LOGIC PROGRAMMING - AN ASSESSMENT OF ITS POTENTIAL FOR ARTIFICIAL INTELLIGENCE APPLICATIONS IN LAW

Robert N. Moles[1]
The Australian National University
Faculty of Law

ABSTRACT

In the area of artificial intelligence and law, a good deal of work has been undertaken with a view to replicating the process of legal decision making using computer programs. This article evaluates the work of one particularly significant group in this area - the Logic Programming Dept at the Imperial College of Science and Technology in London. The techniques involved in "logic programming", and the issue of "isomorphism" are discussed with a view to articulating the assumptions of members of this group concerning the nature of the legal process.

The remainder of the article explains why it is that the assumptions of this group concerning the rule-based nature of law, and the search for the unambiguous meaning of words, are mistaken. Their assumptions are based on a lack of awareness of the dynamics of the legal order and the factors associated with legal interpretation. These erroneous assumptions should have been brought to light in the early stages of their research program, but, as the members of the team explain, they did not consult with lawyers at that stage. They continue to claim that they have either solved (or are solving) significant theoretical problems, whilst at the same time they continue to abandon significant projects without attempting to implement them.

The issues raised have important implications for the funding of research in this area, and demonstrates the need for real interdisciplinary work which will include lawyers and legal theorists as part of the team.

Law is based on rules

Rules may be applied deductively - mechanically

Words have an "unambiguous" meaning

Rules as "atomistic" entities

No need for legal expertise

---

1 ACII. LL.B. (Belfast) PhD. (Edinburgh). Senior Lecturer in Law at the Australian National University.

1 - INTRODUCTION:

The main focus of interest in this article is on a particular technique which has been used in the area of artificial intelligence and law (AI and law). This is the area in which people use computers to try and determine the correct legal outcome regarding some problem with which they are concerned. Some people in this area focus more on working with statutes, others try to work with legal cases (precedents) and yet others try to combine the two approaches. The particular technique which we will be concerned with here has tended to focus mainly (but not exclusively) on working with statutes and regulations. The technique involves the use of "logic programming" and its use has, not surprisingly, been most keenly advocated by the Logic Programming Group at the Imperial College of Science and Technology in London. This group is part of the Department of Computer Science at Imperial College, and for convenience I will refer to it hereafter as "the ICG" - the Imperial College Group. The members of this group have a large share of the responsibility for the development of logic programming in the domain of legal applications, and they have been closely involved with a number of other teams. As Richard Susskind (one of the few legal theorists to have taken an interest in AI and law) points out, there are a number of projects using the 5th Generation logic programming tool PROLOG to represent various legal domains. This includes the running of translations of legislation in a general expert system shell. A shell is a system from which all the "particular" knowledge has been removed, but which retains the basic inference sub-system. Non-specialists can then purchase the shell, and insert the knowledge relevant to their particular project. In this way they may be able to produce an expert system without having to program the inference sub-system. Andrzej Kowalski's article in this edition of the Journal discusses the way in which his students used an expert system shell to assist them in developing their own expert systems. We will also see, later in this article, that the ICG used expert system shells in the development of their work. Insofar as it leads them to accept the separation between the knowledge representation, and the inference sub-system which applies that knowledge, it may have unfortunate consequences for their view of law.

Susskind goes on to point out that the best known of the PROLOG projects have been developed by the team at Imperial College.[2] The significance of the ICG is reflected in the substantial number of articles which they have generated in the emerging field of AI and law. Members of the ICG have also been very much involved in the three International Conferences on Artificial Intelligence and Law, at Boston (1987) Vancouver (1989) and Oxford (1991).[3]

The key figures involved in the initiative to bring logic programming techniques to the legal domain are Robert Kowalski (head of the Logic Programming Department at Imperial College) and Marek Sergot. In fact, they take the view that "within the Logic Programming

---

2 Richard Susskind *Expert Systems in Law* (1987) Oxford University Press p 17. He discusses the use of "shells" at pp 155-57.

3 At the Oxford conference, the "International Association for AI and Law" was formed.

Group at Imperial College the domain of law is seen as a primary source of applications."[4] It is important then to assess the value of this contribution to developments in what is becoming an increasingly important area. Before we attempt to do this, we should perhaps look first at the technique of logic programming and the way in which it works.

2 - THE IMPERIAL COLLEGE GROUP - LOGIC PROGRAMMING

Virtually all of the ICG work has been implemented in the programming language PROLOG. As its name suggests, it is a language particularly suitable for logic programming. In the discussions of their work, the ICG refer to the fact that they have used "Horn Clauses" (I will explain what they are in what follows) extended to allow for negation. One of the leading American researchers in this field, with whom the ICG have a close association is Thorne McCarty, and he has utilised a similar technique. As the ICG explain, the key to their approach is the representation of knowledge by means of definite Horn Clauses.[5] Each clause has exactly one conclusion. It may or may not have a condition (or conditions) upon which the validity of the conclusion depends. Each conclusion (or condition) is regarded by the ICG as an atomic relationship among individuals. For example, the logic programmer might write:

A:- B1...Bn

An alternative form would be: A if B1...Bn

A represents the conclusion while B1...Bn represent a specified number of conditions upon which the validity of that conclusion depends. This means that the conclusion (or goal) - A - is valid if the conditions (or sub-goals) B1...Bn are valid. The Horn clauses (as this form of representation is called) are queried or invoked by means of conjunctions of these atomic relationships. In other words, if the query ?B1...Bn is entered, the computer searches for that pattern of conditions (or sub-goals) in its knowledge base. One can readily see that each condition may, in its turn, be seen as a conclusion, the validity of which depends on yet further conditions. To take our example:

(Conclusion) A: B1...Bn (Conditions)

If we look at the condition B1, it may be the case that before we can determine the validity of that condition, we may have to establish the validity of further conditions C1...Cn upon which the condition B1 depends. This could be represented diagrammatically as follows - the apex of each triangle represents a conclusion, the base represents the conditions upon which the validity of that conclusion depends:

---

4 Robert Kowalski and Marek Sergot "The Use of Logical Models in Legal Problem Solving" (1990) *Ratio Juris* Vol 3 No 2 pp 201-218 at p 201.

5 M.J. Sergot, F. Sadri, R.A. Kowalski, F. Kriwaczek, P. Hammond and H.T. Cory "The British Nationality Act As A Logic Program" (1986) *Communications of the ACM* Vol 29 No 5 pp 370-386 at p 372.

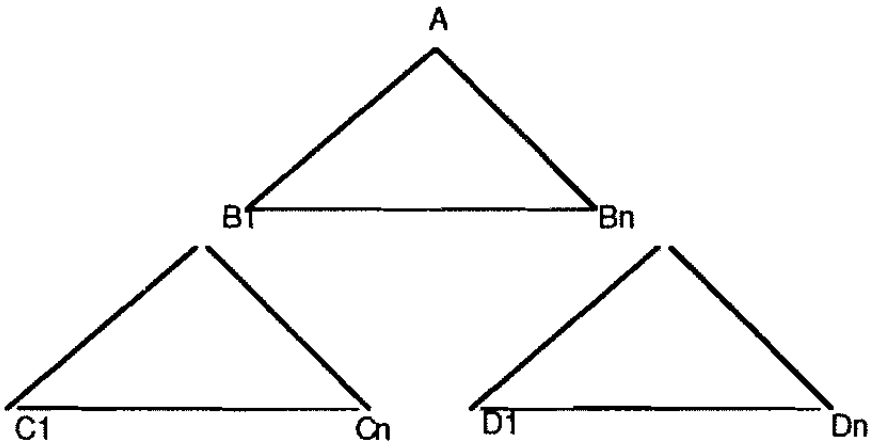

Clearly, we could quickly develop a complex network of these relationships. Yet for all this, there is an important limitation inherent in this approach. As we shall see in what follows, the ICG accept that, whatever its complexity, the approach being adopted here is a purely mechanistic form of "reasoning". What they refer to as the "blind, mechanical application of rules."[6] Thorne McCarty has expressed, as clearly as anyone, the implications which this approach has for the process of legal decision-making:

> Legal analysis, in its simplest form, would then be a process of applying the 'law' to the 'facts'. Put this way, the paradigm seems to be an ideal candidate for an artificial intelligence approach: the 'facts' would be represented in a lower-level semantic network, perhaps; the 'law' would be represented in a higher-level semantic description; and the process of legal analysis would be represented by a pattern-matching routine.[7]

As the ICG team explain:

> The Horn clause form of logic is the basis of the computational paradigm, logic programming, and of the logic programming

---

6 Several members of the ICG said in introducing a recent article, "Later in the article, we will describe how our translations of the Act can be executed as a program by an augmented PROLOG system, so that consequences of the Act can be determined mechanically." M.J. Sergot, F. Sadri, R.A. Kowalski, F. Kriwaczek, P. Hammond and H.T. Cory "The British Nationality Act As A Logic Program" (1986) *Communications of the ACM* Vol 29 No 5 pp 370-386 at p 370. See also Sergot, Kowalski "The use of Logical Models in Legal Problem Solving" (1990) *Ratio Juris* Vol 3 No 2 pp 201-218 at p 205 "Nevertheless, in the day to day practice of law, there are many situations where routine tasks do have to be performed, and where rules and regulations do have to be applied mechanically." "What we have described is a program which operates by blind, mechanical application of its rules." Ibid.

7 L.T. McCarty "Some Requirements for a Computer-based Legal Consultant" (1980) Report LRP-TR-8, Laboratory for Computer Science Research, Rutgers University pp 2-3.

> language Prolog. Every set of definite Horn clauses is a Prolog program.[8]

What then have they been trying to do with Horn clauses? As they sum it up in one of their recent articles:

> Our work has concentrated almost exclusively on providing systems which, presented with a description of some real or hypothetical state of affairs, can be used to determine whether some specific legal consequences would seem to follow.[9]

This would appear to be the standard task of legal analysis - to size up a particular situation and assess what legal consequences might result. If computer programs could be utilised in this way, then, of course, they (and the computer programmers) could dramatically change the nature of the legal enterprise. In order to evaluate this possibility, we should first look to see what the ICG think is involved in this process, and assess the adequacy of their analysis and methods.

The ICG state that they do not wish to get involved with the more difficult problem of simulating legal reasoning:

> .. we should stress once again that we have not addressed the broad and much more difficult problem of simulating legal reasoning. Rather, we have concentrated on the limited objective of implementing rules and regulations with the purpose of applying them mechanically to individual cases.[10] [emphasis added]

They obviously see no problems with regard to the "implementing" "applying" or otherwise giving effect to rules and regulations, without having to engage in legal reasoning. If it were the case that judges regard the process of legal reasoning as the means by which they reach their results, then one could well imagine that they would be attracted to a system which offers an alternative to that difficult (and often contentious) process. The question we have to determine is whether the ICG approach really offers us a viable alternative? In order to answer that question, we need to know a little more about the methodology involved.

We are told that the way in which the ICG systems operate is that:

---

8 "The British Nationality Act as a Logic Program" (1986) *Communications of the ACM* Vol 29 No 5 pp 370-386 at p 372.

9 Sergot, Kowalski "The use of Logical Models in Legal Problem Solving" (1990) *Ratio Juris* Vol 3 No 2 pp 201-218 at p 203.

10 "The British Nationality Act as a Logic Program" (1986) *Communications of the ACM* Vol 29 No 5 pp 370-386 at p 372. Although it appears that they are not always consistent in this regard "we believe that the formalization of legislation and legal reasoning offers potential contributions to computing technology itself." Ibid at p 371. [emphasis added]

> ..the law is modelled by a set of logic sentences (an axiomatic theory) which represent some chosen unambiguous interpretation (sic) of the selected legal sources.[11] [emphasis added]

That is, the ICG construct "logical models which represent statutes or sets of regulations".[12] Once this is done, "an automated theorem prover is used to derive useful consequences [from] the representation".[13]

> Because the formalization of the British Nationality Act is an axiomatic theory, any logical consequence of the axiomatization can, in theory, be derived by means of a complete mechanical theorem prover.[14]

Usually this will involve the application of the law to the "real or hypothetical state of affairs" already referred to. Essentially, "the law is 'applied' to facts by deduction".[15] In what areas then have the ICG attempted to develop their "atomistic" "deductive" "mechanistic" approach? Their main areas of activity have been in connection with the British Nationality Act (U.K.) 1981 and their involvement in the Alvey project. The latter has been a major initiative in the United Kingdom, involving the government, industry and the universities, which was established to develop the new generation of clever computers.[16]

If measured by influence and output, the contribution of this team has clearly been substantial. If other criteria of significance are adopted, such as theoretical adequacy, or practical utility, then the conclusion must be that their contribution has been much less valuable. Let us look first in a little more detail at their approach to the British Nationality Act.

3 - THE WORK OF THE ICG ON THE BRITISH NATIONALITY ACT (U.K.)

We have already seen that the basic tool used by the ICG is the use of Horn clauses as PROLOG programs. As they explained their approach:

---

11 Sergot, "The Representation of Law in Computer Programs" in *Knowledge-Based Systems and Legal Applications* (1990) (ed. Bench-Capon) Academic Press pp 3-67 at p 35. See also Sergot, Kowalski "The use of Logical Models in Legal Problem Solving" (1990) *Ratio Juris* Vol 3 No 2 pp 201-218 at p 201.

12 See also Sergot, Kowalski "The use of Logical Models in Legal Problem Solving" (1990) *Ratio Juris* Vol 3 No 2 pp 201-218 at p 204.

13 Sergot, "The Representation of Law in Computer Programs" in *Knowledge Based Systems and Legal Applications* (1990) (ed Bench-Capon) Academic Press pp 3-67 at p 35.

14 "The British Nationality Act as a Logic Program" (1986) *Communications of the ACM* Vol 29 No 5 pp 370-386 at p 376.

15 Sergot "The Representation of Law in Computer Programs" in *Knowledge Based Systems and Legal Applications* (1990) (ed Bench-Capon) Academic Press pp 3-67 at p 35.

16 *A Programme for Advanced Information Technology* (1982) HMSO. See also Bench-Capon, Robinson, Routen, Sergot "Logic Programming for Large Scale Applications in Law: A formalisation of Supplementary Benefits Legislation" in *Proceedings of the First International Conference on Artificial Intelligence and Law* (1987) ACM Press pp 190-198 at p 191.

> Most of the [British Nationality Act] was translated into Horn clause logic, extended to allow for negation, during the two months of July and August 1983 by a student, without any expert legal assistance.[17]

It was implemented in APES, which is really an "augmented PROLOG system"[18] and which allows the user to insert data in response to queries by the system. Thus APES utilizes the same deduction mechanism used in PROLOG - that of goal-directed pattern-matching.

> The answers it produces [when the Act is embedded in it] are logical consequences of the rules contained in the formalization of the Act together with the information obtained from the user.[19]

APES can also provide information to explain why a given query was generated and how a given solution has been obtained. The ICG feel that an "obvious application [of the formalisation] is to determine whether, in a particular given instance, a particular given individual is or is not a British citizen."[20] We can see from their work, on what has been for them a major project, a number of factors which will prove to be important to our further analysis. The main concern at this stage is to clarify their approach to the question of "knowledge acquisition". That is to say, how "knowledge" or "understanding" is derived from what they regard as the basic textual sources. That they do not regard this as a significant problem is made clear when they say that the basis of their approach to law is as follows:

> The formalisation of legislation by means of rules has almost all the characteristics of an expert system. It differs, however, in one important respect. In a classical expert system, before knowledge can be formalised, it has to be elicited from the subconscious of an expert. Eliciting this knowledge is generally regarded as the main bottle-neck in the construction of expert systems. It is entirely absent, however, in the case of legislation which is already formulated and written down. Thus the use of expert system techniques for representing legislation has virtually all the advantages of expert systems without the attendant disadvantages of eliciting the knowledge.[21] [emphasis added]

---

17 Sergot, Cory, Hammond, Kowalski, Kriwaczek and Sadri, "Formalisation of the British Nationality Act" (1986) 2 *Yearbook of Law, Computers and Technology* pp 40-52 at p 41.

18 Sergot, Sadri, Kowalski, Kriwaczek, Hammond and Cory "The British Nationality Act as a Logic Program" (1986) Vol 29 Number 5 *Communications of The ACM* pp 370-386 at p 377.

19 Ibid at p 377. The standard legal convention of referring to a statute as an "Act" (with the initial capital letter) is often ignored by these writers. I have corrected this in direct quotations without further note.

20 Sergot, Cory, Hammond, Kowalski, Kriwaczek and Sadri "Formalisation of the British Nationality Act" (1986) 2 *Yearbook of Law, Computers and Technology* pp 40-52 at p 46.

21 Ibid at p 49.

In other words, formalising a piece of legislation by way of logic programming does not involve us in any problems of knowledge acquisition. The problem is "entirely absent" in the case of legislation, because it is "already formulated and written down". Clearly these researchers do not distinguish between the writing (which is the legislation) and the meaning of that writing. When they say that the legislation is "already formulated", they are only telling their readers half the story - their computers cannot deal with the writing of the statute in its existing form. Feeding in the whole text of the Act would simply give them a data base, like Lexis. The exercise in which they are engaged requires them to reformulate the legislation so that it can be dealt with by their computers as part of what they call an expert system. While they claim that the reformulated material has "the same structure" as the original, this is clearly not so, as I shall explain in the next section. What they must mean is that the reformulation of the source materials does not affect their meaning. When we see what this reformulation can involve, it is clear that we must treat this claim with considerable scepticism.

4 - THE ALVEY PROGRAMME - ISOMORPHISM.

The Alvey Report suggested that funding of the order of £350m be provided to enable the development of a new generation of computers. This led to the setting up of a number of "Demonstrator Projects" to serve as feasibility studies in a number of areas. The Alvey - DHSS (Department of Health and Social Services) Demonstrator Project ran from 1984 to 1989.

> [It was] supported by the Alvey Directorate of the UK Department of Trade and Industry and the UK Science and Engineering Research Council. The project collaborators were ICL, Logica, Imperial College, and the universities of Lancaster, Liverpool and Surrey.[22]

Bench-Capon, from whose work the following is taken, was originally a member of the ICG before moving to his present position at Liverpool University.[23] Indeed, the present appearance of diversity and range of interest in terms of the contributors to the discussions, reflects an underlying narrowness of perspective.[24]

---

22 *Knowledge-Based Systems and Legal Applications* (1990) (ed. Bench-Capon) Academic Press pp vii-viii. See also pp 69-70 and *A Programme for Advanced Information Technology - The Report of the Alvey Committee* (1982) HMSO.

23 See for example, "Support for Policy Makers: Formulating Legislation with the Aid of Logical Models" in *Proceedings of the First International Conference on Artificial Intelligence and Law* (1987) Boston pp 181-189 which was written whilst Bench-Capon was a member of the ICG.

24 Indeed, we see in the recent article by Kowalski and Sergot that of the 23 citations, 13 are either to other articles by those authors, or to articles by other members of the ICG. Five others refer to MSc theses completed within the Dept of Computer Science at Imperial College. In other words, 18 of the 23 references are "in-house". Kowalski and Sergot,

In his recent work, Bench-Capon (with his co-author Frans Coenen) has laid great emphasis on the need for "isomorphism". In so far as this suggests that there is to be a similarity of structure between the knowledge base and the original source materials, it would appear to be a good thing. However, when we see the number and type of operations which text can be put through and still "retain the same structure", we might wish to think again.

Isomorphism:

An essential aspect of this approach is its claim to have a similarity of structure between the knowledge base and the source material. This similarity is referred to as "isomorphism":

> knowledge based systems in the legal domain will be more disciplined and better engineered if the structure of the knowledge base reflects the structure of the source documents from which the knowledge base is derived.[25]

We are told that the isomorphic nature of the computerised knowledge-base means that it represents groupings of concepts found in legislation and other documents. Bench-Capon provides us with a case-study to explain the procedure which is involved in developing such an isomorphic knowledge base. The project he discusses is one which was developed jointly with British Coal to help them to deal with work related injury claims. I include extracts from his discussion here, only to illustrate the complex changes which are involved in the transition from "documentary sources" to "knowledge base". The reader should bear in mind the question, as we follow this discussion, whether the constant references to isomorphism might not conceal more than they reveal about this process? Bench-Capon states that we start with

> The Mines and Quarries Act 1954
> The H & SE Mines (Safety of Exit) Regulations 1988
> The Claims inspectors Manual (1990)
> A number of significant judgements.[26]

This means that the documentary sources comprise legislation, delegated legislation, administrative guide-lines and precedents.

We then deal with the text in the following stages:

"The Use of Logical Models in Legal Problem Solving" (1990) *Ratio Juris* Vol 3 No 2 pp 201-218.

25 Bench-Capon and Coenen "Exploiting Isomorphism: Development of a KBS to Support British Coal Insurance Claims" in *Proceedings of the Third International Conference on AI and Law* (1991) Oxford ACM Press pp 62-68 at p 62.

26 Ibid at pp 64-65. The following details are taken from pp 64-66 of the same article.

1 The source documents are translated into a form suitable for use in KANT (a computer program, the name for which is an acronym derived from the fact that it is a "Knowledge ANalysis Tool).
2 The documents are browsed and "relevant sections identified". The criterion of relevance is, unfortunately, not explained.
3 The documents are "copied" into a single source "structure" for further analysis. "This structure can be thought of as precis (sic) of the source material". It is a hierarchical tree, which represents groupings of data.
4 The structure of the preceding section is then analysed - to identify Entity, Attribute and Value triples, which are then stored in another structure. He emphasizes at this point that "isomorphism with the source is maintained".

It is amazing to think that text which has been translated, cut up into bits, precised, further analyzed into EAVs and stored in another structure, can still be thought to represent the same structure as the original. However, we are not finished yet.

5 The EAV structure only provides the basis from which "Class Hierarchies" and "Rule Bases" in the "Kant intermediate representation" are formed.
6 Then another program KBB (Knowledge Base Builder) takes the Class Hierarchy and Rule Base and compiles them into rules and objects in the target representation language.
7 We then have a knowledge base in Conjunctive Normal Form.

It may surprise the reader to learn that the intention of all this is that "the rule base should reflect the sources from which it has been developed". One might be tempted to think that if it did, it would be a miracle.

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. We do have some reason to believe that this may well have been the case. We have already seen how the ICG team thought that they could translate the BNA into Horn Clauses without the need for any legal expertise. In the recently published collection of essays edited by Bench-Capon he explains that:

> this book is intended to make available the experience gained, and some of the lessons learnt, from a substantial group of related experiments in the field of knowledge based systems and law performed in the Alvey-DHSS Demonstrator project.[27]

The bibliographic notes which Bench-Capon includes with the essays, indicates the qualifications and experiences of those involved. As to the 10 authors, their qualifications would appear to be as follows:

Bench-Capon:	philosophy, economics,
E.S. Cordingley:	physics, social research
J. Forder:	computer consultant

27 *Knowledge-Based Systems and Legal Applications* (1990) (ed. Bench-Capon) Academic Press p vii.

D. Frohlich:	psychology
N. Gilbert:	sociology
P. Luff:	computer science
C. Portman:	engineering
M. Sergot:	mathematics
G. Storrs:	cognitive psychology
A. Taylor:	philosophy, information systems.[28]

Although the book is called Knowledge-Based Systems and Legal Applications, it would appear from this that there was very little emphasis on involving people with legal expertise. Indeed, we know that Bench-Capon thinks it to be a virtue of isomorphism that people with no previous experience in building legal knowledge bases can, with this methodology, be set to work fairly quickly; "two out of the three developers had not previously built a legal KBS."[29] In his discussion of the interdisciplinary nature of the project, Bench-Capon does point to the "unusually wide spread of educational backgrounds" but then goes on to say:

> One, perhaps surprising, feature of the composition of the team was that there were so few members with a legal background - although there were project members with first degrees in law there were no practising lawyers or academic law departments involved.[30] [emphasis added]

He does continue to say that the justification for this was that "lawyers only rarely get involved in social security" and that adjudication and the first level of appeal are carried out by lay people.[31] However, given the great emphasis which was placed on the size and length of the project (65 researchers, 30 or more at any given time, working over some 5 years) and the fact that this was a feasibility study to explore "the application of KBS [knowledge based systems] to law"[32] - the explanation provided for the lack of involvement of those with legal expertise is far from convincing. This is especially so when we consider that the goal specified by the ICG was not simply to change the order of the wording of legislation, but to come up with something which has meaning - something which would enable them to produce legal outcomes, given a certain state of affairs.

---

28 Ibid pp xi-xiv.

29 Bench-Capon and Coenen "Exploiting Isomorphism: Development of a KBS to Support British Coal Insurance Claims" in *Proceedings of the Third International Conference on AI and Law* (1991) Oxford ACM Press pp 62-68 at p 66.

30 *Knowledge-Based Systems and Legal Applications* (1990) (ed. Bench-Capon) Academic Press p 70.

31 Ibid p 71.

32 Ibid p 69. It is clear from this article generally that there is a similarity to the approach taken by Kowalski and Sergot. The emphasis throughout their work on the legal significance of what they are doing, followed by the acknowledgment that they neither had nor sought any assistance from legal experts. That although they appear to be dealing with "AI and law", they are really only structuring a "layperson's" understanding of the issues.

5 - PROBLEMS FROM THE LAWYER'S PERSPECTIVE

Given that many of the people working in this area have had relatively little involvement with lawyers, it seems to me that it would be useful to indicate the aspects of their work which would appear unconvincing to those with an interest in legal interpretation. I should like to add, in passing, that I regard these problems as being of general application, and that they would apply whether we were dealing with interpretation with regard to legal texts, or with regard to historical or religious documents. In other words, it appears to me that the ICG have little awareness of the issues of epistemology which are involved in questions of interpretation generally. I can only touch upon a few of the issues in this article, but sufficient, I hope, to indicate the nature of my concern.

Law is based on rules

As we have seen, the work of the ICG has proceeded on the assumption that the law is a rule-based system and that the rules can be computerised without any significant problems arising from the process of knowledge acquisition. It is also the case that the ICG have concentrated mainly on statutory provisions. It has appeared to many that legislation is more a matter of rules than is case law, and in this respect the ICG are no exception.[33] Firstly, we should note that the ICG have

> laid particular emphasis on separating the representation of the law from the inference mechanisms that apply this representation for some given purpose.[34]

As mentioned earlier, their insistence on this point probably reflects the use by the ICG of shells to develop their systems.[35] This approach is bound to lead to considerable difficulty, for the knowledge of the purpose for which the representation is to be applied is often an integral factor in establishing the representation itself. Let me illustrate this by the use of an example which I have already developed in some detail elsewhere.[36]

The matter concerns the Domestic Violence and Matrimonial Proceedings Act (U.K.) 1976 S 1 which allowed the County Court to issue an injunction to restrain an aggressive party to a relationship from molesting the other party to that relationship, or a child living with that party. In determining the meaning of the statutory provision in question,

---

33 See also J. Popple "Legal Expert Systems: The Inadequacy of a Rule-Based Approach" (1991) *The Australian Computer Journal* Vol 23 no 1 pp 12-16, who despite the title of the article, still suggests that "rules are appropriate for representing statutory law" at p 15 although he feels they may well be inadequate to represent case law at p 13. .

34 Sergot, Kowalski "The use of Logical Models in Legal Problem Solving" (1990) *Ratio Juris* Vol 3 No 2 pp 201-218 at 208.

35 Sergot, Sadri, Kowalski, Kriwaczek, Hammond and Cory "The British Nationality Act as a Logic Program" (1986) Vol 29 Number 5 *Communications of The ACM* pp 370-386 at p 371.

36 See Moles *Definition and Rule in Legal Theory* (1987) Blackwell especially chap 5 "Rules - Their Application and Development".

the court had to look first at the County Court Act 1959 to ascertain the general jurisdiction of the County Court.[37] They then had to look at the Matrimonial Homes Act 1967, to determine the rights of spouses under that Act. Only when these matters had been examined in some detail, (including an examination of the cases decided under those provisions) could the court determine the scope of the provision in question. The conclusion the court came to was that it was impossible to coustrue S 1 of the Violence Act as altering substantive rights - it only affected procedural matters.

This approach illustrates the way in which the meaning of a statutory provision may well be affected by statutory provisions which appear as parts of other pieces of legislation, as well as by those rules which appear in the same piece of legislation. If one were to look at the provision of S 1 of the Violence Act in isolation, it would appear to alter substantive rights. When seen in the context of the County Court Act 1959 and the Matrimonial Homes Act 1967, it would appear to affect only procedural and not substantive matters.

This might lead one to the view that one simply requires a more extensive rendering of the legislative context in order to establish the meaning of the provision in question. If true, then that fact alone increases to a staggering degree the complexity of representing statutory provisions in knowledge bases. For example, the ICG claim that an incomplete rendering of the BNA would contain at least 500 rules. Susskind points out that the Crystal expert system shell, "is an excellent tool for building rule-based expert systems" and that "Crystal can hold up to 3000 rules."[38] If the interlinking of statutes is necessary in the way that I suggest, then Crystal would be overwhelmed before we got out of first gear. The fact that statutes are interdependent, not independent, raises fundamental difficulties in terms of representing statutes as systems of rules. While knowledge base builders might, perhaps, have realised that there could be a relationship between the Matrimonial Violence and the Matrimonial Homes Acts, would they have appreciated the significance of more fundamental and pervasive legislation such as that affecting the general jurisdiction of the court in which the matter is heard? If we depend upon computer science students or others without legal experience to establish these knowledge bases, then the answer is obvious and the potential complexity of the knowledge bases is being significantly underrated.

However, there is one more level of complexity which has to be taken into account. The court itself would not have appreciated the relationship between these various statutory provisions if it had not been for its awareness of the common law property rights of the parties involved. In discussing the 1967 Matrimonial Homes Act, Bridge L.J. stated that:

---

37 *B v B* [1978] 1 All E.R. 821. The points mentioned in this part of the discussion are taken from the judgment of Bridge L.J. at 823-834.

38 Susskind "Out of the Research Laboratory and Into the Marketplace" in *Proceedings of the First International Conference on Artificial Intelligence and Law* (1987) Boston ACM Press pp 1-8 at p 5.

> S 1 only operated to protect a wife who had no proprietary, contractual or statutory right to remain in the matrimonial home; it did not protect a wife who owned property jointly with her husband, whose rights of occupation in relation to the matrimonial property fell to be determined accordingly at common law, independently of the elaborate legislative code for which the 1967 Act makes provision.[39]

In other words, the key to the relationship between these different pieces of legislation is only provided by an understanding of the nature of the existing position at common law. The "common law" is, of course, the non-statutory part of the law which is evidenced through the reported cases. That this is not an uncommon feature of the law is indicated by L.C.B. Gower in his discussion of the Companies Acts - this is thought by many to be a typical example of an area which is governed by complex legislation and which consists of many detailed rules:

> behind the Acts is a general body of law [law in this context meaning "common law"] and equity applying to all companies irrespective of their nature, and it is there that most of the fundamental principles will be found.[40] [emphasis added]

The fact that most of the fundamental principles will not be found in the statutes themselves, but in the reported decisions of cases, has immense implications for those who would work outwards, as it were, from the detailed rules of individual statutes. How could one begin to encode the meaning of individual statutory rules, for example, without knowing whether they are intended to affect the substantive law or only matters of legal procedure?

The significance of this will, I am sure, not be lost on our knowledge base builders who presently believe that they can build up legal knowledge by an acretion of detailed atomistic rules. Indeed Bench-Capon suggests that:

> the principles of isomorphism results (sic) in a very disciplined and teachable methodology, and one which permits the construction of the knowledge base to be spread over a multi-person team in an organised and sensible manner.[41]

When we remember that perhaps two-thirds of this team will not have had any previous experience with legal systems, and that none of them will have had any experience as legal practitioners or legal academics, one wonders where the connective legal expertise (knowledge of fundamental principles) will come from.

We also have to appreciate, as we have seen, that the people working on these systems also think it a virtue of them that they separate

---

39 [1978] 1 All E.R. 821 at 827.

40 L.C.B. Gower *Modern Company Law* (4th edn 1979) p 8.

41 Bench-Capon and Coenen "Exploiting Isomorphism: Development of a KBS to Support British Coal Insurance Claims" in *Proceedings of the Third International Conference on AI and Law* (1991) Oxford ACM Press pp 62-68 at p 66.

the representation of the statutory provision from the mechanism which applies it for any given purpose. However, this approach fails to appreciate the significance of what is known as the "consequentialist" view.[42] Consequentialism requires an evaluation of any potential outcome in a way which will have regard to its political, moral or social acceptability. In legal cases, to the extent that such an outcome is not regarded as acceptable, it will lead to a reformulation of the law (or rule) involved so as to ensure that it does give rise to an acceptable outcome. In an extreme case it may well lead to the abandonment of the law or rule involved altogether.[43]

This aspect of consequentialism brings a further social dimension into the network of statutory and common law rules. The case of *Stoke on Trent City Council v B & Q Retail* provides a useful example of this.[44] B & Q were large retailers of DIY products and opened their shops on Sundays in contravention of the Shops Act (U.K.) 1950. The maximum penalty for such a contravention was considerably less than the potential profits of a day's trading, and so B & Q were happy to pay the penalty on each occasion. The local authority applied for an injunction to prevent them from opening. The court noted that the injunctive remedy was only available as part of a civil action - yet the breach of the Shops Act only gave rise to criminal penalties. The court also observed that it was a fundamental principle of the law (not, incidentally, stated in any statute) that where Parliament imposes a penalty for an offence, Parliament must consider that the penalty is adequate and Parliament can increase the penalty if it proves to be inadequate. It is not the job of the court to add more onerous penalties to a criminal statute. It follows therefore that:

> the court should be reluctant to grant an injunction which if disobeyed may involve the infringer in sanctions far more onerous than the penalty imposed for the offence.[45]

However, none of these important considerations prevented the court from allowing the injunction in this case. They felt that in balancing the consequences which would flow from the various alternatives, it would be better to issue the injunction to prevent continuing breaches of the criminal statute. It might well be thought that here we are dealing with rather esoteric parts of the law, and that the builders of expert systems could not be expected to account for these exceptional cases in which the courts develop or alter the existing rules. Unfortunately this is not so. As Neil MacCormick points out in opening his discussion of consequentialism:

> am I right in thinking that decisions are 'commonly determined' by such considerations? The answer must be Yes; to dip into the

---

42 For a discussion of consequentialist arguments see Neil MacCormick *Legal Reasoning and Legal Theory* (1978) Oxford University Press chap VI.

43 A number of cases in which this has happened are discussed in Moles *Definition and Rule in Legal Theory* (1987) Blackwell pp 167-172 and pp 254-256.

44 [1984] 2 All E.R. 332.

45 [1984] 2 All E.R. 332 at 341.

> Law Reports is to be confronted at every turn with such arguments.[46]

Far from being an isolated or exceptional feature of legal decision-making, it is in fact pervasive. It brings into clear focus the defeasibility of legal rules. If we do not understand that, then we do not understand their stability. If the individualised rule-oriented basis cannot explain that, then what is its value?

Rules may be applied deductively - mechanically

We have already given one good reason why a legal expert might not be happy with the claim that it is important to separate "a representation of the law" from the "inference mechanisms that apply this representation for some given purpose". Such a strategy would not allow for the feedback and reformulation which consequentialism either requires or allows for. In addition, a legal theorist might want to say that working from the representation of law to the particular application does not involve any "mechanism", inferential or otherwise. To be sure, it is, as Kowalski and Sergot emphasise, "this separation... [which] makes possible the use of logical models of the law in more ambitious systems".[47] But if the use of logical models depends on this separation, and it is this separation which is inimical to the nature of legal reasoning, then the proper conclusion might well be that the law is not a suitable application domain for the use of logical modelling.

One can appreciate why it is that a "one-tool" team based on logical modelling might not be happy with this finding. However, by continuing to work with assumptions that fit their tools, rather than those which are appropriate to the domain of application, we are likely to be the recipients of more talk of progress, but very little evidence of it as we will see below.

Words have an unambiguous meaning

Much of the credibility of the ICG approach depends on the view that words can, and often do, have an unambiguous meaning. As Bench-Capon and Sergot put it, the objective is to:

> represent some chosen unambiguous interpretation of the selected legal sources[48].

In more recent work they have moved from a discussion of the BNA to the logical implementation of the library regulations at the Imperial College library. It is, in my view, difficult to see this as a sign of progress. However, the procedure which they are using, to determine

---

46 MacCormick *Legal Reasoning and Legal Theory* (1978) Oxford University Press p 129.

47 Sergot, Kowalski "The use of Logical Models in Legal Problem Solving" (1990) *Ratio Juris* Vol 3 No 2 pp 210-218 at p 208.

48 Bench-Capon and Sergot "Towards a rule-based representation of open texture in Law" In *Computing Power and Legal Language* (1985) (ed Walter) Greenwood Press pp 39-60 at 42. For a similar statement, see the opening passages of the *Ratio Juris* article.

the "unambiguous meaning" of the words used, is to get people to fill in a questionnaire and indicate which of the possible meanings of "must", "not" and so on, is the "correct" meaning. As the form allows for a range of different interpretations of the words used, this approach is likely to confirm for them what many other people already know - that words do not have a single unambiguous meaning. Unfortunately, the tool they use to determine the meaning of words (questionnaires), is as inappropriate for that task as the tool they use to determine the meaning of law (logic programming).

However, it is an approach similar to that which found favour with Richard Susskind in his attempt to determine the equally difficult question as to whether the law really is a system of rules. He said that he would carry out a survey of the jurisprudential literature.[49] He acknowledged, of course, that it would not be possible to survey the whole of the jurisprudential literature. In fact, he determined that law was a system of rules by "surveying" only those whose avowed position was based on the fact that the law was a system of rules:

> the overwhelming majority of the materials surveyed were British writings of analytical jurisprudence (and philosophy) composed since the mid-fifties and early sixties...the impetus for which was derived largely from the work of H.L.A. Hart.[50]

It is well-known, of course, that it was Hart who perpetuated the idea of the law being "a system of rules."[51] By the use of "the survey", Susskind was able to give this paradigm a new lease of life. "In striving to identify consensus, approximately 50 major texts and 100 leading articles were surveyed."[52]

I would venture to suggest that this is in fact a misuse of the survey technique, whether to "find" consensus, or the unambiguous meaning of words. Susskind was perfectly familiar with the work of Hart and his followers, and was well able, therefore, to find any number of books and articles which supported the "law as rules" view. He then developed his position in Expert Systems, on the basis of what this purported consensus within jurisprudence had to say.

It appears that this view has been accepted without demurer by the AI community. McCarty, for example, states that Susskind has been

---

49 This discussion is taken from his article "Out of the Research Laboratory and Into the Marketplace" in *Proceedings of the First International Conference on Artificial Intelligence and Law* (1987) Boston ACM Press pp 1-8.

50 Ibid at p2.

51 This view was explained by Hart in *The Concept of Law* (1961) Oxford University Press and critieised in some detail in Moles *Definition and Rule in Legal Theory - A Reassessment of H.L.A. Hart and the Positivist Tradition* (1987) Blackwell.

52 "Expert Systems in Law - Out of the Research Laboratory and into the Marketplace" in *Proceedings of the First International Conference on Artificial Intelligence and Law* (1987) Boston ACM Press pp 1-8 at p 2.

"deeply concerned with the jurisprudential foundations of the field".[53] I would like to suggest that one cannot settle theoretical issues by using the techniques of quantitative analysis and that to attempt to do so shows a lack of awareness of basic jurisprudential issues. Of course, Susskind was telling certain sections of the AI and Law community what they wanted to hear, and hence their enthusiasm for it. It was also somewhat self-serving:

> It was thought that the most rigorous of these writings constituted the source materials with greatest potential given the overall purpose of the project.[54] [emphasis added]

Potential for what? What was the overall purpose of the project in which he was engaged?

> The Oxford Project had three goals. The first objective was to design, develop, and implement an expert system in Scottish divorce law.[55]

Susskind was already committed to the development of an expert system before conducting his survey. No wonder he did not survey the vast array of jurisprudential literature which has seriously questioned the adequacy of the "law as rules" approach. One might take the view that a more sophisticated analysis of the issues involved might be better for the AI and law community in the long run.

Similarly with the ICG. Such a survey in connection with the Imperial College library regulations might seem to be a neat idea, but again, they would do better to face the real issues of meaning rather than try to dodge them in this manner. I would much rather know how the judge (or chief librarian) will interpret the meaning of a relevant statutory provision (or library regulation). Whether this happens to accord or not with the outcome of a survey of other people, who do not have responsibility for determining the matter, is neither here nor there.

Rules as "atomistic" entities

It is not surprising to find that the PROLOG programmers view rules as atomistic entities, given what we have already seen of their position. It means, if correct, that the source material can usefully be mapped as a series of atomic relations in PROLOG. It is assumed that a piece of legislation has a relatively autonomous existence as do even sections of an Act. In discussing the work on the BNA, Sergot points out that:

---

53 McCarty "Artificial Intelligence and Law: How to Get There from Here" (1990) *Ratio Juris* Vol 3 No 2 pp 189-200 at p 194.

54 Susskind "Out of the Research Laboratory and Into the Marketplace" in *Proceedings of the First International Conference on Artificial Intelligence and Law* (1987) Boston ACM Press pp 1-8 at p2.

55 Ibid at p 1.

> the system ran a relatively self-contained part of the Act, consisting of approximately 150 rules.[56]

As we have already seen, the ICG took the view that an incomplete rendering of the Act would contain about 500 rules.[57] The parts of the Aet which were omitted from their representation refer to the way in which the Act in question relates to other legislative sources. Their exclusion of this part of the Act indicates their lack of awareness of the legislative connectivity of which we have already spoken. It also demonstrates the way in which this aspect has been ignored by the ICG, even where the connections are relatively explicit.[58]

This atomistic quality of legal rules is referred to as a virtue of isomorphism by Bench-Capon, for it means that when corrections are needed:

> the offending piece of the knowledge base can then be detached, its parent source fragment associated with it, and corrections made without any need to worry that surrounding parts of the knowledge base or source will be jeopardised by the changes.[59]

What he is referring to here is the problem of maintaining the knowledge base. As we have already seen, one has to travel a long and difficult road to get from the basie documentary sources to the aetual knowledge base which is being constructed. If one had to re-do that whole process each time a change occurred, whether as a result of the discovery of error, or because of changes to the legislation or rules, then the problem of change would quickly become insurmountable. However, if the original documentary sources, and the corresponding section of the knowledge base, can be seen as comprised of atomistic entities, then when a change occurs, one only has to alter that 'piece' of the original statutory source and the corresponding 'piece' of the knowledge base. This is why Bench-Capon places such importance on the need for isomorphism which stresses the similarity of structure betwcen the knowledge base and the source material. The atomistic assumption upon which this is based is, of course, extremely convenient. If, however, as I have suggested, the issue of connectivity and context is all important, then a change in one rule, which is itself part of the context in which the other rules occur, will give rise to a change in the meaning of those other rules as well.

Following on from their view regarding the autonomy of legislation, it was also assumed by the ICG that because the legislation was at the time they were commencing their project, fairly recent, that

---

56 "The British Nationality Act as a Logic Program" *Communications of the ACM* (1986) Vol 29 No 5 pp 370-386 at p 383.

57 The ICG's implementation of the Act did not include sections of Part 5 of the Act which referred to amendments to other Acts and repeals. Part 5 also deals with a number of other matters including decisions involving the exercise of discretion by the Secretary of State, ibid p 383.

58 Ibid.

59 Bench-Capon and Coenen "Exploiting Isomorphism: Development of a KBS to Support British Coal Insurance Claims" in *Proceedings of the Third International Conference on AI and Law* (1991) Oxford ACM Press pp 62-68 at p 66.

fact meant that it was therefore free of the influence of case-law.[60] This is an assumption which no decent lawyer would make, as even a new piece of legislation will always involve some of the "fundamental principles" which we looked at earlier. The previous case-law will be of even more importance when we are dealing with legislation such as the BNA which has a long history in prior legislation.[61]

I appreciate that what we are looking at are pilot projects, but concern must be expressed when we see that the prevailing methodological assumptions are likely to fundamentally undermine their viability. The problems which remain are, of course, those of the abstraction of knowledge and of the individualization of rules.[62] Whether the researchers involved are unaware of the problems, or whether they have merely chosen to ignore them, is difficult to determine. There is, as we shall see, some evidence to support the latter view.

The view that legislation consists of a series of discreet propositions, each of which can be manipulated independently of the others, is not a view which is either sound in theory, or which would find any resonance amongst well informed lawyers. Such a view is rather reminiscent of the Humean view of a theory of knowledge in which our understanding is derived from certain "sense-data", which are the neutral originals from which we develop ideas. Immanuel Kant very wisely turned this idea on its head by pointing out the extent to which the knower participates in constructing that which is known - a factor which appears to be entirely ignored by the writers we are considering.[63]

The point can be appreciated in the more specifically legal context by considering the aspect of statutory interpretation. If one were to turn to Sir Rupert Cross's well-known book on statutory interpretation, one finds, in his chapter on "The Basic Rules Stated" a section on "context" in which he makes the following observation:

> it is difficult to believe that the notion of construction in complete isolation was ever taken wholly seriously.[64]

Cross explained what he meant in discussing the case of *Re Bidie*. He pointed out that Lord Greene, in the Court of Appeal, criticised the trial judge's approach to the matter thus:

> I think he attributed too much force to what I may call the abstract or unconditioned meaning of the word...The real question which we have to decide is what does the word mean in the context in which we find it here, both in the immediate

---

60 "The British Nationality Act as a Logic Program" (1986) *Communications of the ACM* Vol 29 No 5 pp 370-386 at p 370.

61 See the discussion which follows regarding the case of *Attorney General v H.R.H Prince Ernest Augustus of Hanover* [1957] 1 All E.R. 49.

62 For a discussion of the theoretical problems associated with this see "Epistemology - The Common Ground" chap 6 of Moles *Definition and Rule in Legal Theory* (1987) Blackwell.

63 A useful discussion of these matters may be found in Marjorie Grene The *Knower and The Known* (1966) Faber and Faber chaps 4 and 5.

64 Sir Rupert Cross *Statutory Interpretation* (1976) Butterworths p 44.

> context of the subsection in which the word occurs and in the general context of the Act, having regard to the declared intention of the Act and the obvious evil that it is designed to remedy.[65]

He also made it clear that this notion of "context" is not just something which has to be considered on the odd occasion. He referred to the statement by Lord Reid that:

> In determining the meaning of any word or phrase in a statute the first question to ask always is what is the natural or ordinary meaning of that word or phrase in its context in the statute.[66] [emphasis added]

Any interpretation based on the "atomistic" approach to the meaning of words or phrases in a statute fails to appreciate the significance of this aspect which Lord Reid said must always be considered - the meaning of the word in its context in the statute.

The following example (utilised by Cross) is even more appropriate, given the fact that a major focus of the work of the ICG is on the BNA. It also illustrates the falsity of the assumption made by the ICG team that if the statute has only recently been passed, then there cannot be any relevant cases concerning it. *Attorney General v H.R.H. Prince Augustus of Hanover* [1957] was a case dealing with the provisions of the British Nationality Act 1948. Although that Act repealed earlier legislation, it was clear from the speeches in the House of Lords, that one could not appreciate the significance of the 1948 Act without a proper understanding of the earlier legislation. As Viscount Simonds pointed out, "the question is to be answered on a consideration of a statute passed just 250 years ago".[67] He then pointed out that:

> it is proper, too, to have in mind what was the state of the law in regard to naturalisation in 1705. By the common law, only those persons who were born on English soil were subjects of the English Crown.[68]

In this case, in order to understand the current legislation, one had also to understand not only the previous legislation which was passed some 250 years ago, but also the state of the common law at the time that the earlier legislation was passed. In addition to the need for an understanding of the preceding statutes dealing with that subject, Viscount Simonds said that one would also have to appreciate a number of other factors:

> words and particularly general words, cannot be read in isolation, their colour and content are derived from their context...and I use 'context' in its widest sense, which I have already indicated as including not only other enacting provisions of the same statute

---

65 Ibid p 45 citing Lord Greene in *Re Bidie* [1948] 2 All E.R. 995 at 998.

66 Ibid p 29 citing Lord Reid in *Jones v Director of Public Prosecutions* [1962] A.C. 635 at p 668.

67 *Attorney General v H.R.H Prince Ernest Augustus of Hanover* [1957] 1 All ER 49 at 50.

68 Ibid at 51.

> but its preamble, the existing state of the law, other statutes in pari materia, and the mischief which I can, by those and other legitimate means, discern the statute was intended to remedy.[69]

His Lordship, in fact, referred specifically to the contention of the Attorney General in this case that where the enacting part of the statute is clear and unambiguous, it cannot be cut down by other factors, such as the preamble.[70] But as he sensibly goes on to point out:

> it must often be difficult to say that any terms are clear and unambiguous until they have been studied in their context... the elementary rule must be observed that no one should profess to understand any part of a statute or of any other document before he has read the whole of it. Until he has done so, he is not entitled to say that it, or any part of it, is clear and unambiguous.[71]

It is not surprising that the judge also felt constrained to reject the argument of the Attorney General that the words be given their "prima-facie and literal meaning". Lord Normand also stated that one had to read the whole Act, inform oneself of the legal context of the Act, including Acts related to it and the factual context. There was no suggestion in his Lordships speech that this was an approach peculiar to this area of law, for he added:

> It is the merest commonplace to say that words abstracted from context may be meaningless or misleading.[72]

It is clear that this line of thinking has had little impact on the writers we have been looking at. In fact Susskind, in common with the ICG, take precisely the opposite view:

> the author devotes considerable attention to precisely this issue in EXPERT SYSTEMS IN LAW and offers a "semantic" theory of clear cases according to which such cases are, roughly, those in which the facts of the case can, in accordance with the conventional and "acontextual" use of legal and ordinary language, be subsumed unambiguously within the terms of valid legal rules.[73]

Susskind continues, by suggesting that:

> Ronald Dworkin introduced the idea of "acontextual meaning" to jurisprudence. It refers to the meaning we would assign to words "if we had no special information about the context of their use or the intentions of their author".[74]

---

69 Ibid at 53.

70 Ibid.

71 Ibid at 55 per Viscount Simonds.

72 Ibid at 56.

73 Susskind "Out of the Research Laboratory and Into the Marketplace" in *Proceedings of the First International Conference on Artificial Intelligence and Law* (1987) Boston ACM Press pp 1-8 at p 3.

74 Ibid at p 3.

Many of us, of course, would not wish to assign meaning to words in such a situation. To suggest that one could do so would be to make the same mistake which we have already referred to and to which the Chief Justice of the Australian High Court referred in a recent address:

> to treat the law as a discrete set of principles in a vacuum and without a context is to misconceive its dynamic and ubiquitous nature.[75]

Context, of course, is something which is integral to meaning, and cannot be incorporated at a later stage, or dispensed with, as Susskind and the ICG team suggest.

No need for legal expertise

The people that we have looked at so far clearly accept that a purely logical model of any legislation can be made, and that this can be done without the help of anyone who knows even the most basic of the approaches taken to the problem of statutory interpretation. The ICG have proceeded to deal with complex legislation, with a long statutory history, on the basis of what either they or their students assume it to mean. On occasion, it is true, they have gone beyond merely using their own intuition to work out what the legislation means. In dealing with Social Security matters, for example, they referred to the S-manual, or administrative guide-lines which indicate how to interpret the legislation.[76] We saw earlier, in our discussion of isomorphism, how Bench-Capon utilised the Claims Inspectors' Manual. Although no expert legal help was enlisted, the ICG acknowledge that this must not be the case when implementing a real system:

> For our project, the accuracy of the representation was not a critical consideration at this stage. Our formalisation could therefore be undertaken with no expert legal assistance (except that the S-Manual does provide some indication of how to read the various provisions). In general, accuracy of the formalisation is, of course, critical, particularly if one were constructing a representation to be used in practice[77].

This is to assume, however, that the problem with regard to "accuracy" is merely a matter of changing the detail of content. It fails to appreciate that an expert may have a great many useful things to say about how one goes about the process of interpretation. The expert advice will therefore have implications for the method being employed and the way in which the knowledge is structured. Although the ICG group

---

75 The Honourable Sir Anthony Mason, A.C., K.B.E., Chief Justice of Australia *Address at the Inauguration of the Faculty of Law University of Wollongong* Australia 19 February 1991.

76 Bench-Capon, Robinson, Routen, Sergot, "Logic Programming for Large Scale Applications in Law" in *Proceedings of the First International Conference on Artificial Intelligence and Law* (1987) Boston ACM Press pp 190-198 at p 192.

77 Ibid.

appear to have considered this possibility, they also appear to have rejected it:

> Access to an expert might well have changed the exact form of the rules in our program, but it would not have changed the method we used to formulate and compute with those rules.[78]

Their claim in this regard would be more persuasive if it were arrived at after considering what an expert has to say. But to arrive at such a conclusion in advance of the consideration of such expert opinion, suggests that the reason for the rejection of the value of expert advice might be based more on reasons which have less to do with philosophical adequacy than with other factors. They do in fact indicate what other factors influenced their approach.

> It is naturally most convenient if there is access to such an expert adviser from the beginning...This is not always realistic given the demands on the time of an expert lawyer and the costs of employing the services of such a person.[79] [emphasis added]

So they set about interpreting the BNA without the benefit of such assistance from an expert. Whilst they state that it would be "most convenient" to obtain expert advice at an early stage, I would like to suggest that it was in fact "more convenient" for them to manage without it. When they say that it would not be "realistic" to obtain expert advice, "given the demands on the time of an expert" and "because of the costs involved" they appear to be assuming that highly paid barristers or solicitors are the only people to have the expertise which was needed and that this could only be made available to the team at great expense. It is, of course, the case that lowly paid legal academics regularly make their skill and expertise available without necessarily demanding a fee on each occasion, as indeed do some solicitors and barristers.

As a result, the ICG proceeded to implement their programming of the BNA without access to the legal advice they so obviously needed:

> Our representation of the British Nationality Act was undertaken with no expert legal assistance. Our model of the Act expresses a layman's reading of the provisions. This in itself renders our British Nationality Act program of limited practical value. We could not use it in its present form for solving problems of British citizenship in actual legal practice.[80]

They try to fudge the issue by saying that the program which they have produced "is of limited practical value" because, as they say, it could not be used "in its present form" to solve problems of actual legal practice. But if the model cannot solve problems of "actual legal practice", then it seems that it would be better to say that it is of NO

---

78 Sergot, Kowalski "The use of Logical Models in Legal Problem Solving" (1990) *Ratio Juris* Vol 3 No 2 pp 210-218 at p 207.

79 Ibid at p 207.

80 Sergot, Kowalski "The use of Logical Models in Legal Problem Solving" (1990) *Ratio Juris* Vol 3 No 2 pp 201-218 at p 207.

practical value. The clear implication is that it might well be of some practical value, in a more developed form.

They suggest that if they were now to bring in an expert adviser, such a person might want to extend or modify the program. Whilst this is clearly a possibility, the more likely scenario is that such an expert would advise them to abandon it altogether. Their hope that "expertise" can be brought in to "further develop" what they are doing, presupposes that they have been on the right track so far. Who would attempt to start building a house themselves - in order to avoid the cost of a bricklayer - in the hope that they could always call in the brickie later on to put things right? The point is, of course, that if the foundations are not laid properly, the only sound advice might be to clear the site and start again.

After several years of research, and a good deal of money having been expended, they conclude that they have produced something which might be seen as a "common sense or a 'layman's' reading of the British Nationality Act". One wonders whether, given their use of symbolic logic and extended Horn-clauses, it even amounts to that. In one respect, the ICG are refreshingly frank about the factors which motivated them to do as they did. It seems to boil down to a combination of naivety and convenience:

> We have stressed in this article the top-down, goal directed development of our formalization of the British Nationality Act. We adopted this approach for purely practical reasons. It allowed us to delay addressing the more complex issues of knowledge representation until it became unavoidable to do so, and it enabled us to avoid considering how to represent the various commonsense knowledge needed to understand the legislation until we discovered what knowledge was required.[81] [emphasis added]

As they point out, they have never suggested that their programs model "the entire process of legal reasoning". In fact, they have stated quite clearly that they have not attempted to model "any" of the process of legal reasoning. As they admit, they have derived the logical consequences of "a model of law" - but this is indeed different from deriving such consequences from "the law itself". It is fair to say that the model from which they derive these logical consequences has not been built using the insights of any person who has knowledge of the way in which such consequences are derived by lawyers, or of anyone who has any wider knowledge of the nature of legal reasoning. They have used logic programming to model the logical consequences of their own untutored assumptions as to how a statute, dealing with an area with which they have no experience, would be read and used by someone about whom they have no understanding.

It would appear that the reason they have pursued this line is because they were, before they started on law, already committed to the use of a particular computing tool, not to the understanding of law. As Susskind pointed out:

---

81 "The British Nationality Act as a Logic Program" (1986) *Communications of the ACM* Vol 29 No 5 p 385.

> Some of the workers on the PROLOG projects were committed to using PROLOG before selecting the law as an appropriate domain. The goal of some of their projects was to represent legislation in PROLOG come what may.[82]

It is not surprising, therefore, that the ICG ignore aspects of legal decision-making which cannot be worked with their very limited tools. In fact, there is good evidence to suggest that they do not really have a serious interest in understanding the workings of the law, but that their main focus is to develop computing technology.[83] Rather than take the opportunity to extend their range of skills, by developing a real understanding of the law, they simply stuck to doing what they know best. As a result, it has to be admitted that very little progress has been made.

In their recent retrospective article, they make the claim that there are no outstanding technical obstacles which need to be overcome to finish off a program in this way, but I am not convinced by what they say.[84] As they explain:

> Experience in other projects, and with applications in other domains, suggests that there are no outstanding technical obstacles which need to be overcome to finish off a program in this way, but that this stage of the process can often involve a considerable amount of work and extra programming effort.[85]

This might lead us to believe that after all the hard years of work and successfully working out all the theoretical problems, we are now within sight of the first prototype of a working program. After all, it is not as though this "extra programming work" was not foreseeable from the outset, and, therefore, it is something which should have been allowed for from the outset. When, however, we see that the final slog along the home straight is given as the reason for the abandonment of the program, in favour of taking up "more ambitious" projects, we may, perhaps, be forgiven for remaining a little sceptical. I would like to suggest that the real reason for the abandonment of the project is that the difficulties of continuing to work without an adequate research protocol are now becoming so obvious that to continue with further work would become an even greater embarrassment. The ICG are not the only group to have found themselves in something of a black hole. In 1977, McCarty made the claim:

---

82 Susskind *Expert Systems in Law* (1987) Oxford University Press p 24.

83 "The British Nationality Act as a Logic Program" (1986) *Communications of the ACM* Vol 29 No 5 "We believe that the formalization of legislation and legal reasoning offers potential contributions to computing technology itself" p 371 and "..the accumulated experience of managing complex systems of law may teach us something about the maintenance of large bodies of software" p 385.

84 Susskind too points to the lack of progress in the field whilst at the same time stating that technical obstacles to further progress have been sorted out *Expert Systems* pp 18, 19.

85 Sergot, Kowalski "The use of Logical Models in Legal Problem Solving" (1990) *Ratio Juris* Vol 3 No 2 pp 210-218 at p 208.

> I would be willing to hazard a guess: if a substantial amount of resources was committed to the project, a sophisticated prototype system of demonstrable utility to a corporate tax lawyer could be developed within approximately 10 years.[86]

Well, more than 10 years have passed, and Taxman I and Taxman II have come and gone and despite the commitment of very substantial resources we still appear to be no closer to obtaining an implementable system. In fact, McCarty pointed out as recently as 1990, "I will argue that we have not advanced very far at all in these past ten years."[87]

So is the constant talk of promise, but no progress, going to be replaced by the real interdisciplinary work which is needed? Well, it does not appear to be likely. In McCarty's recent discussion of his new "language for legal discourse" (LLD) we find mention of the fact that "atomic formulae" the "reification of relationships as objects" and "Horn Clauses" are the more important building blocks of LLD.[88] And despite all the previous years of promises, it is only with the advent of LLD in 1989, that we take the first concrete step towards the realisation of the goal, which is to establish a solid foundation for further theoretical work.

In summary, where could we say we are now? McCarty is only taking his first steps after some 15 years (although apparently he is going to continue to use the same tools as previously). The ICG claim to have sorted out all the theoretical problems, although they too will not be finishing off any of their programs just now - the work would be too boring. Instead they are setting to work on more ambitious programs - like the work on the library regulations at Imperial College? In the meantime, one can imagine that their jurisprudential adviser, Richard Susskind, will continue to reassure them that all is well, because he has counted up some more pages of people who agree with Hart. They don't sound too convincing, do they?

It is clear that if real progress is to be made, then all of those involved will need to approach, with an open mind, the basic question of the suitability of the legal domain for the development of expert systems. In order to determine that issue, they will first need to seriously engage with legal theorists who will provide them with insight into the nature of legal analysis, rather than just tell them what they want to hear. If they cannot bear to face the difficult issues, then we are now producing a new generation of students qualified in both law and computer science, who will pass them by on the way to tackle those interesting issues. These new students will be both able and willing to work with the legal

---

86 McCarty "Reflections on Taxman: An experiment in artificial intelligence and legal reasoning" (1977) 5 *Harvard Law Review* pp 837-893 at p 892.

87 McCarty "Artificial Intelligence and Law: How to Get There from Here" (1990) *Ratio Juris* Vol 3 No 2 pp 189-200 at p 189.

88 McCarty "A Language for Legal Discourse" in *Proceedings of the Second International Conference on Artificial Intelligence and Law* (1989) Vancouver ACM Press pp 180-189 at pp 181-182.

theorists, and together they will increasingly compete for the available research funding.[89]

89 I am much indebted to Dr Neil McLeod of the law faculty at the Australian National University, for his detailed comments on earlier drafts of this article and to Dave Brown and Andrzej Kowalski of the Melbourne law faculty for the discussions we have had on these issues. Surend Dayal has made an important contribution whilst working with me as a research assistant on this material. He is one of the students I had in mind whilst writing the last paragraph of the article.