EXPERT SYSTEMS - LAWYERS BEWARE!*

Ronald STAMPER, James BACKHOUSE and Karl ALTHAUS

Abstract: Two fundamental paradigms are in conflict. Expert systems are the creation of the artificial intelligence paradigm which presumes that an objective reality can be understood and controlled by an individual expert intelligence that can be replaced by machinery. The alternative paradigm assumes that reality is the subjective product of human beings striving to collaborate through shared norms and experiences, a process that can be assisted by but never replaced by computers. The first paradigm is appropriate in the domains of natural science and mathematics but dangerous in social science, business and, especially, the law. Expert systems are constructed on the basis of a number of metaphysical assumptions that are invalid in the legal domain. These assumptions are assimilated through a number of commonplace metaphors that guide the thoughts of the majority of people entering the computing field who are usually trained in first paradigm subjects such as mathematics and the natural science. This inappropriate paradigm hinders our progress in the field of computers and law. We need to adopt a socially orientated view of the nature of reality, of language, of meaning, of intelligence, and of reasoning. It will be easier then to build computer systems to facilitate social interactions in the legal domain and easier to understand why boxes that try to imitate legal expertise are intrinsically fraudulent.

Expert systems as a genre have not escaped criticism. Their supposed virtues are loudly trumpeted by the well-funded[1] community of 'knowledge engineers' who build them. The weight

---

*Este trabajo ha sido presentado como ponencia en el reciente IV Congreso Internacional "LAW AND COMPUTERS", organizado por la Corte di Cassazione (Tribunal Supremo) de Italia y celebrado en Roma del 16 al 21 de Mayo de 1988 y se recoge, con el correspondiente permiso, en este número extraordinario de THEORIA dedicado al tema en la seguridad de que con su selección y rápida difusión prestamos un gran servicio a nuestros lectores interesados en la materia, y en especial a los españoles, quienes, de nos ser así, difícilmente tendrían tan inmediato acceso a tan importante texto.

THEORIA - Segunda Época - Año III
Curso 1987-88, Nº 7-8, pp. 317-340

of uncritical literature in favour of expert systems in almost every conceivable domain deserves a counterpoise, to which this paper will contribute its gram.

Among the most powerful critics of expert systems and its parent domain, artificial intelligence, are Dreyfus and Dreyfus (1986), Weizenbaum (1976), and Winograd and Flores (1986). The Dreyfus brothers demonstrated five levels of skill, ranging from the novice to the expert, of which two, they argue on psychological grounds, are beyond the range of the machine. Weizenbaum questions the epistemological assumptions of the computer-orientated thinkers and Flores and Winograd point to the solution in terms of human knowledge and expertise being rooted in the involvement of agents in an infinitely complex and unanalysable reality of physical and social action. Their arguments are commended to anyone needing to develop resistance to the deluge of expert systems sales literature one tends to receive today.

Undoubtedly, expert systems have roles to play in business, especially where complex routine decisions have to be made at an operational level (see Lee 1985). They provide a means of delegating such decisions in the form of precise rules to be followed in a bureaucratic manner by a technician or a clerk. They can be especially useful in circumstances where there is too rapid a turnover of staff to make it economic to train the technicians or clerks in the relevant skills. But that is not how these systems are presented. Rather, we are expected to accept them as black-boxes containing enough 'knowledge' to make them surrogates for human experts. When we cease to regard it as a mechanical aid to communication between a responsible expert and a cautious assistant, but treat an expert system as a surrogate for some remote person (the original source of the 'knowledge-base') who delegates his responsibility to the machine, then the expert system is a potential danger, no matter what the field of application.

From Factory Floor to the Law

The law is one of the domains most recently to find favour among the community of 'knowledge engineers', as expert-systems practitioners like to be known. They come to lawyers bearing gifts constructed from some tacit assumptions about knowledge, reality, reasoning, information, truth, intelligence, and so on that are harmless in narrowly circumscribed domains of engineering. Outside a micro-world of established consensus about an objective reality, the expert system can embody a dangerouly misleading form of pseudo-science. The law contributes to the resolution of conflicts where consensus certainly does not exist and legal problems are difficult if not impossible to confine to a narrow domain. Legal expert systems should be treated with suspicion for those reasons alone. Here we shall

throw doubt upon their tacit assumptions behind their gifts, their solutions in search of lucrative problems.

My colleagues and I entered the domain of computers and law early in the 1970s. We produced one of the first expert-system shells[2] but in the days before that unjustifiable jargon expression had begun to circulate. We started, not with a mathematical solution in search of problems, but with the problem itself: How, in general, can we represent social behaviour as systems of social norms. The outcome of our work is radically different from the ready-made solutions such as predicate logic in which some enthusiasts believe a legal domain can be expressed as an axiomatic system (Sergot et al 1986). It entails abandoning the normal tacit assumptions of knowledge engineering in favour of quite a different set of metaphysical assumptions. The new approach is not set out here (space does not permit). Emphasis is given to criticising the dominant and (we believe) mistaken paradigm within which expert system builders think. Our results support the development of so-called expert-systems but we prefer to avoid that misleading expression. Above all, we have rejected the simplistic assumptions about reality and meaning in the old paradigm, and we conclude that legal decision-support tools should aid the processes of human collaboration, especially through the sharing of knowledge and the resolution of conflicts (Jacob and Stamper 1986).

The Current Dominating Paradigm

When a computer specialist solves a problem, he brings to it a framework of assumptions that served him well on other occasions. Those assumptions will determine what he sees when he looks at the issues. The history of a community of problem solvers tells us about their customary world-view. The computer community, in so far as such generalisations are justifiable, has a very different intellectual history, individually and collectively from the the legal community. The computing milieu has an outlook and attitudes predominantly drawn from the natural sciences, engineering, mathematics and accounting. The resulting weltanschauung is appropriate for the hardware and software branches of computing, but not for the analysis of non-routine problems in the law and practical business affairs in general.

Anyone who questions an established world-view is likely to be unwelcome and face a daunting task (Kuhn 1962). To question the tacit assumptions of a whole community is to threaten powerful economic interests which have invested in the historically established corpus of analytical tools that embody the assumptions being questioned. We are not so ambitious.

Our object is to call attention to the limits within which the established computing weltanschauung can be justified. The expert-systems movement seems to be in danger of straying past those limits as it attempts to move into the legal domain.

The assumptions to be questioned are seldom discussed explicitly, so how do we uncover them? We shall use, here, the insights provided by drawing attention to the metaphors regularly and unquestioningly employed in the discussion of expert-systems and the related fields of problem analysis. The principal metaphors that we shall examine are

1. the conduit metaphor of language
2. the chemical engineering metaphor of data-processing
3. the set metaphor of reality
4. the correspondence metaphor of meaning
5. the platonic metaphor
6. the information-processing metaphor of mind

all of which tend to support the hyperbole of expert systems sales talk. These metaphors are commonplace in much of our everyday discourse, a fact that makes them seem so much like commonsense that anyone can be made most uncomfortable if asked to abandon them. They are, in fact, blinkers preventing us from seeing to the side and so, thinking laterally.

A Fragment of an Expert System

To illustrate the points to be covered below, a fragment of a rule-system, typical of those embodied in an expert system will be useful. Prolog is the favourite language for expert system building in Europe and the work of Kowalski's team on systems to support social security administration is one of the heavily funded projects of the Alvey Project in the UK. Of their work, we take an extract from a paper that is readily available in the literature (Sergot et al 1986). It deals with a slightly different domain of the law, The British Nationality Act 1981, but it typifies the kind of system that interests us. The choice of this example is not a comment on the specific authors, almost any other expert system could have served our purpose.

A 'knowledge engineer' who builds an expert system has to extract knowledge from a domain expert so that it can be stored in a knowledge-base. The expert expresses his concepts in words and he encapsulates his skills in rules in order to convey them to the engineer. These rules are like those used in the law, especially statute law. They are loaded into the knowledge-base. The 'inference-engine' in the expert system will apply the knowledge in the rules to the raw facts which are the input from a user, and output decisions, advice or other information he needs.

The legal draftsman is a kind of intermediary to the knowledge engineer. He provides a statement of the policy extracted from the political expert. In this fragment, it is the first sub-section of the British Nationality Act 1981:

1.- (1) A person born in the United Kingdom after commencement shall be a British citizen if at the time of birth his father or mother is
(a) a British citizen; or
(b) settled in the United Kindom.

The meaning of this legal norm is then captured by the knowledge engineer in a number of Horn Clauses in Prolog, as follows:

```
x acquires British citizenship by section 1.1 on date y
  if x is born in the U.K.
  and x was born on date y
  and y is after or on commencement
  and x has a parent who qualifies under 1.1 on date y

x has a parent who qualifies under 1.1 on date y
  if z is a parent of x
  and z is a British citizen on date y

x has a parent who qualifies under 1.1 on date y
  if z is a parent of x
 and z is settled in the U.K. on date y

z is a parent of x
  if z is the mother of x

z is a parent of x
  if z is the father of x
```

The expert-system will use this knowledge-base to generate questions from the sentences in the conditions of the rules, such as

Who is the mother of Karamjit Raj?

so that the rules can be evaluated to answer questions about the issue concering the client, Karamjit Raj.

The Conduit Metaphor

Michael Reddy (1979) introduced this expression in an essay on the language we use for talking about language. He shows that our normal English idoms presume that words, sentences, paragraphs, books, libraries (data- and knowledge-bases) are vehicles, containers, pipelines through which meanings, concepts, ideas, emotions and so on can be carried from one place to another. No doubt the same can be said of Italian or any other European language. As Reddy points out, this deeply embedded metaphor leads to some dangerously misleading notions of how language is used in human collaboration. This pathology should be evident in one of the paragraphs above, introducing the fragment of and expert-system. The reader probably found nothing unremarkable in the passage. Care has been taken to use just the kinds of expressions that fill the average page of expert-system literature. Here it is, with the significant expressions emboldened:

"A 'knowledge engineer' who builds an expert system has **to extract knowledge from** a domain expert so that it can be **stored in a knowledge-base.** The expert **expresses his concepts in words** and he **encapsulates his skills in rules** in order **to convey** them to the engineer. These rules are like those used in the law, especially statute law. They are **loaded into the knowledge-base.** The 'inference-engine' in the expert system will apply the **knowledge in the rules** to the **raw facts** which are **the input** from a user, and **output decisions, advice or other information** he needs."

Notice that there would be no dispute that the precise form of the words or signs employed is a secondary issue compared with the facts, concepts and knowledge that they carry.

These high-lighted, metaphor-laden expressions imply that some mysterious substances that we call 'knowledge' and 'information' are carried by words, rules and other symbols used within the computer. This palpable nonsense is the foundation of knowledge engineering seen, as it is by its commercial exploiters if not by everyone in the field, as a means of 'bottling' knowledge and selling it to anyone who is prepared to pay for some software. There is no need, in this view, for anyone to understand the whole structure of the expert system's programme. It is sufficient to accumulate knowledge, the essential oil of experience, bottle it and apply it as a balm to any appropriate problem. Of course, knowledge, an abstract substance extracted by the engineer from the expert, has an independent existence.

This nonsense is dangerous for several reasons. It obscures responsibility. It devalues the role of critical understanding in the maintenance of valid knowledge. It pretends that

knowledge can be detached from its social context. It suggests that signs (words, sentences, drawings, books etc) are carriers of this abstract substance and that the human process of interpreting signs is not essential to the nature of the knowledge represented. What antidotes are there?

The remedies should be second-nature to a good lawyer. Ask who is providing the knowledge and what axe he may have to grind. Be suspicious if the knowledge-base cannot be examined as a whole to help one better to comprehend the domain, and also expose it to criticism. If the knowledge-base is not the constant focus of the critical attention of a community of domain specialist, doubt its validity. Apply the lawyer's normal skills to considering the possibility of multiple and conflicting interpretations to the system's inputs and outputs. But who will apply these remedies?

Unfortunately, the second-rate lawyer, by second-nature, will be only too glad to abdicate thought in favour of a machine that will increase his income. This combination of expert system purveyor and irresponsible user raises another range of issues tackled elsewhere (Göranzon 1987).

The Chemical Engineering Metaphor of Data-Processing

The conduit metaphor appears in a most exaggerated form in the literature of computer applications. Presumably because one knows that electrons flow through the carefully engineered components of a computer, it is natural, indeed almost inevitable, that the computer specialist will think of a business or a legal system functioning through the flow of information through components that generate, transmit, store, process and output information. The care that has been devoted to the experimentally demonstrating of the nature of electrons, does not prevent information from being treated as a mystical fluid that can be taken for granted. The builder of your computer-based information system is almost your information plumber.

In the paragraph quoted above to illustrate the conduit metaphor, you will find the commonly encountered notion that the facts or data, are treated as a kind of "raw" material that the computer can turn into distilled "information". It is even normal practice (in computerise) to use the word 'data' in the singular to suggest that 'it' is a substance processed by the computer. The expert-system is regarded as just a new kind of processing device, perhaps a more powerful distillation column, to be incorporated into that information refinery that we call our judicial system.

This chemical engineering metaphor build further misunderstandings upon the conduit metaphor. Its principal danger derives from its inventing another mystical fluid, called information. The spurious objective reality of information arises from the belief that there is an abstract entity that flows. The commonly accepted model

input -------- process ---------- output

obscures the fact that we have

signs ----- encode ----- process ----- decode ----- signs

nothing flows but there is a cause-and-effect chain linking one lot of signs to another. Nothing is actually carried by the signs, they only (literally) have any significance in the context of the social groups of the people who interpret them:

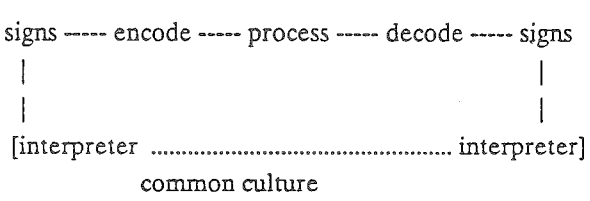

The illusion of flow depends upon there being a common culture shared by the interpreters of the 'input' and 'output' signs. The people involved are able, on account of their shared culture, to influence each other's behaviour through the mediation of the signs, so completing semiologically the circuit of cause-and-effect in which the computer plays a minor physical role. The whole **information system**, as opposed to the **electronic signalling system** is based upon the human cultural system. By suppressing awareness of the cultural system, the impression is given that information exists, as a kind of abstract substance, independently of human beings and their social systems.

The lawyer knows that the cultural system is permanently in tension. Those tensions do not always surface as legal conflicts but, when they do, they generally manifest themselves as differences in the interpretation of words, expressions and rules. The idea of an expert-system as a processing device through which we pass the raw data of cases to filter or distil from them decisions or advice, begins to look dubious. Lawyers, on the contrary, need systems that will handle just those situations that do not support the illusion of information flow.

Formal and Natural Language

The unconscious influence of the conduit and chemical engineering metaphors upon the knowledge engineer is visible in the Nationality Act fragement we have quoted above. The has taken the meaning of the statute and put it into Hórn Clauses of Prolog. The transfer from one 'container' to another has been carried out with as much confidence as he might feel when transferring the contents of a drawer to a briefcase. It is not unnatural, we are all accustomed to paraphrasing sentences. More is happening here than the familiar task of paraphrasing.

A translation from natural language to another form of natural language expression is quite different from the writing of these rules. The knowldege engineer is actually translating into a formal language.

It is important to recognise that expressions such as:
... acquires British citizenship by section 1.1 on date ...
... is born in the U.K.
... was born on date ...
... has a parent who qualifies under 1.1 on date ...
are just single symbols as far as Prolog is concerned. They might equally well have been written in the forms:

A(x,y)
B(x)
D(x,y)
P(x,y)

as far as the computer is concerned. The use of the longer form is confusing. The resemblance to natural language in spurious but it certainly helps to give the expert-system purchaser the impression that the computer understands more than it does.

This unjustifiable faith in the conduit theory has at least three dangers. There is the problem of misleading people by passing off a formal language as though it is a familiar natural language, mentioned above. There is also a likelihood that this practice will lead the expert-system to fall into disarray. And, furthermore, it makes an obscure body of knowledge even more obscure. Let us make a brief note of the other two.

Most practical computer-based information systems are designed, constructed and maintained by many people over a period of time. They may, at various times, translate related items of statute law into Prolog. The predicates they introduce may all be paraphrases of the same text, such as:

.. acquires British citizenship by section 1.1 on date ..
.. starts his British citizenship by sect.1.1 on day ..
British citizenship of .. acquired on .. under sub-section 1.-(1)
..is a citizen from..by virtue of British statute 1981c61, 1.-(1)
..is a natural born Briton commencing on .. (ref: UK 1981c61)

As systems grow in size and age, and as systems are linked together to provide composite 'knowledge-bases' the result will be growing confusion. The knowledge is no longer a shared linguistic structure but a mixture of incommensurate private code-books!

This privatising of knowledge has the even worse effect of potentially destroying knowledge. It is quite difficult to gain a thorough grasp of a body of legislation and related case law. The vocabulary of the draftsman and the elaborate syntactic structures he creates can be made more comprehensible by the provision of commentaries by experts. But the typical, present-day expert-system is not designed to comment and clarify, rather, it is intended to 'capture' and 'deliver' expertise as though it were a commodity. It does so by imposing, upon an already difficult problem, an unnecessary additional layer of complexity in the form of the knowledge engineer's own private formal vocabulary. The expert-system is an opaque black box which makes the domain of expertise more obscure than necessary, thus preventing or at least hindering the processes of comprehension and criticism. Knowledge loses its validity if not subjected to continual criticism.

The Set Metaphor of Reality

Lending support to the notion that knowledge and information are independent of people, is the mathematical concept of sets. This very simple idea has proven indispensible for mathematical analysis. Unfortunately, its use in non-mathematical domains has been legitimated by many great philosophers (Whitehead and Russell 1910/1927, to mention but two). One can make the set of all objects and their relationships the metaphor on which to construct a model of the real world.

It is natural to conceive of the universe as composed of discrete individuals as though they were the tools and components disposed upon some celestial laboratory bench. The individuals can be grouped into sets by their properties (all red things, all things that are people etc). They can be related in pairs (one individual inside another, one individual married to another etc) and into higher order structures. Whole systems for modelling reality, accounting form meanings and defining truth are based on this idea (Carnap 1942, Barwise and Perry 1983). Predicate logic is normally taught on the basis that this metaphor gives an adequate

account of reality. It populates the world with individuals such as Karamjit Singh who would like to be confirmed as a British citizen, as well as his father and mother. It then provides the one-place predicates that classify things (red, person, is born in the UK etc) and the two-place predicates for the relations among individuals (are married, is a parent of, etc) To know the nature of the real world is then just a matter of knowing to which classes and relationships every individual belongs.

As an abstract mathematical theory, the set metaphor can be put to good use. As a concrete representation of a domain of practical human affairs, the metaphor is weak, as we shall see later. In use, in the Prolog illustration we have chosen, the metaphor can lead to arbitrary models of the world.

People and dates are treated as individuals in our Nationality example. The knowledge engineers could also have treated the United Kingdom and Great Britain as individuals, but notice that their appearance in the knowledge-base is spurious because, as far as the formal system is concerned, they are submerged as undifferentiated features of whole predicates such as "is born in the U.K." Despite the mathematical elegance of set theory, when used in practice for knowledge engineering, it reflects as much the psychology of the system designer as it conveys knowledge of the world. The objects on the celestial work-bench keep changing according to the whim of the observer! A change in the law might well force the expert-system to alter its picture of reality, acknowledging individuals that were previously implicit in properties and relations. For legal or other continually evolving problem domains, this instability in the basic frame of reference is potentially dangerous.

The likelihood of progressive confusion is only a minor fault of the set-orientated, 'knowledge-representation' languages such as Prolog. More serious deficiencies lie in their evading the issues basic to legal conflict resolution. By modelling the world in terms of sets, we make some colossal assumptions.

1. The world is divided up, independently of who is observing it, into discrete individuals.
2. Those individuals are uniquely identifiable.
3. Individuals can be grouped into sets, the precise membership of which defines the meaning of the property to which it corresponds.
4. Pairs of individuals can also be grouped into sets in the same way in order to define relations.

These assumptions are usually acceptable in the domains of routine natural science and engineering (but not on the frontiers of those subjects) but certainly they are unacceptable in the legal domain.

The role of the law is to establish boundaries and maintain them, even to allow them to move gradually in a controlled manner. To assume, as in 1 to 4 above, that all kinds of boundaries are fixed, and fixed in an objective way, independently of any human agency, is to evade the central issues with which the law is concerned.

Individuation, the first assumption, causes many a legal dispute. A piece of land may easily be turned into two pieces and vice versa. For the sake of gaining some legal advantage either of these courses of 'action' might be undertaken. In the UK, several meadows of ecological importance, have been subdivided into minute portions and sold individually to different people in order to raise funds for charitable purposes but also to ensure that the land would be very difficult to convey into the ownership of anyone who might be tempted to exploit it commercially. The spatial and temporal boundaries which circumscribe an individual are frequently the subject of heated dispute: consider the current debates about abortion law and the laws governing organ transplant operations, for example. You will find it easy to provide other examples of legal problems associated with determining whether notionally to create or destroy individuals or where to set their boundaries. Individuals do not exist in their own right, it is the function of our social system, aided often by the law, to decide what constitutes an individual.

Identity is also socially determined in the realm of practical business and legal affairs. Take the problem of copyright in a literary work. This can be bought and sold as an individual resource but what constitutes the particular work may not be a trivial issue to deal with. A publisher who buys the rights in a future work from a popular and prolific author may find himself in dispute about the identity of the work in which he believed himself to have an interest; the issue may have to be resolved by the courts. Anyone concerned with manufacturing and marketing mass-produced goods, or dealing with the legal liabilities associated with them, will know that fixing and maintaining a definite identity for each item can be a difficult and expensive task. Administrative and legal systems cannot take identity as a given feature of an objective reality.

Classifications are perhaps the most common focus of attention in legal disputes. The dividing line between guilty and not-guilty is not an objective fact. A responsible agent, usually a magistrate or a jury has the responsibility for deciding. In the Common Market, we Europeans are frequently in hot dispute over the classification of various goods. We all like Italian ice-cream and agree that is what it is, but many non-British and plenty of British people in the community would like the product sold under that classification today to be treated as 'non-dairy, frozen confectionery'.

Reality in the legal domain, and, indeed, in the non-routine aspects of practical human affairs, is subjective. Boundaries dividing individual from individual, their identities and classifications are not given but intrinsincally open to dispute. We actually construct our reality using social mechanisms for resolving disputes, of which the law, politics and armed conflict are examples. Objective reality is a luxurious invention which we can enjoy when the social disputes have been resolved.

These criticisms can be related to the Nationality example. To avoid the oversimplified view that reality is objectively a set of identifiable individuals, the expert-system should at least compel the sources of knowledge about classifications and relationships to be accounted for in delivering its advice. At best, this information is an optional extra.

The Correspondence Metaphor for Meaning and Truth

The metaphors examined so far conspire to support an over-simplified theory of meaning. We have seen that the model of reality based upon the idea of sets impedes the consideration of those key legal problems that account for the majority of disputes - problems of meaning. Earlier, we saw that the conduit and chemical engineering metaphors encourage the belief that words, expressions and whole sentences have meanings of their own. If these signs 'carry' meanings, interpretation is only a matter of 'uncovering' them - there is no suggestion that different meanings are made by different agents in the process of interpretation, a more appropriate view for the treatment of law. An over-simplified notion of meaning is not a sound basis for legal expertise which is preeminantly concerned with the resolution of disputes about meaning.

An additional metaphor, which explains meaning as a correspondence between a sign and what it signifies, adds its weight to the conspiracy. We are accustomed to thinking of words, expressions and sentences as akin to verbal fingers, pointing at the entities in the world they represent. A mathematician would describe this as a function, mapping a linguistic expression onto the real world. The idea is reassuringly simple. It has also been elaborated into the most sophisticated account of the meaning of natural language so far generated by language theorists, which is known as 'Montague semantics'(Montague 1974 and Dowty et al 1981). However, mathematical elegance an unsuitable criterion for assessing a juridical theory of meaning.

The basic correspondences are between the names of individuals and the individuals themselves. Thus 'Karamjit Raj' maps onto the person we know by and who answers to that name. The next level of structure is the set of individuals onto which the name of some property or

common noun is mapped. Thus the set which includes not only Karamjit Raj but Ronald Stamper, Jim Backhouse, Karl Althaus and many others, is the meaning of the word (predicate) 'person'. Other kinds of words have their meanings supplied by mappings onto rather more elaborate set constructs, but there is no need to look at these for our purposes.

So far so good, if we are looking for a mathematical theory, but notice what the mathematical theory expects us to swallow in practice. The precise meaning of person is the extension of the set of all persons, that is to say all those individuals in the set pointed to by the word 'person'. There are good mathematical reasons for preferring a definition of meaning in terms of an extension of a set, rather than in terms of the intension or criterion by which one judges whether an individual is a person or not (see Quine 1960 ch 6). Unfortunately, this makes a correct knowledge of the meaning of 'person' beyond anyone's capability in practice. There is no reason to believe that a good theory in a mathematical sense should also be even adequate in a totally different domain, such as the law.

The problem of establishing a meaning in practice in the manner decreed by the correspondence theory is even more difficult. If the individuals comprising the set of persons changes, then, in the correspondence theory of meaning, the meaning of 'person' changes also. People are born and die, so the meaning of 'person' seems unresolvable in practice. However, for the mathematician, there is an escape. He is prepared to include in the set of persons all those who have been and who ever might be in times past and future. This does not solves the practical problem of coming to grips with the meaning, indeed it makes it worse. In addition, to being acquainted with the collection of all the individuals we originally recognised as comprising the real world, we now have to add intervals or instants of time, and any number of imaginary worlds. The result still looks quite good as a mathematical theory but it seems to us to be in tatters as a theory of meaning for tackling problems in business, the law or other practical human affairs.

The treatment of whole propositions is particularly interesting. These are held to correspond, for semantic purposes to one of the two entities, truth or falsity. By combining the names of individuals with predicate names according to grammatical rules, well-formed sentences are generated. Then, by inspecting the meanings of these components in relation to one another, the meaning (truth or falsity) of the whole sentence can be computed. Thus, when we substitute Karamjit Raj and Parvati Raj for x and z in

z is the mother of x or M(z,x)

to form the sentence

Parvati Raj is the mother of Karamjit Raj

all (!) we have to do is to check whether or not the ordered pair

< Parvati Raj, Karamjit Raj >

is a member of the set that is the meaning of

... is the mother of ...

and if it is, the sentence is true, otherwise it is false. This procedure raises a number of significant issues.

Of course a database could be constructed to hold not all, but all the relevant pairs of < mother, child > names, as surrogates for the actual instances of this relationship. A positive result of a search for such a pair might well be taken as proof of the truth of the proposition

Parvati Raj is the mother of Karamjit Raj

but it would be unwise to treat the absence of the pair

< Parvati Raj, Karamjit Raj >

as demonstrating the falsity of the proposition. In practice it is hardly ever possible to mimic the mathematical function mapping expression onto the real world. One is forced to ask the user of the system to affirm or deny each fundamental proposition occurring in the condition of a rule to be evaluated. Most of the expert systems you will encounter devote most of their operation to posing these questions to the user. This is fine for system with a single user who is the source of all the information - at least he knows its reliability. When, however, the reliability of the source of information is at issue, the simple model of an objective reality comprising sets of individuals breaks down. In general, the need to weigh the evidence according to its sources is essential in expert legal reasoning.

Truth, in this correspondence theory, is treated as a real-world individual onto which propositions are mapped by the 'meaning function'. Truth, in this view is a primitive notion. But lawyers know that truth in the realm of practical affairs represents an agreement among the interested parties within a framework (of which the judicial system is a key part) assigning different powers to them. The truth concerning any matter, in practice, is our name for this kind of weighted consensus about it. Objective truth can be invented in those narrow domains where a consensus has been obtained, such as in uncontroversial applications of the

natural sciences, routine engineering and mathematics. Legal truth is not the basis of meaning but the product of legal reasoning.

The dangers of a theory of meaning based upon the correspondence metaphor are principally
1. that its pretence of objectivity draws attention away from the parties involved and the reliability of the 'facts' they provide;
2. that its meaning functions relating the knowledge-base to reality are mathematical abstractions incapable of being operationalised;
3. that it elevates an abstract notion of truth to a key role without establishing a role for the human agents who are involved in arriving at truth in practical affairs.
This list emphasises the artificiality of the concepts behind the legal expert systems being constructed today.

Plato's Metaphor

Mathematicians, almost without exception, regard their subject as the study of an abstract, platonic reality. They encourage themselves and persuade many others to have faith in their platonic realm as a source of absolute, reliable knowledge. Their linguistic trick of "letting x = ..." encourages the notion that meanings may be arbitrarity assigned and then carried by the chosen symbols. The valuable, but minor, mathematical skill of symbol manipulation certainly demands a high degree of intelligence of a particular kind, leading to the assumption that it is the prototype of all intellegent behaviour. The quest for truth in mathematics and the habit of "seeing" mathematical truth with an inner eye conspires to support the view that truth is a simple, primitive notion.

Lawyers are concerned with a most unplatonic reality. The world of practical affairs contains no ready-made truths. It finds the truth through institutional arrangements such as the courts, where evidence and the providers of evidence are judged in an attempt to reach a reliable consensus upon which to base just decisions. Each court will make its best endeavours to reach the truth in this sense but failure is known to be endemic, so an appeals procedure allows the process to be resumed. Ultimately, as we all know, there is a degree of arbitrariness about the result, but what matters is that there is a clear assignment of responsibility for whatever truth is established and the responsible agent is involved in an informal structure of social and political pressures that will influence him or her. Truth in the world of legal and other practical affairs is no simple matter, the notion of responsibility is a safer primitive on which to base an epistemological theory.

The Information-Processing Metaphor of the Mind

Guiding the thinking of knowledge-engineers, and presumably accounting for their faith in the machine as a surrogage for a human expert, is another metaphor that would have us believe that the human mind is like a computer. A harmless idea until one forgets that it is but a metaphor that should not be pushed too far.

Many psychologists use this idea. They have taken the information-processing metaphor from the computing community as a successor to the telephone exchange metaphor that psychologists found helpful in days before the computer, and to the earlier, 18th century, use of the automaton as metaphor for human behaviour. It is legitimate for psychologists to sieze upon any metaphor that helps them to conjecture hypotheses that will help explain human behaviour. Cognitive psychology, which exploits this metaphor, now overlaps almost completely the domain of computing called artificial intelligence, the home of knowledge engineering. The idea is that the human brain receives input data from the senses and then processes them in order to perceive the world as it is, and subsequently processes the resulting concepts, encoded internally, in order to think and solve problems, make decisions and control the motor mechanisms of the body. However useful it may be, not all psychology can be accommodated within this paradigm.

The brain as computer is an idea matched by the AI community's perception of the computer as a brain. The mutual support that these two metaphors lend each other is less convincing as soon as one realises their incestuous relationship.

Let us examine the consequences of this metaphor for legal expert systems. In the first place, it implies that perception and intelligent thought are achieved through processes of sign manipulation. Therefore a box of mechanisms that can perform any of the possible operations on signs or symbols, with great speed, is likely to be able to be able, when suitable configured, to behave intelligently. In particular, it implies that a suitable symbol manipulating device will be able to reason.

Reasoning by the manipulation of symbols is the province of logic and deduction using the resolution principle is the foundation of Prolog, the knowledge manipulation language used in our chosen example. This is exactly the kind of reasoning employed in mathematics when theorems are proved, step by step. Predicate logic, which Prolog implements, was developed as a mathematical tool, only subsequently has it been applied to legal problems. The information-processing metaphor supports a view that expert thinking is a matter of good deductive reasoning.

Deductive reasoning operates by transforming valid propositions into other valid propositions by executing transformations that guarantee the preservation of validity. Truth is the notion of validity we normally employ. We have already noted the inappropriateness of truth as a primitive concept in the fied of practical human affairs. Such kinds of reasoning that preserve truth values already established are of trivial importance in legal problem solving.

An exercise in mechanical, deductive reasoning of the kind one might give to a student of AI will illustrate its intrinsic triviality:
Given, as a knowledge-base, a stanza of Lewis Carroll's poem, The Hunting of the Snark, write a logic programme to answer questions relating to it.
The solution might well be based on the stanza:

Twas brillig, and the slithy toves
Did gyre and gimble in the wabe;
All mimsy were the borogoves,
And the mome raths outgrabe.

and generate such dialogues as:

Q: What were the toves doing?
A: Gyring.
Q: Were they doing anything else?
A: Yes.
Q: What might make a mome rath outgrabe?
A: Being brillig, the gyring and gimbling of slythy toves, or the borogoves being mimsey.

These syntactical manipulations are quite clever and have their uses. They can generate dialogues about the citizenship of people in relation to the British Nationality Act using exactly the kind of logic programme quoted above.

Legal reasoning has relatively little to do with this kind of closed realm and its symbolic manipulation. Initially, Karamjit Raj's claim to be a British citizen might be examined this way and, if it appears to fail, deductive reasoning could reveal the factors upon which it depends. If that is all you want, a preliminary screening of cases, expert systems can be most helpful. The skill involved, as Dreyfus and Dreyfus argued, is that of the competent beginner or, in this case, a competent bureaucrat. Legal expertise only comes into play when strategies have to be developed to push the meanings of the terms employed in the act in a direction that will favour the client. A lawyer's reasoning will only yield results if it goes beyond the closed symbolic system where truth values are fixed.

Beyond the closed system, the lawyer must exercise his intuitive grasp of the field, the institutions, and the people involved with the problem. Our fragmentary example, taken from the British Nationality Act, probably affords few opportunities to demonstrate this kind of legal expertise. However, the term 'settled' seems to offer a range of interpretations but subsection 1(8) invokes definitions given in three other subsections which, in their turn, call upon the provisions of two other Acts. In addition to all these explicit references, there will be many other implicit ones. By adding to the knowledge-base rules that attempt to capture the experience of lawyers in the field, the expert system has the potential to guide a user through an abundance of possible lines of enquiry. Up to a point this will help but gradually the labyrinth of possibilities embodied in an opaque box will become impossible to follow in a coherent fashion, so that the user can only respond to the stimuli generated by the system. The ideal expert system will be a mechanical aid to the essentially intuitive expert problem solving.

Institutional Buttresses

Unspoken metaphysical assumptions usually have deep historical foundations and their superstructures are usually institutions that help to keep the assumptions in place. The world of computers provides no exception to this rule. Look at the influences that have brought us to our present position.

Mathematics and natural science in the service of defense and engineering are the major traditions at work. Mathematical concern with computation is responsible for the challenge to build machines that can calculate. Among the problems that have justified the early research were the balistics of weapons and the breaking of enemy ciphers. The computer was the response. The natural sciences generated an endless menu of problems for computer science to digest over subsequent years. The leading edge of AI research still remains in the hands of military funding bodies. Computer scientists quite brazenly invented "artificial intelligence" as a means of attracting funds to support the exciting innovations in programming because the concept was easy to sell to the US Defense Department who could see the advantage to be gained from wielding "intelligent" weapons. SDI or Star Wars is the culmination of this line of thinking (the essays by Parnas 1985 on this topic are relevant). We must not be afraid to question the relevance of solutions to problems on the defense agenda to our own.

Part of the same tradition is operational research. This mathematically orientated body of rational problem solving techniques originated in World War II. Subsequently, its practitioners moved out of defense work into industry (one of the authors among them) but

without any marked change of attitude. In its appropriate fields, OR has been very successful. Recently it has gained momentum by climbing on the expert-systems bandwagon. The results have been some excellent systems to support decision-making in narrowly circumscribed areas of production and maintenance - just the same kinds of areas where OR had already a well-deserved reputation. We should view OR as another institutional butress to the expert-system with intellectual attitudes not dissimilar to those in the defense industry.

Of course, the Defense Industry and the Armed Forces are social systems. They have to solve many problems that have less to do with balistics and cryptogaphy, making weapons intelligent and translating intercepted messages, than they have to do with issues akin to those faced by lawyers dealing with fuzzy affairs of human disagreement. They probably have much to gain from the very different approach to knowledge, information, meaning and reality that we hint at in this paper.

Teflon Software

Expert systems are a spin-off from defense research - a counterpart to the non-stick fryingpan. They are undoubtedly of value in some contexts but we have been asking whether the legal domain is one of them.

Legal problems are the architype for social problems encountered in our practical affairs. In business we have to solve our problems quickly, negotiating agreements on endless large and small issues with seldom the need to go to court. The difference is essentially that the norms of business are largely informal but in the legal field they are made explicit in the form of rules. In both, problems of interpretation have to be resolved, but in the law the solutions are explicit precedents for later decisions whilst in business informally accumulated wisdom and (folk)lore suffice for most purposes. Social systems including business systems function because the people involved share a normative structure. In the legal domain the norms and the processes of using them are to a greater extent spelled out in detail and perfomed in a relatively slow and visible manner. The mechanisms of human cooperation and conflict resolution make quite a different starting point for research from the military and engineering problems that have shaped our computing tradition so far (see Stamper 1985b). Our own research aims to understand the general problems of social systems, especially business sytems, via the understanding of legal norms and legal reasoning (see, for example Stamper 1985c).

The hidden contention of this paper is that, starting from the study of legal problems, in particular, and business and other social problems in general, we shall arive at a very different

philosophical position from that of main-stream computing. This is especially so in the case of legal expert systems where we find that the tacit assumptions of the mathematician and natural scientist are not justified.

The LEGOL/NORMA Project is now supported by Digital Equipment (DEC). It has been supported in the past by IBM, the UK Science and Engineering Research and Economic and Social Research Councils, and by the Computer Board for Universities. The main products are a battery of problem analysis techniques under the title MEASUR and the languages NORMA (for knowledge representation) and LEGOL (for knowledge manipulation). The key idea on which these are based is that the informal social system is the foundation of our knowledge, our use of signs and the meanings we give to them. NORMA has been evolved by the study of social norms, especially those to be found in the legal domain. Although it allows us to specify sytems in a formal manner, the specification is never self-contained but carries references to the agents responsible for every norm or fact in the knowledge base. In short, all the methods and tools we have produced are intrinsically subjective in their metaphysics. In arriving at this position, the greatest difficulty proved to be learning how to escape from the metaphysical prison described in this paper[3]

Notes

1. The European ESPRIT Project, the UK Alvey Project and other national research programmes have given rather uncritical backing to expert systems research. The COST-13 Project is one notable exception.

2. The LEGOL-2.1 interpreter, programmed by Clare Tagg was demonstrated in prototype form at the Symposium on Computer Science and Law in 1979. The specification for LEGOL-2.0 was completed principally by Susan Jones in 1977 (see Jones et al 1979). The interpreter for LEGOL-1, programmed by Peter Mason, ran in 1975.

3.Many writers have helped by prompting or confirming our change of position including Barnes 1982, Bloor 1976 and 1983, Berger and Luckman 1967, Kuhn 1962, Lakatos 1976, Lakoff and Johnson 1980, Twining and Myers 1976, Whorf 1956, Wynn 1979.

[This paper was based upon the inaugural address give in April 1986 to the Netherlands Society of Lawyers]

References

Barnes, B. 1982. T.S.Kuhn and Social Science, London, Macmillan.

Barwise, J and J.Perry, 1983, Situations and Attitudes, Cambridge Mass, MIT Press.

Berger,P.L. and T.Luckmann, 1967, The Social Construction of Reality, Harmondsworth, Penguin.

Bloor,D.,1983, Wittgenstein: A Social Theory of Knowledge, London, Macmillan.

Carnap,R.,1942 Introduction to Semantics, Harvard University Press.

Ciampi,C.(ed), 1982, Artificial Intelligence and Legal Information Systems, Amsterdam, North Holland.

Dowty,D.R.,R.E.Wall and S.Peters, 1981,Introduction to Montague Semantics, Dordrecht, Reidel.

Dreyfus, H.L. & S.E.Dreyfus, 1986, Mind over Machine, Oxford, Blackwell.

Feigenbaum,A.E. and P.McCorduck, 1983, The Fifth Generation, Reading Mass, Addison-Wesley.

Gilbert,N. and C.Heath, 1985, Social Action and Artificial Intelligence, Aldershot, Gower Press.

Göranzon,B. (ed) 1987, Proceedings of the Symposium on Culture, Language and Artificial Intelligence (2 Vols), Stockholm, Swedish Centre for Working Life.

Jones,S.,P.Mason and R.K.Stamper, 1979, "Legol-2.0: A Relational Specification Language for Complex Rules", Information Systems, Vol 4, No 4.

Kowalski,R., 1979, Logic for Problem Solving, Amsterdam, North Holland.

Kuhn,T.S., 1962, The Structure of Scientific Revolutions, Chicago, Chicago University Press.

Lakatos,I. 1976, Proofs and Refutations, Cambridge, Cambridge University Press.

Lakoff,G and M.Johnson 1980, Metaphors We Live By, Chicago, University of Chicago Press.

Lee,R.M. 1985, "Bureaucracy as Artificial Intelligence" in Methlie and Sprague, 1985.

Martino,A.A. (ed), 1982, Deontic Logic, Computational Linguistics and Legal Information Systems, Amsterdam, North Holland.

Martino,A.A. and F.Socci Natali (eds), 1986, Automated Analysis of Legal Texts: Logic Informatics and Law, Amsterdam, North Holland.

Minsky,M. 1985, The Society of Mind, London, Heinemann.

Montague,R. 1974, (ed by R.H.Thomason), Formal Philosophy, New Haven, Conn, Yale University Press.

Niblett, B., 1980, Computer Science and Law, Cambridge, Cambridge University Press.

Reddy,M.J., 1979, "Conduit Metaphor - A Case of Frame Conflict in Our Language about Language", in Ortony 1979

Sergot,M.J., F.Sadri, R.Kowalski, F.Kriwaczek, P.Hammond and H.T.Cory, 1986, "The British Nationality Act as a Logic Programme", Communications of the A.C.M. Vol 29, No 4.

Stamper,R.K. 1980, "LEGOL: Modelling Legal Rules by Computer", in Niblett 1980.

Stamper,R.K. 1985a, "Information: Mystical Fluid or a Subject for Scientific Enquiry", The Computer Journal Vol 28, No 2.

Stamper,R.K. 1985b, "Management Epistemology: Garbage In, Garbage Out" in Methlie and Sprague 1985.

Stamper,R.K. 1985c, "A Logic of Social Norms for the Semantics of Business Information", in Steel and Meersmann (eds) 1985.

Stamper,R.K. 1986, "The Processing of Business Semantics: Necessity and Tools" in Proc.IFIP Working Conference Knowledge and Data, forthcoming from North Holland.

Steel,T. and R.Meersmann (eds) 1985. Database Semantics, Amsterdam, North Holland.

Twining,W. and D.M.Myers 1976, How to Do Things with Rules, London, Weidenfeld and Nicolson.

Whorf,B.L. 1956, Language, Thought and Reality, Cambridge Mass, MIT Press.

Weizenbaum,J. 1976, Computer Power and Human Reason, W.H.Freeman.

Whitehead,A.N. and B.Russell, 1910/1927, Principia Mathematica, Cambridge, Cambridge University Press.

Winograd,T. and C.F.Flores 1986, Understanding Computers and Cognition, Norwood, N.J., Ablex Publishing Co.

Wynn,E. 1979, Office Conversation as an Information Medium, (PhD thesis), Berkley, University of California.

Ronald STAMPER:

Department of Information Management

University of Twente, The Netherlands

James BACKHOUSE:

Department of Information Systems,

London School of Economics

Karl ALTHAUS

MTM Systems,

London