

LEADING LAW STUDENTS TO UNCHARTED WATERS AND MAKING THEM THINK: TEACHING ARTIFICIAL INTELLIGENCE AND LAW

Andrzej Kowalski¹
The University of Melbourne
Law School

ABSTRACT

The experience of teaching a new course to law students entitled 'Legal Reasoning, Expert Systems and Artificial Intelligence' is described. The multi-disciplinary nature of the field of artificial intelligence and law is emphasized together with the importance of legal theory, legal reasoning and philosophy. The students, working in groups, constructed rule-based expert systems, using an inexpensive commercial shell called VP-Expert, in domains of case-based law (medical negligence, confidential information, assault, false imprisonment and the defence of fair comment in the law of defamation). With due emphasis given to the theoretical and philosophical presuppositions of the field of artificial intelligence and law, law students in their final year(s) of study are capable of building rule-based legal expert systems of a very high standard during a one semester course.

I. BACKGROUND

II. INTRODUCTION

III. THE COURSE

A. REQUIREMENTS

B. STRUCTURE

C. THEORETICAL AND PHILOSOPHICAL UNDERPINNINGS

D. ARTIFICIAL INTELLIGENCE AND EXPERT SYSTEMS

E. PROBLEMS

F. THE COMPLETED EXPERT SYSTEMS

G. STUDENT FEEDBACK

IV. CONCLUSIONS

¹ B.A., LL.B., LL.M., Lecturer-in-Law. This article is dedicated to the students who took the course. Their patience and enthusiasm in the face of technical difficulties was extremely rewarding as was the high quality of their work and the spirited discussions. I wish to thank Bob Moles, Dave Brown, Dorota Gertig, Sally Walker and Gerry Simpson for their valuable comments and suggestions on drafts.

I. BACKGROUND

In the first semester of 1991 the Law School of The University of Melbourne offered for the first time an optional subject called 'Legal Reasoning, Expert Systems & Artificial Intelligence'.² In selecting materials for the course, I decided to use an article by James Boyle concerning the experience of teaching a torts class for the first time.³ The article was selected for the course not so much for its focus on the teaching experience but for its discussion of some ideas about the indeterminacy of law. It did, however, provide an example to suggest that one might share one's experiences of teaching a new subject. Artificial intelligence ("AI") and law is a relatively new and embryonic field and provides, as a result, virtually uncharted waters concerning teaching in this area.⁴ It might therefore be helpful to others who may be thinking of providing courses in this area to record some of the ideas and considerations gleaned from my experiences. Far from being determinative of a methodology for teaching AI and law, these suggestions are clearly no more than the provision of a framework, which might provide the basis for stimulating further discussion and the suggestion of alternative possibilities.

II. INTRODUCTION

It seems to me that AI and law should be regarded as a multi-disciplinary field of study. AI and law may be seen as interrelating with both the discipline of AI (of which AI and law is a sub-discipline) and with many of the central aspects of the humanities. Thus the field of AI and law is related to computer science and programming, psychology, legal theory, jurisprudence, philosophy, epistemology, linguistics and, no doubt, other disciplines. It was important to me that the course reflect, insofar as it could, this multi-disciplinary nature, especially in light of the tendency in the study of traditional legal subjects to confine the discourse to 'black letter' law.⁵

² I know of no other similar subject offered by a law school in Australia, but there are other such courses in North American and European law schools. In designing the course, I had the benefit of having taken a course in this field at The University of British Columbia, Faculty of Law, which was designed and taught by Professor J.C. Smith.

³ Boyle, James, *The Anatomy of a Torts Class*, 34 (1985) *American University Law Review* 1003.

⁴ On using AI and expert systems as teaching aids see Gregor, S.D., Rigney, H.M., & Smith J.D., *The Applicability of a Knowledge-Based System to Legal Education*, (1991) 23 *The Australian Computer Journal* 17; Ashley, Kevin D. & Alevin, Vincent, "Toward an Intelligent Tutoring System for Teaching Law Students to Argue with Cases", *Proceedings of the Conference, The Third International Conference on Artificial Intelligence & Law*, 25th-28th June, 1991, Oxford, England, The Association for Computing Machinery, New York.

⁵ However, the field of AI and law may already be in danger of losing sight of its multi-disciplinary roots. Although it is still developing its own structure and place within the framework of the study and practice of law, many of those involved in its development are closing their eyes to the

Richard Susskind emphasized the importance of theory and philosophy when he wrote:

"It is beyond argument, however, that all expert systems must conform to some jurisprudential theory because all expert systems in law necessarily make assumptions about the nature of law and legal reasoning. To be more specific, all expert systems must embody a theory of structure and individuation of laws, a theory of legal norms, a theory of descriptive legal science, a theory of legal reasoning, a theory of logic and the law, and a theory of legal systems, as well as elements of a semantic theory, a sociology and a psychology of law (theories that must all themselves rest on more basic philosophical foundations). If this is so, it would seem prudent that the general theory of law implicit in expert systems should be explicitly articulated..."⁶

I therefore wanted the course to convey the theme of the multi-disciplinary nature of AI and law and expose students to the important theoretical and philosophical issues as well as some of the related disciplines. By the same token, emphasis was also placed on the principal task allocated to the students: the modelling of legal domains by the construction of legal expert systems. Building a working legal expert system for the first time requires an in-depth analysis of both the

importance of philosophy and theory upon which it is necessarily and unavoidably founded. The Proceedings of the recent Third International Conference on Artificial Intelligence & Law contain a proliferation of 'technical' papers and almost a complete absence of philosophical or theoretical material, whether on legal theory and the nature of legal reasoning as such or on the relationship between AI and law and its associated disciplines. See Bob Moles' article in this edition of this Journal for a legal theorist's critique of the work of the Logic Programming Group (at the Imperial College of Science and Technology in London) which featured prominently at the Conference. It may well be the case that the field of AI and law has cut itself off as a specialist discipline and established its parameters somewhat prematurely. One hopes that the field does not simply become a form of glorified computer programming, abandoning serious inquiry into its related fields. For if there is a divergence of views among lawyers and jurisprudentialists about the nature of law and legal reasoning, then the AI and law community (which is primarily concerned with modelling these processes) must engage with the range of issues or slide into inconsequentiality. It is interesting in this context to note McCarty's comments about the lack of substantial progress in the field to date, with which many researchers will agree. I would venture to suggest that this state of affairs may be largely attributable to the premature demarcation of the field of AI and law. McCarty's identification of knowledge representation as the central issue facing the field, and where progress must be made if the field is to advance, is a way of acknowledging (albeit indirectly) this very problem. McCarty, L.T., *Artificial Intelligence and Law: How to Get There from Here*, (1990) Ratio Juris Vol 3 No 2 p. 189.

⁶ Susskind, Richard E., *Expert Systems in Law: A Jurisprudential Approach to Artificial Intelligence and Legal Reasoning*, (1986) 49 Modern Law Review 168 at p.183. See Bob Moles article in this edition of this Journal for a critique of Susskind's approach to this topic.

legal domain of the system and the various stages of system design and construction. The second element will be particularly time consuming for novices. Thus the competing goals of a multi-disciplinary review of AI and law and the in-depth analysis required to build a legal system necessitated careful balancing in light of the limited twenty-six hours of teaching time available to a one semester course.

For the students, I felt the best way to settle the issues as to the viability of the field of AI and law was to get them to 'do it themselves' by building their own legal expert systems. Demonstrations of working systems, explanations of approaches to knowledge representation and the like were important components of the course, but these could not adequately substitute for the process of building legal expert systems themselves. There was also the attraction inherent in the field of AI and law that building a legal expert system is an opportunity to 'empirically' test theories of legal reasoning and jurisprudence. Further motivation to ask students to build working expert systems came from my background as a practising lawyer, from where springs my conviction that the field of AI and law has no future if it cannot produce tools and systems of utility to practising lawyers. Other motivations to teach the course were to give final year law students the opportunity to demystify the processes of legal reasoning and argument and to place 'black letter' law in its context as part of a wider theoretical and practical framework. Boyle takes up this point when he comments on the failure of law teachers:

"...to tell students that a lot of the time we were teaching them a standard rhetoric and *not* a set of rules or a body of knowledge at all. If this was part of 'thinking like a lawyer', what would happen when you demystified the process and taught arguments explicitly and overtly? Whatever happened, it would surely be better than expecting the students to pick up these arguments by osmosis while they are still under the illusion that there are *right answers* and that, correctly understood, these arguments will deliver them."⁷

I also wanted to dispel the niggling question I had in my mind: 'Why would law students go to the trouble to construct a legal expert system when there were many other elective law subjects offering more traditional and tried content and perhaps less work?'

III. THE COURSE

A. Requirements

There were no prerequisites for the course. There was an initial enrolment of approximately 40 students which then had to be reduced because a quota of 18 students, for reasons explained below, was imposed on the course. Selection of students was based on seniority, with priority given to those in the later years, primarily because I thought by this stage of their legal careers they would have sufficient:

⁷ *Supra*, note 3, at p. 1009.

- breadth of legal experience to give them a number of choices of legal domains for an expert system; and
- exposure to the nature of legal reasoning in a number of domains to enable them to contribute to discussions on this topic, which was central to the course.

In addition, students in earlier years who missed out would be able to enrol in the course in subsequent years. Information about the final enrolments is summarized in Table 1:⁸

Table 1

Final Enrolments

	2nd YEAR	3rd YEAR	4th YEAR	5th YEAR	TOTALS
LL.B.	1	2	1		4
B.A./LL.B.		1		3	4
B.Sc./LL.B.			2	3	5
B.Com./LL.B.				2	2
TOTALS	1	3	3	8	15

The course prospectus stated clearly that no knowledge of computers was assumed or necessary, but that some knowledge of word processing might be helpful. Although the task of teaching students would no doubt have been easier if they were reasonably computer literate, I reasoned that if I was going to advocate the demystification of black letter law and legal reasoning I ought not drape the process of demystification with a cloak of assumed computer knowledge and terminology. I was nevertheless certain that the course would attract students with computer skills and therefore had to ensure that they did not feel that the course was less demanding for them, if it was set at the level of those without computer experience. The manner in which this tension was resolved will be discussed later.

Interestingly, there were only three students in the course who could be classified as having significant computer literacy:

- one of the B.Sc./LL.B. students had taken computer science and programming subjects towards her B.Sc.;
- one of the LL.B. students operated his own desktop publishing business; and
- another of the B.Sc./LL.B. students was an experienced self-taught MS DOS user.

Most of the other students had some word processing experience, but some of them were not even familiar with basic MS DOS commands such as FORMAT and COPY. Two students were Apple users with virtually no experience of MS DOS machines. There was also one Commodore Amiga user with a little MS DOS experience. Thus, as a

⁸ After the selection process was finalized and student numbers stabilized a total of 15 students were enrolled in the course.

whole, the class generally did not have substantial computer experience or knowledge.

A computer laboratory was necessary and we were fortunate that the Law School completed installation of a new facility with 18 networked IBM compatibles immediately before the commencement of the course. A quota of 18 students was therefore imposed on the course because approximately one third of the classes were held in the laboratory and I also wished to keep numbers down to facilitate seminar-style discussions. The remaining classes were held in a seminar room. Software initially installed on the network consisted of WordPerfect 5.1,⁹ Microsoft Word 5.0¹⁰ and dBASE IV 1.1.¹¹ Other important equipment included a Kodak Datashow, overhead projector and screen to allow computer screen images to be displayed to the class.

The most important piece of software was the shell to be used by the students to build their expert systems. The only viable choice was the student edition of VP-Expert given that the students could purchase it for approximately AUD \$80.¹² No other shells could compete in terms of price and features. Although the student version has limitations imposed on it, such as the size of the knowledge base, I was confident that the students would not exceed its capacities in the space of a one semester course. However, as will be seen later, this proved to be a miscalculation.

B. Structure

The students were divided into groups of three to five students at the start of the course for the purpose of constructing their legal expert systems. I reasoned that this approach was necessary because no one student would have the capacity to process 40 or 50 cases during a one semester course being the minimum number of cases (as a rough rule of thumb) a reasonably sophisticated legal expert system requires in its database. Had the students prepared individual expert systems, then 18 systems may have been produced, none of which could probably have functioned with any real degree of sophistication. Steps were taken to ensure that each group had at least one person with some computer experience. I had some trepidation that a person with computer skills might be saddled with an unequal amount of work but it worked out that because of the many stages to expert system design and construction the groups were more than able to divide the work load equitably while ensuring that skills of members were put to the best use.

One student was very reluctant to work in a group and expressed a desire to build an expert system by himself. He explained it was not from a desire to work alone as such, but he felt he had the necessary computer skills and motivation to build a system by himself. Despite my explanations that I would expect a system of similar standard to those built by the groups and that he would be taking on a huge amount of work, he nevertheless chose to go it alone. I had reservations, but there

⁹ A product of WordPerfect Corporation, Orem, Utah, U.S.A.

¹⁰ A product of Microsoft, Seattle, Washington, U.S.A.

¹¹ A product of Ashton-Tate Corporation, Torrance, California, U.S.A.

¹² A product of Paperback Software International, Berkeley, California, U.S.A.

was no other alternative in light of his enthusiasm than to let him go ahead.

An interesting aspect of dividing the students into groups for the production of assessable work is that it is for the most part an unknown experience in Australian law schools. Students do participate in study groups and (to the chagrin of some teachers) assist one another with assignments, but collaborating to produce a joint piece of assessment is rare. Yet from my experience, I knew that the ability to work effectively as a member of a team (eg. in large litigation matters or as part of a specialist group within a law firm) is an important aspect of the practice of law.

Another structural decision was whether the groups ought to produce individual expert systems or whether they ought to work on modules for one large expert system project. The modular approach has the advantage of allocating each group very well-defined tasks but no group then directly experiences each phase of the construction of the system. This experience is available on the individual system approach, but the concern I had was that at some point the students would encounter a problem or obstacle which might appear insurmountable and thus abandon the system altogether. The task of building only one module of a large system must intuitively feel more manageable than the task of building all of the modules of a system, albeit a smaller system. I opted for the individual system approach, despite the risks involved, because I thought that the process of building an expert system from start to finish would be more instructive and rewarding for the students.

The students were requested to select domains from case-based law rather than statute-based law. This request was a reflection of my opinion that unstructured areas of purely case-based law present the most interesting and challenging domains for the construction of legal expert systems. It has been suggested that it is easier to develop knowledge representation models for legal expert systems from statutes because the statutes themselves suggest models;¹³ in purely case-based law no such preliminary model exists and the model must be designed entirely by the knowledge engineer after an in depth review of the cases. The domains chosen by the students in consultation with me were: medical negligence, false imprisonment, civil assault and battery, the defence of fair comment to a defamation action and breach of confidential information.

A traditional method of assessment for the course might have consisted of a substantial research essay, an examination or a combination of the two. However, this approach did not appear appropriate in light of the quite sizeable task being asked of the students - the construction of substantial legal expert systems - combined with the traditional propensity of many law students to leave things until the last minute. With this project, unless the students started very early and continued throughout the course to develop their systems progressively, there was every likelihood that there would be a collection of half-baked systems produced in last minute fits of panic. A four stage assessment process

¹³ Sergot, Cory, Hammond, Kowalski, Kriwaczek & Sadri, *Formalization of the British Nationality Act (1986)* 2 Yearbook of Law, Computers and Technology, p. 40 at p. 49.

was therefore opted for which resulted in a piece of assessment falling due for submission every three weeks or so during the semester with the final systems due at the end of the first semester examination period. Each piece of assessment was designed to be a module of or contribute to the construction of the expert systems:

- (1) Summaries of cases in the domain to be available for display by the system to the user - 10%.
- (2) Development of the structure and field names of the database and an analysis of the cases in the domain in terms of the structure and field names - 15%.
- (3) Doctrinal level or work-in-progress expert system - 15%.
- (4) Completed expert system - 60%.

The final expert systems were only required to be rule-based, as more sophisticated systems, such as case-based reasoners, would have been too time consuming.¹⁴ At any rate, the choice of shell virtually mandated this approach. I originally planned to set students an essay on some of the theoretical issues in the field but this idea was abandoned due to some of the practical and technical difficulties faced which will be discussed below.

C. Theoretical and Philosophical Underpinnings

This part of the course was, in my opinion, central and essential to the students' understanding of the issues and viability of the field of AI and law. I readily acknowledge that theoretical and philosophical knowledge imparted to students will inevitably reflect the teacher's own orientation. Rather than attempting to overcome this, I simply tried to draw to the students' attention the judgmental aspects of the more blatant mindsets or opinions offered. Another caveat is that the survey of legal theory presented here as relevant to the field of AI and law is not meant to be exhaustive in any way, but merely instances of paradigms of debate.

My approach to teaching legal theory and reasoning in the context of AI and law may essentially be described as 'from deconstruction to reconstruction'. This is based on the premise that the field of AI and law is primarily concerned with modelling the processes of legal reasoning and decision making with rule-governed and predictive models. Thus I take the view that in formulating theories of legal reasoning and decision

¹⁴ The actual requirements for the final expert system were that the system: (1) ask the user questions without relying on AUTOQUERY; (2) reach a conclusion based on the user's input and display that conclusion to the user; (3) use the MENU command to display headings of relevant cases in the database; and (4) allow the user to select a case heading and display to the user the text of that case summary. AUTOQUERY is a VP-Expert statement and MENU is a VP-Expert clause.

The AUTOQUERY statement causes VP-Expert to automatically address a question to the user if the inference engine requires a value for a variable for which there is no associated question. The question posed is, "What is the value of *variable*?" I considered that this 'rough' approach was inappropriate for legal expert systems where carefully worded questions are vital. Thus the students were required to draft their own questions to associate with variables, rather than relying on AUTOQUERY.

making on which we can base our models we must address those theories of law which deconstruct and bring into doubt the rule-governed nature of law. The deconstruction aspect involved a consideration of some theories of legal reasoning and theory that have sought to deconstruct legal reasoning, the legal system and the language of the law. My foray into deconstructionism began with the late Professor Fred Rodell of the Yale Law School who asserted:

"Legal words and concepts and principles float in a purgatory of their own, halfway between the heaven of abstract ideals and the hell of plain facts and completely out of touch with both of them. And that is why, in the last analysis, the language of The Law is inherently meaningless."¹⁵

Rodell may be described as a rule sceptic and was part of the school known as American Legal Realism. Another form of scepticism within the American Legal Realist movement may be described as fact scepticism. This distinction was devised by the late Jerome Frank who classified himself as a fact sceptic.¹⁶ Whereas rule sceptics are concerned with deconstruction of formal legal rules, according to Frank:

"...'fact sceptics' [are also] engaging in rule scepticism...But the fact sceptics go much further. Their primary interest is in the trial courts. No matter how precise or definite may be the formal legal rules,...no matter what the discoverable uniformities behind these formal rules, nevertheless it is impossible, and will always be impossible, because of the elusiveness of the facts on which decisions turn to predict future decisions. The fact sceptics, thinking that therefore the pursuit of greatly increased legal certainty is, for the most part, futile - and that its pursuit, indeed, may well work injustice - aim rather at increased judicial justice."¹⁷

Moving to the present, I then considered the Critical Legal Theory Movement which seems to have subsumed the bifurcation of deconstructionism contained in American Legal Realism.¹⁸ I found it

¹⁵ Rodell, Fred, *Woe Unto You, Lawyers!*, Pagent-Poseidon Ltd, Brooklyn, New York, 1959, at 136 as cited by Smith, J.C. & Deedman, C., "The Application of Expert Systems Technology to Case-Based Reasoning", *Proceedings of the First International Conference on Artificial Intelligence and Law*, (Boston), Association for Computing Machinery, New York, 1987, p. 84 at p. 86.

¹⁶ Frank, Jerome, *Law and the Modern Mind*, Preface to Sixth Printing, Peter Smith, Gloucester, Massachusetts, 1970, p. x.

¹⁷ *Ibid*, at p. xi.

¹⁸ The Critical Legal Theorists ("CRITS"), who might be described as the inheritors of the American Legal Realist legacy, share a similarly pessimistic view of the indeterminacy of law combined with a critical treatment of the legal process and system. A related deconstructive school is Radical Feminist Theory which focuses its deconstructive efforts on patriarchal and misogynist structures, institutions and attitudes within society. See, for example, MacKinnon, Catharine, *Towards a Feminist Theory of the State*, 1989, Harvard University Press. Whereas the American Legal Realists were somewhat anecdotal in their critiques, Radical Feminist Theory and Critical Legal Theory have

disappointing from the point of view of teaching the course that legal theorists interested in the field of AI and law, such as Susskind,¹⁹ appear to have overlooked Critical Legal Theory which I would regard as one of the most significant schools of legal theory of our day. Our concern with Critical Legal Theory in the field of AI and law should be primarily with its rule scepticism component and, particularly, the nihilist movement which holds that law is fundamentally indeterminate.²⁰ However, the field of AI and law appears to have largely ignored the implications of Critical Legal Theory and nihilism for the construction of legal expert systems.²¹ The course therefore examined closely rule scepticism and nihilism, and also touched very briefly on other critical aspects of Critical Legal Theory. This follows from my views about the multi-disciplinary nature of the field of AI and law and my belief that it is important that we are able to view the field in a context wider than that of the traditional boundaries of the study of law or computer science.

In the course I used the exchange between Joseph Singer and John Stick (and the commentaries of spectators²²) as a paradigm of the nihilist debate. Singer explains the indeterminacy thesis as: "The claim that a legal doctrine is indeterminate means that the doctrine allows choice rather than compelling it"²³ or that "legal theory is infinitely manipulable".²⁴ He defines a legal theory or set of rules as "completely determinate if it is comprehensive, consistent, directive and self-revising"²⁵ and then sets out to demonstrate how legal doctrine does not satisfy his definition. By comparison, Boyle expresses the indeterminacy thesis as "law is...totally manipulable, a grab bag of arguments sources, and examples of flip-flop rhetoric."²⁶

Without entering into the detail of the debate, it may be said that the focus of Stick's critique is on Singer's misrepresentation of the philosophers, particularly Rorty, on whom he relies for support.²⁷ Stick also suggests that Singer errs in implicitly ascribing to legal reasoning a deductive and axiomatic model when "no satisfactory theory

progressed to be more incisively analytical and far ranging in their deconstructive efforts.

19 eg. Susskind, Richard, *Expert Systems in Law*, Oxford University Press, 1987.

20 This focus does not seek to diminish or marginalize the valuable deconstructive contributions of Critical Legal Theory and Radical Feminist Theory to which I have briefly referred. These are at the forefront of contemporary social and legal theory and, I believe, have significantly altered the perspectives of those who have followed the scholarship.

21 For an exception, see Smith & Deedman, *supra*, note 15.

22 See, *infra*, note 30.

23 Singer, Joseph W, *The Player and the Cards: Nihilism and Legal Theory*, (1984) 94 Yale Law Journal 1 at p. 11.

24 *Ibid*, at p. 10.

25 *Ibid*, at p. 14.

26 *Supra*, note 3, at p. 1005. Boyle subsequently proceeds to teach his torts students how to manipulate and 'flip' legal arguments to support different positions.

27 Stick, John, *Can Nihilism be Pragmatic?*, (1986) 100 Harvard Law Review 332 at 332 to 345 and 392 to 401.

of legal reasoning exists".²⁸ He accuses the nihilists and Singer of "selective parsing of the language of a few theorists and judges (neglecting the hundreds of thousands of practicing attorneys)".²⁹ Other commentators have criticized Singer and the CRITS for merely 'trashing' liberals without seriously addressing the tenets of liberalism.³⁰ Obviously this debate is of some complexity which cannot be addressed in detail in this article; this brief description is simply to illustrate some of the central issues that were discussed in classes.³¹

I then wanted the students to consider the implications of nihilism and rule scepticism for the construction of legal expert systems. For if it is the case that legal argument, rules and doctrine are inherently 'mush', then the task of building a legal expert system with *legal* rules, using the language of the law, ought to be impossible. If legal rules are mushy then we must look for other sources of rulebases. As Smith & Deedman point out, one could construct an expert system based on the political beliefs, prejudices, preferences and prior decisions of judges.³² One might simply input the name of the judge and some information as to the theoretical, political or social issues raised by the case at hand and the likely decision would be the output. Smith and Deedman dismiss this approach for practical reasons because the identity of a judge is usually not known until very shortly before, or the actual day, of the trial. They state that even an expert system of this sort would need to be able to produce relevant cases and statutes of assistance to lawyers to support its predictions.³³ However, I do not think this disposes of the matter entirely. One could envisage how such an expert system might contain a database of cases previously decided by judges profiled according to the political or social issues raised in the case and the judges' decisions. It might be said that such a system would be founded on a Critical Legal Theorist's or Legal Realist's perspective of the legal system and legal reasoning. The expert system would therefore, as mentioned earlier,

²⁸ *Ibid.*, at p. 348.

²⁹ *Ibid.*, at p. 334.

³⁰ Williams, Joan C., *Critical Legal Studies: The Death of Transcendence and the Rise of the New Langdells*, (1987) 62 New York University Law Review 429 at p. 486; Hunt, Allan, *Law Confronts Postmodernism*, (1990) 35 McGill Law Journal 507 at p. 528.

³¹ In discussing this debate in classes, I drew on my experiences as a practising lawyer to which I compared Singer's assertion that "the claim that legal theory is infinitely manipulable expresses a universal experience of lawyers." See *supra*, note 23, at p. 10. My anecdotal account suggested that law is only partially indeterminate at any moment in time. Partial indeterminacy arises from the very nature of legal practice and the concept of the advocate in the adversarial system of justice. Lawyers inevitably represent different parties with different interests over a period of time and must adapt their arguments and submissions accordingly. Despite this feature of legal practice, judicial decision making still operates within a constrained framework as evidenced by the ability of experienced practitioners to predict legal outcomes and give sound advice in *legal language* to their clients on which their clients can, for the most part, rely in conducting their affairs. See Kowalski, Andrzej, "Case-Based Reasoning & the Deep Structure Approach to Knowledge Representation", *Proceedings of the Conference*, *supra*, note 4, at p. 21.

³² *Supra*, note 15, at p. 86.

³³ *Ibid.*

present an opportunity to 'empirically' test a jurisprudential theory. A practical problem of building such a system would be in defining a theory of relevancy for the domain; for when one adopts a critical and deconstructive approach to these issues, the curtain which traditionally separates what is law from what is not law is torn down. Nevertheless, I believe there is great potential for the field of AI and law to build systems based on a variety of theories of legal reasoning. The course therefore considered the appropriateness of other legal theories, such as positivism, natural law theories and sociological jurisprudence, as foundations for the construction of legal systems.³⁴

The next step was to consider, in light of the philosophical and theoretical issues, the desirable features of a methodology for the design of legal expert systems. This is the 'construction' aspect of my approach to the teaching of legal theory and reasoning. When developing a methodology, one must bear in mind that one of the principal goals of the field of AI and law is to model with computers the processes of legal reasoning and decision making. Further, the models must be capable of producing accurate advice and predictions with supporting case and statutory authority. One of the fundamental underlying assumptions of the field of AI and law is that legal decision making and reasoning is a predictable and rule-governed process. The quest is to find the rules. Rules are after all, in one sense or another, the fundamental building blocks of knowledge representation. Even sophisticated expert systems, such as case-based reasoners, are founded on a rule-based structure. The attempt to build a legal expert system, whether rule-based or case-based reasoning, necessarily involves the knowledge engineer basing the system on his or her view of the structure of the domain.³⁵ Even though the rules of the system may not be explicitly stated at the structural level, they are readily discernible from the structure. This becomes particularly evident if the knowledge engineer uses an object-oriented approach or frames to represent the domain.³⁶ In enumerating our desired features of a methodology for expert system construction, we considered Deedman's three criteria:

- (1) Law must be an autonomous system of rules.
- (2) Judges must not have a discretion in deciding hard cases; and
- (3) Law must be teleologically neutral.³⁷

I then turned to the 'deep structure approach' as a methodology for constructing legal expert systems. The deep structure approach to knowledge representation has been well documented elsewhere.³⁸ Suffice it to say that it involves a search for patterns of *factual* attributes of cases which may explain judicial findings. The thesis of the deep structure

³⁴ Smith & Deedman, *supra*, note 15 at pp. 84 to 89. They describe positivism as "a mild form of rule skepticism" (p. 87).

³⁵ Kowalski, *supra*, note 31, at p. 23; Kowalski, Andrzej, *Beyond Rule-Based Legal Expert Systems: Using Frames and Case-Based Reasoning to Analyze the Tort of Malicious Prosecution*, LL.M. Thesis, The University of British Columbia, 1990, at p. 37.

³⁶ See, *infra*, notes 44, 45 and 46.

³⁷ Deedman, Cal, *Building Rule-Based Expert Systems in Case-Based Law*, LL.M. Thesis, University of British Columbia, 1987, at p. 41.

³⁸ Smith & Deedman, *supra*, note 15; Kowalski, *supra*, note 31. See also McCarty, *supra*, note 5, advocating a deep structure approach.

approach is that judges and lawyers reason in a rule-governed and predictable manner, even though they may not necessarily be consciously aware of the process. It eschews reliance on, or reference to, the language of the law, legal concepts, rules or doctrine where the area or point of law is characterized by indeterminacy. By identifying factual attributes which show a pattern of significance throughout a body of cases, one can develop a structure and *meta* rules for the domain which may be used to analyze the domain and the cases. The deep structure approach thus provides a theory of relevancy for the domain.³⁹

The students were asked to undertake an analysis of their domains in terms of a deep structure approach. Despite being an obvious predilection of the teacher, it was rationalized on the basis that the limited time available to teach the course combined with the goal of producing working expert systems necessitated adopting and learning thoroughly one methodology of knowledge representation in terms of the students' expert systems. It was made clear to the students that the deep structure approach is merely one method among many of knowledge representation.⁴⁰

D. Artificial Intelligence and Expert Systems

The students were given a brief version of the history of AI from the Dartmouth Conference through to the present.⁴¹ More time was devoted to legal expert systems as a subset of the field of AI. Demonstrations of working systems were important aspects of the early classes to give the students an idea of how their own finished systems might look.

Time being of the essence, the students were expected to begin the process of designing their systems as soon as possible. While researching the case law in their domains, I wanted them to bear in mind the deep structure approach and focus on the search for a set of factual attributes. It was thus important from the outset to explain in detail the concept of a database and how the databases would be linked to and be used by their expert systems. It was also vital that they comprehend the relationship between the deep structure approach, the rulebase and the field names of the database.

Important aspects of expert systems such as backward chaining and forward chaining were discussed. It was interesting to note how intuitively familiar the process of backward chaining was to them. They were very much at home with the concept of a goal eg. the client wants to sue, and working backward to attempt to prove the goal. The process

³⁹ On the issues of relevancy and context, we looked at briefly some linguistic debates on the nature of legal language: Levinson, Sanford, *Law as Literature*, (1982) 60 *Texas Law Review* 373; Graff, Gerald, "Keep off the Grass," "Drop Dead," and Other Indeterminacies: A Response to Sanford Levinson, (1982) 60 *Texas Law Review* 405.

⁴⁰ On the importance of knowledge representation in the field of AI and law see McCarty, *supra*, note 5.

⁴¹ The term 'artificial intelligence' was coined by John McCarthy at a Summer Research Project held at Dartmouth College in 1956. This Conference is generally viewed as the official birth date of the field of AI.

of backward chaining through a set of rules required only a brief description and they treated the process as self-evident in terms of legal reasoning and their own legal experience.

The students' intuitive rapport with backward chaining meant that they could readily grasp the computer programming aspects of this process. They were therefore in a position to fairly quickly commence drafting and experimenting with sets of rules. More technical programming aspects such as database retrieval and case display were taught by providing copies and explanations of source code from other legal expert systems. Since the students did not have to write this code from scratch, they had more time to concentrate on the more important issues such as logic flow and knowledge representation. Notwithstanding the availability of pro forma code, there was a great deal of variance in the approaches adopted by the students for matters such as case retrieval and display, help screens and the display of conclusions. Each approach was quite unique and there were no slavish duplications of the pro forma code.

Other aspects of legal expert systems were considered including:

- case-based reasoning and comparison of my approach to the approach of other researchers such as Ashley,⁴²
- comparing the benefits of rule-based systems with case-based reasoning systems.⁴³
- frames,⁴⁴ object-oriented programming,⁴⁵ Schank and Abelson's theories of scripts and memory structures⁴⁶ and the process of legal classification.⁴⁷
- developments in natural language processing systems to enable the bottle-neck of data acquisition, database construction and knowledge representation to be overcome.⁴⁸

E. Problems

⁴² Ashley, Kevin D., *Modelling Legal Arguments: Reasoning with Cases and Hypotheticals*, M.I.T. Press, Cambridge, Mass., 1990.

⁴³ Kowalski, *supra*, note 35, at p. 34; Popple, J, *Legal Expert Systems: The Inadequacy of a Rule-Based Approach*, February 1991, *The Australian Computer Journal*, Vol 23, No 1 at p. 11.

⁴⁴ Kowalski, *supra*, note 35 at p. 15; Parsaye, Kamran & Chignell, Mark, *Expert Systems for Experts*, 1988, John Wiley & Sons Inc, pp. 48 to 67, 161 to 210.

⁴⁵ Vossos, Dillon, Zeleznikow & Taylor, *The Use of Object Oriented Principles to Develop Intelligent Legal Reasoning Systems*, 1991, Vol 23, No 1, *The Australian Computer Journal*, at p. 2.; *OOP in the Real World: A White Paper from Borland International*, P.C. A.I., September/October 1989, at p. 37.

⁴⁶ Schank, Roger C. & Abelson, Robert P., *Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures*, 1977, Lawrence Erlbaum Associates, New Jersey.

⁴⁷ Leff, Arthur Allen, *Contract as Thing*, (1970) 19 *American University Law Review* 131.

⁴⁸ Gelbart, D. & Smith J.C., "Beyond Boolean Search: FLEXICON, A Legal Text-Based Intelligent System", *Proceedings of the Conference, supra*, note 4, at p. 225.

Almost all of the practical problems experienced during the course may be traced back to the fact that at that time the Law School had not appointed a system administrator for the student computer laboratory. We experienced network and software crashes, viruses and printer problems and the eternal mystery of disappearing files. In the absence of a system administrator, I was the first person to be called upon to solve problems and when they were beyond my limited technical skills we had to call in the University's Information Technology Services department where response time was sometimes a problem. Students complained about the awkwardness of formatting and merging text files of case briefs for use by the VP-Expert SHOWTEXT command. It would have also been desirable to offer a short seminar in basic MS DOS techniques at the start of the course, especially considering there was no prerequisite of computer skills for the course, and to have somebody available on a daily basis to deal with simple questions about copying disks and other administrative matters. Now that a system administrator has been appointed, we expect future courses to proceed with minimal technical hiccups and the system administrator will play an active role in teaching and assisting students. From my experience, I would strongly recommend against teaching an AI and law course without the assistance of a system administrator.

One problem which was not anticipated was that most students quickly exceeded the limits of the student edition of VP-Expert. The most irksome limits for the students were the limit on the size of the knowledge base requiring them to chain knowledge bases together⁴⁹ and the undocumented limit on the SHOWTEXT command of a text file of approximately 64K. Students commented that they suspected there were other undocumented limitations built into VP-Expert eg. levels of nesting, which caused their systems to randomly crash sometimes with an error message "The limitations of this edition have been exceeded". There appears to be no ready solution to this problem since there are no other viable shells at a similar price to VP-Expert.

Implementing the deep structure approach proved to be a major obstacle for the students. I found myself returning to explanations of the deep structure approach on many subsequent occasions and assisting the students in developing a deep structure in their domains. This may be partially attributed to the nature of the students' legal education in that they had acquired over time (probably not without difficulty) an understanding of traditional methods of legal analysis. They were then being asked to discard this recently acquired knowledge for a new set of analytical structures and principles. This task was made even more difficult by the students' history of exposure to the traditional reluctance in the teaching of 'hard law' subjects to articulate "the arguments explicitly and overtly" and identify the theoretical presuppositions on which the area of law is *really* based.⁵⁰ Whether we like it or not, the field of AI and law, in its attempts to model legal processes and reasoning, is inextricably concerned with deconstruction and identification of the true constructs of legal domains.

⁴⁹ The Fair Comment group was forced to chain seven knowledge bases for a total knowledge base size of 64K.

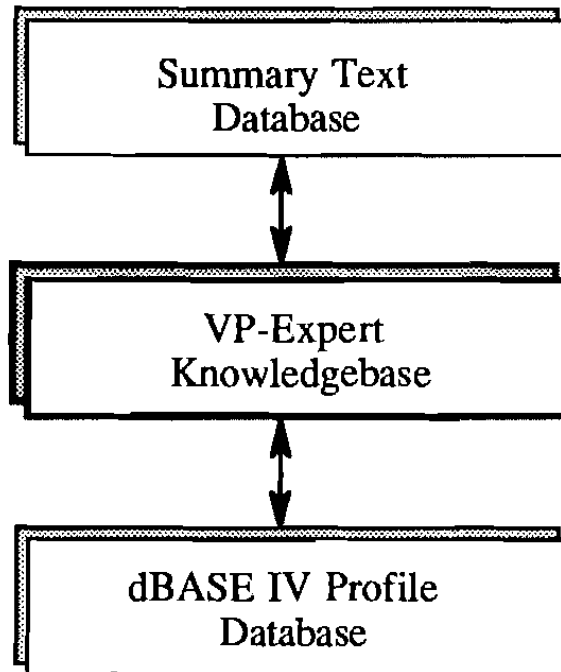
⁵⁰ See Boyle, *supra*, note 7.

F. The Completed Expert Systems

The completed expert systems were based on the following design architecture:

Figure 1

Design Architecture of Students' Expert Systems



The Text Databases contain the summaries of the cases prepared by the students.⁵¹ A user may thus view a summary of any case displayed by the expert systems during a consultation. The dBASE IV Profile Databases contain 'profiles' of cases.⁵² The schemata of these databases was based on the following model:

Figure 2

Schemata Model of dBASE IV Profile Databases

NAME:
CITATION:
YEAR:
COURT:
JURISDICTION:
HELD_FOR:
ATTRIBUTE1:

⁵¹ Preparation of these summaries was the first assignment referred to in Section B, *infra*.

⁵² Preparation of these databases was the second assignment referred to in Section B, *infra*.

|
|
|
|
|

ATTRIBUTE_n:

The fields NAME through to HELD_FOR contain the formal descriptive details of the cases. The remaining fields represent the factual attributes of the domain considered by the students to be significant after undertaking a review and analysis of the case law by the deep structure approach. The VP-Expert knowledgebases contain the rulebases and questions developed by the students.⁵³ The rulebases are based on a backward chaining strategy which seeks to prove the goal of establishing a cause of action or a defence, depending on the domain. In conventional expert system style, the user is lead through a series of questions, based on the factual attributes derived from the deep structure analysis, designed to elicit the user's fact situation. Whenever the expert systems reach a conclusion, they substantiate the conclusion with reference to relevant cases. The retrieval of relevant cases was basically by pattern matching between the user's answers and profiles of cases in the dBASE IV Profile Databases.

As discussed in Section E, some students had difficulties in applying the deep structure approach. However, for the most part, students were able to develop alternative structures for difficult or indeterminate areas or points of law by identifying significant factual attributes of cases. The defence of fair comment to a defamation action is a good example of a difficult area of law which the Fair Comment Group was able to successfully model. Some of the difficult issues the students had to contend with are:⁵⁴

- The defamatory material must contain sufficient reference to the facts on which the comment is based.
- The person to whom the material is published reader of the material must be able to identify the material as a comment.
- There must be a sufficient basis of fact in the material to support the comment or the material must be protected by 'absolute privilege'.
- The comment must not have been affected by malice.
- The matter must be of public concern.

The students' approach was to develop a deep structure set of factual attributes to 'objectively' analyse each of these issues. During a consultation with the Fair Comment Advisor, the user is given the option of relying on his or her own subjective opinion about the above issues or have the expert system determine them 'objectively' by a deep structure analysis. From purely factual questions, the Fair Comment

⁵³ This was the second and third assignments referred to in Section B, *infra*.

⁵⁴ Walker, Sally, *The Law of Journalism in Australia*, Law Book Company, 1989, pars. 3.9.10 to 3.9.17.

Advisor is able to determine these issues with a reasonable degree of accuracy.⁵⁵

There are other difficult issues in this area of law. For example, can the defence operate where the defendant publisher was not actuated by malice, but the person who supplied the information to the publisher was actuated by malice? The answer to this question may depend upon whether the supplier of the information was an employee of the publisher, such as a journalist, or a stranger.⁵⁶ One can see how difficult it would be for a lawyer to determine these issues, let alone a legal expert system. The Fair Comment Advisor takes account of these complicated issues and provides guidance as to the likely outcome of a fact situation dealing with these issues.

The Medical Negligence Group adopted an interesting approach, which provides an excellent example of deep structure research, in attempting to define a set of deep structure factual attributes for a negligence action based on the failure of a medical practitioner to disclose to the patient a known risk inherent in a medical service. The students started from the premise that a medical practitioner will be justified in not disclosing the risk if it can be shown that the patient/plaintiff was 'incapable' of using the information rationally or if the medical practitioner believed on reasonable grounds that disclosure would have been harmful to the patient. In terms of the latter scenario, they identified a number of factors from the cases which appeared to be relevant:

- The chance of success of the proposed treatment.
- The risk to the patient of not proceeding with the proposed treatment.
- Whether the risk of proceeding with the treatment exceeded the risk of not proceeding.
- Whether there was an alternative to the proposed treatment.
- The gravity of the unexplained risk if it eventuated.
- The probability of the unexplained risk eventuating.

The students went on to explain:

"All of these factors seem to be weighed up by the courts. However, they give rise to extremely complicated rules, in which only some of the possible combinations of factors are covered in the available cases...The solution to the dilemma here was to identify a pattern within the 'harmful' factors."

"Initially it seemed that a defendant medical practitioner could justify withholding information when *chance_success* = high and *risk_not_proceed* = high, but not if *chance_success* was low and *risk_not_proceed* was medium. This left other combinations unclear, until a case was found where the plaintiff succeeded where *chance_success* was low and *risk_not_proceed* was medium, and

55 Members of the Law School specialising in media law have evaluated the Fair Comment Advisor and were impressed by the accuracy of its conclusions and the supporting authorities which it cites.

56 Walker, *supra*, note 54.

another where the plaintiff succeeded where *chance_success* was high and *risk_not_proceed* was low. This allowed a simple rule to be written along the lines that if *chance_success* was high and *risk_not_proceed* was high, then 'harmful' was satisfied and the medical practitioner was justified in not giving the patient information. This may be a simplification, but it seems to work and it accords with commonsense, i.e. that a medical practitioner would have a 'good reason' to negate a patient's right of choice."⁵⁷

These examples indicate the quality of the analysis that the students applied to their domains. All of the final expert systems were of a very high standard and well beyond my expectations and requirements. Some were exceptional in the amount of work, the knowledge representation analysis and finished product. Whereas only a rule-based system was required of the students, one student in the Assault Group implemented a case-based reasoning override so that the system would check for inconsistencies between the rulebase and the cases in the database and then resolve the inconsistency by assessing the precedential weights of cases. The solo student (Breach of Confidential Information) produced a first class system with an excellent multi-level help facility. Members of faculty have been evaluating some of the systems and the reports as to the validity of the legal analysis embodied in the systems have been very encouraging.

G. Student Feedback⁵⁸

Many of the students commented that the course workload was quite substantial for a one semester subject and suggested that the course be taught over two semesters. This is a reasonable suggestion as such, but it should be remembered the technical difficulties discussed above tended to increase the students' workloads. Further, to the students' credit, the technical difficulties did not prevent them from submitting expert systems of a very high standard. With a system administrator and a fully functional computer laboratory, I believe the course can still be taught over one semester, although a two semester course would certainly allow more time to be spent on various sections of the course. It should also be remembered that a one semester course is attractive to many students from the point of view of finishing it quickly and 'getting it out of the way'. However, given the encouraging number of enrolments and the obvious interest in this field, I would not rule out running the subject over two semesters. The primary advantage of a two semester course would be extra time to consider some of the theoretical issues in depth before commencing construction of the expert systems.

There was a great variety of personal reactions to the theoretical component of the course. This is to be expected in any course which deals with theory or philosophy. Perhaps the significant point is that there were certainly many spirited discussions and debates during the course. As mentioned earlier, the students required on going assistance

⁵⁷ Medical Negligence Group.

⁵⁸ One student from each group was requested to write down their impressions of the course including comments, criticisms and suggestions. I wish to thank Kaye Hargreaves, Theo Zervas, Lewis O'Brien, Jason Harrop and Matthew Read for their time and efforts.

in order to understand and apply the deep structure approach. The range of comprehension or acceptance varied from excellent to not at all: "despite my persistent questioning the concept of 'deep structure' was never articulated to my satisfaction."⁵⁹ This was an inherent difficulty with the course.

Having completed the course, it is interesting to note some of the students' reactions to the future and viability of the field of AI and law:

- "...the use of an expert system specific to an area of law that can be used in practical situations with a relative degree of confidence was indeed satisfying and thought provoking with respect to the various applications and domains of law that it can be applied to. The commercial possibilities that arise are outstanding and fellow students who tried the systems related their astonishment and amazement."⁶⁰
- "Whilst one may be able to go some way towards closing the system by resolving questions raised in academic commentary on the doctrine, the work of a court in a democratic technically advancing community requires continuing flexible consideration of pre-existing doctrine. How can a computer hope to dynamically re-frame rules in this manner, when it is divorced from the essential experience of being alive and part of a society?"⁶¹
- "Computers are like dogs. You can play with them and make them do things but their intellectual capability remains limited. EXPERT SYSTEMS, as revealed in the course, consist of human EXPERTise applied to computer SYSTEMS. As yet the artificial intelligence is not in evidence. The exciting part is the feeling of being present at the beginning of great new developments. For that privilege, any amount of frustration at system crashes is worthwhile."⁶²

IV. CONCLUSIONS

The quality of the expert systems produced by the students clearly demonstrates that legally trained people with little or no computer skills can successfully build legal expert systems of a reasonable degree of sophistication within a relatively short time frame and do so using a very inexpensive shell. I believe that a major factor contributing to these results was the due regard paid to the importance of legal theory and reasoning during the course. The quality of the work is good evidence of the students' enthusiasm for this new discipline and certainly dispelled any questions I entertained about the students opting for safer and more traditional subjects and forms of assessment.

If law students are able make such advances in the field of AI and law, then the same ought to be true for practising lawyers, although their formal training in legal theory may be buried in the recesses of their legal

⁵⁹ Kaye Hargreaves, Fair Comment Group.

⁶⁰ Theo Zervas, Fair Comment Group.

⁶¹ Jason Harrop, Breach of Confidential Information.

⁶² Kaye Hargreaves, Medical Negligence Group.

education. However, I suggest that their real world experiences provide an equally valid theoretical basis. Indeed, as discussed earlier, there is a nice congruence between the deep structure approach and the experiences of practising lawyers. Perhaps the greatest barrier to the acceptance of legal expert systems by practising lawyers will be the traditionally technophobic mindset of the legal profession.

While the students' expert systems are only research projects and not commercial products,⁶³ the skills developed by the students should stand them in very good stead for more substantial future applications. As Bob Moles suggests in his article in this edition of this Journal, "we are now producing a new generation of students qualified in both law and computer science...[who] will be both able and willing to work with the legal theorists". Even if some of them may have resolved never to build another legal expert system, they should all, nevertheless, have benefited from the enquiries into the nature of legal theory and reasoning and be well prepared for the future impact of AI on the practice of law. In this latter capacity they should be able to make informed decisions and judgements on expert systems and AI tools, which will no doubt proliferate in the legal market.

⁶³ With some notable exceptions there are very few legal expert systems available commercially or used to any great extent in commercial markets. For an exception see Mead, David & Johnson, Peter, "Legislative Knowledge Base Systems for Public Administration - Some Practical Issues", *Proceedings of the Conference, supra*, note 4, at p. 108.