprise of AI, not only because it has been the main driver of the current wave of commercialization of AI, but because it has set the major scientific problems for AI for the next few years.... [1, p. 385]

We think of DENDRAL as the "grandfather of expert systems". (Though, in the modern style, perhaps we should think of it as "the mother of all expert systems".) In 1968, we wrote a paper summarizing the first three years of the DENDRAL experiments, in which we used the DENDRAL results to challenge the "generality" paradigm. We stated and defended our knowledge-is-power hypothesis of problem solving. That paper, in several places, links the system's behavior with that of "the expert", and the felicitous phrase "expert system" came into use in our project thereafter. Much more important than the coining of a term was the fact that we helped to set in motion a shift in paradigm in AI from one based on generality to one that was *knowledge-based.* The knowledge-based paradigm is today the main operating paradigm of AI. For example, one of the major textbooks of AI (Rich and Knight [12] concludes the book with this:

> If there is one single message that this book has tried to convey, it is the crucial part that knowledge plays in AI programs. [12, p. 579]

DENDRAL was not the only agent that brought about the shift of paradigm (at MIT, Moses and the Mathlab/Macsyma group were influential supporters of and early contributors to the expert systems viewpoint), but it was one of the most significant agents.

By 1967, the DENDRAL project faced a crisis of knowledge representation. The amount of new knowledge (represented as LISP code) that was pouring in via the knowledge acquisition interactions with the chemists produced a complexity of the knowledge base that we could neither manage nor sustain. Inspired initially by the Newell–Simon use of productions as an architecture for problem solving, we conceptualized productions as modular situation–action "rules" in terms of which we could represent the knowledge of mass spectrometry. By early 1968, we had re-represented the entire knowledge base of DENDRAL, had provided a clean representation for the remainder of the DENDRAL Project, and had given ourselves the right representational leverage for the soon-to-happen MYCIN work. This contribution of DENDRAL to AI has been one of the most robust. As expert systems moved into industrial and commercial use, most of the implementations were rule-based systems.

Well, DENDRAL, what have you done for us lately?

Riding on the maturation of time-sharing technology and the birth of the ARPAnet, Lederberg and Feigenbaum established the SUMEX facility, a national computational resource for applications of AI to Medicine and Biology. DENDRAL was ported to SUMEX's PDP-10 and made available throughout the 1970s and early 1980s to a wide national community of academic and industrial chemists. They used DENDRAL primarily to gain the advantage of its superb structure elucidation methods (not, however, for its mass spectral analysis expertise).

In the original article that is the subject of this note, we pointed to a problem of technology transfer at that time: the absence of "satellite engineering firms" for AI that could "map research programs into marketable products" that would benefit the chemical industry. In this simplistic wish of our younger selves was a stunning naiveté: that such firms, if they existed would want to "harden" our software; and that buyers for our software existed in the industry. Also in the original article, as a way to improve the chances for technology transfer, and with some prescience, we called for a small computer to come into being, hopefully inexpensive, that would run advanced symbolic manipulation languages (we asked for INTERLISP)!

Several factors, not least among which was the impending end of federal funding for the DENDRAL Project, led Stanford to license the DENDRAL programs to a company specializing in software for chemical structure manipulation, synthesis planning, and literature searching. The key players of the DENDRAL Project in the 1980s were all of a breed that might be called "computational chemists". They were hired by the company to do the technology transfer, and more importantly for the company to guide and develop other projects within the firm. (Our former chemist collaborators regard the DENDRAL Project's training of many of the first generation of computational chemists as one of the most significant contributions of the Project).

DENDRAL as a software entity with a unique identity was "de-constructed". Its structure manipulation algorithms have been used in the various products of the firm, including three chemical database management systems.

These systems, using DENDRAL ideas, algorithms, and chemical structure representations, are used (according to the company) by the "overwhelming majority of the world's chemical and pharmaceutical industries to manage their chemical information".

The focus on mass spectrometry was apparently not a marketable focus. Nor apparently was DENDRAL's unmatched capability to do systematic structure elucidation—even when coupled with excellent modern interactivity that would allow a chemist to shape and control the search for structures.

As AI researchers, we seriously underestimated the problems of technology transfer and the nature of the barriers to technology diffusion. "Underestimate" is charitable: we really didn't have the foggiest idea. This same lack of understanding was to plague the embryonic AI software and applications industry throughout the 1980s. While the small, cheap symbol-manipulating computers we asked for did indeed help to lower acceptance barriers, we largely ignored the social, psychological, and business aspects of reluctance to try new tools. We comment on all of these, and give the subject extensive discussion in our recent case study paper [9]. We also failed to appreciate until 1988 [5] the crucial role played by champions of the technology in industry. A technology does not transfer itself; it is transferred on the strong shoulders of champions of the technology. None of the early industrial users of DENDRAL became an industrial champion for DENDRAL.

What about Meta-DENDRAL? What did it do for AI? For industry?

Around the time of Meta-DENDRAL's birth (circa 1970), work in the machine learning area was at a low ebb (exceptions were Waterman's work at our lab on learning of production rules, an ancestor of Meta-DENDRAL; Michalski's work on variable-valued logic and its application to learning; and Winston's thesis work on concept acquisition). Meta-DENDRAL was the stimulus that led to the resurgence of the machine learning area. This was due to three factors:

First, Meta-DENDRAL focused on knowledge: the learning of knowledge, not process. In addition, it took a knowledge-based approach to the learning task. Meta-DENDRAL demonstrated in a concrete way that "knowledge acquisition is itself a knowledge-based task". In the early 1970s, these were powerful ideas.

Second, it turned out that an important part of Meta-DENDRAL's learning algorithms were generalizable. The generalization—to Version Spaces—was done by Mitchell [10] in his thesis, spawned much research, and was very influential.

Third, Meta-DENDRAL had demonstrably significant results. A paper we published in 1976 [4] reported new mass spectral fragmentation rules for certain subfamilies of the chemical family called androstanes. As far as we know, it was the first paper in the literature of science that reports the discovery of new scientific knowledge (albeit of a routine kind) by a computer program (there is now another). Perhaps in the future that will be viewed as a landmark event.

We characterized Meta-DENDRAL's task as an induction problem with not many examples to start with and no teacher to pre-classify them. Some of the design considerations implied by this goal are still being addressed by the machine learning community. Here are several examples:

- *Noisy data*: We could not assume that the empirical data given to the program were complete and correct. The data were known to contain spurious (noisy) data points and to omit data points that the theory predicted should be present.
- *Multiple concepts*: Meta-DENDRAL had to learn the preconditions (LHSs) for more than one concept (mass spectral process), but did not know how many concepts needed to be learned.
- *Unclassified data*: Meta-DENDRAL was given sets of x–y points without having those points labeled as positive or negative instances of a concept. Thus we first had to generate possible explanations of each x–y point before we could consider positive and negative evidence associated with each explanation.

This is the stuff of excellent AI science. But did Meta-DENDRAL find any application in any industrial setting? No. Nor have any of the other complex machine learning procedures (however, machine induction based on algorithms of Quinlan have had a marginal success). The industry of AI applications is still awaiting the dawn of an era of knowledge engineering significantly aided by machine learning.

References

[1] D.G. Bobrow and P.J. Hayes, Artificial Intelligence—where are we? *Artif. Intell.* **25** (3) (1985) 375–415.

[2] J. Brinkley, R. Altman, B. Duncan, B.G. Buchanan and O. Jardetzky, The heuristic refinement method for the derivation of protein solution structures: validation of Cytochrome-b562, *J. Chem. Inf. Comput. Sci.* **28** (4) (1988) 194–210.

[3] B.G. Buchanan and E.A. Feigenbaum, Dendral and Meta-Dendral: their applications dimension, *Artif. Intell.* **11** (1978) 5–24.

[4] B.G. Buchanan, D.H. Smith, W.C. White, R. Gritter, E.A. Feigenbaum, J. Lederberg and C. Djerassi, Applications of artificial intelligence for chemical inference, XXII: automatic rule formation in mass spectrometry by means of the Meta-DENDRAL program, *J. Am. Chem. Soc.* **98** (1976) 6168.

[5] E.A. Feigenbaum, P. McCorduck and H.P. Nii, *The Rise of the Expert Company* (Times Books, 1988).

[6] P.D. Karp, Hypothesis formation and qualitative reasoning in molecular biology, Knowledge Systems Laboratory Tech. Report 89-52, Stanford University, Stanford, CA (1989).
[7] D.B. Lenat and E.A. Feigenbaum, On the thresholds of knowledge, in: *Proceedings IJCAI-87*, Milan, Italy (1987) 1173–1182; also: *Artif. Intell.* **47** (1991) 185–250.
[8] R.K. Lindsay, B.G. Buchanan, E.A. Feigenbaum and J. Lederberg, *Applications of Artificial Intelligence for Organic Chemistry: The DENDRAL Project* McGraw-Hill, New York, 1980).
[9] R.K. Lindsay, B.G. Buchanan, E.A. Feigenbaum and J. Lederberg, DENDRAL: a case study of the first expert system for scientific hypothesis formation, *Artif. Intell.* **60** (1993) (to appear).
[10] T.M. Mitchell, Version spaces: an approach to concept learning, Doctoral dissertation, Department of Electrical Engineering, Stanford University (1978).
[11] A. Newell, H.A. Simon and C. Shaw, Chess playing programs and the problem of complexity, E.A. Feigenbaum and J. Feldman, eds., *Computers and Thought* (McGraw-Hill, New York, 1963).
[12] E. Rich and K. Knight, *Artificial Intelligence* (McGraw-Hill, New York, 2nd ed., 1991).