

Expertise and expert systems: emulating psychological processes

Mildred L. G. Shaw and Brian R. Gaines

Knowledge Science Institute
University of Calgary
3635 Ocean View Crescent, BC, Canada V0R 1L1
{mildred, gainses}@cpsc.ucalgary.ca

Introduction

The role of personal construct psychology in computer research and applications concerned with the development of 'expert systems' and their beginnings in 'artificial intelligence' and 'cognitive science' are covered in this chapter. Research on expert systems led to the identification of the 'knowledge acquisition bottleneck,' that it was generally extremely difficult to make overt the presumed knowledge of human experts in order to program it for computers.

The history and reasons for the adoption of repertory grid methodologies and tools to overcome the knowledge acquisition bottleneck are described. Then a more fundamental analysis is made of why expert systems to date have had only limited success, and the merits of a personal construct approach to emulating human expertise in greater depth than has been achieved with existing cognitive science models are presented.

In conclusion, it is noted that the techniques developed to emulate human expertise are essentially ones for modeling and emulating any person's psychological processes, not just those of people valued by others as 'experts'. PCP-based expert systems methods and technology have wide relevance, for example, in clinical and educational research and applications.

Research on programming computers to think

The arrival of the first commercial digital computers in the 1950s led to widespread interest in the potential applications of computing. The use of the term 'giant brains' became common in the press although it was clear that the precise, logical operations of computers had little in common with the human brain. However, interest in simulating human thought processes was common among the early computer pioneers. Alan Turing, a brilliant Cambridge logician who had helped develop computers to break enemy message encryption during the second world war, wrote a paper on "Computing machinery and intelligence" for the journal *Mind* in which he considers the question "Can machines think?" (Turing, 1950). He answers it in behavioural terms, proposing what has come to be known as the 'Turing test,' that if a person communicating with the computer and with another person through the same medium (such as communicating teleprinters) cannot distinguish them correctly then the machine, for all practical purposes, can be said to think.

Research on programming computers to think became widespread, and McCarthy, Minsky, Rochester and Shannon (1955) proposed to the Rockefeller Foundation that it fund a study of *artificial intelligence* (See Jack Adams-Webber on this topic in Chapter X) to be carried out during the summer of 1956 at Dartmouth College, Hanover in the USA. The year of this proposal also saw the publication of Kelly's seminal work on personal construct psychology but, as discussed in the following section, the pioneers of artificial intelligence and cognitive science never became aware of this work. The next decade was also the era of the development of

'computer science' as a new field of study when computers were very expensive, and university and funding agency budgets were hard-pressed to supply the demand for computer facilities. In Britain the competition between those wishing to undertake research in computer science and in artificial intelligence was so intense that the UK Science Research Council commissioned Sir James Lighthill to report on the state of the art in "machine intelligence." His report (Lighthill, 1973) was damning about both the achievements and the prospects for such research and had a strong negative influence world-wide on funding for research to program computers to think (Fleck, 1982).

Embattled AI researchers focused on specific goals to develop programs that emulated human expert performance in fields of obvious practical value such as mineral exploration and medical diagnosis, and in the mid-1970s announced a number of breakthroughs in what came to be called "expert systems" (Michie, 1979). The first successful expert systems were DENDRAL (Feigenbaum et al., 1971) for reconstructing molecular structures from mass spectrometer data and MYCIN (Shortliffe, 1976) for diagnosing microbial infections from medical data. The systems were programmed as collections of 'production rules' that expressed a relationship between a premise and a conclusion such that if the conditions of the premise were satisfied then those of the conclusion could be drawn. For example, a rule from MYCIN is:-

- If
 - 1) the infection is primary-bacteremia, and
 - 2) the site of the culture is one of the sterile sites, and
 - 3) the suspected portal of entry of the organism is the gastro-intestinal tract,
- Then there is suggestive evidence (.7) that the identity of the organism is bacteroides.

Such rules are obtained from specialists in microbial infections and their application to particular data is fairly simple data processing. The rules are validated through their application to many cases and revised when they fail to give the correct diagnosis. MYCIN was designed to interact with a clinician in order to make a diagnosis and suggest therapy for a particular patient with a suspected microbial infection. It first gathers data about the patient and then uses this to make inferences about the infections and their treatment.

The success of the early expert systems attracted industrial and research attention, and a major industry developed in the early 1980s. The research objectives were then defined by one of the commercial AI pioneers, Hayes-Roth (1984), in a workshop on "AI Applications for Business" in May 1983. He enumerated some situations appropriate to expert systems, such as: the organization requires more skilled people than it can recruit or retain; job excellence requires a scope of knowledge exceeding reasonable demands on human training and continuing education.

As a modern example of the success of expert systems technology, the April and July 2000 issues of InTech Magazine published by the Instrumentation, Systems and Automation Society, has a two-part paper from Eli Lilly on the use of an expert system in its fermentation plant (Alford et al., 2000). The evaluation in 2000 corresponds well to Hayes-Roth's predictions in 1983. Within a few weeks, the expert was satisfied that the expert system reliably came to the same conclusions he would have by looking at the same data. The expert system then took over this part of the expert's job, freeing up 40 hours per month of his time for other work.

The knowledge acquisition bottleneck

Expert systems appeared at first to be a major validation of the possibility of digital computers being able to emulate human thinking, and there is continuing evidence of some successful applications. However, the industry has not grown to the extent predicted, largely because programming such systems has been very much more difficult than expected. Feigenbaum (1980), one of the pioneers of expert systems, termed this the “knowledge acquisition bottleneck”. Hayes-Roth, Waterman and Lenat (1983) in their book on Building Expert Systems noted that, since the programmer has far less knowledge of the domain than the expert, communication problems impede the process of transferring expertise into a computer program. The vocabulary initially used by the expert to talk about the domain with a novice is often inadequate for problem-solving; so that the programmer and the expert must work together to extend and refine it. One of the most difficult aspects of the programmer’s task is helping the expert to structure the specialist knowledge, to identify and formalize the expert’s concepts.

From a personal construct perspective, the task of the expert system programmer is to reconstruct the conceptual and operational framework that an expert in a domain uses to solve problems in that domain, noting that the terminology used may be highly idiosyncratic, that is, personal to the expert. However, the expert is, by definition, someone who is effective at problem-solving in the domain and, hence, her or his knowledge is valid in some practical sense. The expert’s knowledge has been acquired by some mix of processes, such as trial and error, mimicking others, reflection on personal experience, reading text written by other experts, conversations with others, and so on. That corresponds to Kelly’s notion of an individual as a *personal scientist* (Shaw, 1980) about which he asks:-

“Might not the individual man, each in his own personal way, assume more of the stature of a scientist, ever seeking to predict and control the course of events with which he is involved? Would he not have his theories, test his hypotheses, and weigh his experimental evidence?” (Kelly, 1955, p.5).

Kelly merges the notions of prediction and control into the unitary notion of *anticipation* and hence his fundamental postulate:-

“a person’s processes are psychologically channelized by the ways in which he anticipates events.”

Thus, from a personal construct perspective, the task of the expert system programmer is to model the personal construct system of the expert in operational form as a computer program such that the program is able to anticipate events in the same way as the expert. It was suggested in the early years of expert systems (Gaines and Shaw, 1980) that new methods for rule extraction made Kelly’s *repertory grid* a suitable tool for eliciting knowledge from experts. Existing computer programs for interactive elicitation of repertory grids were rapidly modified to support knowledge acquisition for expert systems (Shaw and Gaines, 1983; Boose, 1984). The approach proved successful in industrial applications (Boose, 1986), and a framework based on personal construct psychology became accepted as the foundation for developing knowledge acquisition techniques and tools (Ford et al., 1993; Gaines and Shaw, 1993).

A personal construct alternative to rule-based cognitive models

While repertory grids were widely used as knowledge acquisition tools in the 1980s and 1990s, expert systems themselves failed to achieve as much as had been expected, and a large-scale artificial intelligence industry did not materialize. Various writers have speculated on the reasons for that failure, the deepest analysis being that of Dreyfus and Dreyfus (1986). They see the problem as a manifestation of Wittgenstein's (1953) argument that the notion of human behaviour following a rule is paradoxical because, as he showed, by a suitable interpretation every course of action could be made to accord with the rule.

The pioneers of cognitive science had modeled the human mind as a repository of so-called 'production rules' (Anderson, 1983) and the designers of expert systems had followed this model in their 'knowledge representation' schemes. Dreyfus and Dreyfus argue that such representation is a major weakness and that systems based on it could never fully emulate human expert behaviour. In the artificial intelligence literature, Clancey (1989) has criticized approaches to expert system development based the assumption that expertise can be captured in overt knowledge, and has noted that all processes of behaving, including speech, problem-solving, and physical skills, are generated on the spot, not by mechanical application of scripts or rules previously stored in the brain. He argues that knowledge is a capacity to behave adaptively within an environment; it cannot be reduced to representations of behaviour or the environment.

Repertory grid-based knowledge acquisition tools had of necessity delivered knowledge in the form of rules so that it could be utilized by existing expert system knowledge representation tools. However, the analysis leading to the rules is not part of the construction process and may be regarded as an artifact of the need to use rule-based expert systems technology. Kelly developed personal construct psychology from a perspective that was consistent with that of Wittgenstein, and did not introduce rules in his psychological model. For him construing was all that was necessary to account for human behaviour, and anticipation was a by-product of construction. That is, construction intrinsically supported anticipatory processes without the storage of anticipatory 'rules' but, at a particular stage in the construction of experience, these anticipations might have a regularity that an observer could ascribe to 'rules of behaviour.' In research on the philosophy of mind the term 'supervenient on' (Kim, 1993) is used to describe a phenomenon which is a by-product of another phenomenon but not essential to it; anticipations are supervenient on constructions. However, as the person construes more experience then the supervenient anticipations might change and the observer could construe this in terms of the person 'learning new rules.' Kelly also emphasized that personal construct psychology does not need a notion of 'learning' on the part of the person or personal scientist. Construction alone is sufficient to account for the person's mental processes and behaviour, and it could also account for the models being produced by observers or psychologists. The Wittgenstein paradox presents no problems to personal construct psychologists because there is no assumption that human behaviour is rule-governed.

It is unfortunate that the development of cognitive science in the mid-1950s became dominated by those whose background was in mathematical logic. Kelly published his major work on personal construct psychology in 1955, and it could easily have become adopted as the foundation for what became called 'cognitive science' and provided foundations for artificial intelligence and expert systems. In the few years until his death in 1967 he made a number of presentations to wider audiences that might have triggered recognition of the far-ranging

implications of his work. In April 1961 he presented personal construct psychology to Luria and other members of the Moscow Psychological Society in Moscow as “a mathematical approach to psychology” (Kelly, 1969) paralleling the development of mathematical psychology in the USA by Miller, Mosteller and others (Hirst, 1988). In June 1962 he was an invited commentator at a conference on “the computer simulation of personality” held at Princeton University and stated:-

“In this connection I would like to make a plug for the psychology of personal constructs. Not only is it a system built upon the notion that scientists and human beings, alike, approach truth by erecting simulation devices—called constructs—but is a theory deliberately formulated in a language system which is based on binary elements and which does not accept the so-called subject-predicate error of the Indo-European language system.” (Kelly, 1963).

However, Kelly’s work was not recognized in the 1950s by computer and cognitive scientists.

Personal construct psychology as a foundation for modeling human expertise

The models of human thought processes derived from personal construct psychology and from mathematical logic can be contrasted through a simple example. Suppose a child has three constellations of experience:-

- child is well-behaved; mother is attentive; mother smiles;
- child is naughty; mother is attentive; mother frowns;
- child is passive; mother is inattentive.

A machine learning program might derive the rules:-

- child is well-behaved **implies** mother smiles;
- child is naughty **implies** mother frowns;
- child is passive **implies** mother is inattentive;
- child is well-behaved or naughty **implies** mother is attentive.

So that a child who is well-behaved might infer that her mother will smile, but how does the child know when she is well-behaved?

A PCP model would be that the child construes her three sets of experience in terms of the constructs: *well-behaved—naughty*, *attentive—inattentive* and *smiles—frowns*. Supervenient on the construing of the three constellations of experience are *all* the compatible anticipations, that is, those above plus:-

- mother smiles **implies** child is well-behaved;
- mother frowns **implies** child is naughty;
- plus others.

These reverse implications will be used to give meaning to the construct *well-behaved—naughty* in novel situations. To act to make the mother smile the child will chose situations where the child is well-behaved and the mother smiles. If the child wants the mother’s attention then the child may chose situations where the child is naughty, the mother frowns but also pays attention to the child. There are no ‘rules of behaviour’ but there is the choice of situations in a rather more flexible way than would be entailed through sets of rules. In addition there is an increasing

repertoire of behaviour as the child construes new situations in terms of her behaviour and the mother's smile or frown. One might say the child is 'learning' but there is no specific mechanism for learning, only for construction. One might say that the mother's smile 'reinforces' good behaviour but there is no reinforcement, only construction and choice. Kelly's view is that construction provides a complete account of human behaviour and can also model the constructs of different schools of psychology.

Now apply that model to human expertise. It models the expert as a construing agent not as a 'knowledge base' of rules. The model automatically updates as more experience is construed, that is, as the expert system attempts to solve more problems. The experience can be used in a variety of ways to solve problems and to give meaning to new situations, for example, the availability of a new drug or treatment. 'Knowledge acquisition' is intrinsic to personal construct-based expert systems and does not need to be treated as a separate phase. Expert knowledge can be transferred to the system not only through exemplary problem solving but also by commenting on the system's problem solving and by choosing problems for the system which are at the limits of the system's current capabilities. That is, experts can make their behaviour available to be mimicked and can also act as coaches commenting on performance and setting tasks, all major strategies in supporting human development.

An example repertory grid-based expert system development and application tool is WebGrid which is freely available on the World Wide Web (<http://repgrid.com/WebGrid/>). To use WebGrid an expert enters exemplary situations and, once some have been entered, can enter test cases to see how the system performs (Gaines and Shaw, 1997). If the system is incorrect, the expert can change the result and enter the corrected test case as an additional example until the system is generally correct. The system retains only the repertory grid of constructions as its knowledge base. WebGrid can produce sets of rules at any stage that characterize and explain its model of expertise at that stage, but these are not stored just produced on request and are truly supervenient on the expert's construction.

Conclusions

Expert systems were recognized as a breakthrough in artificial intelligence, in programming computers to emulate human thinking. However, they were based on a form of cognitive science that took mathematical logic as its foundations and was not well-suited to modeling the full richness of human behaviour. Personal construct psychology developed over the same time period but was not recognized by those working on artificial intelligence and cognitive science as a complete psychological system providing more effective foundations for cognitive science and expert systems. Repertory grid elicitation *was* recognized as a valuable knowledge acquisition technique with which to develop rules for expert systems, but the knowledge transferred in the form of rules was static and brittle, and did not lead to the systems being open to experience. It would be timely to adopt personal construct psychology as the foundations of cognitive science and use it to build expert systems that fully emulated the capabilities of human experts, not only to solve problems but also to be effective in dealing with new problems as they arise.

As a final comment, it is noteworthy that while the expert system community has focused on emulating the capabilities of those with expertise of value to industry, the technology developed is useful for modeling the psychological processes of any person. Kelly noted that all people may be construed as 'scientists' in their processes of modeling their worlds and validating those

models. Similarly, we construe everyone as 'experts' in being themselves and living their lives in their own way, whether or not the capabilities involved in doing this are singled out as being of special value by others. In therapy or in education for example, emulation of the person in the computer may provide a cognitive mirror (Shaw, 1980) in which an individual can view their psychological processes and come to understand them better. If there are problems arising from these processes, the increased understanding may help the individual to develop alternative constructions to address them. One by-product of research on the application of personal construct psychology to expert system development is that it has advanced our capabilities to model and emulate a person's psychological processes in a way that may be useful to that person. The motivation for the research may have been to emulate the expertise of value to industry, but the outcomes have far wider significance.

Acknowledgments

Financial assistance for this work has been made available by the Natural Sciences and Engineering Research Council of Canada.

URLs

WebGrid can be accessed at <http://repgrid.com/WebGrid/>

Related papers on WebGrid can be accessed through <http://repgrid.com/reports/>

References

- Alford, J., Cairney, C., Higgs, R., Honsowetz, M., Huynh, V., Jines, A., Keates, D. and Skelton, C. (2000) 'Real rewards from artificial intelligence.' *InTech*, , 52-55.
- Anderson, J.R. (1983) *The architecture of cognition* Cambridge, MA: Harvard University Press.
- Boose, J.H. (1984) 'Personal construct theory and the transfer of human expertise' In *Proceedings AAAI-84* pp.27-33. California: American Association for Artificial Intelligence.
- Boose, J.H. (1986) *Expertise Transfer for Expert Systems* Amsterdam: Elsevier.
- Clancey, W.J. (1989) 'Viewing knowledge bases as qualitative models.' *IEEE Expert*, 4, 9-23.
- Dreyfus, H.L. and Dreyfus, S.E. (1986) *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer* New York: Free Press.
- Feigenbaum, E. (1980) *Knowledge Engineering: the Applied Side of Artificial Intelligence*. STAN-CS-80-812 Department of Computer Science, Stanford University.
- Feigenbaum, E.A., Buchanan, B.G. and Lederberg, J. (1971) 'On generality and problem solving: a case study using the DENDRAL program' In D.Michie, (Eds) *Machine Intelligence 6* Edinburgh: Edinburgh University Press.
- Fleck, J. (1982) 'Development and establishment in artificial intelligence' In Elias, N., Martins, H. and Whitley, R., (Eds) *Scientific Establishments and Hierarchies* pp.169-217. Holland: D.Reidel.
- Ford, K.M., Bradshaw, J.M., Adams-Webber, J.R. and Agnew, N.M. (1993) 'Knowledge acquisition as a constructive modeling activity.' *International Journal of Intelligent Systems*, 8, 9-32.
- Gaines, B.R. and Shaw, M.L.G. (1980) 'New directions in the analysis and interactive elicitation of personal construct systems.' *International Journal Man-Machine Studies*, 13, 81-116.
- Gaines, B.R. and Shaw, M.L.G. (1993) 'Basing knowledge acquisition tools in personal construct psychology.' *Knowledge Engineering Review*, 8, 49-85.

- Gaines, B.R. and Shaw, M.L.G. (1997) 'Knowledge acquisition, modeling and inference through the World Wide Web.' *International Journal of Human-Computer Studies*, 46, 729-759.
- Hayes-Roth, F. (1984) 'The industrialization of knowledge engineering' In Reitman, W., (Eds) *Artificial Intelligence Applications for Business* pp.159-177. Norwood, New Jersey: Ablex.
- Hayes-Roth, F., Waterman, D.A. and Lenat, D.B., Ed. (1983) *Building Expert Systems* Reading, Massachusetts: Addison-Wesley.
- Hirst, W. (1988) *The Making of Cognitive Science*. Cambridge, UK: Cambridge University Press.
- Kelly, G.A. (1955) *The Psychology of Personal Constructs* New York: Norton.
- Kelly, G.A. (1963) 'Aldous: the personable computer' In Tomkins, S.S. and Messick, S., (Eds) *Computer Simulation of Personality* pp.221-229. New York: Wiley.
- Kelly, G.A. (1969) 'A mathematical approach to psychology' In Maher, B., (Eds) *Clinical Psychology and Personality: The Selected Papers of George Kelly* pp.94-113. New York: Wiley.
- Kim, J. (1993) *Supervenience and Mind: Selected Philosophical Essays* Cambridge: Cambridge University Press.
- Lighthill, J. (1973) 'Artificial intelligence: a general survey' In *Artificial Intelligence: a paper symposium* UK: Science Research Council.
- McCarthy, J., Minsky, M.L., Rochester, N. and Shannon, C.E. (1955) A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence. Dartmouth College. <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>.
- Michie, D., Ed. (1979) *Expert Systems in the Micro Electronic Age* Edinburgh: Edinburgh University Press.
- Shaw, M.L.G. (1980) *On Becoming A Personal Scientist: Interactive Computer Elicitation of Personal Models Of The World* London: Academic Press.
- Shaw, M.L.G. and Gaines, B.R. (1983) 'A computer aid to knowledge engineering' In *Proceedings of British Computer Society Conference on Expert Systems* pp.263-271. Cambridge: British Computer Society.
- Shortliffe, E.H. (1976) *Computer-Based Medical Consultations: MYCIN* New York: Elsevier.
- Turing, A. (1950) 'Computing machinery and intelligence.' *Mind*, 59, 433-460.
- Wittgenstein, L. (1953) *Philosophical Investigations* Oxford: Blackwell.