

Artificial Intelligence, Expert Systems and Law

*Richard E Susskind**

This paper introduces and explains the various branches of Artificial Intelligence and Law, summarises past and present activities within both Europe and the rest of the world and, on a practical level, considers the feasibility and the benefits of these systems as well as the problems to which they give rise. It then points to the keys to successful development.

I ARTIFICIAL INTELLIGENCE AND LAW

Artificial Intelligence (AI) is concerned with the design, development and implementation of computer systems that can perform tasks and solve problems of a sort for which human intelligence is normally thought to be required. For example, AI programs have been written to understand the spoken word, to translate from one language into another, and to recognise images and objects in the physical world.

Artificial intelligence as applied in the legal field can be sub-divided into two categories: expert systems and knowledge-based systems; and enhancements to legal information retrieval systems.

Expert systems and knowledge-based systems

The broadly agreed goal of workers in the fields of expert systems and knowledge-based systems is to use computer technology to make scarce expertise and knowledge more widely available and easily accessible. Expert and knowledge-based systems are therefore computer applications that contain knowledge and expertise which they can apply – much as a human expert does – in solving problems, offering advice and undertaking a variety of other tasks. In law, these systems should be able to apply their legal knowledge in guiding users

* This paper is based on a report submitted to the Council of Europe in January 1990, entitled "Systems Based on Artificial Intelligence in the Legal Field".

Richard E. Susskind is Special Advisor on Law and Information Technology at Masons Solicitors, London. He is Visiting Professor at Strathclyde University's Centre for Law, Computers and Technology and is Chairman of the Society for Computers and Law.

through complex legal issues; in identifying solutions to problems; in planning tasks; compiling documents and managing the flow of cases; and in offering advice and making specific recommendations.

If there is any distinction at all between expert systems and knowledge-based systems, it is that the former are more powerful than (and are a subset of) the latter in that the former hold expertise and not just knowledge. There is a more fundamental set of distinctions, however, and that is between types of expert systems (the term hereafter used to refer both to expert systems and knowledge-based systems). There are currently five identifiable classes of expert systems as applied in law.

1. *Diagnostic systems*

These systems offer specific solutions to problems presented to them. From the facts of any particular case, as elicited by such a system, it will analyse the details and draw conclusions, usually after some kind of interactive consultation. These systems are analogous to the medical diagnostic systems which make diagnoses on the basis of symptoms presented to them. An example of a diagnostic system in law would be a taxation system that could pinpoint the extent to which and why a person is liable to pay tax, doing so on the basis of a mass of details provided to it.

2. *Planning systems*

In a sense, planning systems reason in reverse. For these systems are instructed as to a desired solution or outcome and their purpose is to identify *scenarios*, involving both factual and legal premises, that justify the preferred conclusion. In tax law, a planning system could recommend how best a taxpayer should arrange his affairs so as to minimise his exposure to liability. The knowledge held within planning systems can be very similar to that held within diagnostic systems; what is quite different is the way that that knowledge is applied.

3. *Procedural guides*

Many complex tasks facing legal professionals require extensive expertise and knowledge that is in fact procedural in nature. Expert systems as procedural guides take their users through such complex and extended procedures, ensuring that all matters are attended to and done within any prescribed time periods. An example of such a system would be one that managed the flow of a complex tax evasion case, providing detailed guidance and support from inception through to final disposal.

4. *The intelligent checklist*

This category of system assists in auditing or reviewing compliance with legal regulations. Compliance reviews must be undertaken with relentless attention to detail and extensive reference to large bodies of regulations. Intelligent checklists provide a technique for performing such reviews. They formalise the process. In

taxation, an intelligent checklist approach could be used to assist in the review of a company's compliance with corporation tax.

5. *Document modelling systems*

These systems – also referred to as document assembly systems – store templates set up by legal experts. These templates contain fixed portions of text together with precise indications as to the conditions under which given extracts should be used. In operation, such a system will elicit from its user all the details relevant to a proposed document. This is done by the user answering questions, responding to prompts and providing information. On the basis of the user's input, the system will automatically generate a customised and polished document on the basis of its knowledge of how its text should be used.

Enhancements to legal information retrieval systems

Although legal information retrieval systems such as *Lexis* have dominated the field of computers in law, these systems nevertheless have serious shortcomings. Often they retrieve an excess of irrelevant documents, or alternatively not all pertinent documents within the data-base are located during every consultation. There is much more to legal problem-solving than searching for the occurrence of words within documents and, in recognition of this, researchers have sought to improve the performance of legal information retrieval systems by using AI techniques. Two approaches can be adopted here: the introduction of a “front-end”; and the deployment of conceptual retrieval techniques.

1. *Front-ends*

The idea here is that guidance on the use of an interactive system is offered to users before searching within the database commences. Such a front-end system will help the user actually formulate his search request. It will ask the user questions, help to pinpoint relevant terms and concepts, and help ensure that all but only the relevant documents are retrieved. Front-ends can be built using classical “knowledge engineering”/AI techniques – they can embody the knowledge of experts in the use of legal information retrieval systems and make that expertise available to others.

2. *Conceptual retrieval*

Those who argue that there is more to legal problem-solving than searching for keywords will often suggest that legal reasoning and research involves familiarity and manipulation of legal *concepts*. Accordingly, researchers have developed methods of augmenting the traditional approach with conceptual retrieval techniques which will allow users to search through massive bodies of legal data, not just on the basis of the occurrence of keywords but in terms of the fundamental

concepts relating to any problem at hand. This could mean, for example, that systems will search not only for words expressly articulated by the user, but also for terms conceptually implicit in such requests.

The above analysis focuses on the functions of AI systems in law but says little of the enabling technologies. In developing expert systems or enhanced legal information retrieval systems, researchers have drawn and will continue to draw on a wide range of AI techniques. Over and above the techniques and methods normally associated with expert systems, workers in artificial intelligence and law increasingly make use, in particular, of natural language processing and neural computing.

II HISTORICAL ANALYSIS

The purpose of this section of the paper is to put the field of AI and law in historical context. It proposes a four stage model in terms of which activities in the field can be analysed and classified.

The four stages of evolution in the field of AI and law

Activities in the field of law and AI tend to progress through four key stages: negligible activity; preliminary research; extensive research and development; and commercial exploitation. Although these stages overlap, each has characteristics unique to it.

1. Stage 1 – Negligible activity

During this stage, there are no sustained attempts to carry out serious scientific investigations into the topic. Nevertheless, and even in the days prior to the coming of the computer, in this first stage there may still be considerable speculation about what might be called the mechanisation of the legal problem solving process. Some contributions in this connection may be no more than fictional, but others show remarkable foresight of the potential and the dangers of computerising legal reasoning.

2. Stage 2 – Preliminary research

Eventually, speculation and fiction give way to the desire to inquire into the field more thoroughly and rigorously. In this second stage, the preliminary research will lay the foundations for later and larger initiatives, but at this stage work is confined largely to exploring AI and law from a theoretical perspective. Researchers may be from law as well as from the world of computing. Perhaps surprisingly, Stage 2 is often dominated by the latter tradition – it is computer scientists' basic research that will establish the technical feasibility of building AI systems in law and will stimulate lawyers into further inquiry.

3. Stage 3 – Extensive research and development

The potential of the field having been established during Stage 2, there will follow a period of great activity, involving many research projects, largely devoted to the development of prototype systems and demonstrators. At the same time, fundamental, theoretical research will continue and will have direct impact on the research and development projects. In this stage, as in the previous two, work will largely be confined to academic establishments. The findings of Stage 3 will progress AI in law from the research laboratory into the marketplace.

4. Stage 4 – Commercial exploitation

In this final stage, commercial organisations explore the technology with a view to profitable development. These organisations will draw heavily from Stage 3 activities, not simply in borrowing ideas but also in recruiting staff. The stage of commercial exploitation itself can be subdivided into several phases, each representing varying degrees of success and financial gain. The key feature of Stage 4 activities is their commercial orientation – there is little concern for theory, although it must be stressed that Stage 3 research and development will still continue (at increasingly advanced levels) in parallel with Stage 4 activity.

A brief history of artificial intelligence and law

The four stage model set out above can be used in analysing the worldwide history of the field.

Stage 1 would correspond to the time period before 1970. During that period, visionaries, science fiction writers, and technologically oriented lawyers wrote on the topic of computers engaging in legal reasoning and even replacing judges. Interesting though these contributions were, they were often detached from the technical realities and from the nature of the legal process.

It was not until 1970 that work began in earnest. This was the beginning of about five years' preliminary research – Stage 2 activity. Buchanan and Headrick, a computer scientist and lawyer from Stanford University, published the first detailed analysis of the field in 1970. Shortly afterwards, a man who is now considered to be the father of the field, Thorne McCarty, began his TAXMAN project, which was to run well into the 1980s. Other key projects during that period were carried out by Meldman in the United States, Popp and Schlink in Germany, Sprowl in the United States and Stamper in England. Collectively, these workers and a few others, undertook the preliminary research that has served as the foundation for more ambitious work over the last fifteen years.

During the decade between the mid-1970s and the mid-1980s there emerged about 30 sustained research and development projects throughout the world. These projects built upon the Stage 2 work, lending further credence to it by developing operational systems that demonstrated the potential of the field so much more clearly than abstract reports could ever have done. Important projects

during this stage were conducted at the Rand Corporation in California, Imperial College in London, the Norwegian Research Centre for Computers and Law in Oslo, and in England's Open University and Oxford University. It is interesting to note that the work of these projects – during the world's Stage 3 – emanated more from institutions than from individuals. The findings and products of these research and development programs attracted interest not only within the computing and legal communities but also in the popular press. In turn, the commercial world began to take greater interest.

From 1985 onwards, Stage 4 commercial exploitation came about. The findings of earlier research projects were combined with the teachings of traditional data processing and constrained by the pragmatic demands of profit-making organisations. It must be stressed, however, that even on a world-wide basis, commercial exploitation of AI in the legal field is still, in 1990, at a very early phase, for returns on investments so far have been low (with the exception of the use of document modelling systems in the United States).

It should be said, finally, in this brief history of AI and law, that there have now emerged, quite clearly, two types of worker in the field. On the one hand, there is the 'pragmatist' whose overriding aim is to develop and implement commercial systems that can actually assist in the solving of legal problems. On the other hand, there is the 'purist', for whom completion even of modest prototypes is not always necessary for success. The major goals of the purist are to clarify the nature of legal reasoning, of human and artificial intelligence and of computational models of law. Live systems are not necessary for this. Generally, pragmatists operate in the commercial world, while purists can be found in research establishments. Pragmatists are at Stage 4 of the evolutionary path outlined earlier; while purists remain at Stage 3, often with no intention of being involved in Stage 4 activities.

In assessing contributions to the field of AI and law, it is important to bear in mind this distinction between pragmatist and purist approaches. It is crucial to appreciate that workers in these camps have in mind quite different goals and their orientations often diverge radically. There must be room in the world of AI and law for both pragmatists and purists. Indeed it is essential that contributions emerge from both camps. Equally crucial is that neither claims superiority over the other. Above all, perhaps, collaboration between the two is desirable.

III PRACTICAL ISSUES

This section seeks to offer answers to four questions often asked by persons interested in exploiting the potential that AI and expert systems techniques offer for the law. (Note that hereinafter the term "expert system" is often used in place of "AI" as it is this branch of AI that is attracting the greatest commercial interest.) The four questions are as follows:

- Is it feasible to build artificially intelligent systems in the legal domain?

- What are the benefits of this approach?
- What problems are there for workers in this field?
- What are the keys to success?

This section deals with each of these questions in turn.

Questions of feasibility

People unfamiliar with the field will often say that it is not possible or not feasible to build expert systems in law. Careful analysis of this assertion, however, reveals that there are really five dimensions to this feasibility issue; that there are really five questions to be answered. These questions ask whether building expert systems in law is technically possible, jurisprudentially sound, commercially viable, organisationally suitable and strategically appropriate.

1. Technical possibility

The key issue here is whether hardware and software in the field of AI has developed to such an extent that sound, reliable and robust systems can be designed, developed, implemented, tested and maintained. It is now widely accepted that expert systems have matured to such an extent that technologies and techniques are indeed now available, in 1990, to support the development – at least – of what might be termed ‘first generation’ systems. Generally, such systems will be stand-alone, will operate in small problem domains, will not be capable of solving all problems that human experts can, but nevertheless will deliver appreciable business benefits.

2. Jurisprudential soundness

Any expert system in law necessarily makes assumptions about the nature of law and of legal reasoning. Accordingly, there is a growing literature on the philosophical implications and presuppositions of building such systems. Some theorists have suggested that building such systems is to misrepresent, distort and oversimplify the legal problem-solving process. Others have argued there to be a direct match between what the computer can do and what legal reasoning is all about. However, the view that has been most widely supported, and is now gaining even greater acceptance, is that there are no fundamental objections from the point of view of legal philosophy to building expert systems in law of limited scope. These limitations in scope refer to the category of proper user (the lawyer or legally informed person); the way in which such a system should be used (as an ‘intelligent assistant’ and not a replacement for a legal expert); and the limitations of solving problems on the basis only of legal *rules* (so that these systems currently cannot reason on the basis of legal principle, policy or purpose).

3. *Organisational suitability*

Even if it is technically possible and jurisprudentially proper to build these systems, it must always be asked whether such a system could actually function effectively in any target organisation. Some users of the systems might feel de-skilled; others may feel uncomfortable about interacting with computer systems rather than human beings; while still others will lack the confidence to operate any computer system of whatever sort. This question of feasibility is often overlooked by enthusiasts and champions of the field. Yet, from a purely practical point of view, many systems of the future will be regarded as failures precisely because they will not have been integrated with and accepted by the organisation for which they were developed.

4. *Commercial viability*

The costs associated with developing expert systems in law are considerable. Not only are skilled computer professionals required, but if possible extensive time, effort and therefore cost of human experts will also be expended. For an expert system in law to succeed in the commercial world, the benefits that accrue from its deployment must outweigh these substantial outlays. A major difficulty here, however, is the notorious difficulty of quantifying the benefits of this technology: for many of the benefits, as shall be seen from the next part of this report, are of an intangible sort. A major challenge for this field, therefore, is to provide guidance on the compilation of cost/benefit analyses of systems.

5. *Strategic relevance*

Complete commitment to this emerging technology may be further inhibited by a perception that even if these systems can be built, their operation would not sit comfortably with the nature of the organisation for which they are being developed. It may be, for instance, that a firm of solicitors decides not to be heavily involved with any sort of technology, but to offer a distinctively 'human' service. For such a firm, even if systems could profitably be developed, they may be deemed to be strategically inappropriate. The strategic question associated with these systems is, therefore, whether their use supports the wider strategic and business objectives of the organisation considering their introduction.

Significantly, since 1987, the first two questions – about technical and jurisprudential feasibility – have been asked less frequently. Operational systems have themselves, in effect, responded positively. These two questions are the key concerns of those in Stage 3 of the evolutionary path. In the world arena, the focus now is on the third and fourth questions – whether these systems can offer sufficient payback and fit into organisations considering their introduction. Frankly, it is too early to be able to answer these questions with confidence. Early indications of Stage 4 are very encouraging; so what seems certain is that further investment and investigation is necessary. The fifth question, the one relating to strategy, will be the dominant question of the mid- to late 1990s.

The benefits

There are two main categories of benefits. First, there is the benefit from a human resource perspective. Second, there is the impact on quality.

1. *Human resources*

The central idea of expert systems technology is to allow human knowledge and expertise to be distributed more effectively and efficiently. This dissemination of scarce expertise will give rise to a range of improvements in the way human resources can be managed. Using the technology, complex tasks, which in the past would have required the attention of experts, can now be reliably delegated to less experienced persons. These users will benefit from access to first-rate expertise that would have been possible previously only through direct interaction with human experts (whose time would inevitably have been limited). This possibility of delegating tasks and activities could go some way to overcoming the anticipated skills shortages that are likely to prevail in the 1990s. There are training implications here too: in operating these systems, users will gain considerable insight into the knowledge and techniques necessary and sufficient for first class performance in the legal problem-solving arena.

Additionally, expert systems will allow the expertise of many experts to be synthesised and preserved. More than this, they will also perform a liberating function. For it is likely that expert systems will be used largely to assist in the performance of many tasks that for experts are mundane or routine, although in the past required their attention. If expert systems can assist in such tasks, then they will free experts to focus on what they do best – and what are likely to be beyond the scope of computer systems for some time yet – namely, the tackling and solving of problems that are complex and difficult even for experts.

2. *Quality*

Expert systems will also enhance the quality of legal work. By preserving and making widely available scarce legal expertise and in essentially codifying that knowledge, the technology can promote a uniformity of approach to similar problems, a consistency of disposal, and an in-built quality control regime. Moreover, computer systems will not suffer from “off-days” that so often inhibit the performance of human beings.

It is difficult to quantify with precision the human resource and qualitative benefits just noted. Nevertheless, the major dimensions of financial benefit can easily be identified. For the profit oriented organisation, expert systems will allow greater leverage, that is, a higher proportion of lower paid to higher paid employees, while retaining the same quality and quantity (at least) of workload. The technology will also facilitate “value-billing”, whereby the charge for legal services will be based *not* on the time spent on a task multiplied by some hourly rate; rather, the task itself will have been automated and so undertaken far more

quickly or at far lower cost. Value-billing gives rise to charges equivalent to those that would be incurred under conventional billing systems. Yet the task for which the charge is being made will have incurred far less cost for the provider.

For the public service based organisation, such as a government legal department, the financial benefits of using the technology result from the possibility of substantially reducing the cost of the services delivered by the department. The technology will allow tasks to be undertaken by less costly staff in fewer hours while increasing the quality of the work.

While it is easy to identify the sources of profit derived from expert systems generally, it is far more difficult to quantify these benefits in particular cases. As mentioned earlier, a key challenge for those involved with the commercial exploitation of AI and law is the articulation of techniques for precisely analysing the financial costs and benefits of building systems.

Problems and obstacles

A number of problems and obstacles have faced and will continue to confront those wishing to build expert systems in the legal area. Six major problems deserve mention.

1. Lack of knowledge engineers

The person whose task it is to “mine the jewels” from human experts’ heads, then articulate that knowledge and finally embody it in a computer system, is known as the “knowledge engineer”. To be effective in this task, the knowledge engineer must have considerable knowledge both of computer technology and of the law. Unfortunately, there are few people with training in both these disciplines. This lack of potential knowledge engineers will continue to inhibit the number of systems that are developed.

2. Lack of domain experts

The source of knowledge and expertise for any expert system is the human expert himself. This domain expert must work closely with the knowledge engineer and this is a very time consuming process. It is not easy to convince an expert to devote extensive periods of time that could otherwise be used for chargeable work or more direct legal service. Yet, without the commitment of the expert to a project, no system can be developed. Understandably, few experts so far have shown the degree of commitment required.

3. Lack of method

Modern data processing is characterised by the use of so-called “methodologies”; that detailed sets of carefully and clearly formulated standard practices and procedures to guide those developing systems. In contrast, no such methodology exists for the development of expert systems in the legal area and so designers of

these systems are required to face the same obstacles and hurdles that others have faced and overcome in the past.

4. *Lack of tools*

Other than for the development of document modelling systems, there are, as yet, no commercially available software packages devoted to the development of expert systems in law. Today, designers of systems are compelled to use tools not ideally suited to the task and inefficiency and inaccuracy inevitably result.

5. *Quality control difficulties*

A major problem facing the AI community generally is the difficulty of controlling the quality of systems under development. It is enormously difficult to test the reliability and accuracy of the knowledge held within these systems, to test the soundness of the underlying code and to audit their performance. A fully articulated methodology, of course, would address this issue of quality control.

6. *Legal implications*

There can be no doubt that some reluctance to develop fully operational systems has its roots in concern over the legal implications of expert systems in law giving rise to loss. The expert system as a source of information and advice seems to sit somewhere between the professional adviser on the one hand and the text book on the other. However, there have been no decided cases on the issue of liability for expert systems, and while this uncertainty prevails, so too will some reluctance to develop the technology further.

Keys to success

The fourth and final question often asked of expert systems in law looks for guidance on the successful development of systems. Drawing on international findings, there follows a listing of ten keys to success.

- There must be a management or business problem requiring a solution – too often expert systems are ‘solutions looking for problems’.
- A rigorous feasibility study must be undertaken prior to any major development project.
- Quick, inexpensive prototyping of systems can, at an early stage, greatly enhance the understanding of management, experts, and projected users.
- Where possible, existing expert systems developments, methods and techniques should be used rather than inventing new ones.
- It is crucial for the purposes of development and maintenance, to record and document the knowledge held within a system in some schematic form, usually as “knowledge base maps”.
- Conventional data processing skills can and should be regularly used in the

course of developing expert systems in law.

- A fundamental, theoretical, jurisprudential understanding of the law on the part of developers greatly increases the likelihood of the development of reliable systems.
- Once in operation, the performance and organisation fit of systems must frequently be monitored.
- The expectations of management, experts and users must be managed most carefully and developers must not fall foul of 'the fallacy of the successful first step'.
- A system will only be used profitably if there is the commitment of senior management, of domain experts and of the end users.

In conclusion, it is generally anticipated that the 1990s will be the decade during which artificial intelligence technologies begin to deliver substantial business benefits. Given careful and sufficient investment, the administration of the law – both in the public and private sectors – is an area of particular promise. The public administration of the law and the provision of private legal services is inherently *knowledge based* and so is especially amenable to AI and expert systems treatment; for these technologies strive precisely to support and enhance knowledge processing tasks.