Chicago-Kent Law Review

December 1997

Machine Intelligence and Legal Reasoning - The Charles Green Lecture in Law and Technology

J.C. Smith

MACHINE INTELLIGENCE AND LEGAL REASONING

JC SMITH*

INTRODUCTION

Has any one seen HAL, the computer in the Arthur C. Clarke and Stanley Kubrick film, *2001: A Space Odyssey*, "that could think, talk, see, feel, and occasionally go berserk"? HAL was supposed to have come into consciousness and be operational on January 12, 1997.[1] While Clarke and Kubrick, "failed to predict the biggest advance of the past 20 years: miniaturization and microelectronics,"[2] nevertheless we do not even have computers which can manage natural language at the level of a two year old infant. Is Clarke and Kubrick's, (and many others') vision of a HAL-like intelligent computer an impossible dream, or was merely their "technological time line . . . woefully inaccurate"?[3] At the present time we simply do not know for certain. Nevertheless, I will give reasons why, for the foreseeable future at least, the traditional artificial intelligence ("AI") paradigm is not only misconceived, but results in a misdirection of effort and scarce resources. An alternative, albeit more modest paradigm will be offered, which leads research and development in a different and more fruitful direction. These issues, as will be shown, are of fundamental importance to the field of artificial intelligence and law.

The area of artificial intelligence and law has a special interest for both computer scientists and the legal profession. Most legal information, at the present time, is originally created in electronic form, thus providing large quantities of data for experimentation. Many computer scientists in the field of artificial intelligence, therefore, are drawn to the discourse of the law as a fertile field for research, more so than other types of text, because of the increasing reliance of the legal profession on information in electronic form, the fact that the content has structure, and the tantalizing notion that law is based on a

1. *See* Simson Garfinkel, *Happy Birthday HAL*, WIRED, Jan. 1997, at 120.
2. *Id.* at 124.
3. *Id.* at 121.

logical foundation. From the perspective of the lawyer, we have the concept of *the rule of law*, as contrasted with *the rule of persons*; thus, in some sense separating the legal conceptual process from the human. It is this contrast that seduces us into envisaging judges as reasoning machines and computers as judges. Professor Anthony D'Amato wrote an article with a title in the form of the question, *Can/ Should Computers Replace Judges*?[4] In the course of this article D'Amato provides "A Modest Proposal for Dehumanized Decision-making."[5] He concludes his article with the statement that "[b]y re-moving a large area of unpredictable 'judgment' from the law, society may benefit from a sharply reduced number of litigated cases," thereby resulting in fewer judges, fewer attorneys, and fewer cases.[6] He calls this vision, "a determinable legal system."[7]

The very title of such an article entails certain presuppositions about the nature of legal discourse and reasoning, and about the na-ture of what we call artificial intelligence. Such a title presupposes that legal discourse is fairly lucid and stable, and that legal reasoning is rule governed. The very question that constitutes D'Amato's title must implicitly entail these two assumptions for the question to make any sense, as a high degree of lucidity, clarity, and rule governedness would be essential in order that a computer could render an accepta-ble judgment on a legal issue. The question further assumes that the term artificial intelligence is more than a mere metaphor, in that it is, at least in part, real or genuine intelligence, except that it is artificial in the sense that it is not human.

There are, however, a prior set of questions about the nature and structure of human intelligence, legal discourse, legal reasoning, and machine intelligence which ought to be considered before we even ask questions such as whether computers can, and if so, ought to function as judges. The answer to these prior questions ought to determine the kind of research and development we undertake in the field of com-puters and law because the objectives we choose and the methods we use will be different according to the positions we take on these pre-liminary issues.

The question, *Can/Should Computers Replace Judges?* is based on two common delusions. These two delusions, together, result in a mis-understanding of the nature and structure of intelligence. The first of

4. Anthony D'Amato, *Can/Should Computers Replace Judges,* 11 GA. L. REV. 1277 (1977).
5. *Id.* at 1288.
6. *Id.* at 1300-01.
7. *Id.* at 1301.

these has been called "Descartes' Error."[8] Descartes' error was in concluding that intelligence can exist independently of the human body. Another name for Descartes' error is mind/body dualism. The second delusion I will designate as "Leibnitz's Fallacy." Leibnitz's fallacy was in believing that all human thought could be expressed through a universal language which could function with mathematical precision, and that social problems could consequently be solved through computational means. Leibnitz's motto for resolving disputes and difficult issues was, "*Calculemus!*," "Let us calculate."[9] He said that, "[t]hings are like numbers,"[10] and he once wrote a thesis entitled *New Method of Teaching and Learning Jurisprudence*,[11] so he most certainly, had he been asked, would have responded to D'Amato's question with an enthusiastic, "Yes!"

According to the Church-Turing thesis, machine intelligence and human intelligence are essentially equivalent.[12] While human intelligence is far more complex than machine intelligence, there nevertheless is no qualitative difference between the two. If the Church-Turing thesis is valid, then the human brain can be usefully viewed as a neural computer and human languages can be viewed as corresponding to software, and the computer can be viewed as a mechanical brain.[13] The Church-Turing thesis leads to a reverberating set of metaphors. Human intelligence becomes a metaphor for the computer, and the computer becomes a metaphor for the human. The conception of the human as an organic machine, and the computer as a mechanical entity, has not only profoundly affected the field of artificial intelligence, but it has shaped our own view of ourselves.[14] The use of computer discourse to describe human mental processes is a common and widespread phenomenon. People now speak of programming and reprogramming themselves. It is not uncommon to hear forms of mental illness referred to as hardware and software problems. Others speak of memory lapses as problems of information retrieval.[15]

In 1965, Herb Simon predicted that by 1985, (12 years earlier than the predicted birth of HAL), "machines will be capable . . . of

8. ANTONIO R. DAMASIO, DESCARTES' ERROR: EMOTION, REASON, AND THE HUMAN BRAIN 247-48 (1994).
9. Z.A. MELZAK, BYPASSES: A SIMPLE APPROACH TO COMPLEXITY 228 (1983).
10. 4 THE ENCYCLOPEDIA OF PHILOSOPHY 422 (Paul Edwards ed., 1967).
11. *Id.* at 423.
12. *See* RAYMOND KURZWELL, THE AGE OF INTELLIGENT MACHINES 117, 152 (1990).
13. *See* WARREN S. McCULLOCH, EMBODIMENTS OF MIND (1965).
14. *See* SHERRY TURKLE, THE SECOND SELF: COMPUTERS AND THE HUMAN SPIRIT (1984).
15. *See id.* at 285.

doing any work that a man can do."[16] Predictions such as this are implicitly based on the validity of the Church-Turing thesis. If, however, that thesis is invalid, then it is highly unlikely that machines will ever "be capable of doing any work that a man can do." These predictions have failed to materialize. This failure may, on the one hand, be due to the fact that Turing, Simon, and others were overly optimistic as to the time frame, or, on the other hand, the failure may be due to the possible invalidity of the Church-Turing thesis. If these predictions were to come true in some distant future, then it would be very likely that at least some legal disputes could be adjudicated by machines. In order to address Professor D'Amato's question, however, we must examine the Church-Turing thesis more closely, from the perspective of Descartes's Error and Leibnitz's Fallacy.

II. DESCARTES'S ERROR

It is one of the supreme ironies of human understanding that very often the things we take most for granted are the most difficult to explain or define.[17] This is certainly true of the idea of intelligence. We have no theory of the mind or psyche, and therefore, even though we can recognize it, we cannot explain intelligence. It is not the case that we have a number of theories which are, as yet, not proven, and we merely have to wait until more evidence comes in to resolve the theoretical issues which remain in dispute. Rather, we simply have no true theories at all. All that we have are models, and models are not theories. A true theory must explain the nature and origin of a phenomena. A model has little or no explanatory power.

A. The Ingredients of Intelligence

Human intelligence has four necessary conditions: (1) consciousness, (2) the self, (3) perception, and (4) language.[18] Human intelligence entails consciousness.   Cognition takes place within

16. HERBERT SIMON, THE SHAPE OF AUTOMATION FOR MEN AND MANAGEMENT 96 (1965).

17. Martin Heidegger spent an academic life time exploring such ordinary everyday concepts such as *being, time,* and *thinking. See, e.g.,* MARTIN HEIDEGGER, BEING AND TIME (John Macquarrie & Edward Robinson trans., 1962); MARTIN HEIDEGGER, WHAT IS CALLED THINKING (Fred D. Wieck & J. Glenn Gray trans., 1968).

18. T.J.M. Bench-Capon defends the traditional view of AI, but states that "[i]n discussing the notion of intelligence as applied to computer systems, however, we should be aware that what counts as an intelligent computer system is not the same as what would count as intelligent for a human being." T.J.M. BENCH-CAPON, KNOWLEDGE REPRESENTATION: AN APPROACH TO ARTIFICIAL INTELLIGENCE 3 (1990). Thus the defenders of the Church-Turing thesis and the traditional view of AI want to have it both ways.

consciousness. There is no outside of consciousness from which we can view it. Therefore, "[s]cience's biggest mystery is the nature of consciousness."[19] We have developed no theory of consciousness because it is consciousness that is doing the theorizing. Consciousness is self-reflective. It is the perception of perception, and the awareness of awareness.

Human intelligence assumes the existence of a self. Yet where and what is the self? The sentence, "I should learn to spend more time with myself," makes perfect sense, yet who is the I, and whose is the self? Who is doing the thinking, and what is the object of the thought? This paradox demonstrates that language presupposes the subject. "I am thinking about myself." "I am now thinking about myself, thinking about myself." The self as subject always recedes outside of language as the thinker or the speaker. This is the reason we say such things as "my body" even though there can be no self without the body, and there can be no body without the self.[20]

Our perceptual world consists of vision, sound, taste, heat, cold, pain, pleasure, touch, and smell. We have no explanation for even the simplest of sensory experiences. We know red when we see it. We understand the physics of light. We understand how the eye works as a biological sensory system. We do not, however, have any explanation whatsoever for how the processes going on in the retina, the optic nerve, and the brain produce within an individual the experience of redness. Biology and physics alone cannot explain the phenomena of *seeing* red.

Language and consciousness presuppose each other in that each is a necessary but not a sufficient condition for the other. The structure of language entails a syntactical distinction between subject and object and a related conceptual distinction between self and other, which presupposes consciousness. It is impossible to conceive of an animal that had language and yet did not have consciousness. Animals with complex communication systems, such as those possessed by ants, bees, whales, or dolphins do not have language. Animal communication is not structured like a language, but rather like a code. Language consists of very much more than merely the transfer of information. The fundamental difference between plant and animal life is that plants have no internal representation of an outside environ-

19. NICK HERBERT, QUANTUM REALITY: BEYOND THE NEW PHYSICS 249 (1985).
20. *See* Richard Warner, *Is the Body a Physical Object?, in* NATURALISM, A CRITICAL APPRAISAL 255 (Steven J. Wagner & Richard Warner eds., 1993).

ment. Animals, unlike plants, have an inner perceptual representation of not only the external world, but the external and internal parts of their own bodies. Animals, therefore function with an instinctive perceptual system that spans the external and the internal, converging at points that mark the inside from the outside. This perceptual representational system, which models both the animal's body and the external world, permits the organism to modify its behavior and to learn to distinguish between the beneficial and the harmful. The human brain, on the other hand, has evolved the capacity to create a system of symbols to represent the perceptual system, which permits individuals to formulate goals and plan actions to achieve them. Within the human there is no clear boundary between the perceptual system and the symbolic system. The perceptual and symbolic systems merge into a conceptual-perceptual representation of the external world. As we perceive, we filter our perceptions through language, and probably our perceptions and language impose a form on the development of the neural network of the brain, which in turn filters perception and language.

Consciousness functions like a meta-representational system. Thus we have the knower and the known, the self, and the other. When the knower knows itself, the knower unavoidably objectifies itself. We say such things as "I have the right to do what I want with my own body." Since one's own body is an object of observation, mind/body dualism is implicit in consciousness. We thus speak of our minds, our psyches, our souls, and our bodies. An adequate theory of the mind would require an understanding of the relationship between the psyche and the brain. We know that there is a relationship between the two, but we don't know what that relationship is. No one as yet has transcended mind/body dualism, nor has bridged the phenomenological gap between the two. Even though consciousness is central to intelligence, we must nevertheless keep in mind Freud's admonition not to overvalue the property of being conscious. Much of the mental processes which go on in the use of language take place at the unconscious level. When we read the sign "USED FOOD EQUIPMENT FOR SALE," it is completely open as to whether the word *used* is to modify the word *food*, or the word *equipment*. Our mental process almost instantaneously recognizes that we use the adjective *used* far more often to modify the word *equipment*, than we ever use it to mod-

ify the word *food*. We generally reach the appropriate interpretations without consciously thinking about them.[21]

When we view machines as thinking entities, we project our own mind/body dualism onto them. When in a science fiction story or movie such as *2001: A Space Odyssey*, we project intelligence onto an imaginary computer such as HAL, we implicitly endow it with some perceptual mechanisms that would permit the machine to have a dynamic internal representation or model of the external world. The machine would be able to differentiate what was internal to it from what was external to it, and thus would evolve a conception of a self. The machine would further have a symbolic representational system, constituted in language, which would represent the perceptual system. The machine could then be said to be conscious. Thus we project a homunculus, or a little human psyche, inside a mechanical body.

B.   Mind Without Body

Descartes's error is to conceive of the possibility that the mind is independent of the human body.[22] Implicit in the Church-Turing thesis is the assumption that intelligence can exist independent of the human body. While it is true that machines can calculate, calculation is only one of many particular functions of intelligence, and it is not the same thing as intelligence. Intelligent entities can calculate, but it does not follow from this that calculating entities are intelligent. The essential characteristic of intelligent agents is the capacity to act in the fullest sense of action. Human intelligence is manifested in the capacity to have goals, plan and evaluate alternative strategies, and carry them out.

Antonio R. Damasio, a professor of neurology, argues in his book, *Descartes' Error: Emotion, Reason, and the Human Brain*, that emotions are a necessary condition of human action.[23] According to Damasio, emotions and feelings are not "intruders in the bastion of reason" but are "enmeshed in its networks, for worse *and* for better."[24] He gives examples where the emotional part of a person's brain has been damaged, but the reasoning part is left untouched. These people are able to formulate and select the best plan of action, but still, nevertheless, remain unable to act. Action requires emotion. Damasio writes:

---

21. *See* STEVEN PINKER, THE LANGUAGE INSTINCT (1994).
22. *See* 1 THE PHILOSOPHICAL WORKS OF DESCARTES 101 (1972).
23. *See* DAMASIO, *supra* note 8.
24. *Id.* at xii (emphasis in original).

> I believe that, relative to the brain, the body proper provides more than mere support and modulation: it provides a basic topic for brain representations.
>
> . . . .
>
> In the perspective of the above hypothesis, love and hate and anguish, the qualities of kindness and cruelty, the planned solution of a scientific problem or the creation of a new artifact are all based on neural events within a brain, provided that brain has been and now is interacting with its body. The soul breathes through the body, and suffering, whether it starts in the skin or in a mental image, happens in the flesh.[25]

Freud, who wrote in the German language, used the German word for *soul* which the English translators translated as psyche, a Latin term, because the English word "soul" has religious connotations, which the German word does not.[26] Freud did not believe that mind could exist independent of the living human body, nor, for that matter, does Damasio.

Damasio goes on to say that:

> My view then is that having a mind means that an organism forms neural representations which can become images, be manipulated in a process called thought, and eventually influence behavior by helping predict the future, plan accordingly, and choose the next action; Herein lies the center of neurobiology as I see it: the process whereby neural representations, which consist of biological modifications created by learning in a neuron circuit, become images in our minds; the process that allows for invisible microstructural changes in neuron circuits (in cell bodies, dendrites and axons, and synapses) to become a neural representation, which in turn becomes an image we each experience as belonging to us.[27]

If emotions are necessary for action, and the capacity to act is necessary for intelligence, then we can safely assume that emotion is a necessary aspect of human intelligence.

Descartes's error was in conceiving mental phenomena as independent of the human body. The traditional paradigms of artificial intelligence commit the same error. Terry Winograd, one of the pioneers of the field of artificial intelligence chose for the title of an essay the question, *Thinking Machines: Can there be? Are we?*[28] The title of D'Amato's article, *Can/Should Computers Replace Judges?*, while

---

25. *Id* at xvii.

26. *See, e.g.*, BRUNO BETTELHEIM, FREUD AND MAN'S SOUL, at xi, 11-12 (1983).

27. DAMASIO, *supra* note 8, at 90.

28. Terry Winograd, *Thinking Machines: Can there be? Are we?*, in THE FOUNDATIONS OF ARTIFICIAL INTELLIGENCE 167, 167 (Derek Partridge & Yorick Wilks eds., 1990) [hereinafter FOUNDATIONS].

somewhat more specific, raises the same issues. Winograd's conclusion bears equally on D'Amato's question. Winograd concludes:

> In the tradition of artificial intelligence, we project an image of our language activity onto the symbolic manipulations of the machine, then project that back onto the full human mind.
>
> But these projections are like the geometric projection of a three-dimensional world onto a two-dimensional plane. We systematically eliminate dimensions, thereby both simplifying and distorting. The particular dimensions we eliminate or preserve in this exercise are not idiosyncratic accidents. They reflect a philosophy that precedes them and which they serve to amplify and extend. . . . [W]e all too easily dismiss the concerns of human meaning that make up the humanities, and indeed of any socially grounded understanding of human language and action. . . . [W]e lose sight of the tacit embodied understanding that undergirds our intelligence.[29]

A classic example of Descartes's error in the context of artificial intelligence is the concept of intelligent "Software Agents." One web site proclaims that "[s]oftware agents are on-line pseudo-people."[30] A software agent is defined as a piece of software which acts to accomplish tasks on behalf of its user. Etzioni and Weld ascribe the following set of properties to software agents.[31] Software agents would be autonomous, goal-oriented, collaborative, flexible, self-starting, temporally continuous, character-based (defined as "a well-defined, believable personality"[32] and emotional state), communicative, adaptive, and mobile. If the agent is mobile then it is able to "transport itself from one machine to another and across different system architectures and platforms."[33] Thus, a software agent is just as independent of the machine as Descartes's soul or mind is of the human body.

III. LEIBNITZ'S FALLACY

Leibnitz's dream of a universal, mathematical-like language that would permit human thought to be ideally carried out as a process of logical calculation has turned out to be highly flawed. Language is much more complex than had previously been realized. Gödel has demonstrated that logical systems, in and of themselves, are either not provable, or they are incomplete.[34] Heisenberg's uncertainly principle

---

29. *Id.* at 188-89.
30. Michael Coen, *What's an Agent?* (visited Feb. 24, 1998) <http://www.ai.mit.edu/people/mhcoen/agents/section2_1_1.html>.
31. *See* Oren Etzioni & Daniel S. Weld, *Intelligent Agents on the Internet: Fact, Fiction, and Forecast,* IEEE EXPERT, Aug. 1995, at 44-45.
32. *Id.* at 45.
33. *Id.*
34. *See* KURT GÖDEL, ON FORMALLY UNDECIDABLE PROPOSITIONS (1962).

has revealed some of the limitations of mathematical models of the physical universe.[35] Our most profound philosophers such as Nietzsche, Heidegger, and Wittgenstein have clearly demonstrated that language is far more complex than what was assumed to be the case by the traditional rationalism and logical empiricism which, in the past, has inspired the field of artificial intelligence.[36] Leibnitz's view of language is wrong on at least two counts. First, it assumes that precision is better than imprecision, and secondly that the relationship between words and what they signify is necessary or fixed. Both of these presuppositions run counter to the nature and function of legal discourse. Many of the concepts of law such as "reasonable doubt," "malice," "proximate cause," "responsibility," etc., are functional precisely because of their imprecision. The issue of what is the most appropriate legal doctrine to apply to a set of facts lies at the very heart of "hard cases," and the very practice of using relevant precedents assumes that the same legal doctrine can be applied to different factual situations. Legal reasoning thus exhibits a marked mobility in the relationship between the words and concepts of legal discourse and the material facts of life they can be used to signify.

A. The Value of Imprecision

In learning about human intelligence through the process of simulating intelligence in the machine, we are haunted by an uncertainty about our own thought processes. It would seem that we are not clear about what kind of thinking is ideal and what kind is second best. The issue may be reduced to the simple but difficult question as to whether mathematical and logical reasoning is better than less formal thought processes such as practical reasoning, and the somewhat related question as to whether a deductive system of knowledge is in some way superior to a body of knowledge that is neither systematic or logical, but is practical and intuitive. Poets and novelists would favor the latter, but scientists, mathematicians, and logicians would likely prefer the former. As far as human thought is concerned, we need not privilege one form over the other because humans can do both and can appropriately choose which form is best for which purposes. For machines, however, it is different. The capacity of the machine to carry out mathematical and deductive computations is far

---

35. *See* LEE SMOLIN, THE LIFE OF THE COSMOS 245-47 (1997).
36. *See* JOSEPH WEIZENBAUM, COMPUTER POWER AND HUMAN REASON 221 (1976).

greater than that of humans, while, on the other hand, the facility of the machine with natural language is very limited.

The way in which we comparatively evaluate mathematical–logical–deductive reasoning and less systematic and structured discourse has critical implications for artificial intelligence. The most characteristic aspect of machine intelligence is that it is computational, and because it is computational it is precise. Propositions are either true or false and premises either follow deductively or they do not. Consequently, precise forms of human thought such as logical deduction or mathematical reasoning, where everything is reducible to two states representable by a binary system, can be represented or simulated very well within an intelligent machine system. By far the greatest part of human intelligence is not computational, and consequently it is not precise. Words and corresponding concepts such as high, low, hot, warm, cold, far, very far, near, very near, etc., are all imprecise. One can always substitute a precise mathematical measure in terms of kilometers, meters, and centimeters, or in terms of the gradations of a temperature scale on a thermometer. Such precision, however, has a cost: the cost of measurement or of getting exact information.

Human intelligence is efficient because we consciously or unconsciously reduce transaction costs by maximizing the degree of imprecision with which we can function before the costs of imprecision outweigh the savings gained by being inexact. We say that a person is young, but we could equally say that the person is ten years, eight months, three weeks, four days, and five and one-half hours old. Seldom, however, do we need to go beyond the age in years. We almost never act on the basis of complete information. We can state a general rule of action that humans reduce the transaction cost of acting by finding the appropriate degree of imprecision of information by measuring the costs of acquiring the information against the costs of imprecision.

There is an inverse relationship between information costs and the costs of imprecision. The higher the degree of precision or accuracy in human action, the greater the information costs, and conversely, the higher degree of imprecision we can tolerate in our actions, the greater the degree that we can reduce information costs. At the same time, there are costs to imprecision. Human action generally takes place at the point of economic maximization as between the costs of acquiring the information necessary for precision, and the costs of acting with less than maximum information. We might repre-

sent this with a graph having two axes, cost and precision, as seen in Figure 1.

FIGURE 1

COST

Cost of Imprecision

Minimization of
Costs

Information
Costs
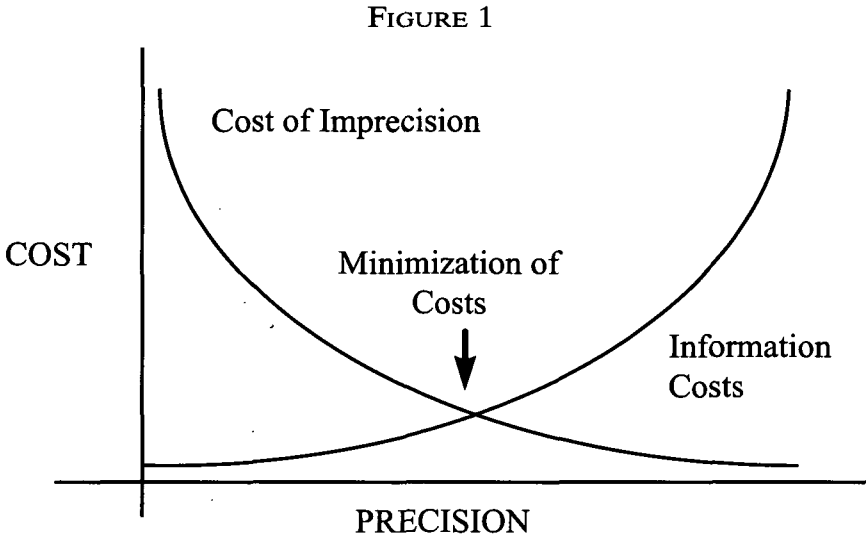
PRECISION

The relationship between the costs of imprecision of knowledge and the costs of the acquisition of knowledge furnishes the primary justification from an economic perspective for the existence of rules. The optimum time for an eight-year-old child to go to bed will differ from day to day depending upon a large number of variables such as their activities of that day, what time they got up in the morning, and the activities awaiting them the next day. It is costly for the parent to argue and debate these various factors every evening to find the optimum time. It is more economic to set a general time, for example, of nine o'clock each night. Thus, rules save the cost of finding the optimum action or point of action on each occasion by tolerating a degree of imprecision. Rules should never be absolute, however, as there will always be occasions when the costs of imprecision will be too high. Thus, parents will waive the rule for a special television program or a visit of a close relative such as a grandparent.

The discourse of the law is replete with imprecise concepts and terms. Reasonable belief, proximate cause, the standard of care of the reasonable person, good faith, malice, balancing of interests, etc., are to name but a few. The law often deals with very fundamental issues in extremely imprecise ways. Legal discourse is often so imprecise

because the information costs to achieve a greater degree of precision are generally too high. Take, for example, the issue of the criminal responsibility of children. Some children have the mental capacity to be responsible for their actions by the time they reach five, while others achieve it at a much later age. The law, however, arbitrarily sets the age of criminal responsibility at a fixed point such as the age of seven. This saves the cost of an intensive testing program in each individual case where a young child has committed a crime.

Some children reach a sufficient degree of maturity to properly exercise a franchise in their early teens. Some adults, on the other hand, fail to demonstrate sufficient maturity to be entitled to vote. To save the extensive information costs for ascertaining the maturity of each citizen, the law sets the arbitrary age of eighteen. For some it is not soon enough; for others, it is too soon; but for most, it is an appropriate age.

One way in which the indeterminacy of human language and concepts can be dealt with in the computer is by using "fuzzy sets" or "fuzzy variables'" that is, "classes of objects in which the transition from membership to non-membership is gradual rather than abrupt."[37] According to Professor Lofti Zadeh, the pioneer of what has become known as fuzzy logic, "Fuzziness . . . is a concomitant of complexity. . . . In fact, it is the capability to manipulate fuzzy concepts that distinguishes human intelligence from the machine intelligence of current generation computers."[38]

B.   *The Metaphorical Nature of Language*

The Swiss linguist, Ferdinand De Saussure, who is generally credited with being the father of modern linguistics, demonstrated that the relationship between the linguistic sign and what it signifies is arbitrary.[39] One of the essential problems of language is the nature of the relationship between words and concepts and the material world of sensed perception. One of the most common misperceptions which people have about how language functions is the belief that words get their meaning by pointing to things. "There is an absolute non-equivalence between discourse and pointing,"[40] in that the word *cat,* for

---

37. Lofti A. Zadeh, *The Birth and Evolution of Fuzzy Logic*, 17 INT'L J. GEN. SYS. 95, 99 (1990).

38. *Id.* at 100.

39. *See* FERDINAND DE SAUSSURE, COURSE IN GENERAL LINGUISTICS 67 (1959).

40. JACQUES LACAN, THE SEMINAR OF JACQUES LACAN: BOOK III, THE PSYCHOSES 1955-1956, at 137 (Jacques-Alain Miller ed. & Russell Grigg trans., 1993).

example, does not get its meaning by pointing to a perceived animal in the external world, but rather in terms of a discourse which extends out into a series of contrasts between the organic and the inorganic, between plants, animals, and humans, which are different forms of life, and between cats, dogs, rodents, etc., which are different forms of non-human mammals. The idea that words directly represent something is an illusion, albeit a common one.[41] The word *yellow,* for example, derives its meaning in terms of a system of an already existing discourse of color in which it is contrasted with other colors such as blue. The meaning of *yellow* is formulated within the discourse of colors from which it cannot be separated as a single unit and still retain its meaning.

FIGURE 2
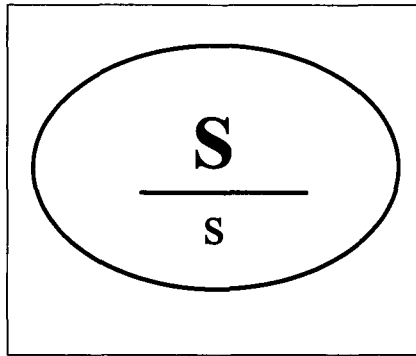
The approach to language developed by Lacan from the pioneering work of Saussure furnishes a linguistic foundation upon which to construct a more robust understanding of how legal discourse functions in the legal process. At the center of this approach is the relationship between the signifier and the signified which Lacan diagrammed as "the signifier over the signified," where the bar signifies the separation between the two such that signifiers slide over what is signified as seen in Figure 2.[42] Signifiers get their meaning in relationship to other signifiers, thus forming chains of meaning.[43] According to Lacan, "[f]or language to be born, it must always already be grasped as a whole."[44] "The word can in no way be regarded as a unit of language, even though it constitutes a privileged elementary

41. *See id.*
42. JACQUES LACAN, ÉCRITS 149 (1977).
43. *See id.* at 150-54.
44. LACAN, *supra* note 40, at 228.

form."[45] Thus "language is a system of positional coherence . . . ."[46] Language, according to Lacan, "is at its most effective when it manages to say something by saying something else. . . ."[47] "[M]etonymy exists from the beginning and makes metaphor possible."[48] Metaphor, however, is opposed to metonymy and operates at a different level.[49]

Saussure classifies the relationships between different terms in a language as syntagmatic and associative.[50] The relationship between signifiers is syntagmatic.[51] Signifiers form chains of signification which take the form of regular or orderly collections of statements, propositions, or doctrines.[52] Some of the syntagmatic relationships are that of opposition between concepts, derivation of one concept from another, class and sub-class, and between the universal and the particular (quantification), the logical relationships between propositions, and other logical patterns of inference. Saussure points out that the notion of syntagm applies not only to words but to groups of words and to complex units of all length and types whether consisting of compounds or derivatives.[53] A chain of signification that constitutes a syntagma generally bears a syntagmatic relationship to other syntagma. "Between the syntagmatic groupings . . . there is a bond of interdependence; they mutually condition each other."[54] As stated by Lacan, "no signification can be sustained other than by reference to another signification."[55]

Syntagmatic groupings of language at one level of discourse are associated with other levels of discourse as the signified.[56] The relationship between signifiers and the signified is thus, metaphorical. An example of the metaphorical characteristic of legal, political, and moral arguments, and the arbitrary nature of the relationship between language that functions as signifiers and language that functions as the signified can be shown by comparing the different signifiers which have been applied to the fetus as the signified. Each of the signifiers is

45. *Id.* at 225.
46. *Id.* at 226.
47. *Id.* at 224.
48. *Id.* at 225.
49. *See id.* at 227-29.
50. *See* DE SAUSSURE, *supra* note 39, at 122-27.
51. *See id.* at 127.
52. *See id.* at 70, 115.
53. *See* DE SAUSSURE, *supra* note 39, at 124.
54. *Id.* at 128.
55. LACAN, *supra* note 42, at 150.
56. *See id.* at 157.

a part of a separate chain of signification which will serve a particular interest. The signifiers parasite, property, person, soul, and baby all have been applied to the fetus as the signified. They belong to the chains of signification of individual autonomy, property rights, human rights, Christianity, and domestic relations. Each signifier is chosen to be applied to the signified according to what interest the speaker seeks to serve, whether that of the mother, the father, the state, the church, or the fetus itself. Thus in the discourse of disputes which surround the fetus, the signifiers float without any right, true, or necessary connection. They function metaphorically. The concept of the fetus, which functions as the signified in moral, political, religious, and legal discourses, functions as a signifier in *biological* discourse.

The nature of the relationship between the discourse of legal doctrine (the Law) and the discourse of the observations, information, and data of everyday life (the Facts) is one of the most difficult issues facing the practitioners of AI and law, working in the area of case-based reasoning and legal expert systems. If the relationship were fixed and determinate, our task would not be nearly as difficult as it is. Because the relationship is indeterminate and shifting, simulating legal reasoning in the computer presents an overwhelming challenge. Viewing the discourse of legal doctrine as signifiers and chains of signification, and the discourse of the material facts of life as the signified furnishes us with a useful conceptual tool for examining the relationship between what we call the law and the facts.

At times the lawyer starts with a set of facts as a given, and attempts to find sets of legal doctrine which will, when applied to those facts, produce the result desired by the client. Take, for example, a set of facts in which a hospital wrongfully posts information that a particular nurse has had a highly infectious disease, and people assumed that she had AIDS rather than infectious hepatitis. The cause of action could be drafted in any one or more of the following legal actions: negligence, breach of contract of employment, defamation, or invasion of privacy.[57] The law of contracts, the law of defamation, the law of privacy, and the law of negligence—each are doctrinal sets which are made up of chains of signification. This discourse is normally fairly abstract at the purely doctrinal level, and doctrinal arguments are generally fairly concise because they are frequently circular within their own frame of reference. The discourse that would be used to describe the material facts of the situation involving the nurse and the

---

57. *See* Peters-Brown v. Regina Dist. Health Bd., [1995] W.W.R. 337, 340-41.

damage she suffered constitutes the signified. The essence of a legal argument is to persuade the judge to select as the chains of signification, a particular doctrinal analysis which, when associated with the facts of the case, a conclusion favorable to her client will follow as a matter of course.

On other occasions, the lawyer has located the appropriate legal doctrine that might produce the desired result, and seeks alternative fact situations where that same doctrine has been applied in order to use as precedents. For example let us take the case where a young female university student is returning to her car in the university parking lot, and while walking down an unlit path bordered by shrubs and bushes is sexually assaulted. The legal issue is whether or not an occupier of land owes a duty of care to invitees or licensees to take precautions to protect them from the intentional wrongful acts of third parties. There are a wide range of possibilities that could give rise to this same legal issue, and consequently would be relevant cases. The nature of the premises, the occupier's business, the location of the wrongful act, and the nature of the wrongful act are all variables that can change and still remain within the range of relevancy. The separation of the discourse of legal doctrine from the discourse of the material facts of life, therefore, not only has a theoretical justification, but a practical application regarding how we structure and represent legal knowledge in the computer.

The way in which legal discourse is structured and operates is incompatible with the traditional paradigm of AI. Like Descartes's view of the autonomy of the mind, the conventional notion assumes that language can function independently of the perceptions of sight, sound, smell, touch, sensations, and emotions that are deeply-rooted in the human body. By equating human and machine intelligence, the traditional perspective implicitly entails the patently false assumption that the mere manipulation of symbols in the computer can adequately represent languaging processes rooted in perceptions and emotions. The precise binary and Boolean logical function of the computer are incongruous with the way humans reduce information costs through imprecision. The separation between the discourses of legal doctrine and the discourse of the material facts of life, and the equivocal and indeterminate nature of the way the one is related to the other in the processes of legal reasoning, is inconsistent with the logically based inference models of the prevailing conceptions of AI. An investigation of the potential for simulating legal thought processes in the machine requires an understanding of the structure of

legal doctrinal discourse and how it relates to the facts. In addition, we will need an alternative view of artificial intelligence that would permit us to represent the discourse of legal doctrine separately as signifiers in chains of signification which can then be associated with the discourse of everyday life in terms of relationships whose structures can be represented in the machine.

IV. THE STRUCTURE OF LEGAL DISCOURSE

The significant feature of signifiers is that they have no meaning standing by themselves, but only get their meaning in relationship to other signifiers. Signifiers thus remain as meaningless abstractions until they are associated with something else as a signified. The signified, itself, is likely to function as a signifier in the context of another level of discourse. Signifiers thus may be seen as the form of more familiar things, the signified, so long as we conceive of form as a syntactical structure rather than a physical shape. Pure signifiers, so long as they are not related to a particular signified, tend to be abstract by nature or form, empty of content. Thus, in many ways signifiers function as variables that can be associated with alternative specifics when given a concrete reference.

A. The Fundamental Signifiers of Legal Doctrine

1. Legal persons

The concept of the legal *person* furnishes us with a paradigm example of a fundamental signifier of legal doctrinal discourse. A legal person is not the same thing as a human being as not all human beings are legal persons and not all legal persons are human beings. Legal persons have no specific meaning in the abstract, but are empty of content. F. H. Lawson states of the legal person that "[l]egal personality and legal persons are, as it were, mathematical creations devised for the purpose of simplifying legal calculations."[58] David Derham, using similar language, writes that the "legal person ... is just as much a pure 'concept' as 'one' is in arithmetic. It is just as independent from a human being as one is from an 'apple.'"[59] The signifier, or legal person, does not get its meaning by pointing at or standing for an individual, but rather gets its meaning from its function and position in the syntagma, *legal obligation*. Legal persons have syntactical

---

58. F.H. Lawson, *The Creative Use of Legal Concepts*, 32 N.Y.U. L. REV. 909, 915 (1957).
59. David P. Derham, *Theories of Legal Personality*, in LEGAL PERSONALITY AND POLITICAL PLURALISM 1, 5 (Leicester C. Webb ed., 1958).

meaning in that they are either right holders or duty holders in the fundamental relationship of legal obligation. The right holder may be associated with a particular individual, or a business, or an association of people seeking some common purposes such as a trade union or an association, providing there is a legislative scheme that permits the group or the association of people to be granted *corporate* standing. There are specific kinds of legal persons such as landlords, tenants, trustees, testators, beneficiaries, etc. each of which gets its meaning in terms of the bundle of legal relations that constitute the particular legal interest. Thus the legal interest *reversion* defines the entity, *landlord*; the legal interest *leasehold estate in land* defines *tenant*; and *trust* defines *trustee*.

2. Actions

The meaning that links signifiers into chains of signification or syntagma is multi-relational. Chains of signification are in themselves related to each other in matrices or webs of meaning. The concept of *personhood* derives its meaning in part from the concept or idea of *action*, and *action* gets its meaning in terms of *actors* or *agents* who are particular kinds of *persons*. Only *agents* possess full *personhood*, and *agents* are those who have the capacity to perform *actions*, and *actions* are those things which *agents* can perform. Thus *action*, in this sense, is a signifier, and as such is a fundamental variable in the chain of signification which links legal persons in terms of the relationship of legal obligation.

Full fledged or standard actions entail certain properties such as intention. If an essential element of a standard action is missing, then we have a case of a non-standard act, or an adjusted act. We can assume if someone says that "John shot the donkey," that the action of shooting the donkey had all of the key elements of a standard act. If, however, someone says that John shot the donkey *unintentionally, accidentally, mistakenly, carelessly, negligently, inadvertently, or unconsciously*, then the action of shooting the donkey is non-standard. Each of these *adjusters* negate the presence of an element of a standard act. When we identify all of the elements which are negated, we will have the list of characteristics which constitute a standard act.[60] The concept of an action in legal doctrinal discourse does not get its meaning from the various bodily movements people make, but rather from a

60. *See* S.C. COVAL & J.C. SMITH, LAW AND ITS PRESUPPOSITIONS: ACTIONS, AGENTS AND RULES 1-28 (1986).

theory of the autonomy and responsibility of agents. The various defenses in the criminal law and the law of torts thus relate to the presence or absence of these elements that constitute a standard action.

3. Things or objects

The actions that constitute the subject matter of legal obligations are often in relationship to things or objects. Thus the object of the action constitutes a variable in the chain of signification which makes up the relationship of legal obligation. The object term of a legal relation, or a set of them, is not necessarily a physical thing, but may be immaterial such as a form of intellectual property or other bundle of legal relations. Again, rather than pointing to a physical object, this variable term in a legal relation, or set thereof, gets its meaning from the set of relations within which it functions as a term.

4. The fundamental syntagmatic relation

One of the fundamental signifiers of all normative discourse, whether moral, legal, or political, is *obligation*. What one has an obligation to do, one both ought to do and is obliged to do.[61] It is the element of necessity which distinguishes what we ought to do from a moral point of view, from what we are morally obligated to do. The meaning of obligation in the abstract comes from a matrix or web of chains of signification which underlies and furnishes the ideological foundations for Western social order. While we can maintain a clear distinction between legal and moral obligations, the signifier *obligation*, as it functions in each context, nevertheless shares many chains of meaning such as those relating to individual responsibility and the nature and structure of agency and action.

B. Some Fundamental Syntactical Properties of Legal Relations

1. Universalizability

The essential logical property of normative (as contrasted with merely descriptive) propositions about the existence or non-existence of a legal obligation is that it is universalizable. The universalizability of legal judgments about obligations underlies a series of legal doctrines such as the doctrine of precedent that like cases shall be decided alike, equality before the law, equality under the law, justice as equal

---

61. *See* J.C. SMITH, LEGAL OBLIGATION 34-59 (1976) [hereinafter LEGAL OBLIGATION].

treatment, the duty of judicial impartiality, and others.[62] A normative statement that a particular person ought to do something is equivalent in meaning to a proposition that there are good reasons for that person to do that action. Reasons for action entail assumptions about cause-effect relationships between actions and desired states of affairs. Thus, in the final analysis, the universalizability of normative judgments is founded upon the consistency of cause and effect in that like causes have like effects. It logically follows that if, prescriptively, it is the case that a person has a legal obligation to do a particular act, then it is the case that that person ought to do that act. In which case, it follows that: *Any judgment made in regard to a particular situation, that a particular person is or is not legally obligated to do a particular act, logically entails that the judgment instances a rule of law such that anyone in a relevantly similar situation is or is not legally obligated to do the same act.* All criteria of relevancy are teleological. Universalizability functions in normative legal discourse to maintain a teleological consistency within the legal system.

2. Opposition

The relationship of opposition between concepts is an important element of meaning by which signifiers are differentiated in relationship to each other. Opposition is not the same thing as negation. Bachelor and spinster are oppositional terms, but neither is the negation of the other. While it is the case that a bachelor is not a spinster, it does not follow from the statement, *It is not the case that a particular person is a bachelor*, that the person is a spinster. An oppositional relationship between two signifiers entails that they are opposed to each other within the framework of a common relationship, that of marriage. While both negate the condition of marriage, they are oppositional vis-À-vis each other.

The fundamental relation of law is that of legal obligation. Legal obligation is a three-term relation connecting two legal persons in regard to a particular act or pattern of behavior. In regard to any two legal persons and any particular act, one person will have the obligation to perform the act, the *duty*, and the other person will have the claim that the act be performed, the *right*. A right is equivalent to the three term relationship of legal obligation viewed from the perspective of the particular person having the benefit of the relationship. A duty is equivalent to the same three term relation viewed from the

62. *See id.* at 88-108.

perspective of the person having the burden of the relationship. Thus, rights and duties are oppositional.

3. Negation

Any proposition can be negated in some form or other. If the proposition expresses that something is the case, then it can be put in a negative form that something is not the case. Or if the proposition expresses that something is true, then the proposition can be negated by adding the term "not." Propositions can also contain negation within. Actions can be done or not done. Thus, negation can function within a proposition about legal obligations at two points, acting and not acting, and the existence or non-existence of the relationship itself. The relation of legal obligation can be expressed as the propositional form: *A particular person has a duty to another particular person to do a particular act.* The same relationship can be expressed in its negative form as: *It is not the case that a particular person has a right that another particular person do a particular act.* We can also express the negation of action as: *It is the case that a particular person has a legal obligation to refrain from (not do) a particular act.* Thus through using negation at these two points, we can create four kinds of propositional forms relating to the existence or non-existence of an obligation to do, or not to do, a particular act.

4. Logical implications

These four propositions can be placed at the four corners of the traditional Aristotelian Square of Opposition:[63]

---

63. *See* R.E. Robinson et al., *The Logic of Rights*, 33 U. TORONTO L.J. 267, 274-75 (1983).

FIGURE 3
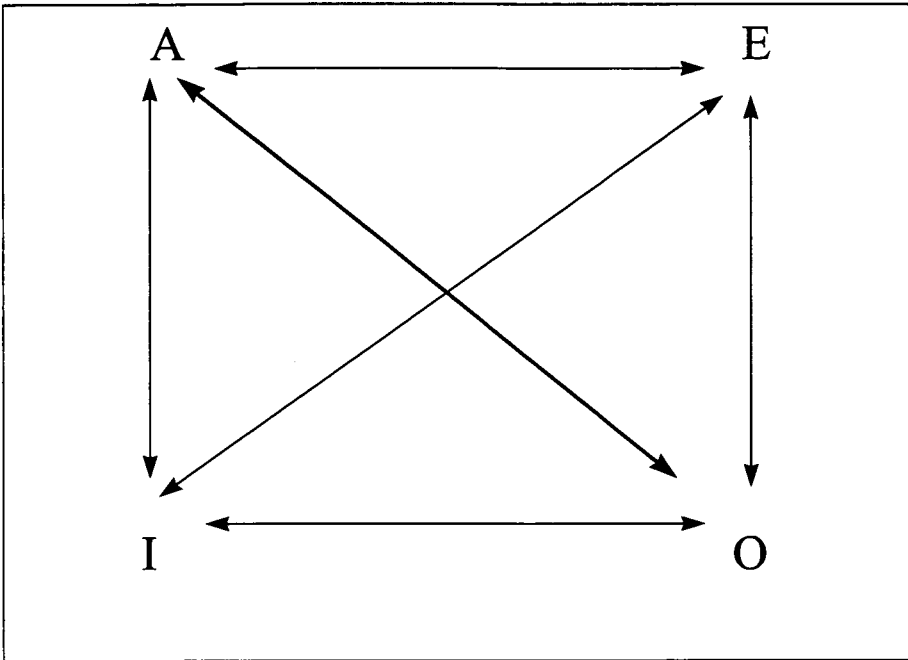
A. The duty holder is legally obligated to the right holder to do a particular action.

I. It is not the case that the duty holder is legally obligated to the right holder to do a particular action.

E. The duty holder is legally obligated to the right holder to refrain from doing a particular action.

O. It is not the case the duty holder is legally obligated to the right holder to refrain from doing a particular action.

By placing these propositional forms at the appropriate corners of the square of oppositions, the logical relations that hold between these propositional forms can be shown.
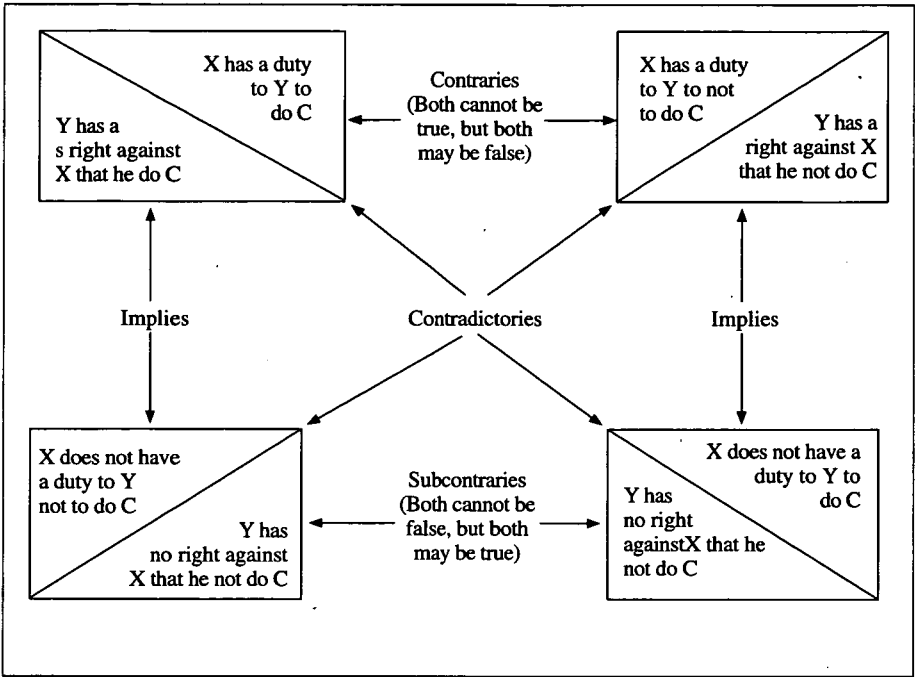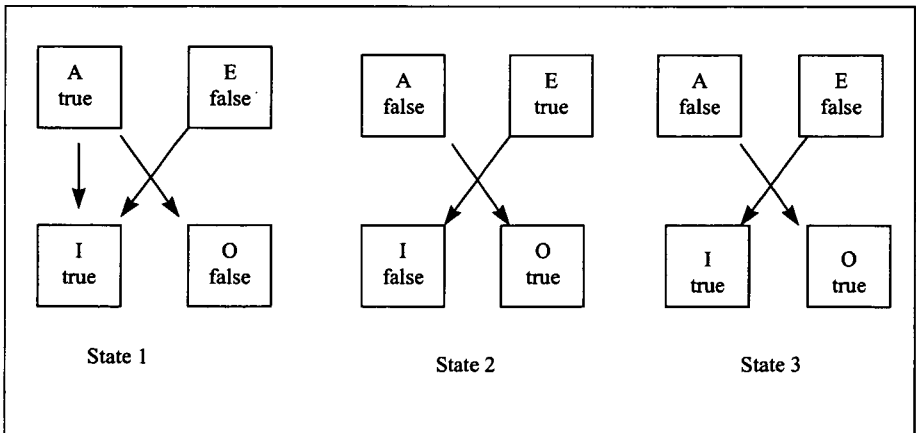
FIGURE 4

FIGURE 5

Depending upon the truth values of A, E, I, and O, there are three possible states for the square of opposition: in state 3, the duty holder has neither a duty to do the particular act nor a duty to refrain

from doing it. In that case, she can be said to be free or to have a choice to do or not do it, thus having a legal liberty or privilege. A liberty is equivalent to the negation of a legal obligation to do a particular act and the negation of a legal obligation to refrain from doing the act.

5. Quantification

The entity variables for any relation or set of relations can be either particularly or universally quantified. Every entity variable has a universe of discourse in terms of which the variable will be given a concrete value. Thus if the variable is particularly quantified, the relationship will only be true for some of the specific values of the universe of discourse. If the variable is universally quantified, then the relationship will hold true for every instance of the entities which make up the universe of discourse. The classical distinction between rights in rem and rights in personam reflect the syntactical property of quantification. Legal relations can be grouped in sets to create legal interests such as contracts. We can differentiate between four different kinds of legal relations or sets of legal relations in terms of how the entity variables of right and duty holders are quantified. They are:

a. Contract — The duties and rights are between particular legal persons.
b. Property — The duties fall on everyone but only particular persons have the rights.
c. Public Law — The particular public officials have the duties but everyone has the rights.
d. Tort Law — Everyone has the duties and everyone has the rights.

6. Classification

Legal relations such as rights, duties, privileges, and liberties are generally grouped together into sets such as contracts, property, trusts, mortgages, patents, copyrights, etc. From the existence of practices where we are free to create, terminate, or transfer legal relations, we postulate the existence of powers. Legal powers have their origin in the performative rules which prescribe how the practices are to be invoked.

Sets of legal relations may be made more complex by adding further variables. By adding the concept of an object to a set of legal relations we can construct the bundle of rights that constitute property. The legal relations, whether rights, duties, liberties, privileges, or

powers, that make up the set which constitute property contain at least four terms: the right holder, the holders of the duties, the patterns of behaviors or acts that are permitted or prohibited, and the object to which the patterns of behavior relate. Property rights can be made even more complex by adding the variable of time. Estates in land are bundles of rights relating to realty or funds, which are viewed as extending on a plane of time. Thus the presence of the time variable can be used to distinguish the sets of rights that constitute personal property from the sets of rights that constitute an estate in realty or funds. The time element in estates will have one of each of four sets of properties:

1. Infinite or Finite
2. Definite or Indefinite
3. Present or Future
4. Absolute or Defeasible (conditional)

7. Hierarchy

Hierarchical structures are found throughout law. Most legal concepts have a variety of hierarchical relationships to other legal concepts. The classification of estates, for example, reflects the following hierarchical configuration:

FIGURE 6
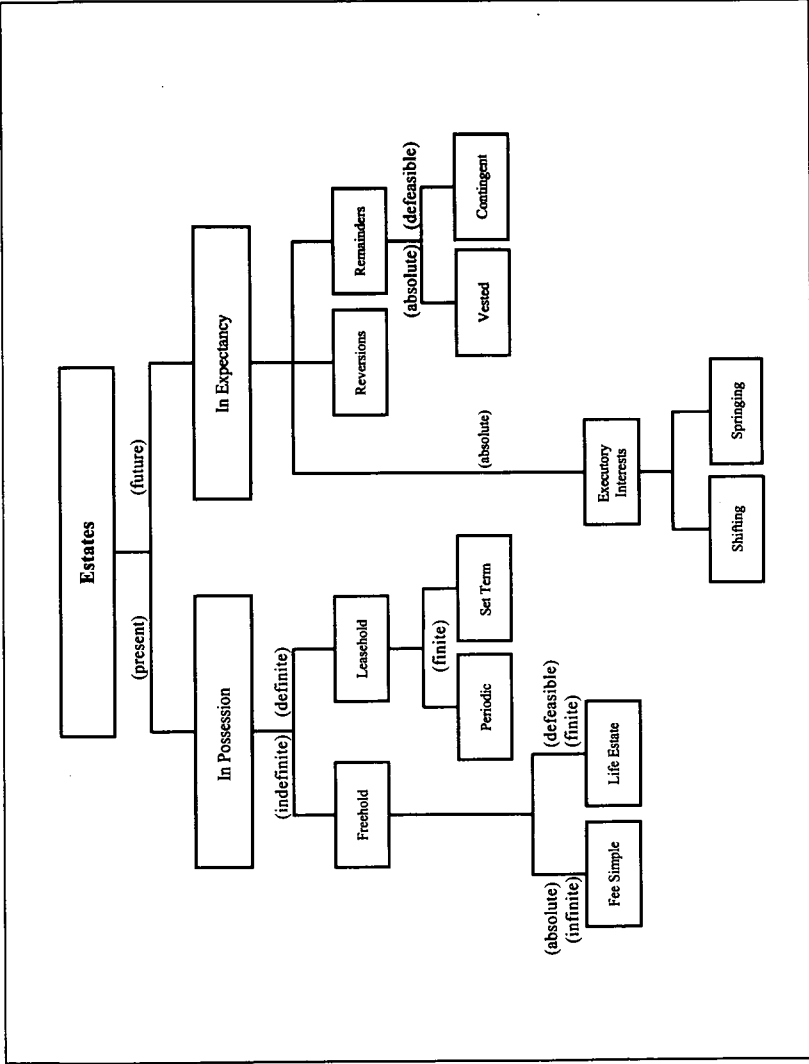
8.   Chaining

A further kind of complexity is possible in that the object of property relations can include sets of legal relations as well as material land and objects. We can therefore have a set of legal relations as the object of a further set of legal relations. The situation of the ordinary householder constitutes at least four levels of sets of legal relations between the human beings and the actual land and home. Thus, Mary Jones (1) *owns* (2) a *tenancy-in-common* in (3) an *equity of redemption* in (4) a *fee simple.* Joint tenancies and tenancies-in-common are sets of relations which allow more than one person to co-own a piece of property. They function to regulate the rights, duties, liberties, and powers as between the co-owners. Mortgages constitute a security interest whereby the mortgagor has certain rights to the property as security for a debt, while the person holding the equity of redemption has certain rights of use and enjoyment, including the right to pay off the indebtedness and remove the mortgage. Thus people can have almost all of the benefits of ownership without the full capital outlay. At the end of these sets of doctrinal chains of signification is the material and visual piece of earth and the buildings that stand upon it, which is the material factual content that functions as the final signified at the end of this set of complex chains of legal relations.

V.   THE RELATIONSHIP BETWEEN LEGAL DOCTRINE AND THE MATERIAL FACTS OF LIFE

The functions of the lawyer may be roughly divided into two branches–legal advising, drafting, and conveyancing, and litigation and dispute settlement. The former looks to the future and plans behavior, and the latter looks to the past, to correct what has gone wrong. The former attempts to shape the material facts of life for the future, while the latter deals with the past where the facts are fixed. The former is a matter of creative planning and falls within the practice of the solicitor, while the latter deals with pathology or what has gone wrong, and how best to correct it, and falls within the traditional role of the barrister. These two functions are closely interrelated in the sense that if something is not done correctly or some unseen event arises in the former, then litigation may follow. Work in the field of AI and law is relevant to both kinds of function. Much has been done in the field of automatic document assembly, even though this area of law and computers has not caught the attention or interest of those

working in the field of AI and law as much as has case-based reasoning.

A. Advising, Drafting, and Conveyancing

When a group of business people wish to pool their assets in order to create a business, the function of the lawyer is to construct the legal framework that will accomplish their desires. Because they wish to raise more capital, the lawyer constructs a public corporation rather than a partnership or some alternative form. A great many contracts would be entered into. It would be the function of the lawyer to become, as far as possible, aware of all the risks so that the parties could agree on where they were to lie. Shares would be sold, and further capital might be raised through the sale of corporate bonds. The corporation itself, when the incorporation was complete, would purchase and own assets, and enter into contracts in its own right. The articles and memoranda of incorporation would set out the various legal relationships holding between a variety of the legal entities involved. If the lawyer has accomplished the task well, that is to say if the lawyer is fully knowledgeable regarding the legal rules and selects the appropriate legal instruments and drafts the documents clearly, no problems should arise.

It is the function of the lawyer to select the appropriate legal practice (doctrinal area of law) to accomplish what the client desires. For example, let us assume that the client is a wealthy married woman in her late seventies, lives in a jurisdiction that imposes an inheritance tax on death, and requests her lawyer to draft her will. The lawyer will explain that if she leaves everything to her husband, who then leaves everything to their children, the estate will pay a heavy succession duty on the death of each parent. If, however, the survivor is given a life estate and the children vested remainders, the estate is only taxed once as the children will take a vested remainder on the death of the first parent.[64] There is a functional or teleological history behind the evolution of the law of property reflecting the interest of families in accumulating property and controlling it after death, and a public interest in the devolution and taxation of property and the limitation of the control of the dead hand.

If, in the above example, the wealthy client is concerned about her spouse's capacity to manage her estate if she predeceases him, the lawyer will likely suggest a trust where the surviving spouse can have

---

64. *See, e.g.*, 26 U.S.C. § 2056 (1994) (providing for federal taxation of survivor's estate).

the rights of beneficial enjoyment without the rights and powers of control. The set of relations that constitute a property interest can, in addition to being divisible on the plane of time, be split into two sets — those which constitute the rights, liberties and powers that relate to control, and those which constitute the rights and liberties of beneficial enjoyment. The former are held by the Trustee, and the latter are held by the *Cestui Que* Trust.[65]

The distinction between legal doctrine and the material facts of life underlie nearly every aspect of the practice of law. A contract, in the abstract, as a signifier, is an empty form to which we can associate different persons in regard to different patterns of action. There is no necessary or fixed relationship between any particular set of legal doctrinal structures and any particular material facts. We construct and structure the future by selecting the appropriate doctrinal structures that will accomplish our goals. The relationships between the concepts that constitute a particular legal doctrine such as Property or Contract are syntactical, but when we call a particular transaction a lease or a license, the relationships between the concepts and the particular set of facts are associative according to our teleological needs.

B.   Adjudication

Contrast the above with the following: an elderly gentleman is crossing the street at an intersection having a traffic light and is struck by a car. Both the elderly pedestrian and the driver experience certain perceptions. Both parties to the accident will commence to reconstruct the events so as to minimize their responsibility. When a client has been involved in a traffic accident, lawyer and client seek, if possible, to conceptualize the event in the form of a story script which is about cause and fault, blame, and responsibility. The cause-blame script is associated with particular kinds of events, in order to obtain a particular result. Legal reasoning is reiterative as between the facts and the law as the lawyer seeks to conceptualize the facts into a script which can be associated with a legal doctrine which, if applied, will achieve the desired result, and at the same time shape the legal doctrine to fit the script, as limited by the events. A decided case consists of a script which is then associated with a doctrinal chain of signification. These scripts, in terms of which we think as lawyers, determine the boundaries of relevancy.

---

65. *See, e.g.,* BLACK'S LAW DICTIONARY 208 (5th ed. 1979).

Just how do we as lawyers think when we approach a legal prob-
lem? The discourse of legal doctrine is generally independent of any
particular factual situation, but instead is often conceived of in rela-
tionship to broad factual themes or *scripts*. A script is a kind of story
line, where the main items are variable. We conceive, for example, of
the tort of assault and battery to the person in broad abstract terms
that can take a wide variety of specific instances, such as a blow with
the fist, striking with a club, or simply giving an unwanted and unsolic-
ited kiss. The lawyer's function in adjudication is to select a legal doc-
trinal structure that can be related to a factual script or story that can
match the factual reconstruction of the events through the testimony
of witnesses, as shaped by the rules of evidence and court procedure.

The typical legal dispute generally starts with an event or set of
events of some kind or other. People then conceptualize the events in
the form of stories, which are descriptions of what took place. Natu-
rally the stories of the two participants in a dispute will give different
versions of the events. Lawyers conceptualize facts in the form of
scripts, or story patterns, which are shaped by the interests and objec-
tives of their clients. As an example take the case where a women has
entered into and proceeded with a surrogate mother contract, and
then changes her mind after the birth of the baby. Assuming that the
court refuses to enforce the contract on the grounds that it is against
public policy, the courts still have permitted the biological father to
make a claim to the infant.[66] The issue could be conceptualized as
being similar to an adoption contract where a pregnant women agrees
to deliver up the child at birth to an infertile couple for adoption. The
case would be similar in that a contract to surrender the child after its
birth exists, but it is different in that the other contracting parties have
no biological relationship to the child. If this doctrinal conceptualiza-
tion was used, the mother would likely win as courts are reluctant to
enforce such contracts as they discourage unwed mothers to make
prior arrangements for giving up a child for adoption. It could be con-
ceptualized as a biological father seeking custody of a child born out
of wedlock. This conceptualization would be similar in that the dis-
pute would be between two biological parents, but it would be differ-
ent in that there is an absence of a contract. The father would be
unlikely to win using such a conceptualization as the courts do not
look with sympathy upon males who impregnate women to whom
they are not married. The issue could be conceptualized as a custody

66. *See* In re Baby M, 537 A.2d 1227 (N.J. 1988).

dispute in which case it would be similar in that the contest would be between two biological parents, but it would be different in that the biological parents would have never been married to each other. With this form of doctrinal conceptualization it is likely that the biological father will win given that the test being what is likely to be in the best interest of the child, because the father, being more wealthy than a woman who is likely driven to contract out her womb out of financial necessity, will be in a better financial position to provide for the child. Thus there are three possible ways in which the dispute can be conceptualized, each in some ways similar, and each in some ways different, two which would favor the mother, and one which would favor the father. The issue is as good as decided as soon as the judge adopts the particular form of conceptualization.

VI. MACHINE LANGUAGE AND REASONING

According to Terry Winograd, the field of artificial intelligence "has not achieved creativity, insight and judgment," nor have we "yet been able to construct a machine with even a modicum of common sense or one that can converse on everyday topics in ordinary language."[67] He goes on to point out that "[t]he source of the difficulties will not be found in the details of silicon micro-circuits or of Boolean logic," but that the "basic philosophy that has guided the research is shallow and inadequate,"[68] drawn as it is from the outmoded traditions of rationalism and logical empiricism rooted in the ideas of mechanized reasoning that Descartes and Leibnitz, among others, found so attractive. The computer is not intelligent. Machines do not think. Yet they are the product of intelligence. Information in digital form is used to symbolically represent the perceptual-conceptual system, and eventually that information is transformed back into the form of visual patterns and natural language. It would never occur to us to think of a book as being intelligent, no matter how profound its content. No one speaks of book or print intelligence. We store information in books and we store information in computers. Is the difference between the two (the fact that information can be processed in the computer), so significant that we are justified in imputing intelligence to one and not to the other?

Most who are familiar with the field of artificial intelligence would agree with the theorist who, in discussing the paradigmatic con-

---

67. Winograd, *supra* note 28, at 167.
68. *Id.*

fusion surrounding AI, states that "[i]n spite of what I regard as AI's significant achievements . . . the not so well-kept secret is that AI is internally in a paradigmatic mess. There is really no broad agreement on the essential nature or formal basis of intelligence and the proper theoretical framework for it."[69] Machine or artificial intelligence is a metaphor. We use metaphors to refer to something new which is different from everything else, but has sufficient similarities to sustain a metaphorical reference. The following is the best definition of artificial intelligence that I can think of. Artificial intelligence is what ever we would like to make the computer do which humans can do, but we do not quite know how to make the computer do it, but on the other hand, we think that there is a distinct possibility that we can make the computer do it. Yet if and when we actually do succeed in making the computer do it, then it is no longer artificial intelligence but pattern matching or something else equally mundane. The Holy Grail always becomes a tin cup once we have obtained it.

A.   The Logical Structure of Machine Language

Machines have no consciousness, no self, and no perception, but they do have language. Machine language, however, is very different from human language in that machine language is not a symbolic representational system for the machine, but is a symbolic representational system for human language. At the most fundamental level of its symbolic representational system, the computer can have only two states, the presence or absence of an electromagnetic phenomenon. Thus any information that a computer can process, whether words, graphics, or sound, is represented in the computer in terms of binary numbers.

Boolean logic, with its three fundamental operations, AND, OR, and NOT, will process two kinds of entities such as true or false, yes or no, open or closed, on or off, or 0 or 1. The convergence of Boolean logic with binary mathematics thus forms the underlying structure of computer design. When on or off switches are arranged according to Boolean principles, circuits can be created that can perform both mathematical and logical operations. Simple AND, OR, and NOT gates can be combined in a large number of complex ways permitting the performance of complex logical and mathematical calculations. By representing letters with binary numbers, the computer can repre-

---

69. B. Chandrasekaran, *What Kind of Information Processing is Intelligence?*, in FOUNDA-TIONS, *supra* note 28, at 14.

sent and manipulate human language. Information passes between the human and the machine through at least three levels of representation. Each level has its own kind of language or languages. The first level is that of binary machine language which is machine readable only. The third level is natural language which is human readable only. Connecting the two are a set of programming languages which are readable by both the human and the machine.

Machines can perform complex mathematical and logical calculations at a speed and complexity far beyond that of any human being. Computers can outperform even the world class chess masters.[70] Yet a child of three can far outperform the computer in the manipulation and use of natural language. If intelligence entails consciousness—the self, perception, as well as language—then there is only human intelligence. If we take the Church-Turing thesis seriously, then we must look upon the computer as an idiot savant, so far as intelligence is concerned. It can barely manage to use natural language, but can brilliantly solve mathematical and logical problems. Language is constituted by complex webs, chains, and matrices of semantic and syntactic structures which do not remain fixed. Some working in the field of AI recognize that "even science is not the paradigm of literal language it was once considered to be; rather, metaphor is vital to the modeling processes that result in advances in science," and that "metaphor may be the very heart of how we think and learn."[71]

B. The Limits of Machine Reasoning

Artificial, mathematical, and logical languages are infinitely too simplistic to deal with the complexity and indeterminacy which surrounds human action and conflicts of interest. Anne Von der Leith Gardner in her book, *An Artificial Intelligence Approach to Legal Reasoning,* attempted to simulate the legal reasoning process around the issue of offer and acceptance in the law of contracts.[72] The basic conclusion one has to reach after reading the book is that even a single issue of law was too complex to be effectively simulated as a decision-process in the computer.[73] The machine could handle only straight-forward, clear, obvious cases. The value of the book lies in its

---

70. *See* Bruce Weber, *Computer Defeats Kasparov, Stunning the Chess Experts,* N.Y. TIMES, May 5, 1997, at B2.

71. EILEEN CORNELL WAY, KNOWLEDGE REPRESENTATION AND METAPHOR 8-9 (1991).

72. *See* ANNE VON DER LIETH GARDNER, AN ARTIFICIAL INTELLIGENCE APPROACH TO LEGAL REASONING 24-26 (1987).

73. *See id.*

demonstration of the difficulties of representing the process of legal reasoning in case-based law even for a single legal issue. Kevin Ashley, in his book entitled *Modeling Legal Argument,* describes HYPO, his case-based reasoner.[74] HYPO, however, handles only a relatively small number of cases in an artificial situation.[75] Ashley's pioneering work in case-based reasoning reflects just how difficult it is for the computer to handle issues of relevancy in the simplest of situations.

The reality of legal reasoning is far too complex to describe and comprehensively represent. The best that one can accomplish in the field of artificial intelligence and legal reasoning is to create models which can simulate legal reasoning in a small and limited area. Any model of a knowledge-based system of the law will be subject to the limitations of selection which are required because of the complexity of the law and of the legal reasoning process as carried out by human beings.

VI. How Human Intelligence Functions

The mathematician Zdzislaw Melzak refers to a measurement made some time ago by certain scientists of a change of dimension of forty microns, about the width of a human hair.[76] What he points out as interesting was not the minuteness of the change, as much finer measurements had been made much earlier, but that the measurement was of a change in the diameter of a neutron star at a distance of several thousand light years away in the Crab nebula.[77] This certainly furnishes us with a magnificent product of human intelligence. He further points out, however, that the method used was based on principles developed by Thales over twenty five hundred years ago to measure the height of an Egyptian pyramid.[78] It is a difficult task to directly measure the height of a pyramid because one would have to drill a hole from its top and drop a measuring line or rope down to its base. Thales merely used his similarity theorem which provides that triangles with the same angles are similar.[79] According to Melzak, Thales managed to avoid the difficult method of measurement by:

---

74. *See* K.D. Ashley, Modeling Legal Argument: Reasoning with Cases and Hypotheticals (1990).

75. *See id.*

76. *See* Zdzislaw A. Melzak, *Thales Revisited,* 83 Eos 203 (1995).

77. *See id.* at 205.

78. *See id.* at 204-05.

79. *See id.* at 204.

the use of his theorem which employs the similarity structure and allows the crucial introduction of proportion: the unknown height H of the pyramid is to the length of its known shadow S as the known height h of a vertical stick is to the stick's known shadow-length s. Thus the unknown height H is obtained from the simple proportion H:S = h:s.[80]

One of the more striking manifestations of intelligence is the capacity of humans to carry out a difficult task by creatively transforming it into the form of a number of simpler tasks, and then reconstructing it in a more complex form in order to achieve a desired objective.[81] Melzak calls this process the *bypass principle* which "is a way of dealing with complexity or with difficulty by means of a bypass which promotes a transport or a passage or the solution of a problem in a three-stage reduction process whose first and last stages are each other's inverses."[82] Speech is transformed by being represented by letters which can then be easily carried far beyond the range of the human voice. Words can be represented in the form of clicks of a telegraph key, which in turn become electrical impulses, which after passing through a wire are turned back into clicks, which are interpreted as words. Many of the more complex actions that best reflect human intelligence follow a pattern whereby a very difficult task is broken down into or transformed into a representation in a much simpler form.[83] In this form, a set of easy tasks can then be carried out. When, through a process of inverse transformation, the representation is returned to its original form, the original task is completed. Through the use of our intelligence, we accomplish the original hard task, and thus a complicated process makes a hard task easy. This gives us the formula:

> *DIFFICULT TASK to SIMPLIFIED TRANSFORMATION to a set of EASIER TASKS to an inverse SIMPLIFIED TRANSFORMA- TION to the COMPLETED TASK.*

Before writing existed, passing down the history and culture of a group was a difficult task. In each generation, people had to acquire the information from the elders, memorize it, and pass it on to the next generation. With the evolution of writing, the concepts and corresponding sounds of the language were encoded into a visual symbolism, recorded, and then decoded in the process of reading. When language first evolved, it existed first as an oral tradition. Representa-

---

80. *Id.* at 205.
81. *See* MELZAK, *supra* note 9, at 3-4.
82. *Id.*
83. *See id.* at 4 (discussing Melzak's bypass principle).

tions of memory could be passed on from generation to generation. With the evolution of writing, visual symbols could be used to represent the systems of sounds. These symbols could be given a more lasting form by inscription in material ranging from stone to papyrus to paper. The literate tradition took a revolutionary form with the development of the printing press. With the more recent development of digital technology, we enter a new era in the symbolic representation of human thought.

VIII. CAN COMPUTERS RESOLVE LEGAL ISSUES?

Expert systems are the closest approximation to human reason that have yet been devised in the field of artificial intelligence. Expert systems are computer programs based on various artificial intelligence techniques, which perform generally specific and difficult tasks at the level of a human expert.[84] Expert systems often are constituted by three parts: a knowledge base, an inference engine which manipulates the knowledge base, and explanatory and informational material supplied to the user.[85] Expert systems can be constructed in special programming languages or they can be created through the use of commercially available shells.[86] They can be written in either a logic-based or an object-oriented language, and they can function as a rule-based or case-based reasoner. If computers can be made to resolve legal issues, expert systems and case-based reasoning technology are likely to furnish us with the best tools.

Legal knowledge is technical, yet it does not have the conceptual clarity of the natural sciences. Lawyers think axiomatically,[87] yet the law cannot be put into the form of a deductive system. Law is a systematic body of knowledge, yet at the same time it is not necessarily consistent. Law is said to be rational, but not all legal arguments can be put into the form of logical syllogisms. The complexity of legal discourse is directly related to the nature and variety of the functions which it must serve, ranging from shaping the future to correcting the past, and to the complexity of the relationship between the doctrinal discourse of the law and the material facts of life to which that discourse is applied.

---

84. *See, e.g.*, DONALD A. WATERMAN, A GUIDE TO EXPERT SYSTEMS 1-60 (1986).

85. *See* Richard Forsyth, *The Architecture of Expert Systems*, in EXPERT SYSTEMS 9 (Richard Forsyth ed., 1984).

86. *See* WATERMAN, *supra* note 84, at 339-65.

87. *See* Juan Miquel, *Stoic Logic and Roman Jurisprudence*, in THE WESTERN IDEA OF LAW 333 (J.C. Smith & David N. Weisstub eds., 1983).

Legal reasoning is universalizable in that relevantly like cases are to be decided alike. When one case is like another, there will also be differences, but the differences will not be considered to be relevant. Legal reasoning, therefore, has a metaphorical structure in that when a term is used as a metaphor in a different context, the new reference must have similarities but also differences. Metaphor is not simply a peripheral device, but is central to language in that human cognition moves from the known and the familiar to the new and the different. Case-based reasoning, precedent, and reasoning by analogy are basically metaphorical by nature because one is moving from the known and familiar to the new, and somewhat different. Metaphorical reasoning requires criteria of relevancy in order to identify the differences which are not relevant, and the similarities which are. Relevancy is measured in terms of what is significant or important as affecting an outcome or a result.

Case-based reasoning has low information costs. If an actor can follow a prior procedure to get the same desired results, the individual need not understand the workings of the process, nor obtain new information in order to obtain the desired ends. Following a precedent justified in doctrinal terms is a very economical way of reasoning because it avoids the necessity of an extensive study of the impact or potential results of the alternative choices. Precedential reasoning requires little external information beyond the specific facts of the case. This method of argumentation attempts to find the best match as between the facts of a case and the facts of a precedent within the confines of doctrinal structures that require no external factual justification.

If each litigant in each case had to directly justify their claim in terms of economic, moral, or public policy considerations, the information costs would be considerable. Argument by precedent, however, avoids just this very type of argument which, although highly relevant, would add tremendous costs to the litigation process. Case-based reasoning entails the assumption that if a new dispute is similar to a correctly decided case in all relevant respects, then the resolution of the new dispute will be right and just if the same resolution is adopted for the new case. One significant characteristic of case-based reasoning in law is that the criteria of relevancy are rarely fully articulated. One might well infer from this feature that legal reasoning is based on widely shared assumptions about what is important, and that the information costs of arguing these assumptions in each case would outweigh any benefit to be gained.

A.   Hard Cases

If the relationship between legal doctrine and the material facts of life were fixed, or strictly rule-governed, then the task of simulating legal reasoning in the machine would be feasible, albeit complex. Unfortunately, however, that is not the case. The relationship of legal doctrine to facts in the context of litigation is such that it appears to many that the judge can exercise a discretionary choice as between several possible forms of doctrinal classifications or conceptualizations, and the doctrine is often bifurcated between opposing sets of concepts, rules or norms, and conflicting policies.

On the other hand, most members of the legal profession can tell the difference between good and bad legal arguments and forms of legal reasoning. Only a small percentage of disputes are litigated, and lawyers generally know when and when not to settle when a legal point is in issue. We can safely assume, therefore, that the relationship between legal doctrine and the material facts of life is not simply random. And if not random, then it must follow some kind of structure. While it probably is the case that there is no such thing as one and only one right answer to every legal issue or dispute, we can probably assume that of any two possible alternative answers leading to different legal results, one generally can justify one alternative as a better answer than the other. If all the above is true, while the courts are generally faced with hard cases, we can assume that they will seldom be faced with impossible cases to decide. Hard cases, therefore, are an inevitable aspect of law.[88]

There are many sources of indeterminacy in legal reasoning, and several different methods might be used to describe and categorize them. The nature of the indeterminacies will probably differ from one field of law to another while other patterns of indeterminacy may prevail throughout all parts of the law. What is common to all, however, is that the greater the degree of indeterminacy, the less rule-governed the area of law will be, and consequently the less will be the degree of predictability of outcome for any set of facts raising a legal issue, and the more difficult it will be to program a machine to resolve such a legal issue.

---

88. *See* S.C. Coval & J.C. Smith, *Rights, Goals, and Hard Cases*, 1 LAW & PHIL. 451 (1982).

B. *Deep Structure*

While in many areas of the law there is a significant amount of indeterminacy within legal doctrine itself, the substantial area of indeterminacy is found at the point where the legal doctrine is associated with the material facts of life which gave rise to the dispute. Traditional legal theories can all be conceived in terms of their presuppositions regarding language and how legal discourse relates to the material facts of life. Natural law and rights theories, for instance, presuppose that the relationship between legal discourse and the discourse surrounding the material facts of life, both of the human and of nature, is a fixed one.[89] A law is unjust if the legal discourse of rights, duties, and liberties is applied to the material facts of life in a way contrary to the true nature of humans and their natural environment. Legal Positivism, on the other hand, views the relationship as fixed by the conventions which determine political-legal jurisdictional validity, and where the rules on this are not clear, the judge has a discretion.[90] American Legal Realism, on the other hand, assumes that the relationship is more random, leaving judges with a significant amount of discretion as to what legal doctrines they use to conceptualize a set of facts.[91]

The University of British Columbia ("UBC") Faculty of Law Artificial Intelligence Research ("FLAIR") Project had its origins in a three year cooperative project between UBC and IBM (Canada). FLAIR, from the beginning, pursued a theoretical vision, having as its objective the integration of a particular paradigm of artificial intelligence with psycholinguistics and legal theory. Our first project was to test the validity of a particular theory of the nature and structure of the relationship between legal doctrine and the material facts of life, which we have characterized as teleo-analytic jurisprudence. Teleo-analytic jurisprudence assumes that the relationship between the technical doctrinal discourse of the law and the material facts of life is teleological, and that the validity of a legal judgment can be measured in terms of the consistency of the ordering of goals with that already embodied in the law.[92] This ordering, in turn, is evaluated as good or just according to how well it reflects the shared ordering of goals of the particular society.

89. *See* LEGAL OBLIGATION, *supra* note 61, at 4-7.
90. *See id.* at 7-13.
91. *See id.* at 13-21.
92. *See id. See generally* COVAL & SMITH, *supra* note 60.

Our first project was the development of a methodology for simulating case-based reasoning in the computer through representing the ordering and relationships of the underlying teleological or goal structures of the law. We were able to construct successfully a number of expert systems which can give an opinion on a difficult legal issue in the particular areas of law covered by the system. The legal knowledge contained within the system is represented by taking legal doctrine and doctrinal rules, and through the process of simplified transformation, representing them in terms of the goals and their ordering which underlie the law. The goals and their ordering are, through further simplified transformation, simplified into multiple sets of single more precise individual factual rules which reflect the underlying teleological structure which gave rise to the legal ordering. The user is asked a series of questions which require only answers of a factual nature relating to the specific problem, after which the system responds with a predicted legal outcome.

In order to create case-based reasoners and expert systems which can resolve disputes, the patterns of consistency by which legal doctrine (chains of signification) are associated with the material facts of life (the signified) must be identified and represented in the system. The doctrinal rules of law and legal concepts cannot be used as the knowledge base of the system. At the purely doctrinal level, legal rules are circular because the relationship between the concepts is syntactical and formal. There can be no liability in the law of negligence, for example, unless there is a duty of care. The existence of a duty of care, however, is always considered to be a question of law. If an interest in land gives the right to exclusive possession, then it is a property rather than a contractual interest; and if an occupier has a property interest, then they will have the right to exclusive possession. If, on the other hand, legal rules are expressed in terms of both facts and law, the number of rules and exceptions would be totally unmanageable. It would be impossible, for example, to have a set of rules which would be decisive for all of the factual situations in which the law would recognize or not recognize a duty of care. One can state a doctrinal rule that intent is a necessary element for a killing to constitute first degree murder. One cannot, however, derive a set of rules which can decisively determine for each case of a homicide whether or not there was an intent. On the one hand, if a legal rule is purely doctrinal, then any justification it furnishes to a factual situation will be circular; on the other hand, if the justification for the decision is factual, it will not be rule-like or rule-governed. Applying our para-

digm of artificial intelligence, we have sought to represent cases in terms of a simplified transformation which can be represented in the machine and handled in a computational manner. Where it is possible to recognize the patterns of consistency which constitute the associative relationship between the legal doctrinal structures and the material facts of a case, these patterns can function as a deep structure which can be represented in the knowledge base of an expert system or case-based reasoner.

1.   Remoteness of damages

The law of remoteness of damages in negligence is an example of one of the most indeterminate, ambiguous, and bifurcated areas of legal doctrine.[93] The issue of remoteness of damages relates to the limits of liability for damages resulting from a negligent act. The problem of remoteness of damages arises from the fact that a wide variety of consequential damages and losses can follow from an act of negligence in addition to those readily foreseeable. The harmful consequences of a negligent act are often multiple rather than single. One harmful consequence can often lead to another. As a result of an accident caused by the defendant's negligence, the plaintiff might require medical treatment which may produce further harmful consequences. Physical injury may cause further economic loss. Each particular harm must be dealt with separately in ascertaining whether, for purposes of recovery of damages, it is or is not too remote.

Because there has always been two sets of conflicting tests, one being defendant-oriented, favoring a narrow range of recovery, and the other being plaintiff-oriented, favoring a wider range of liability for causally related damage,[94] it would appear that judges have a complete discretion as to where they draw the line of recovery down the chain of events which follow a negligent act, beyond which a particular loss or injury will be too remote. If the judge wishes to find the damage too remote then she will use a strict foreseeability test.[95] On the other hand, if she wishes to give recovery for the damage, she will use a causation test such as "proximate" or "direct" cause.[96] The former favors freedom of action by placing narrower limits on the range

93. *See* J.C. SMITH, LIABILITY IN NEGLIGENCE 91-175 (1984) (discussing why ascertaining a rational basis for prediction of judicial decisions in remoteness cases is problematic).
94. *See id.* at 93-100.
95. *See id.* at 111.
96. *See id.*

of risk, while the latter favors compensation for loss by extending the boundaries of liability.

The point of balance between these two conflicting goals of the law can be stated in the form of a deep structure rule:

> Damages resulting from a negligent action are not too remote if they are reasonably foreseeable in the particular, or, if not, are one of a foreseeable class of damages.[97]

Although often a particular damage may not be reasonably foreseeable, it is reasonably foreseeable: 1) that dangerous activities when carried out negligently create a wide variety of particular kinds of risks of harm; 2) that injury to persons can also result in further damage from particular susceptibilities or medical complications to the injured person or to others; and 3) that the creation of a risk invites rescue.[98] From the ultra-hazardous category it follows: 1) that no physical injury or property damage caused by a motor vehicle accident, explosives, highly inflammable or toxic substances, fire, high-voltage electricity, weapons, or dangerous machinery will be too remote; 2) that no increased physical or emotional injury resulting from an unusual or particular susceptibility of a person suffering damages as a result of a negligent act will be too remote; and 3) no medical complication resulting from an injury to a person is too remote. No physical or emotional injury suffered by a rescuer will be too remote.[99] While any specific instance of any one of the above may not have been reasonably foreseeable, they are all members of a class of the kind of damages which often do occur.[100] An examination of 334 reported cases in Canada, Great Britain, Australia, and New Zealand raising an issue of remoteness shows that approximately 92% of these cases were decided consistently with the above underlying deep structure rule.[101] Thus, one of the areas of law most often considered to be the least rule governed actually turns out to be highly structured when examined from the perspective of the underlying teleology of that area of the law. Remoteness issues reveal a conflict between the goals of deterring negligent behavior and compensation for harm caused, and freedom of action. Strict liability would inhibit action, and no liability would encourage reckless conduct. Analysis of the underlying teleological structures of the law (teleoanalytic jurisprudence) reveals

---

97. *See id.* at 138, 159-60.
98. *See id.* at 136-37, 160.
99. *See id.* at 160.
100. *See id.* at 132-33.
101. *See id.* at 138-40.

a high degree of consistency in the way in which legal doctrine is associated with the discourse of the material facts of life, which traditional legal theories have failed to recognize.

2. The Nervous Shock Advisor

The Nervous Shock Advisor ("NSA"), developed at FLAIR, is an expert system which will give advice as to whether nervous shock caused by the negligence of another is recoverable in tort law, or is precluded from recovery on the grounds that the shock to the nervous system is too remote.[102] The system will advise, in other words, how far down the chain of strict causality the law will go in compensating victims of negligently induced nervous shock. Nervous shock is a term coined by lawyers within the common law tradition. Its closest medical counterpart is posttraumatic stress disorder. Most U.S. jurisdictions have a similar cause of action called "emotional distress." In law, the classic case is the situation where the claimant is a bystander who witnesses something terrible happening to someone else and suffers shock as a result, such as where a mother sees her child killed by a negligent driver. The NSA contained every case on this legal issue from Canada, Great Britain, Australia, and New Zealand.

In the common law systems of the British Commonwealth, the plaintiff can recover for posttraumatic stress disorder if they are the party injured or endangered. Where, however, the posttraumatic stress disorder is suffered by someone other than the person at risk of physical injury, the law of nervous shock will be found to be one of the most difficult, obscure, and confusing areas of legal doctrine that one could possibly find. In the first place, there is no general consensus in the common law systems of the British Commonwealth as to whether nervous shock is an issue of remoteness in the law of negligence or a substantive tort in its own right. In spite of the doctrinal ambiguity surrounding this area of the law, an examination of the cases reveal that there are two necessary conditions for recovery for posttraumatic stress disorder where the posttraumatic stress disorder is suffered by someone other than the person at risk of physical injury. There must be a certain degree of relationship between the parties, and a certain degree of exposure to the traumatic event. Relationship is a fuzzy class which can be "defuzzified" by giving a precise numerical value to

---

102. *See* Cal Deedman & J.C. Smith, *The Nervous Shock Advisor: A Legal Expert System in Case-Based Law, in* OPERATIONAL EXPERT SYSTEM APPLICATIONS IN CANADA 56 (Ching Y. Suen & Rajjan Shiṅghal eds., 1991) [hereinafter Deedman, *The Nervous Shock Advisor*].

each different situation on a scale ranging from close to distant. In a similar manner the degree of exposure can be broken up into a series of specific situations where each member of the set can be given a precise numerical value. The reasons for judgment in all of these cases, however, are given in purely doctrinal form. Legal justification is doctrinal and not factual.

Figure 7, a diagram of the knowledge structure of the Nervous Shock Advisor, reflects an entirely factual structure, which is related to the doctrinal structure through patterns of consistency in legal decisions which reflect the underlying teleology of this area of the law.

FIGURE 7
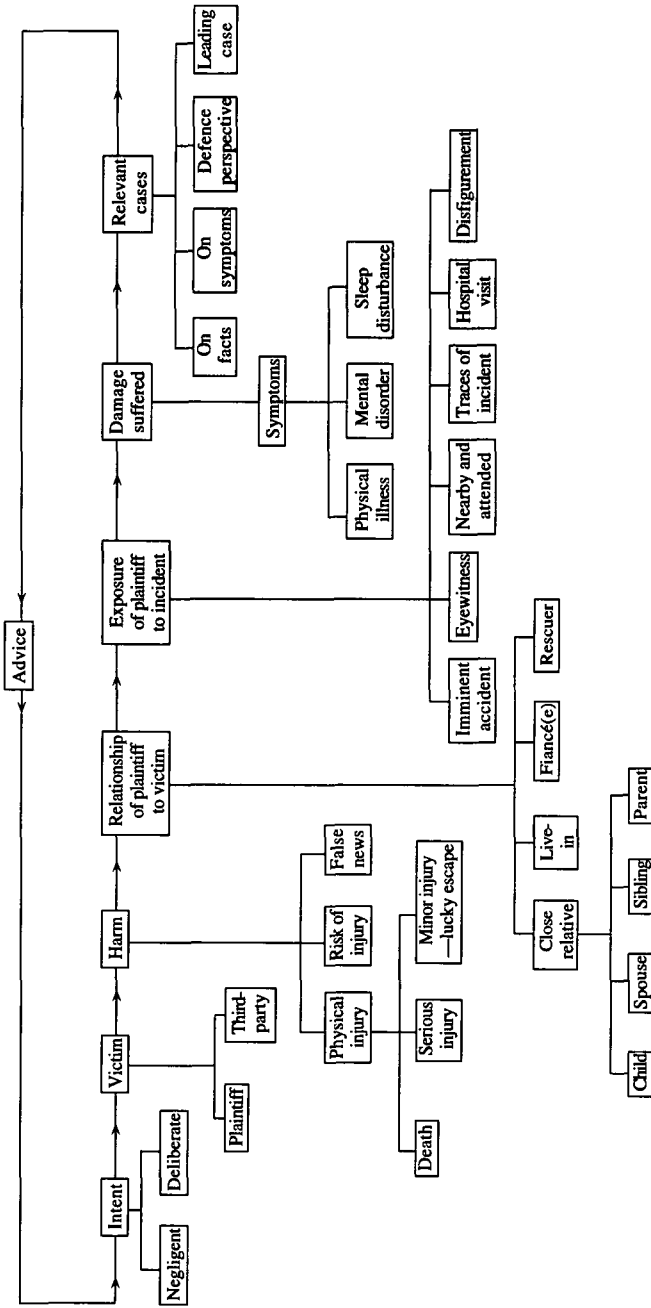
The following hypothetical set of facts can be used to demonstrate how the system will respond to the answers of a set of factual questions by rendering a legal judgment:

The plaintiff is the mother of a six year old child. As a result of the defendant's negligence, the plaintiff's young daughter suffered extensive burns about the face. The accident took place at school, and the child was rushed immediately to the hospital. When the plaintiff arrived there, she found her daughter resting under sedation, with her face entirely covered with bandages. After several weeks of hospitalization, the bandages were removed, and the plaintiff saw her daughter's badly scarred face for the first time, and as a result suffered posttraumatic stress disorder or nervous shock. As a result of the posttraumatic stress disorder the mother is unable to sleep at nights, and consequently has had to give up a high paying job. The defendant admits liability to the child, and the amount of damages has been agreed upon by the parties through the negotiations of their lawyers, but takes the position that there is insufficient proximity of exposure to the accident to create a duty of care so far as the mother is concerned.

The system will ask the user a series of factual questions to form a profile of the problem in terms of the deep structure of the system, whereupon it will give the following response:

I am pleased to inform you that your client has a cause of action for nervous shock
The following is a breakdown of the various elements involved in this decision. If any element has a certainty factor of less than 100%, you should regard it as a weak link in the chain and focus on bolstering your case in that area.
Plaintiff exposure to the incident: 60%.
Judgment: An action lies 60% percent.

The cases are returned in terms of three categories: *on point, relevant by analogy,* and *contra.* The system returns the case of *S. v. Distillers Co. (Biochemicals), Ltd.,*[103] where a mother recovered for nervous shock sustained upon giving birth to an infant who suffered severe deformities during embryonic development due to taking of the drug thalidomide by the mother during her pregnancy. As well as the closest *contra* case, the system makes available a list of cases where there was insufficient exposure to the incident to justify recovery. The system also reports a list of the leading cases and a list of cases relating to the symptoms that the plaintiff might have suffered, along with a summary of the facts and the express statements of the judges relating to the symptoms.

103. 3 All E.R. 1412 (Q.B. 1969) (U.K.).

3.   The limits of expert systems

As a result of a decade of research and development at the UBC FLAIR Project, we feel we can say with some confidence that the answer to the question of whether machines can be made to think like lawyers is a very qualified *"Yes."* Nevertheless, however, very few true expert systems exist in the field of law. There are a number of reasons why this is the case. Expert systems in law are extremely time consuming, and therefore expensive to create and maintain. Lawyers are reluctant to trust the decision of a machine, unless they fully understand the rule structure by which the decision was reached, and if they did understand that, they would not need the expert system query. Commercial programming shells, which facilitate the creation of expert systems, have not proven to be reliable, in that they come and go from the market. If they are not maintained by the software company, the system cannot be maintained, but must be re-created in a different shell, or written in a suitable programming language.

The purpose of creating expert systems in law at the FLAIR Project was to learn more about the nature and structure of legal reasoning by attempting to simulate it in the machine. From this perspective, we consider our research, whereby we have developed expert systems through the representation of the teleologically deep structure which underlies legal doctrinal rules in their application to factual situations, to have been highly successful. We have found, as have many others, that the simulation of legal reasoning in expert systems can be a highly effective way to teach the nature and structure of legal reasoning. Every year for the past six years in the legal theory course, *Legal Reasoning and Artificial Intelligence,* which we offer at the UBC Faculty of Law, students have used the methodology in conjunction with the commercial shell, VP Expert, to create small, but interesting and effective expert systems in restricted domains.

The expert systems developed at the FLAIR project represent the relationship between legal doctrine and factual situations in terms of the underlying patterns of consistency in the ordering of interests and goals when they come into conflict.[104] Thus the underlying rule

---

104. *See* J.C. Smith & Cal Deedman, *The Application of Expert System Technology to Case-Based Law, in* THE PROCEEDINGS OF THE FIRST INTERNATIONAL CONFERENCE OF ARTIFICIAL INTELLIGENCE & LAW 84, 85 (1987); *see also* Deedman, *The Nervous Shock Advisor, supra* note 102; Andrzej Kowalski, *Case-Based Reasoning and the Deep Structure Approach to Knowledge Representation, in* THE PROCEEDINGS OF THE THIRD INTERNATIONAL CONFERENCE OF ARTIFICIAL INTELLIGENCE & LAW 21 (1991); Marilyn T. MacCrimmon, *Expert Systems in Case-Based Law: The Hearsay Rule Advisor, in* THE PROCEEDINGS OF THE SECOND INTERNATIONAL CONFERENCE OF ARTIFICIAL INTELLIGENCE & LAW 68 (1989); J.C. Smith, *Action Theory and Legal*

structure of these expert systems represent the relationship between legal doctrine and the material world of facts to which they are associated by judicial decisions, rather than representing the legal doctrine itself. Following our "creative transformation" paradigm we take an area of law where the cases appear to be decided inconsistently in terms of the application of legal doctrine, and analyze them in terms of the underlying conflicting interests and the likely effect which alternative orderings of those interests will have on the ordering of other interests and goals reflected in the law. Where we can find patterns of consistency of ordering in terms of the underlying teleology reflected in the law (a teleological deep structure) we will then try and create rules in terms of factual components which reflect these patterns of consistency. Thus cases are reduced to goal orderings (the simplified transformation). Factual questions are asked of the user in order for the system to formulate a representation of the problem in terms of a conflict of goals or interests. The system then seeks to find the closest corresponding conflicts and orderings in its knowledge base, and match it to that of the problem (the set of easier tasks). The system then, on the bases of this match, suggests a solution to the issue that is consistent with the teleological structures of the cases in its knowledge-base (the inverse transformation) legal doctrine.

C.   *Putting the Judge into the Machine*

Let us return to the question I used to initiate this examination of the two alternative views of the nature of artificial or machine intelligence, *Can/Should Computers Replace Judges?*, and see what would be entailed in constructing a judicial machine that could actually adjudicate cases. This comparison will permit an evaluation of the Church-Turing view of AI with that of AI as creative transformation.

1.   The Cycic Judge

The most comprehensive artificial intelligence project thus far which is based on the assumptions of the Church-Turing thesis is Cyc. Cyc is a very large, multi-contextual knowledge base and inference engine started in 1984 at the Microelectronics and Computer Technology Corporation ("MCC") in Austin, Texas and transferred to Cycorp,

---

Reasoning, in TORT THEORY 104 (Ken Cooper & Elaine Gibson eds., 1993); Cal Deedman, Developing Conceptual Systems for Structuring Legal Knowledge to Build Knowledge-Based Systems (1993) (unpublished Ph.D. thesis, University of British Columbia) (on file with author).

Inc. as a spin-off company from MCC.[105] Consistent with the Church-
Turing thesis, Cyc is a massive knowledge base containing an enor-
mous amount of fundamental human knowledge in the form of facts,
rules, and heuristics for reasoning about the material facts of human
existence and everyday life. It now contains "more than 2 million bits
of assertions."[106]

Assuming the Church-Turing thesis, a computer-based system
that could actually replace judges would have to be constructed on top
of Cyc with the addition of a full knowledge of case law and statutes
and regulations. We would thus have the Cycic Judge. This would be
one way to attempt to construct a computer that could resolve legal
issues. Typical of the statements of the critics of the Cyc Project is the
opinion that "[w]e believe this is a mistaken approach because it
wastes precious human resources and makes dubious theoretical
assumptions . . . ."[107]

2. A more modest proposal

Judges, lawyers, and academics can be conceived of as function-
ing within the context of a vast information system that reflects a com-
plex matrix of goals and their orderings.[108] Each case decided in favor
of a plaintiff or a defendant resolves a conflict of interest by hierarchi-
cally ordering the goals pitted against each other in the dispute. The
doctrine of precedent produces trends of consistency, with corrections
taking place where different circumstances result in conflicting sets of
ordering. Legal doctrines are constantly being restated, refined, and
modified. Legislation makes more prominent and massive corrections
to the goal structure, which in turn will filter their way through the
individual cases. When the practice of following precedent is applied,
and there are no express contradictions in the legal doctrine, legal de-
cisions will be consistent. If one understands the teleological structure
that is isomorphic for the doctrine and the decisions, one will be able
to maintain a fairly high degree of predictability as to how like cases
will be decided in the future, or the probable outcome of a dispute,
given a particular set of facts.

---

105. For online information on the Cyc Project, see *Cycorp: Company Overview* (visited Feb.
24, 1998) <http://www.cyc.com/overview.html>.

106. Garfinkel, *supra* note 1, at 187.

107. YORICK A. WILKS ET AL., ELECTRIC WORDS: DICTIONARIES, COMPUTERS, AND MEAN-
INGS 7 (1996).

108. *See* COVAL & SMITH, *supra* note 60; *see also* Coval & Smith, *supra* note 88.

One case is relevantly similar to another to the degree that it shares the same teleological structure. A case shares the same teleological structure when the same goals or interests are at stake, with similar causal relationships to other sets of like goals or interests. An examination of the law will show that the decisions of the courts and the effects of legislation result in a fairly consistent ordering of our values. The prevention of physical harm, for example, is ranked higher and more important than the prevention of economic loss. The interest of employees to carry out effective collective bargaining generally takes priority over private property rights, if the negative impact is minimal. From the fact that truth is a complete defense in the law of defamation, we can infer that the dissemination of valid information is more important than, and thus takes priority over, the protection of reputation.

We have found that we can represent the goal matrix of the law in the form of a few basic deep structure rules for establishing priority. In particular, an examination of the goal matrix of the law will show that whenever one ordering of a pair of conflicting goals will maximize only the first goal at the extreme expense of the second, while the converse ordering will maximize the second goal and produce only a minimal interference with the first, the law will generally prefer the second ordering. If one particular ordering of goals as a result of a legal decision will produce a loss of a more highly valued social goal, while the opposite ordering will not, the law will tend to favor an ordering which will not have a negative impact on other goals within the law.

An alternative to the methodology of the Cyc Project would be to use the creative transformation methodology to represent the entire body of case and statute law in terms of goals, their orderings, and their cause-effect relationships on each other in terms of alternative hierarchical priorities. The logic of the system would be based on a set of deep structure teleological rules such as the following:

Rule One

If, when goal 1 is ordered over goal 2, the result will be a small encroachment ($\sim$) on goal 2, and if, when goal 2 is ordered over goal 1, the result will be a substantial encroachment ($\sim$) on goal 1, then goal 1 should be given priority over goal 2.

$$\text{If } \frac{g^1}{g^2} \rightarrow g^1 \ \& \ \sim g^2 \text{ and if } \frac{g^2}{g^1} \rightarrow g^2 \ \& \ \sim g^1 \text{ then } \frac{g^1}{g^2}$$

Rule Two

If, when goal 1 is ordered over goal 2, the degree of encroachment on goal 2 is the same as the degree of encroachment on goal 1 when goal 2 is ordered over goal 1, and if goal 1 is consistently ordered over (is more important than) goal 2, then goal 1 should be given priority over goal 2.

$$\text{If } \frac{g^1}{g^2} \rightarrow g^1 \ \& \ \sim g^2 \text{ and if } \frac{g^2}{g^1} \rightarrow g^2 \ \& \ \sim g^1 \text{ and if } g^1 > g^2 \text{ then } \frac{g^1}{g^2}$$

Rule Three

If, when goal 1 is ordered over goal 2, the result will be a small encroachment on goal 3, and if, when goal 2 is ordered over goal 1, the result will be a substantial encroachment on goal 3, then goal 1 should be given priority over goal 2.

$$\text{If } \frac{g^1}{g^2} \rightarrow g^1 \ \& \ \sim g^3 \text{ and if } \frac{g^2}{g^1} \rightarrow g^2 \ \& \ \sim g^3 \text{ then } \frac{g^1}{g^2}$$

Rule Four

If, when goal 1 is ordered over goal 2, the result will be a small encroachment on goal 3, and if, when goal 2 is ordered over goal 1, the result will be a substantial encroachment on goal 4, then goal 1 should be given priority over goal 2.

$$\text{If } \frac{g^1}{g^2} \rightarrow g^1 \ \& \ \sim g^3 \text{ and if } \frac{g^2}{g^1} \rightarrow g^2 \ \& \ \sim g^4 \text{ then } \frac{g^1}{g^2}$$

Rule Five

If, when goal 1 is ordered over goal 2, the degree of encroachment on goal 3 is the same as the degree of encroachment on goal 4 when goal 2 is ordered over goal 1, and if goal 3 is consistently ordered over (is more important than) goal 4, then goal 1 should be given priority over goal 2.

$$\text{If } \frac{g^1}{g^2} \rightarrow g^1 \ \& \ \sim g^3 \text{ and if } \frac{g^2}{g^1} \rightarrow g^2 \ \& \ \sim g^4 \text{ and if } g^3 > g^4 \text{ then } \frac{g^1}{g^2}$$

Rule Six (transitivity)

If goal 1 is consistently ordered over (more important than) goal 2 and goal 2 is consistently ordered over (more important than) goal 3, then goal 1 should be given priority over goal 3.

$$\text{If } g^1 > g^2 \text{ and } g^2 > g^3 \text{ then } g^1 > g^3$$

All of the above rules would be subject to an overriding *ceteris paribus* clause—all other things being equal. The competing goals themselves may not be on an equal status. Certain goal values such as those enshrined in a bill of rights take precedent over nearly all other goals, particularly those not given constitutional priority. The complexity of judicial reasoning is far beyond what is possible to represent in the computer. Nevertheless, if one were to attempt the process, the latter alternative would be easier to apply than the Cyc approach, and probably more effective.

IX. MACHINE INTELLIGENCE AS CREATIVE TRANSFORMATION

Computer scientists who approached artificial intelligence from the perspective of the Church-Turing paradigm generally assumed that "coded knowledge of the real world was required if machines were to understand language in any serious sense."[109] The vocabularies of such systems, however, remain primitive and extremely limited. Significant developments in the field of natural language processing generally have taken place where the researchers are working from alternative paradigms. Some of the most interesting and promising areas of artificial intelligence and language have come from research which combines psycholinguistics with computational linguistics. Psycholinguistics is a branch of linguistics that examines language in terms of how it functions in human thought and memory.

The idea or concept of the lexicon is the element where psycholinguistics and computational linguistics can be made to converge because the lexical component of language is common to both the human and the computer.[110] Most of the lexical properties of language can be represented in computer-based lexicons. As humans we are used to using lexicons or dictionaries which are ordered alphabetically. Alphabetical organization is a widely used computer lexicon

---

109. MELZAK, *supra* note 9, at 1.
110. *See, e.g.*, R. Beckwith et al., *WordNet: A Lexical Database Organized on Psycholinguistic Principles, in* LEXICAL ACQUISITION: EXPLOITING ONLINE RESOURCES TO BUILD A LEXICON (U. Zernik ed., 1991); GEORGE A. MILLER & PHILIP N. JOHNSON-LAIRD, LANGUAGE AND PERCEPTION (1976).

function. We are also familiar with ordering information hierarchically. Hierarchies are widely used in computational linguistic technology to organize large databases of information. Computer scientists use hierarchies to create inheritance systems whereby the subordinate can be assumed to have inherited or to possess the properties of the superordinate. Thus *dogs*, a subordinate category of *mammals,* can be assumed to possess the defining properties of mammals. Synonyms can be linked in computer-based lexicons. The use of lexicons permit us to create lexical matrices in the computer. The contents of books, documents, and databases can, through the process of creative transformation, be simplified in the form of lexicons which then can be processed to accomplish tasks which would be extremely difficult, or not even possible to accomplish in their original natural language form. The computer permits us to do a variety of creative transformations which can enable us to combine the best aspects of research in the traditional print libraries with the advantages of electronic databases, resulting in new forms of computer-based information systems.

A.   *The Lexicon as a Creative Transformation*

FIGURE 8

Private Law
 Tort Law
   Negligence
    Duty
      Non-feasance
        Occupiers Liability
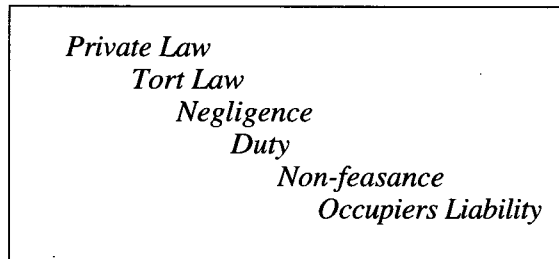
Obtaining information from a book by reading through the entire book would be a difficult task. One need not read through an entire book, however, in order to ascertain the nature of its content. Almost all books contain a table of contents which represents the subject matter of the book. Legal knowledge is brought to the lawyer structured from the general to the particular, from what is fundamental to what is derivative. If the lawyer was dealing with the legal issue of the liability of an occupier of land for a wrongful act such as an assault, committed by a third party on an invitee or licensee of the occupier, which might have been prevented if the occupier had taken some specific safety measure, the lawyer would organize her information need in

the form of a hierarchical string such as that shown in Figure 8. The lawyer would then attempt to find a corresponding match in the table of contents of a legal text.

Doctrinal hierarchy is not the only form in which legal knowledge is presented in the traditional library. Most books have an index as well as a table of contents. Legal concepts and relevant terms will also be found listed alphabetically. Without these two forms of content representation, and the corresponding page numbers which returns the reader from the representation to the actual content of the book itself, the retrieval of information from a book would be much more difficult. This form of organization is not only used in the alphabetical index, but also in research tools such as the legal encyclopedias and dictionaries. Books facilitate both top-down (from the general to the specific) and bottom-up (from the specific to the general) methods of legal research. Tables of contents and indexes are simplified transformations of the contents of the book which enable the reader to go quickly to the portions of the book likely to contain the information relevant to their information needs. The table of contents is a representation of the hierarchical semantic structure of the content of the book. An index is a representation of the content of the book in terms of the important words presented in the form of an alphabetical lexicon.

Another important way in which legal knowledge is presented to the lawyer in the traditional library is in the form of abstraction. An abstract gives a short summary of a legal document or case to help the researcher know what the document is about and whether or not it is likely to be relevant, and therefore whether more time is needed to examine it further or to read it. Abstraction seeks to give the fundamental content in as short as possible form. Lawyers use many forms of abstracts such as headnotes, summaries of the law, and various forms of condensed versions of cases. Abstractions are also creative transformations that make a difficult task easier to perform. Rather than reading through a large number of cases in full text, the full texts are represented by short abstracts which permit the lawyer to peruse quickly, select those which appear to be relevant, and inversely transform them by returning to the full text of the selected abstracts.

There are certain problems with traditional library sources of information that electronic databases do not face, the most important being the element of time. There are thousands of cases in any single jurisdiction, and the law of several jurisdictions may be highly relevant to a problem. Seldom can any lawyer afford to spend the kind of time

on a legal problem which would be required to discover all of the relevant cases. The computer, on the other hand, can go through thousands of documents in an electronic database in a matter of a few seconds. Furthermore, a law library takes up a great deal of room. A lawyer cannot put a library of books in her briefcase and take it home at night. An entire library in electronic form, however, can be carried in one's briefcase on a few CD-ROMs. Cases now are composed on computers with word processing programs. In electronic form they can be immediately made available on disks, through bulletin boards, over the Internet, or through a modem connected to the telephone. The cases exist in electronic form long before they appear in print. Therefore, we rely on electronic data to obtain the most recent decisions. The traditional library can never provide us with the most current decisions. Furthermore, because the cases start in electronic form, and carry no added value such as headnotes, legal information in electronic form is much cheaper than when it is released in print. A good deal of legal information, such as the most recent cases, can be obtained free in many jurisdictions over the Internet.[111]

B. Difficult Tasks

We at the University of British Columbia FLAIR Project considered our research in the area of expert systems to be highly successful. We nevertheless abandoned a fruitful direction of research to turn our resources and efforts in an entirely different direction. We did so because we had become fully aware of the problematic nature of the traditional AI paradigm which requires massive knowledge bases of encoded information and/or strict formal logical inference patterns. At the same time, we recognized that the success of our expert systems was based on a form of creative transformation whereby cases were represented in a simplified transformation in terms of goals and their ordering. When our research led us to adopt an alternative AI paradigm in the form of creative transformations, we sought to undertake a different approach whereby we could solve practical problems through creative transformations rather than by simulating legal reasoning.

111. *See, e.g.*, Cornell Law School, *Legal Information Institute* (last visited Feb. 24, 1998) <http://www.law.cornell.edu/>; Chicago-Kent College of Law, *Chicago-Kent Law Review Online Research* (last visited Feb. 24, 1998) <http://www.kentlaw.edu/lawrev/libraries.html>; Emory University School of Law, *Emory Law Library Federal Courts Finder* (last visited Feb. 24, 1998) <http://www.law.emory.edu/FEDCTS/>.

While it is now possible to store enormous amounts of reported decisions in electronic databases, the retrieval of relevant cases remains an extremely difficult task. The area of information retrieval presented us with the kind of challenge that would permit us to further explore the potential of our alternative AI paradigm using lexicons as the form of simplified transformation. When we started this change of direction back in 1989, lexicology was considered a backwater of natural language processing. Since then, it has become a major field of research, thus justifying our earlier intuitions.[112] With these thoughts in mind, we set out to develop a new approach to document retrieval in large domain-specific databases such as those we find in law. The method we chose was to represent the databases, and the documents which constituted them, in terms of lexicons organized alphabetically, hierarchically, and in terms of frequency.

1.   Making the invisible visible

Information in the form of books has significant advantages over information in electronic form. The content of a book is observable on its face through tables of contents, indexes, and summaries, and one can skim the contents by quickly paging through and glancing at each page. Moving through the information, however, takes a great deal of time. Information retrieval in an electronic database is a much more difficult task than the retrieval of information from books because the information in books is visible on the page, but the information in an electronic database is not visible until it is actually retrieved. Information must first be retrieved and then transformed into the visual symbols of the alphabet. The difficulty lies in that you cannot see the information until you locate it, and you cannot locate it unless you can see it. Books, nevertheless, can be put into an electronic form with tables of contents and indexes, which can easily be brought to the screen. If, however, one wanted merely to access a book, there is no significant advantage to be had in transferring the content of the book to an electronic database so far as information retrieval is concerned. In fact, much is lost, in that skimming through the book by quickly flipping the pages and looking for relevant concepts or significant phrases is not possible in an electronic database. Merely putting books into electronic databases does not take full advantage of the

---

112. *See, e.g.*, WILKS ET AL., *supra* note 107; COMPUTATIONAL APPROACHES TO THE LEXICON (B. T.S. Atkins & A. Zampolli eds., 1994); MACHINE TRACTABLE DICTIONARIES: DESIGN AND CONSTRUCTION (Cheng-Ming Guo ed., 1995); COMPUTATIONAL LEXICOGRAPHY FOR NATURAL LANGUAGE PROCESSING (Bran Boguraev & Ted Briscoe eds., 1989).

lightning speed with which information can be processed in an electronic database.

The process of information retrieval, whether from a book or an electronic database, follows the same pattern. The information needs of the researcher must be simplified in the form of a transformation which will represent the information need. This creative transformation, whether a member of a conceptual hierarchy or a list of words, is then matched with the corresponding representation of the content of the book or of the electronic database. When the relevant match is found, the simplified representation is transformed back into the relevant text.

Information retrieval becomes more efficient in databases rather than books because of the huge amount of information which can be inputted and stored in databases, and the speed with which that information can be accessed, manipulated, and retrieved. Yet tables of contents and indexes are manually created, and thus are expensive and time-consuming to create. Simple tables of contents and indexes such as those found in books have diminishing returns as the size of databases increase. The difficult task in doing research in electronic databases is to design an information system that will automatically generate hierarchical conceptual structures, alphabetical indexes for words and phrases, and sets of lists of significant factors out of the content of each document, in a domain specific database. These would permit the user to examine the content of the database and locate the words and concepts that match the information needs. The Flexicon technology consists of a set of libraries and tools which can be used to create these lexicons in specific domains such as law or medicine, or any other areas that lie within specific subject fields. Thus the invisible is made visible.
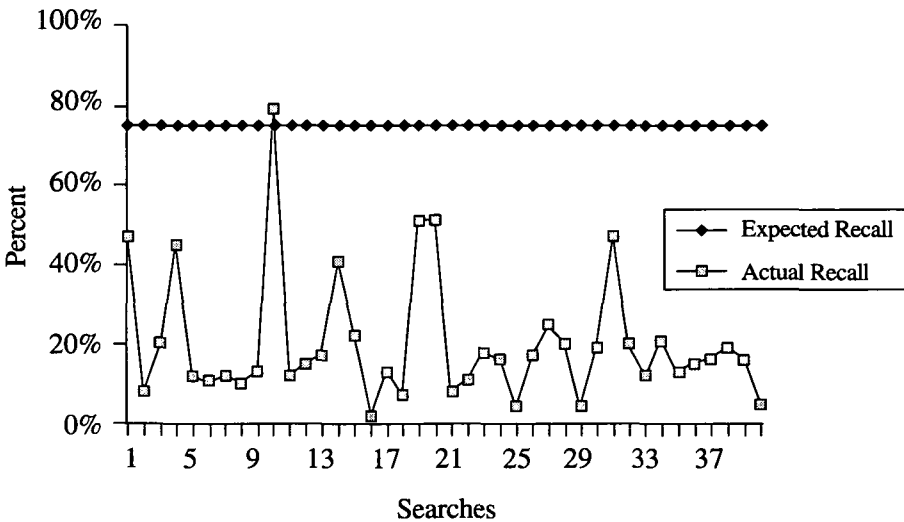
2. Retrieval

Nearly all databases, whether online or CD-ROM, either come with search engines attached, or are accessible through available search engines. Search engines are generally one of two different kinds: either exact-match or best-match. Exact match search engines are generally Boolean. A Boolean search is an exact-match engine in that a Boolean search engine will only return documents that exactly match the query, and the documents will be returned in no particular order. A Boolean search engine requires that the terms of the query be joined with Boolean logical operators; i.e., with any one or more of

AND, OR, or NOT. If AND is used, then the engine will retrieve only documents which contain every term so joined. Such queries generally return too little. If OR is used, then the search engine will return any and every document which contains any one or more of the so joined terms. Such queries generally return too much. If the Boolean operator NOT is used before a term, then the search engine will not return any document with that word in it, even if there are other matching terms of the query. Most such search engines permit proximity searches which enable an experienced researcher to form more complicated queries by stipulating that certain terms must be within a certain distance of each other.

In best-match search engines, the documents do not have to exactly match the query, but are returned in a ranked order according to their similarity with the query. A best-match search engine that permits one to form the query in the way one normally writes or speaks, is often referred to as a natural language search engine because it permits the use of natural language queries.

FIGURE 9

Expected Recall v. Actual Recall
Blair & Maron

One of the most difficult problems in information retrieval is the problem of knowing whether or not one has retrieved all of the relevant documents which are to be found in the database. The more common method of retrieval of information from an electronic

database is through the use of a Boolean search engine. One of the few studies made regarding this problem was the Blair & Maron study carried out on a large database of legal documents using a Boolean exact match search engine.[113] As can be seen in Figure 9, the Blair & Maron report concluded at the end of an extensive series of tests carried out by experienced lawyers on a legal database, that while the lawyers thought that they had retrieved approximately 75% of the relevant documents in the database, they in fact had only recovered approximately 20%.[114]

Most experienced legal researchers use a combination of both Boolean and best-match search engines because each have certain advantages and disadvantages, depending upon the nature of the information requirements. A Boolean search engine permits the use of multiple word groupings in the search query, but does not return the documents ranked as to how well they match the information need which the query represents. For example, if a search returns sixty documents, the most relevant case may be the very last. The lawyer must examine each of the sixty cases retrieved to discover which are relevant to the legal problem, and thus conform to the information need. Non-Boolean best-match search engines, on the other hand, return the documents with some semblance of relevance ranking but they do not permit the use of multiple word groupings and phrases in the search query. One is limited to using single words. This is a very serious limitation because a large number of legal concepts and query items consist of several words.

Because both the capacity to use multiple word combinations in the search query and the capacity to return cases ranked in terms of relevancy are essential features for efficient research in an electronic database, it would seem to make sense to combine the two kinds of search engines into one. This, however, is not possible because of the unique and fundamentally different nature of their structures. Legal publishers have added on to the two kinds of search engines, so far as their structures will permit, some limited functions of the other kind of engines. The commercial legal Boolean search engines offer some limited ranking in terms of an algorithm that statistically measures the frequency of the terms of the query as compared to the frequency of their occurrences in the database as a whole, and ranks the cases ac-

---

113. *See* David C. Blair & M. E. Maron, *An Evaluation of Retrieval Effectiveness for a Full Text Document Retrieval System*, 28 COMMUNICATIONS OF THE ACM 289 (1985).

114. *See id.* at 293.

cordingly. Best-match search engines can be given a limited multiple word facility by attaching a built-in lexicon of frequently used word groupings which permits the system to recognize and treat a phrase as if it were a single word. If a multiple word item is not in the attached lexicon, however, each word will be considered a separate item.

The Flexicon system is a lexically designed and structured search engine for domain specific vertical information markets such as law and medicine which combines both relevance ranking and multiple term queries, permitting a precision of recall far exceeding any existing search engine on the market. The Flexicon information system permits the user to form a much more detailed representation of the information need than is possible using Boolean or exact-match queries, by the use of the lexicons which reveal the content of the database. Thus, a profile of the kind of document that will fit the information need can be formulated which will be used to retrieve the best matches in the documents comprising the database. While a query normally will only contain anywhere from a few to a dozen or two dozen terms, a profile of an information need, or a document profile can contain a hundred or so terms, if appropriate. Each term in a Flexicon search profile can be given a weight of high, medium, low, or negative weight, and each of the different lexicons can be given different weights. This permits a high degree of precision in document retrieval, and one can, depending upon how far one has time to go down the list of returned documents, obtain a far higher percentage of relevant documents than is possible using the traditional forms of search queries.

3.   Automatic abstraction

The database in a Flexicon information system exists electronically in three forms. The first is the lineal text of each original document, the second is the entire database in the form of the lexicons, and the third is each document in the database individually in the form of the lexicons. Each document is retrieved in the form of a FlexNote which provides a profile of the case in terms of legal doctrinal concepts, facts, cited cases, and cited statutes which are presented to the user in the form of four lexicons which are each ordered according to the importance of the listed items. The ranking is based on term frequency, including the synonyms of legal concepts and the alternative forms of case and statute citations, normalized over the database. One can then glance at the profile of the document on the screen and quickly decide if the document is likely to be relevant. The FlexNote,

therefore, constitutes a machine-created abstract of the document. One can then hypertext from each item in the document profile to each of its occurrences in the document or go directly to the full text.

The objective of the Flexicon technology is to have the machine, as far as possible, process the raw data, and create added value such as FlexNote abstracts, automatic subject classification, machine recognition, and header or title formation. The aim of our research is not to seek to replace judges and lawyers, but to, as much as possible, replace by the machine the experts who give the raw data added value.

C. *Making Research Easy*

The Flexicon system takes an entire database of cases in electronic form and separates the text into five groups. It first eliminates the *noise words* such as *the, and, or, in, at*, etc., but it does not eliminate them altogether. Through parsing functions, it retains them where they function as a part of a phrase. The system attempts as much as possible to recognize and retain whole phrases. The remaining text is divided into legal concepts (including single terms and phrases), factual terms and phrases, the names of cases, and statutory references. The concepts are organized both hierarchically as well as alphabetically, and statute citations and references are organized in the familiar way, alphabetically by statute within each jurisdiction, with the section references in the numerical ordering of the sections and subsections. The factual and case lexicons are in alphabetical order. The quadrant of concepts, cases, facts, and statutes is the organizational structure of the legal knowledge as represented in the machine.

Legal concepts and their synonyms are recognized and extracted by a master-hierarchically ordered-legal-doctrinal-concept lexicon. Each concept in the master-concept lexicon is classified by extraction rules as stemmable, moveable, and separable. Stemmable means that suffixes can vary. Separable means that words can come in between the primary words of the concept phrase, and moveable means that the words can appear in a different order. Thus "attorney's negligence" is separable, stemmable, and moveable because "attorney's negligence" matches "attorney was negligent" and "negligence of the attorney." Cases and statutes are recognized by an extraction strategy using the form of citation plus fuzzy matching (recognizes errors in spelling or citations) and a cited cases grammar which extracts full case names and citations, case names without citations, citations with-

out case names, abbreviated case names, and abbreviated citations, all of which are matched to the full case name and citation. The system selects an ideal form for the citation based upon the case, and then upon the database as a whole. Statutes are recognized by the citation form, the statute grammar, and fuzzy matching. The statute grammar recognizes and extracts the full statute citation, isolates sections, and matches them to the appropriate statute. The current Bluebook furnishes the ideal form. The fact lexicon is a default lexicon with the noise words removed except where they form a part of a fact phrase. The extraction strategy uses a grammar to join adjacent fact words and modifiers to form fact phrases, and facts and fact phrases are equated using stemming.

Not only is the entire text of the database organized into these sets of lexicons, each individual case is also represented in terms of quadrants of concepts, cases, facts, and legislation. The organization within each case, however, is different than in the database as a whole where the lexicons are hierarchical, alphabetical, or both. In the individual case, the four sets are organized in terms of frequency. The four types of profile terms are weighted and sorted by factors reflecting their distribution in the processed document and the data collection. Hypertext links are automatically placed which permits the user to hypertext from the terms in the quadrant to each reference in the case. The quadrant thus functions as a FlexNote abstract of the case providing the user with a compact and structured representation of the text, excluding only noise words. From the concepts *felony murder* and *robbery,* the factual terms *clerk, heart attack,* and *death,* the statute reference, *Section 189 of the California Penal Code*, which sets out the crime of felony murder, and the case, *People v. Dillon*, the leading California decision on this offense, one can immediately deduce that the case is about a charge of murder brought against an accused who is alleged to have caused the death of a store clerk who died of a heart attack during the course of a robbery.

The following is an example of a typical legal problem drawn from the law of torts.

> *The plaintiff, a university student enrolled in an evening class, had to walk down an unlit path bordered by shrubs and bushes in order to reach the parking lot where her car was parked. When proceeding to her car, after her class, an assailant dragged her into the bushes and sexually assaulted her.*

This statement of the problem is also a version of an information need. A good Boolean representation of it would be:

university college /p student and ((sexual! /2 assault) or rape! and (foliage or shrub! or bush! or tree) and security and negligen!

A best-match natural language query would look something like the following:

> *Does an occupier owe a "duty of care" to an invitee or licensee to provide security or other safeguards on the occupier's premises against assault by a "third party" or otherwise provide a safe environment?*

A best-match search engine allowing the use of a natural language query eliminates the noise words, thus creating a simple list in which multiple word items contained in its attached concept lexicon are grouped together as single items.

The Flexicon system consists of two basic kinds of lexicons: static and dynamic. The static lexicons are fixed, while the dynamic lexicons are created during the information retrieval process. The major function of the dynamic lexicons is to reveal the content of the database in a lexical form centering around a particular word or concept. The user enters the word or concept, and the system then creates a lexicon of terms and phrases which include that word or one of its forms, either at the beginning, at the end, or inside the phrase, and returns the lexicon alphabetically ordered.

A Flexicon information-need profile is constructed by selecting items from the lexicons. In most cases, a user would generally start with the law. The concept-hierarchical lexicon will permit the user to select the appropriate areas of the law and highlight in the corresponding dialogue box the items that the user wishes to use to form the profile of the information need. The system will also provide the user with a dynamically created lexicon of all forms of that term, or phrases in which it is to be found, and present them to the user in alphabetical order. Appropriate items would then be highlighted in the list for inclusion in the information-need profile. Behind the scenes, Flexicon automatically includes all of the synonyms of that concept. The default weighting of medium can then be altered for any term where it is appropriate.

The next step would be to consider the factual script to which the legal doctrine will be applied. The user can build up a profile of the factual elements of the script by entering a core factual term in a dialogue box. The system will immediately create a dynamic lexicon of all of the forms and phrases that appear in the database containing a form of that word, or a phrase in which a form of it occurs. The plaintiff, in the above set of facts, will be a student, so one can enter the

word "student," and the system will return an alphabetically ordered lexicon of multiple word items from which one might select "attacks by non-students," "college students," "female student," and "protect students." Similar dynamic lexicons can be created around key concepts such as "security" and "campus."

The capacity of the system to create dynamic lexicons around specific terms permits one to enrich the profile of the information need to cover cases which raise the same legal issues but in a different factual context. The lawyer generally creates a script out of the set of facts or story which gives rise to the legal issues in order to locate cases that may be somewhat factually different, but raise relevantly similar legal issues. The profile can be broadened to cover similar situations, such as other kinds of educational institutions, other kinds of wrongful acts, and other kinds of failures to remove different kinds of risks. Thus, a lookup on the term "parking" will produce a dynamic lexicon from which one can select items such as "restaurant parking lot," "parking garage," "shopping mall parking area," etc.—all places where a criminal assault or a robbery might take place. Thus the user can form a profile of not only the particular set of facts, but also a range of factual situations giving rise to the same legal issues.

The user now would be ready to perform the initial search. The Flexicon system would then retrieve a large number of cases, ranked in terms of relevancy in regard to the problem profile. If one went through these cases in order of relevance, by looking at their FlexNotes, viewing the text and paging down to the facts, one would find that almost all are relevant to our legal problem. In the case quadrant of each FlexNote, the user will see the names of a number of cases, appearing more or less frequently. If one hypertexts down on each occurrence of these cases, one would soon discover that they are considered by the judges as leading authorities for these kinds of problems. As the user locates these cases and highlights them in the FlexNote of a case, the user can add them to the problem profile. About eleven or so cases will be found more frequently cited in the list of returned relevant cases, and so now ought to be added to the problem profile in the case quadrant as leading and frequently cited cases. A similar examination of the statute references would be carried out next to locate any statutory provision cited in a number of the relevant cases. If one knows the name of one party to a case and wishes to look at the case or add it to the problem or information-need profile, the user would enter the name in the dialogue box and

perform a Lookup that creates a dynamic, alphabetically ordered lexi-
con of the cases having that name somewhere within the citation.

Relevance ranking can be significantly improved by weighing in-
dividual items (the default position being medium) as high or low, or
by reducing relevance. Key legal concepts unique to the particular
legal issue can be marked as high, while concepts that appear fre-
quently in several areas of the law can be marked as low or given a
reduction of weight function. The present Flexicon system gives each
quadrant a particular weight. These weights were formulated after ex-
tensive testing on eighteen most promising algorithms.

FIGURE 10
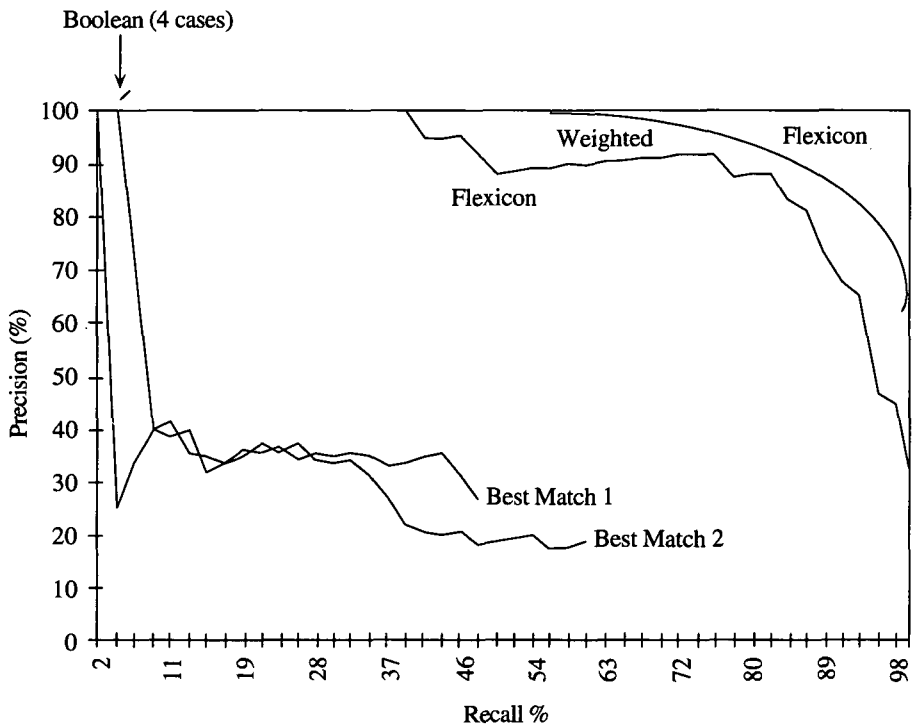
The Flexicon system is still in the development stage. The largest
database upon which it has been tested is approximately thirty three
thousand cases consisting of the California Third and Fourth Series.
Retrieval effectiveness or retrieval performance is measured in terms
of a comparison between how many of the relevant documents are
returned (*recall*) and the proportion of relevant documents to non-

relevant documents (*precision*) at any point of recall.[115] Performance is often portrayed by the use of a graph (see Figure 10, for example) with the vertical axis showing precision and the horizontal axis showing recall. A system that returns more relevant documents than another system and returns a higher number of documents near the top of a ranking is considered to perform better than one that retrieves a smaller number of cases, with less of the relevant documents at the top of the ranking. A perfect performance would be to return all of the relevant documents first, then the semi-relevant documents, where there is a minor modification from the factual script, then the first non-relevant document. The straighter the line, the more efficient is the system. The graph in Figure 10 compares the performance of two Flexicon searches on the above problem with that of two leading best-match search engines using the above natural language query on the same set of California cases. The first Flexicon search was performed with all terms of the problem profile at the default weight of medium, while the second was performed with some terms uniquely specific to this kind of legal issue such as *non-feasance* weighted high, and some terms likely to occur in several different doctrinal areas such as *rape* and *assault* weighted low. All of the initial non-relevant cases that were returned dealt with an invitee or licensee who had suffered harm caused by a third party while on the premises of the defendant occupier, except the wrongful act of the third party was negligent rather than intentional. The term *negligence* in the second Flexicon problem profile was given a negative weight in order to reduce the relevancy ranking of these semi-relevant cases. In addition, the weight of the case quadrant as a whole had been doubled as further testing had proven that case citations give better results than any of the other three quadrants. The precision for the first eighteen documents returned by the unweighted Flexicon profile search is 100%, and then it dips slightly with the first non-relevant case, gradually drops slightly more up to the fortieth case, but then drops substantially. This drop indicates that a substantial number of non-relevant cases are now being returned. The precision for the first forty five cases returned by the weighted Flexicon profile was 100%. Weighting, therefore nearly tripled the degree of perfect relevance ranking. Both best-match searches did substantially better than the Boolean search, and performed very close to each other so far as the results were concerned.

115. *See* Howard Turtle, *Text Retrieval in the Legal World*, 3 ARTIFICIAL INTELLIGENCE AND LAW 5, 10-16 (1995).

In all cases, the test of relevancy was the same; conformity with the factual script. Semi-relevant cases were treated as non-relevant. Only cases that involved the failure of an occupier of land or premises to prevent a wrongful act of a third party to an invitee or a licensee on the property itself were considered to be relevant.

There is another element that must be taken into account, particularly where recall is concerned. The lower the precision in terms of increasing numbers of non-relevant decisions, the greater the amount of time it takes to ascertain which of the returned cases are relevant and to estimate what percentage of the relevant documents have been retrieved. Theoretically one can, in fact, recover all of the relevant decisions in a database of cases if one is willing to spend that amount of time. Almost all systems, however, cut off the number of cases returned where relevance ranking becomes so bad that it is a waste of time to go further down a list of returned cases. A comparison of the results of the best-match, the Boolean exact-match, and the Flexicon system illustrates the capacity of a lexically based search engine for combining relevance ranking with the capacity to use multiple word items in forming the search query. All of the cases returned by the well-formed Boolean query were relevant, but the query found only four out of the approximately sixty relevant cases to be found in the database. If the Boolean query was broadened to return a larger number of relevant cases, it would return a substantial number of non-relevant cases at the same time. The best-match searches, while they retrieved approximately half of the relevant cases, returned them with a substantial number of non-relevant cases intermingled with the relevant ones. The unweighted Flexicon search returned forty five relevant cases up to its point of cut-off. The weighted Flexicon search returned sixty cases up to its point of cut-off and alternative forms of research found no other relevant cases.

The Flexicon system is reiterative in that it facilitates relevance feedback by permitting the lawyer to continue to improve the match between the search profile of the problem and the information need. Rather than simulating what lawyers do, the Flexicon system becomes a partner with the lawyer. The Flexicon system, therefore, is an example of the kind of AI development possible using the creative transformation model of artificial intelligence as contrasted with the more traditional paradigm that would take research in the direction of the Cyc Project and the quest for the "artificial judge."

Conclusion

The Church-Turing thesis is based on premises which are no longer persuasive, at least for the foreseeable future. Our understanding of language and cognitive science has left it behind. At the same time, technology itself has gone in a different direction. The decline in the use of mainframe computers and the proliferation of the desktop computer and workstations; the development of the Internet; communication technology whereby the computer and the television are converging; and the use of optical cable, all are pointing in the direction of huge information networks or matrices, rather than standalone intelligent machines. Thus, the computer can be more usefully viewed as an extension of the human rather than as a simulation of the human. This massive information network will bring enormous changes to the human race, but they will be very different from those projected by the false prophets of the Church-Turing thesis. A more useful metaphor might be to view this as a materialization of a collective unconscious which will shape the consciousness of those who are able to tap into it.

There is only one kind of intelligence, and that is human. As stated by one AI theorist, "[i]ntelligence, as we know it, is (so far) a biological phenomenon, rather than a logical or mathematical phenomenon."[116] If we continue the metaphor of artificial intelligence, we will continue to misconstrue both the nature of ourselves and this marvelous technology which humankind has produced. Computers are the most recent and outmost layer of a set of symbolic representational systems. We are born into language and within it we formulate the human perceptual-conceptual system that constitutes our world. Rather than viewing the computer and the human as two different intelligent systems struggling to communicate with each other, a more useful way to view the relationship between the human and the machine is to view the machine as an extension of the human information system.

Some time ago, many of us saw on the television news a female gorilla tenderly pick up an unconscious three year old who had fallen into the pit that confined the gorilla, and gently place the child by the door where the zoo keeper could retrieve him.[117] It occurred to me then, that if we wish to find intelligence analogous to that of the

---

116. Chandrasekaran, *supra* note 69, at 44.
117. *See Gorilla at an Illinois Zoo Rescues a 3-Year Old Boy*, N.Y. Times, August 17, 1996, at A5.

human, we should look to our fellow species with whom we share over 98% of our genetic code and ponder the wonder and complexity of life, rather than to computers which, after all, are our own creations. In an ancient myth, immortalized in one of the greatest dramas ever written,[118] the Sphinx posed a riddle to Oedipus. *What goes on four in the morning, two at noon, and three at night?* The question, *What is intelligence?*, is just another form of that riddle. Oedipus, according to the story, thought that he had solved the riddle with his answer, "Man!" | but he had not. The Sphinx, herself, represented the answer to the riddle | life-producing female, animal, and something beyond the animal, which we call intelligence. Oedipus's failure to recognize the essence of being led to his tragedy. The answer to the riddle of the Sphinx will be found in our animality, and not in our technology. The idea of intelligent and thinking machines that can perform the range of tasks of which the human is capable is, in itself, a form of Oedipal blindness, and a reflection of our own denial of our animality.

The question posed by Professor D'Amato, *Can/Should Computers Replace Judges?*, is not the right question to be asking as it implicitly assumes the yet unproved Church-Turing Thesis. At this point, we have absolutely no idea whether or not at some distant future time it will be possible to replicate a perceptual-conceptual system in terms of non-organic materials. As stated by one computer scientist:

> AI is based on a mistaken theoretical assumption: the idea that we know what kind of computing thinking is, which in turn mistakenly draws support from the theory of computation. . . . [T]he matter can only be settled by waiting to see what happens. But given AI's confusion about what computer programs are for, its confusion about its theoretical foundations, and its dogmatic refusal to take its failures seriously, it is time that we at least considered the possibility that current AI is incapable of sustaining an inquiry into the nature of intelligence and intelligent thought.[119]

I personally believe that the human race will never celebrate the birthday of HAL or any kind of artificial intelligence system with the capabilities imagined by Clarke and Kubrick in *2001: A Space Odyssey*,[120] because top-down creative and inventive human planning can never begin to match the creative potential of the lengthy complex

---

118. *See* Sophocles, *Oedipus the King*, in THE THREE THEBAN PLAYS 155 (Robert Eagles trans., 1982).

119. Eric Dietrich, *Programs in the Search for Intelligent Machines: The Mistaken Foundations of AI*, in THE FOUNDATIONS OF ARTIFICIAL INTELLIGENCE 223, 229 (Derek Partridge & Yorick Wilks eds., 1990).

120. *See* HAL'S LEGACY: 2001'S COMPUTER AS DREAM AND REALITY (D. G. Stork ed., 1997).

adaptive evolutionary processes[121] which have produced we sentient *homo sapiens*, one of the five species of ape.[122] Time, however, will tell.

For the time being, at least, our time and our resources can be more profitably spent by seeking methods of solving DIFFICULT TASKS related to legal information through the process of CREATIVE TRANSFORMATIONS, permitting the computer to carry out a complicated series of SIMPLE TASKS, which can then be INVERSELY returned in a form wherein the DIFFICULT TASK has been completed. If we consciously pursue the CREATIVE TRANSFORMATION methodology, rather than seeking to replicate or simulate human intelligence in the machine, I think that we will find that we will continue to produce results which will continue to stagger the imagination.

121. *See, e.g.*, RICHARD DAWKINS, THE BLIND WATCHMAKER: WHY THE EVIDENCE OF EVOLUTION REVEALS A UNIVERSE WITHOUT DESIGN (1987); JOHN HOLLAND, HIDDEN ORDER (1995); MURRAY GELL-MANN, THE QUARK AND THE JAGUAR: ADVENTURES IN THE SIMPLE AND THE COMPLEX (1994); M. MITCHELL WALDROP, COMPLEXITY: THE EMERGING SCIENCE AT THE EDGE OF ORDER AND CHAOS (1992).

122. It is now recognized that there are five rather than four species of apes. In addition to the orangutan, gorilla, human, and chimpanzee, there is the bonobo, once thought to be a species of chimpanzee until genetic tests proved conclusively that it is a separate species, closer to the human in terms of the similarity of the genetic code than even the chimpanzee, which is closer to the human than it is to the gorilla. *See* RICHARD WRANGHAM & DALE PETERSON, THE DEMONIC MALES: APES AND THE ORIGINS OF HUMAN VIOLENCE 203-219 (1996).