THE ROLE OF EXCEPTIONS IN MODELS OF THE LAW

Thomas F. Gordon

Gesellschaft für Mathematik und Datenverarbeitung FS-INFRE

ABSTRACT

Existing knowledge representation languages designed for building models of legal domains have required that exceptions to a general rule be collapsed into a single, detailed rule. Although the logic of such rules is clear, they are useless when the complexity of a model grows to realistic proportions. To support this claim, the nature and purpose of models of the law is discussed, and the role exceptions play in furthering this purpose is outlined.

l. Introduction

In this brief report I hope to outline the motivations for our Oblog project at the Research Center for Information Law (FS-INFRE) of the Society for Mathematics and Data Processing (GMD). Oblog is an experimental knowledge representation language designed expressly for use by lawyers in modeling legal domains. It can therefore be compared, e.g., to the TAXMAN project at Rutgers' and the LEGOL project of the London School of Economics . Oblog is strictly experimental at this point; we do not, expect the language to be either expressive or efficient. enough for commercial applications. Rather, we have chosen a modest. subset of the many difficult knowledge representation issues upon which to focus our attention.

Others have addressed various parts of the problem of modeling the law and legal reasoning. McCarty, for example, has concerned himself with the difficult problems of deontic logic and of constructing models from case law precedents. Rissland, to provide another example, has studied the role of hypothetical cases in legal reasoning, and some of the pragmatic aspects of legal argumentation, which she characterizes as various kinds of tactical and strategic 'moves' . It is not possible here to due justice to the range and depth of activity in this interdisciplinary field of law and computer science. The projects just mentioned should, however, suggest something about the kind of research taking place. (McCarty surveyed the field as of 1982 .)

In the Oblog project, we have chosen to focus for the time being on an aspect of legal reasoning which, to our knowledge, has received very little prior attention – the utility of exceptions to general rules. In our effort to introduce a greater degree of precision into the law, we have tended to regard exceptions as an aberration obscuring the logic of the law. For example, we have usually rewritten legal rules of the form

If A and B then C, unless D or E.

to something like

IF A and B and not D and not. E then D.

Exceptions, however, serve an indispensable function in legal models. To support this view, allow me to first describe the character and function of legal models. I will then try to show how exceptions further the utility of these models.

2. Nature of Models of the Law

There are many kinds of models. They can be broadly categorized in terms of the nature of the entities being modeled, the purpose of the model, and the tools used to construct the model. (See Vol. 5, pg. 354-359, of The Encyclopedia of Philosophy for an overview of the various kinds of models used in science.) For example, let us consider models of formal systems, or 'languages', where by 'formal system' we mean a set of symbols; a set of formation rules (i.e., grammar rules) which delimit the possibly infinite set of well-formed formulas in the language; a set of transformation rules for constructing formulas from other formulas, and finally, a set of axioms (i.e. a starter set of well-formed formulas). Together these components make up a syntactic system. Only in relation to a model for the system does the language acquire meaning, or a *semantics.* The entities to be modeled are the components of the formal system, and the purpose of the model is to clarify the meaning and implications of these components. In the case of formal systems, set theory is usually the tool of choice for constructing models of the system, but there are other possibilities.

The use of set theory, which itself may be expressed in a formal system, to explain another formal system may appear to be a form of circular reasoning or infinite regress. Do we not require another mechanism with which to formally explain the meaning of set theory, and so forth? The answer is no, not if we accept that the function of the model is to explain unfamiliar constructs in terms of concepts which are better, if not perfectly, understood.

But. what about modeling the law? Here the entity to be modeled is the law itself, which is clearly not a formal system in the sense just described. Moreover, although expressed *in* language, the law itself is not a language, so it is clearly not our purpose to define the 'semantics' of the law. Rather, our purpose is to *rationally reconstruct* the law. The law is normally represented in natural, albeit technical, language: the language of statutes and cases. These sources of law are not the law itself, but one possible representation of the law. It is clear that these documents are not themselves the law from the fact, that we must first *interpret* statutes and cases to get at the law which they represent, and from the fact that reasonable persons can disagree as to just what the law is, although there is rarely disagreement as to what, words make up the statute or case in question. It is the meaning of the statute or case which is the law, not the text of the document itself.

Before moving on to discuss the utility of models of the law, and tools for implementing such models, let us nail down the subject, matter of interest by listing some domains which, although related to the law and perhaps potential candidates for computer modeling, are not. the object of models of the law per se. Many of these subjects are of current interest to others involved in the rich field of computers and law, but it, may be helpful to distinguish these subjects from the law proper, so as to avoid possible confusion when we begin to evaluate various potential tools and methods.

Michael Dyer of UCLA, for example, is principally interested in modeling how lawyers organize their memory of laws and eases' . Dyer's work has a strong cognitive science flavor; he is not interested in modeling the law per se, but the mental representations of the law constructed by proficient attorneys. Diagrammatically, Dyer's models stand in this relation to the law:

Model -> Mental Representation -> Law whereas
models of the law are direct:

Model -> Law

Prof. Dyer is only tangentially concerned with providing the legal community with useful tools. Nonetheless, work of this kind lays a foundation for eventual legal expert systems. An 'expert system', as originally conceived, is a computer system capable of emulating the reasoning behavior of an expert, in some field. The term 'expert system' has acquired a somewhat broader meaning, but if we restrict ourselves to the narrower original meaning, models of the law are not expert systems, as there is no intent to emulate the reasoning of lawyers. This may be a distinction of little practical importance, however.

Some of McCarty's work also appears to be more concerned with model-ling legal reasoning than in modeling the law. His "prototypes and deformations" work is of this character. The prototypes and deformations model concerns the "construction" and "modification" of legal concepts. Sequences of versions of the law are represented in a prototypes and deformations model, whose purpose is to help us understand the process by which acceptable legal arguments are constructed in "hard" cases and, eventually, to assist attorneys with the construction of such arguments.[1] Thus, the long-term goals of the prototypes and deformation project are much more ambitious than providing tools for representing the law at. a given point in time. This entire discussion, of course, presupposes a definition of the concept of law. McCarty might argue, and persuasively so, that the law includes the principles and policies embedded in a long history of cases, and that a static representation of the law which fails to include a representation of the chain of relevant cases can be of little help in deciding hard cases. It is not necessary to delve into a discussion of the age-old issues of jurisprudence lying just beneath the surface here. I would be merely playing the devil's advocate to dispute this comprehensive view of law. The models of law we propose can only be approximations to the law in this broad sense. They are nonetheless very useful. In fact, as we will argue in the section on the utility of legal models, the are *necessary* for a variety of everyday legal tasks.

The broad definition of "law" recognizes the "open texture" of legal concepts. This quality of the law should not be taken to mean that legal concepts cannot be precisely defined. Rather, it means that our precise definitions are, at best, provisional and that we should be flexible enough to modify them when justice requires. This is the view supported by H. L. Hart, In his essay on Jhering's "heaven of concepts". Hart writes:

> "... we have no way of framing rules of language which are ready for all imaginable possibilities. However complex our definitions may be, we cannot. render them so precise that ... for any given case we can say definitely that the concept either does or does not. apply to it.... Hence there can be no final and

[1] See MacCormick's book on legal reasoning' for a discussion on the distinction between "hard" and "clear" cases.

exhaustive definitions of concepts, even in science.... We can only redefine and refine our concepts to meet the new situations when they arise."

This view of the open texture of the law should not be confused with the "prototype" theories of concepts developed by computational linguists for the purpose of understanding natural language. These theories were developed to explain the lack of precision with which individuals understand and use concepts. We need not address this imprecision in models of the law, as it is not our purpose to represent the range of meanings which an individual may ascribe to a legal concept. Rather, models of the law are models of the *conventional* or *normative* meanings of legal concepts. Conventional meanings may be precisely defined. This is not. to say that we will always want to precisely define such concepts, We may wish to delay decision on an issue in the law by being imprecise. There are ways of being imprecise without abandoning definitions and rules as the principle constructs of the model. Some terms can simply be left. undefined, for example.

2.2. Predictive Models of Judicial Behavior

Although there is a strong correlation between the law and the way judges decide concrete cases, it is obvious that the law is not the only relevant factor influencing a particular judge's decision in a particular case, and lawyers would be foolish to make arguments or plan the affairs of their clients without taking these factors into consideration. (But. see pages 32-34 of MacCormick for a discussion on the duty of judges to decide cases according to law, and the sociological factors constraining judges, to a large degree, to do their duty.)

Despite the importance of these extra-legal considerations, models of the law themselves are not designed to deal with them. Rather than attempting to extend the functionality of computer models of the law to deal with these difficulties, it appears more sensible to consider these separate applications.

8.3. Models of the Non-Legal 'Real' World

Any legal task requires a great deal of general knowledge having no special legal character. A criminal defense lawyer must know, to begin with, that murder is a serious crime. To prepare a deed of appointment, of new trustees of a will trust, an attorney needs to collect a wide range of information, such as the dates of death of the prior trustees and the addresses of the persons concerned. To know what is required, the attorney must of course have some notion of the meanings of 'death', 'date' ’address' and so forth. These examples are due to Nicolas J. Bellord. In his article on the information needs of attorneys, Bellord suggests that *most* of the knowledge required by practicing attorneys is essentially non-legal. Nonetheless, models of the law are not principally intended for representing everyday knowledge about the world. It. should not be surprising that a language designed for constructing models of the law *may* be quite useful for representing some types of

See, e.g.. Cohen and Murphy's article "Models of Concepts".

See pages 200-201 of Nida's "Componential Analysis of Meaning" for a discussion on the conventional meanings of concepts" .

2.4. Document Retrieval Systems and Models of Legal Texts

As mentioned above, it is important to distinguish legal sources, such as statutes and cases, from the law which they represent. In the Oblog project., we are principally interested in directly representing the law. Legal language itself is also a subject. of active interest. Tools designed principally for the retrieval of legal texts must include some representation of the documents to be retrieved, or of the language of the documents. Carol Hafner's conceptual retrieval systems, e.g., includes a computer language for summarizing cases. Legal thesauri, on the other hand, organize the terminology of the law and may be used to semi-automatically construct key-word queries for full-text retrieval systemic . Finally, systems which automatically construct models of the law from original sources will require a thorough model of natural language, including technical legal language.

There is a close relationship between models of the law and systems designed for retrieving relevant legal authorities or understanding legal language. To approach human capability in these tasks, such systems must *include* a model of the law as one their subsystems. (See Sowa's book on conceptual structures for a detailed explanation of the relationship between conceptual models and natural language comprehension.) But it, is important, not to equate these various tasks, as the conceptual models required for natural legal language comprehension are orders of magnitude more complex than those required to represent the law itself. Models of the law are useful in their own right., independently of their possible utility for natural language understanding and document retrieval.

Choosing to focus on a direct representation of the law offers several benefits. The most important of these concerns the function of a model as a rational reconstruction of an area of law, Models developed for document retrieval purposes tend to be *translations* of the original sources, rather than reconstructions. That. is, there tends to be a one-to-one correspondence between the rules of the model and statutes of the original code. This structure is valuable in document retrieval systems, but runs counter to the goals of a rational reconstruction: to offer an equivalent but *simplified,* more elegant representation of the area of law of interest. (See Hart's discussion of Kelsen's notion of rules of law in a "descriptive sense" for one explanation of the role of rational reconstructions.)

Now it is clear that models of the law do not have the authoritative weight. of the original cases and statutes. When it. is time to present an argument before a court, it will be necessary to substantiate the elements of the argument by reference to these original sources. For this purpose a document retrieval system of some kind will be necessary. So rational reconstructions are in no sense a complete solution to the information requirements of attorneys. But, as we will discuss next, they do serve a function not, satisfied by legal thesauri and document. retrieval systems.

Reconstructions can become a kind of "secondary authority" through wide acceptance and use in practice; but we assume here that it will be same time before computer models receive this kind of acceptance.

It. is not. possible to discuss the utility of a tool without reference to the intended setting and circumstances of application. There are many potential kinds of users of computer models of the law: attorneys, judges, secretaries and paralegals, and

perhaps, clients (or would-be clients) themselves. I will restrict our discussion here however, somewhat arbitrarily, to the use of these models by attorneys in advising clients, conducting legal research, and preparing legal documents such as pleadings and briefs. Our starting point is Jon Bing's analysis of the process of recognizing and evaluating the "legal problem" of a client during the initial interview, distilling the legally relevant. facts from the client's telling of his story, finding the legal sources potentially containing (or representing!) the legal norms required to construct arguments in the client's favor, and finally, the actual construction of legal arguments from these norms' . As Bing points out., knowledge of the law is required in each phase of this process. Not all problems are legal problems, and some sort of understanding of the content and limits of the law is required to recognize that a problem is potentially capable of a legal solution. And knowledge of the law is required to sift, out, the relevant facts from those facts which, although important to the client, have no bearing on his legal rights or remedies. We should halt. here to note that this "sifting" analogy is not quite appropriate. The attorney is not a passive listener to the client's version of his tale. Rather, the attorney *creates* a version of the facts compatible with his conceptualization of the law and attempts to *illicit* facts from the client which may be relevant in terms of this conceptualization.

Similarly, knowledge of the law is required to guide the search for relevant. legal texts. A reciprocal process occurs here. The attorney's knowledge of the law guides his selection of cases and statutes, but. his subsequent reading of these sources may extend, modify or update his knowledge of the law (which, in turn, may require him to modify his characterization of the facts or discover additional information). This process repeats until the attorney believes he can make out a satisfactory case for his client., or that no case can be made. Furthermore, later negotiations with opposing parties may require further investigation and research which in turn require him to modify his arguments or characterization of the facts.

As for the construction of legal arguments, they are constructed provisionally throughout this process, using the model which is built up from the attorney's previous understanding of the law together with the modifications resulting from his interpretation of the legal sources he has consulted. Only by constructing and evaluating these provisional arguments can the attorney decide whether further legal research and factual investigation is necessary.

This account of the process of recognizing and analyzing legal problems is admittedly sketchy, but suitable for our purposes here. What. is the potential role of computer models of the law in this process? We have been discussing the importance of knowledge of the law at each stage of this process as if this knowledge were represented only internally, within the attorney. Obviously this is not the case. Due to the sheer complexity of the law, attorneys require a great deal of external support for representing and recovering this knowledge. Traditionally this support takes the form of hornbooks, law summaries, form books, civil practice codes and the like, in addition to the attorney's own notes and records. Computer models of the law can be viewed as an additional source of this kind of support.

In connection with the related problem of tending the ratio decidendi of a case, MacCormick, on pg. 117 of "Legal Reasoning" includes a good discussion of the issues involved in characterizing the legally relevant. facts of a case.

4. The Efficiency Requirement of Computer Models of the Law

Now that we have discussed at some length the nature of models of the law and suggested some uses for these models in the law office context, we reach the kernel of this article. What. properties must a model of the law have in order to serve the various functions outlined above? It would be too much to try to answer this question in full here. I want to focus on a requirement which appears to have been somewhat neglected: the utility of these models must. outweigh the costs involved in building, maintaining and using them. Stated so bluntly, this requirement is obvious and not. particularly illuminating. Certain corollaries to this fundamental requirement are perhaps more interesting. All cases are not of equal value; the defendant in a murder case clearly has more at stake than the plaintiff in a breach of contract suit. concerning the sale of a bicycle.[2] Unfortunately, there is no direct relationship between the legal complexity of a case and its value to those involved. The purely legal issues involved in the dispute over the bicycle could very well be much more interesting and complex than those involved in the murder case. In fact, as Bellord points out, the law of homicide is relatively simple. Moreover, as Bing reminds us: "Proving facts demands resources". Not merely the costs of proving a fact in court must be considered here. Just convincing oneself that a fact. is or is not. the case during the early phases of constructing provisional arguments demands resources. The client may need to be called, witnesses interviewed, records searched, etc..

Ideally then, regardless of the complexity of the legal issues involved, a legal model should assist an attorney in reaching a preliminary analysis which is, in some sense, *probably* correct.. This preliminary, rough analysis should be inexpensive. Then, depending on the importance of the outcome of the case to the client., further legal research and factual investigation should offer the promise of a deeper analysis that is more likely to be accurate than the preliminary analysis. The results of this later analysis may, of course, contradict the results of the earlier analysis. Simply put, additional effort should hold out the promise of compensating benefit.

It may be asked, "What possible value can an admittedly incomplete legal analysis have for a client?" To answer this question, we must remember the adversarial context of legal disputes. The opposing party or parties must also operate under economic constraints, if not necessarily the same constraints. Their arguments too, by necessity, will be incomplete. Each party will evaluate the arguments of the other, and then decide whether to bear the costs of another round of legal research and fact gathering, or to attempt to reach a settlement of the dispute. In the United States, at least, most cases do not. proceed to trial. (This is true as well in criminal cases, in those jurisdictions which permit plea bargaining.)

Having briefly examined the efficiency requirement of legal models, we can begin to appreciate the inappropriateness of requiring all logically necessary conditions for some legal conclusion to be explicitly stated together in a single rule.

---

[2] I will be simplifying the economic issues here considerably. It is certainly possible, although highly unlikely, that the defendant values his life less than the plaintiff values his bicycle. For an entry to the "new" economics of the law, see Richard Posner's "Utilitarianism, Economics and Legal Theory" .

Although this strategy makes the logic of the law transparent., it ignores the important distinction between broadly applicable general principles and more technical rules designed to handle anomalous cases. By requiring exceptions, no matter how often relevant in practice, to be collapsed into the general rule, attorneys are confronted with the choice of either analyzing an issue in great. detail or abandoning the use of the model altogether. Simply put,, collapsing exceptions into the general rule obscures the forest. for the trees.

Although exceptions and other related mechanisms, such as presumptions and burden of proof rules, are only beginning to attract the attention of those interested in computational models of the law or legal "expert. systems", the problems involved are old, if unsolved, ones in the field of artificial intelligence. Common sense reasoning, not surprisingly, has much in common with legal reasoning. The problem of adequately representing the simple fact. that birds fly *except* such birds as penguins and ostriches, is archtypical in artificial intelligence.

The goal here has been to outline the purpose of our current. line of research, not to discuss in detail the technical issues involved or possible solutions. We are at work on a new version of our knowledge representation language, Oblog which, along with several other improvements, will support. reasoning with exceptions to general rules. The knowledge representation issues involved, and our proposed solution, will be described in depth in a subsequent report.

References

1. Bellord, Nicolas J., "Information and Artificial Intelligence in the Lawyer's Office," in *Artificial Intelligence and Legal Information Systems,* ed, C. Ciampi,North-Holland (1982).

2, Bing, J., "Uncertainty, Decisions and Information Systems," in *Artificial Intelligence and Legal information System,* ed. C. Ciampi, North-Holland (1982).

3. Cohen, Benjamin and Murphy, Gregory L., "Models of Concepts," *Cognitive Science 8* pp. 27-58 (1984).

4. Dyer, Michael and Flowers, Margot., "Toward Automating Legal Expertise," in *Proceedings of the First Annual Conference on Law and Technology,* West Publishing Co. (1985).

5. Edwards, Paul, *The Encyclopedia of Philosophy,* MacMillian (1972).

6. Gordon, Thomas F., "Object-Oriented Predicate Logic and its Role in Representing Legal Knowedge," iVr. 135, Gesellschaft. für Mathematik und Datenverarbeitung, Sankt Augustin, West Germany (February, 1985).

7. Hafner, C. D., *An information Retrieval System Based on* a *Computer Model of Legal Knowledge,* UMI Research Press, Ann Arbor (1981).

8. Hart, H. L. A., Essays *in Jurisprudence and Philosophy,* Clarendon (1983).

9. MacCormick, Neil, *Legal Reasoning and Legal Theory,* Clarendon Press (1978).

10. Martino, A. A., *Deontic Logic, Computational Linguistics and Legal Information Systems,* North-Holland (1982).

11. McCarty, L. T. and Sridharan, N. S., "A Computational Theory of Legal Argument.," LRP-TR-13, Laboratory for Computer Science Research, Rutgers University (1982).

12. McCarty, L. T., "Permissions and Obligations," in *Proceedings of the Eighth International Joint Conference on Artificial intelligence,* Karlsruhe, West Germany (1983).

13. McCarty, L. T., "Intelligent. Legal Information Systems: Problems and Prospects," *Rutgers Computer k Technology Law Journal* 9(1903).

14. Nida, Eugene A., *Componential Analysis of Meaning,* Mouton (1975).

15. Posner, R., "Utilitarianism, Economics and Legal Theory," *Journal of Legal Studies* 8(104) (1979).

16. Rissland, Edwina L., "Examples in Legal Reasoning: Legal Hypotheticals," pp. 90-94 in *Proceedings of the Eighth International Joint Conference on Artificial Intelligence,,* Karlsruhe, West. Germany (1983).

17. Rissland, Edwina L., "Argument. Moves and Hypotheticals," in *Proceedings of the First Annual Conference on Law and Technology,* West Publishing Co. (sess).

18. Sowa, J. F., *Conceptual Structures – Information Processing in Mind and Machine,* Addison-Wesley (1984).

19. Stamper, Ronald; Tagg, Clare; Mason, Peter; Cook, Sandra, and Marks, So, "Developing the LEGOL Semantic Grammar," pp. 357-379 in *Artificial Intelligence and Legal information Systems,* ed. C. Ciampi, (1982).