

THE LOGIC PROGRAMMING DEBATE

by

Alan L. Tyree*

I. Introduction

One of the most popular participation sports amongst lawyers interested in artificial intelligence must surely be Imperial College bashing. It is easy to play, quite rewarding and of near universal appeal. I have indulged in the sport myself and can attest to a warm inner glow¹. And it is ever so much easier than the hard work necessary to build and explore actual AI systems.

I do not wish to argue that the Imperial College approach is correct. I do not believe that it is. What I do wish to argue is that the arguments against it are incorrect or, at best, premature. The basis of my objection is that the performance, or non-performance, of a legal expert system is an empirical fact. Most of the arguments presented by the anti-Imperial College approach are a priori arguments. It is easy to illustrate that the a priori arguments have been irrelevant in other areas.

II. Making a chess machine

Suppose we consider the problem of building a chess playing machine. We must pretend that we know nothing of the machines that have actually been built over the last thirty years. The question is: what methods and models should we use to build the machine?

Very plausible arguments suggest that we need "deep models" of chess. It must be impossible to build a machine that plays at even a "good" level of chess without incorporating the understanding of generations of chess players. Chess is a subtle game, and it calls for a psychological understanding of ones opponent. Although chess, like law, has rules, good chess playing cannot be reduced to a set of rules. For this reason we expect a good chess player to make moves which are different from those made in identical positions in different games. In other words, the "context" is important. Good chess playing is intuitive, not a deductive mechanical process.

There are no doubt other "philosophical" arguments which point to the need to develop very sophisticated knowledge representation methods and even more sophisticated models of "chess reasoning". In other words, most of

* Landerer Professor of Information Technology and Law, University of Sydney.

¹ Tyree A, Greenleaf G and Mowbray A, "Expert Systems in Law: the DataLex Project" Proceedings of the First International Conference in Artificial Intelligence and the Law, Boston, May 1987.

the comments made about the ICG work apply to the construction of a chess machine.

The trouble with this line of argument is, as we now know, that it is incorrect when applied to the construction of chess machines. Machines which rely primarily on "brute force" methods have proved to be very competitive with the machines which rely primarily on "smart" methods. The yearly machine chess competition has produced no clear answer to which basic approach will ultimately be successful. It may be that both will be.

III. The lesson for the "rule" debate

These thoughts were, of course, prompted by the major article in the last issue of the *Journal of Law and Information Science*,² the comments by Brown in his report of the Third International Conference on Artificial Intelligence and Law,³ and, to a lesser extent, the note "Implications of a Logical Paradox for Computer-Dispensed Justice"⁴. The first two appear to be very concerned that a certain approach to legal expert systems cannot succeed. Yet, these arguments seem very similar to me to the arguments presented above, ie, they are arguments from an a priori perspective about how, or even if, we can achieve a certain result.

We must begin by acknowledging that there are no machines which can reason in any domain at a level comparable with a human domain expert. Chess machines are probably the most advanced in this regard, yet they are regularly beaten by the very best chess players. On the other hand, we have to acknowledge that there are some very good machines in a number of domains. Chess machines can regularly beat all but the very best chess players. Medical systems can regularly give more accurate diagnosis than all but the most expert of the human doctors⁵.

What is the current situation in law? Unfortunately, we don't have a very good idea. There is nothing like the yearly chess competition to measure performance of legal expert systems. There have not, to my knowledge, been any major comparisons of legal machines with advice given by the general legal practitioner. I suspect, but cannot prove, that there are a number of legal machines operating in specialised areas of law which, while not approaching the performance of an expert in the area, will give advice which is better than that of the general practitioner.

If that is correct, then the critics of rule based systems must at least pause and consider the relevance of their arguments, particularly when those

2 Moles R N, "Logic Programming - An Assessment of its Potential for Artificial Intelligence" (1991) 2 *JLIS* 137.

3 Brown D, "The Third International Conference on Artificial Intelligence and the Law" (1991) 2 *JLIS* 233.

4 Fulda JS, "Implications of a Logical Paradox for Computer-Dispensed Justice" (1991) 2 *JLIS* 230.

5 Buchanan B G and Shortliffe E H, "Rule-Based Expert Systems", Addison-Wesley, 1984.

arguments seem to call for a cessation of research into rule-based legal machines. After all, most (not all) current legal expert systems function on a rule-based knowledge base. In spite of a number of arguments for a more "jurisprudential" approach to the building of legal expert systems, no such machines have been built.

If the argument is that rule-based mechanisms can never achieve the level of expertise of a human, then I have to say that the arguments are not convincing. The theoretical model of rule based reasoning is the production rule system explained by Post in his 1943 paper⁶. Post showed that production rule models were theoretically equivalent to Turing Machines. Since it is generally accepted that a Turing Machine can (theoretically) reproduce any computer, it follows that if it is possible at all to produce a machine which achieves human expert performance, there is no theoretical reason why that machine could not be rule-based.

Maybe the arguments are intended to prove that no such machine can ever be built. In that case, my answer is that we must know a great deal more about the limits of human reasoning and the limits of machine reasoning before such arguments could be convincing⁷.

IV. Some particular observations

A. Programming language

Even if the critics are correct about rule based systems, it would be a pity if legal AI systems reject experimentation with logic programming. Although logic programming statements are (generally) in the form of rules, the rules need not be directly related to the subject matter. The rules of a program may be rules about how to manipulate complex data structures or other entities. In other words, logic programming is intended to be a general purpose programming method. The hope of logic programming research is to produce programming languages which are easier to use and debug.

B. Isomorphism

Both of Moles and Brown attack the notion of "isomorphism". Moles is particularly concerned about the amount of translation and selection required before legal material can be transformed into executable code. There is no doubt that each operation may be expected to lose some desirable aspect of isomorphism, but surely the problem there is simply a technical one. It is hard to understand an objection to isomorphism as a desirable goal even though it may be some time before it can be perfectly realised.

Moles and Brown also both criticise the ICG for failure to use legal experts in the construction of their systems. Again, I and others have also

⁶ Post E, "Formal reductions of the general combinatorial problem" 65 *American Journal of Mathematics* 197.

⁷ If that is the intention, then the argument was better made and answered by Turing in Turing A, "Computing Machinery and Intelligence", (1950) 59 *Mind* 433.

made this observation. Yet it is possible to feel some sympathy with the ICG. Bench-Capon says that the intention of the isomorphism process is that "the rule base should reflect the sources from which it has been developed". This seems to me to be a very desirable goal, no matter how imperfect the procedures. Moles says of this intention: "It is clear that material such as this could only have been written by people with little or no appreciation of the nature of legal interpretation."

The ICG might be forgiven for believing that people with legal training are not going to be too helpful if the only advice is "give up, the problem is beyond you".

C. Results/Methods

Part of the problem with the debate is captured in Brown's puzzled remarks concerning the possibility of reproducing the results of legal reasoning without modelling the method. He suggests that this is not possible, or at least that it would be most unusual. The FINDER system, utilising the PANNDA statistical approach to examples, is very far from any model of legal reasoning that I know⁸. I attach a sample output from the system and note the following: it is "correct" in that it reproduced a significant amount of the reasoning of the English Court of Appeal. Indeed, the opinion is better than we ordinarily get from first year students who are familiar with the subject.

On the other hand, it does not enter into the debate as to whether *Bridges v Hawkesworth* was correctly decided. Does that disqualify FINDER as a legal expert system? This leads to my final comment.

D. How good does a legal expert system need to be?

The answer to the question must depend upon use and cost. FINDER is cheap to build and cheap to maintain. It gives good advice in its (narrow) area of expertise, although it is not as good as the best human advice from experts in the same area. One can imagine that Parker may have been quite reassured to receive the FINDER advice, at least if the advice was priced at less than \$5. He may have then felt more confident in seeking more expensive human advice.

I would like to argue further that FINDER could, provided appropriate procedural safeguards were in place, be used to actually decide disputes. I am definitely NOT arguing that FINDER performs at human level. But, and here

⁸ For a complete description of the method, see Tyree A, Greenleaf G and Mowbray A, "Legal Reasoning: the Problem of Precedent" Proceedings of the Joint AI Conference, Sydney, November 1987. A version of this paper has been included in Gero J and Stanton R (eds), *Artificial Intelligence Developments and Applications*, North-Holland, Amsterdam, 1988, pp231-247. See also the development of SHYSTER by James Popple which develops and improves the basic FINDER algorithm; see Popple, J "SHYSTER and the Authorization of Copyright Infringement", Technical Report TR-CS-92-08, Department of Computer Science, ANU and the references cited therein.

is the crux of the entire matter, that is simply not the issue. Most disputes within the area of FINDER's expertise are not solved by human judges, they are solved by brute force as a result of the high costs of human justice.

So, the choice may not be a choice between human justice and machine justice, but a choice between machine justice and no orderly justice at all. Do we dismiss machine justice as being inhuman in such a case? Do we seriously argue that it must be as sophisticated as human justice? Do we seriously argue that a legal expert system is of no value if it does not perform as well as a human being?

V. Back to Imperial College

The important question then is whether the Imperial College approach can build machines which are good enough at a price which we can afford. I do not know the answer to this question. It is an empirical question and one that requires the assistance of the Imperial College Group. Let's see what kind of problems their system can solve. Let's see what kind of problems give it a nervous breakdown. Does their system recognise when it has a problem that requires more sophisticated attention?⁹

I do not wish it to seem that I think the "big" questions are not interesting. They obviously are, although I personally find that little has been added to our understanding of the "big" question since Turing's 1950 paper¹⁰. It is interesting to note that Turing's suggestion there of how to actually approach the construction of an "intelligent" machine has not so far proved productive. It would be a pity if a priori arguments had prevented more limited and, according to some, more prosaic research in the search for machine intelligence. And, although the "big" questions are undoubtedly interesting, they should not be used to discredit the incremental work which has been the only work to date which has produced useful functioning systems. In my opinion, these functioning machines will teach us more about the way forward than all of the a priori arguments that lawyers can muster.

⁹ This is an important issue in building useful legal systems. See Tyree A, Greenleaf G and Mowbray A, "Generating Legal Arguments", (1989) 2 *Journal of Knowledge Based Systems* 46.

¹⁰ Turing A, "Computing Machinery and Intelligence", (1950) 59 *Mind* 433.