



<https://theses.gla.ac.uk/>

Theses Digitisation:

<https://www.gla.ac.uk/myglasgow/research/enlighten/theses/digitisation/>

This is a digitised version of the original print thesis.

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study,
without prior permission or charge

This work cannot be reproduced or quoted extensively from without first
obtaining permission in writing from the author

The content must not be changed in any way or sold commercially in any
format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author,
title, awarding institution and date of the thesis must be given

Enlighten: Theses

<https://theses.gla.ac.uk/>
research-enlighten@glasgow.ac.uk

Towards a Computational Model of a Theory of Legal Reasoning

by

Alastair Morrison LL.B.(Hons.), Dip.Comp.Sci., M.App.Sci.

A thesis submitted to the
Faculty of Science,
University of Glasgow,
for the degree of **Master of Science**

Department of Computing Science
University of Glasgow
13 May 1992

ProQuest Number: 10992048

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 10992048

Published by ProQuest LLC (2018). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code
Microform Edition © ProQuest LLC.

ProQuest LLC.
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 – 1346

*Thesis
9272
copy 1*

GLASGOW
UNIVERSITY
LIBRARY

Acknowledgment

A large part of my research period was funded by a grant from the Science and Engineering Research Council, to whom I am grateful.

Abstract

This thesis examines the extent to which one can model, computationally, some of the ideas expressed in a specific jurisprudential theory of judicial legal reasoning; that of Neil MacCormick. To introduce the work I explain, in chapter 1, the reasons for undertaking it, and justify the choice of legal theory.

In chapter 2 I rehearse MacCormick's example of deductive justification in a clear case (i.e. one in which there is no, or very little, judicial discretion available) in order to show that rule based reasoning can be used to model the way in which he claims some cases are deductively resolved. I then suggest some simple hypothetical variations of the factual scenario which is the subject of that example, in order to show how easily a user's responses to a rule based system's queries can become uncertain. I propose the use of what I term generic factual rules, as a means to increase a user's confidence in his responses, and the usefulness of the advice a rule based system can provide. Then, in chapter 3, I examine a case based reasoning approach to a different aspect of this problem of user uncertainty. Case based reasoning attempts to emulate reasoning by analogy, which MacCormick, as well as many other theorists, acknowledges as a central part of legal reasoning.

In chapter 4 I discuss how MacCormick claims judges resolve a dispute when deductive reasoning, and reasoning by analogy, prove inadequate for the task, i.e. when 'hard' cases arise. This involves appealing to legal principles, and I discuss the possibility of identifying such principles, particularly previously unarticulated ones, for the purpose of incorporating them into an expert system. I find support for the claim that principles can be framed by a domain expert, in a manner which parallels the way in which MacCormick claims judges formulate them, in the work of J C Smith. I discuss the process by which, according to MacCormick, judges decide whether or not to support a proposed principle, and suggest that it is valid to propose that an expert should try to anticipate their assessments. I conclude chapter 4 by noting that the expert's appraisal of these judicial deliberations, and other matters, would need to be expressed in certainty factors attached to each principle he proposed.

In chapter 5 I examine the extent to which a major part of the

principle formulation task which MacCormick identifies, that of reconciling a proposed principle with those domain cases which are arguably subsumed by it, can be performed using artificial intelligence techniques. I propose the use of a distinct (from the expert system) 'principle-cases reconciliation system', which would be used by the domain expert, for this reconciliation task, and at the end of the chapter I compare the concept of a principle with that of the generic factual rule discussed in chapter 2.

In chapter 6 I discuss the representation and use in the expert system of the principles produced from the reconciliation process. Also in this chapter I discuss the representation and use of domain cases, used by both the principle-cases reconciliation system and the expert system, and the possibility of the expert system user, rather than the domain expert himself, formulating principles. In chapter 7 I note how the user would supply the expert system with the details of his problem, and discuss the problem of grain size which any computational representation of the facts of a case must face.

In chapter 8 I bring together the main points made in the preceding chapters in order to outline how an expert system, rooted in MacCormick's theory, could be built, and how it would function. To conclude I summarise, in chapter 9, the extent to which MacCormick's ideas can be computationally expressed, as revealed by the thesis, and suggest that principles, which are central to his theory, are essential for all future expert systems in law.

Contents

Chapter 1	Introduction	1
Chapter 2	Deductive Reasoning	4
2.1	Introduction	4
2.2	Rule Based Reasoning	4
2.3	Generic Factual Rules	11
2.4	Generic Factual Hierarchy	15
Chapter 3	Reasoning by Analogy	18
3.1	Introduction	18
3.2	Approximate Factual Match	18
3.3	'Factors' and HYPO like Case Based Reasoning	20
3.4	Partial Matching of Cases and Combinations of Factors	23
3.5	Representation (and Utilisation) of Knowledge in HYPO	26
3.6	Generic Factual Rules and Factors	27
Chapter 4	Hard Cases and Principles	30
4.1	Introduction	30
4.2	Acquisition of Principles	31
4.3	Deep Structure Analysis and HYPO	35
4.4	Framing Principles at various Levels of Generality	35
4.5	The need for a Certainty Factor (and MacCormick's contribution to its assessment)	36
4.6	Consequentialist Argument	37
4.7	Conflicting Principles and Consequentialist Argument	40
4.8	Consequentialist Argument and Artificial Intelligence	41
Chapter 5	Principles and Relevant Cases	43
5.1	Introduction	43
5.2	Reconciling Principles with Case Law	43
5.2.1	Identifying Relevant Cases	44
5.2.2	Distinguishing Factually Dissimilar, Unfavourable Cases	45
5.2.3	Reconciling Factually Indistinguishable, Unfavourable Cases	47
5.2.3.1	Wrongly Decided	47
5.2.3.2	Basis of Decision	54
5.3	Using the Same Cases to Deny a Proposed Principle	55
5.4	Obiter Dicta	56
5.5	Unreconcilable Cases	58

5.6	Procedural Context	58
5.7	Parallel between Framing a Principle and Generating a GFR	60
5.8	Lack of Distinction between 'Argument by Analogy' and 'Argument from Principle'	60
5.9	GFRs and Principles	61
Chapter 6 Representation and Utilisation of Principles and Cases		64
6.1	Representing a Principle	64
6.2	Satisfying the Components of a Principle	64
6.2.1	Facts	64
6.2.2	Legal Concepts	65
6.2.3	Open Textured Phrases	66
6.3	Identifying all the Principles applicable to a NFS	68
6.4	Case Profiles Construction and Utilisation	69
6.5	Use of the Principle Based Reasoner	70
6.6	User Projection of Principles	71
6.7	Other Domain Principles	73
Chapter 7 Factual Descriptions of Cases		74
7.1	New Fact Situation entry	74
7.2	Problem of Grain Size	74
Chapter 8 A Principle-Cases Reconciliation System and Expert System based on MacCormick's Theory		77
8.1	Modelling MacCormick's Ideas	77
8.2	Stage One	77
8.3	Stage Two	77
8.4	Stage Three	78
Chapter 9 Conclusions		82
9.1	Recapitulation	82
9.2	The extent to which MacCormick's Theory can be modelled	82
9.3	The need for Principles in any Expert System	82
References		85
Other Material		87

Abbreviations

Main Text

AI	Artificial Intelligence
CABARET	CAsE-BAsed REasoning Tool
CBR	Case Based Reasoning
CKB	Cases Knowledge Base
CRL	Case Representation Language
FLEXICON	Fast Legal EXpert Information CONsultant
FP	Factual Predicate
GF	Generality Factor
GFH	Generic Factual Hierarchy
GFR	Generic Factual Rule
GREBE	GeneratoR of Exemplar-Based Explanations
HYPO	A case-based reasoning program which, amongst other things, can pose HYPOthetical variations of a problem (or NFS).
LCF	Legal Case Frame
MOP	Most On Point
NFS	New Fact Situation
PKB	Principles Knowledge Base
RBR	Rule Based Reasoning
RKB	Rules Knowledge Base
SGA	Sale of Goods Act 1893
UBC	University of British Columbia

Footnotes

<i>AAAI-91</i>	<i>Proceedings of the American Association for Artificial Intelligence</i>	Anaheim, CA (July 1991)
<i>ESL</i>	Susskind, <i>Expert Systems in Law</i>	(1987)
<i>ICAIL-91</i>	<i>Proceedings of the Third International Conference on Artificial Intelligence and Law</i>	Oxford (June 1991)
<i>IJMMS-1</i>	<i>International Journal of Man-Machine Studies</i>	Special Issue: AI and Legal Reasoning Part 1 Vol. 34 No. 6 (June 1991)
<i>IJMMS-2</i>	<i>International Journal of Man-Machine Studies</i>	Special Issue: AI and Legal Reasoning Part 2 Vol. 35 No. 1 (July 1991)
<i>LRLT</i>	MacCormick, <i>Legal Reasoning and Legal Theory</i>	(1978)
<i>MLA</i>	Ashley, <i>Modelling Legal Argument: Reasoning with Cases and Hypotheticals</i>	(1990)

Chapter 1

Introduction

The purpose of this work is to examine the extent to which one can model, computationally, part of a jurisprudential theory of judicial legal reasoning. There has been much work recently on modelling aspects of legal argumentation, however, to my knowledge, nobody has yet attempted to apply this work to a specific jurisprudential theory.

This is not surprising as there is still much debate amongst legal theorists as to what, if any, are the logical thought processes underlying and influencing the way judges reach decisions, and whether these can be articulated. However, I do not believe this should dissuade one from attempting to model any well received theory which does appear amenable to some degree of computational representation. One benefit of such an exercise would hopefully be that it would be easier to evaluate such an 'animated' theory, by testing it on real decided cases, as well as specifically constructed hypothetical cases. Also, if found to have some predictive value, the model could serve as a useful tool for lawyers, and perhaps take us a step nearer to the widespread use of expert systems in law, the potential benefits of which have been well documented. (For example, making specialised advice, available to many people, at low cost.¹)

For the above reasons I believe the ensuing examination should be of interest to lawyers, in spite of the obvious objection which can be levelled against its objective, from the point of view of legal theory, as well as to computer scientists.

I will discuss how one might use current artificial intelligence (AI) techniques, particularly those developed in the work on modelling legal argument, and some proposed techniques of my own, to attempt to build a model of one particular jurisprudential theory of how judges decide cases. I will also suggest how one should handle those aspects of the theory which cannot be computationally simulated.

¹ Berman, D. H. & Hafner, C. D. The Potential of Artificial Intelligence to Help Solve the Crisis in Our Legal System. *Commun. ACM.* Aug. 1989, Vol.32, No.8, pp. 928, 932-3, 937

The theory selected is that presented by Neil MacCormick in his book 'Legal Reasoning and Legal Theory'. It was chosen for the following reasons. It is a well respected treatise. It's author is a contemporary leading jurist, and so is able to comment on this work should he choose to do so. He is very interested in, and actively contributes to, research in AI and law, and so would hopefully be available for consultation should a model actually be built.

The theory describes certain features of legal argumentation which the author claims are operative in certain legal systems, specifically those of the common law tradition, and explains why he believes they ought to be fundamental features of such argumentation. It is thus presented as being both descriptive of operative norms, and as being normative in its own right.² The author draws his example from Scottish and English law, and claims that his theory is descriptive of judicial legal reasoning in general in the Scottish and English legal systems, not just in the domains from which illustrative examples are drawn. Thus if a model is developed, and performs satisfactorily, it would constitute a significant contribution towards the creation of a generally applicable expert system shell.

The theory belongs to the positivist school of jurisprudence. In brief, the central tenets of positivism are as follows.³ Legal rules, found primarily in statutes and cases, are sufficient for the (deductive) resolution of some cases, called 'clear cases'. However there are gaps in the rules. Thus other cases, 'hard cases', arise, which the rules do not cover, or in which the application of the rules is unclear, and such cases require the exercise of judicial discretion for their resolution.

The fact that MacCormick's theory is positivistic in nature is conducive to the goal of describing how a practically useful model based on his theory might be developed, for the following reasons. Firstly, because legal positivism is the theory of legal reasoning most widely accepted today, at least by practising lawyers in common law jurisdictions.⁴ Even if there are opposing schools of thought regarding how law works in theory (the arch exponent of which is Dworkin), it is the positivist school (most closely

² *LRLT*, p. 13

³ a) Gordon, T. F. An abductive theory of legal issues. *IJMMS-2*, p. 96

b) ref. 2, p. 100

⁴ ref. 3a, p. 96

associated with Hart), which seems best to accord with practice.⁵

Secondly, because most of the computational work from which I will draw (i.e. that on modelling legal argument), is grounded in a positivist philosophy. It might also be noted that most of this source work, though applicable to both civilian and common law systems, has been developed with the common law tradition primarily in mind, and MacCormick's examination is of the common law method of judging.⁶

It must be stated that my work in no way claims to be an up to date reflection on MacCormick's views on legal reasoning (although the author has assured me that such modification as his thinking has undergone does not amount to a radical abandonment of the ideas expressed in his book). The aim of my work is not to synthesise the writings on legal reasoning of a particular theorist, but to see how well one can model a cogent, coherent and complete theory of judicial legal reasoning (even if it no longer totally reflects its author's thoughts).

⁵ Sergot, M. The representation of law in computer programs: A Survey and Comparison, in: *Knowledge Based Systems and Legal Applications*, Bench-Capon, T.J.M., Ed.

⁶ *LRLT*, p. 9

Chapter 2

Deductive Reasoning

2.1 Introduction

A large part of MacCormick's theory is concerned with explaining how judges reason when they clearly have discretion in the resolution of a dispute (he claims that from an examination of the case law one can identify factors that influence the exercise of that discretion, and how they influence it). However, before commencing his examination of judicial reasoning in such hard cases (where such clear discretion is involved), MacCormick asserts, in Chapter 2 of his book, that the justification of a legal decision can be (and sometimes is) purely deductive in nature. He illustrates this assertion by showing how the decision by Lewis J., in *Daniels and Daniels v. R. White & Sons and Tarbard* ([1938] 4 All E. R. 258), could be deductively justified using clear rules of law expressed in statutes and cases. This deductive account can be computationally expressed simply using rule based reasoning (RBR), which I will do shortly. This will demonstrate that this first part of MacCormick's theory is amenable to computational simulation. However, given a slight variation of the facts of *Daniels*, matters of uncertainty (and consequently judicial discretion) arise, with which basic RBR is inadequate to deal. This reflects the reality that deductive reasoning, using clear rules of statute and case law, is often inadequate to deal with a legal problem. I will suggest some techniques which would allow an expert system to continue making sound inferences, and providing useful advice, when simple deductive reasoning (implemented by RBR) faces such problems.

2.2 Rule Based Reasoning

To set the context, a summary of the facts of *Daniels*, and the area of law it involved, is required. The facts of *Daniels* are, briefly, as follows.

Mr. Daniels went into Mrs. Tarbard's pub and said 'I want a bottle of R. White's lemonade', and was sold one. He took the bottle home where

he and his wife drank some of the lemonade. They both experienced burning sensations and became ill. The cause of their sickness was subsequently established as being the fact that the lemonade which they had consumed was heavily contaminated with carbolic acid. Examination of the remaining contents of the lemonade bottle showed the lemonade contained a large admixture of carbolic acid.

The plaintiffs, Mr. and Mrs. Daniels, subsequently sued the manufacturer of the lemonade bottle, and the publican who sold it to them, for damages in compensation for their illness, treatment expenses, and loss of earnings when ill. The manufacturer was absolved from liability, the publican was held liable and ordered to pay damages.⁷

The law relevant to the successful action concerned the contractual relationship between a seller and a purchaser; specifically, the implied condition imposed on the seller in a sale by description (by the Sale of Goods Act 1893 s.14(2)) that the goods shall be of merchantable quality. Breach of this condition was the ground on which the liability of the second defender, the publican Tarbard, was established.

I will now suggest how one might develop the rule based part of an expert system whose purpose was to advise on breach of the implied contractual condition just noted. RBR on this module would be sufficient to handle a case like *Daniels* in the deductive manner MacCormick⁸ describes. The knowledge represented would be statutory rules, and 'rules' derived from cases. The sources of these would be primary and secondary sources, and expert opinion.

As proposed by Susskind, individuation of the rules, or 'legal productions', and building the rules knowledge base (RKB), would be primarily a task for the legally trained knowledge engineer; he would immerse himself in the sources and require minimal help from the expert until tuning the system.⁹ Thus from the 1893 Act the knowledge engineer would select the following sections. Section 1(1), establishing the criteria for a 'contract of sale', i.e., that if one party transfers the property in goods to another for a money

⁷ *LRLT*, pp. 19-20

⁸ ref. 7, pp. 30-2

⁹ *ESL*, pp. 59-60

consideration then a contract of sale exists between those parties. ('Property' in goods means the right to possess, use, and dispose of those goods.) Section 14(2), imposing on the contract, where the sale is a sale by description, and the seller is a person who deals in goods of that description, the implied condition that the goods shall be of merchantable quality. Sections 11(1)(a), and 53(1) and (2), stating the buyer's right to recover damages from the seller if he (the seller) breaks a condition of the contract.

These provisions can clearly be re-expressed in the propositional form required of a rule based system (as in fact MacCormick's outline on pages 30-31 of his book clearly shows), and this would be a straightforward task. For example -

If a person (the seller) transfers property in goods and another (the buyer) pays money then there is a contract of sale. SGA s.1(1)

If there is a sale by description and the seller deals in goods of that description then there is an implied condition that the goods shall be of merchantable quality. SGA s.14(2)

If there is an implied condition that the goods shall be of merchantable quality and they are unmerchantable then there is a breach of contract (the buyer is entitled to damages).

Furthermore, through his examination of leading cases in the domain, and commentary on them, the knowledge engineer would identify case rules, express or implicit, which supplement the statutory rules. In effect this is what MacCormick does when he suggests,¹⁰ that the proposition - If a contract of sale exists between buyer and seller, and if the goods in question are a form of bottled drink, and if the buyer requests and receives a certain named beverage, then the goods are sold by description - is implicit, though not expressly stated, in *Morelli v. Fitch & Gibbons* ([1928] 2 K.B. 636). In that case the plaintiff specifically requested a bottle of Stone's ginger wine, and this was held to constitute a sale by description.

Susskind provides jurisprudential support for the idea of deriving rules from cases, he calls them 'case law-statements',¹¹ and concrete evidence of the feasibility of identifying implicit case rules can be seen in, for example, the CABARET¹² project. One should not be too

¹⁰ *LRLT*, p. 30

¹¹ *ESL*, pp. 84-90

¹² Rissland, E. R. & Skalak, D. B. CABARET: rule interpretation in a hybrid architecture *IJMMS-1*, pp. 863, 865

ambitious in trying to divine rules from cases; Deedman for one warns against trying to extract rules from case decisions, when there are simply none present, purely in an effort to make the case suitable for computer manipulation.¹³ However, simply identifying specific groups of facts which clearly seem to satisfy a particular statutory term, and thereby articulating implicit rules, as in *Morelli*, or the CABARET project; and identifying express propositions of law which were clearly intended to clarify such terms, for example, Lord Wright's proposition in *Grant v. Australian Knitting Mills* ([1936] A.C. 85 at p. 100), that a thing 'is not merchantable ... if it has defects unfitting it for its only proper use but not apparent on ordinary examination', would not seem to involve reading anything unwarranted into the case law.

These statutory rules, and case derived rules, would form a RKB, implemented as a standard decision tree. See Figure 1. RBR on this module would follow the conventional practice of querying the user about satisfaction of terminal node conditions, and forward chaining when possible towards the goal.

¹³ Deedman, C. Exposing some Myths about Law and dispelling some Myths about Computers. (paper) pp. 3-4

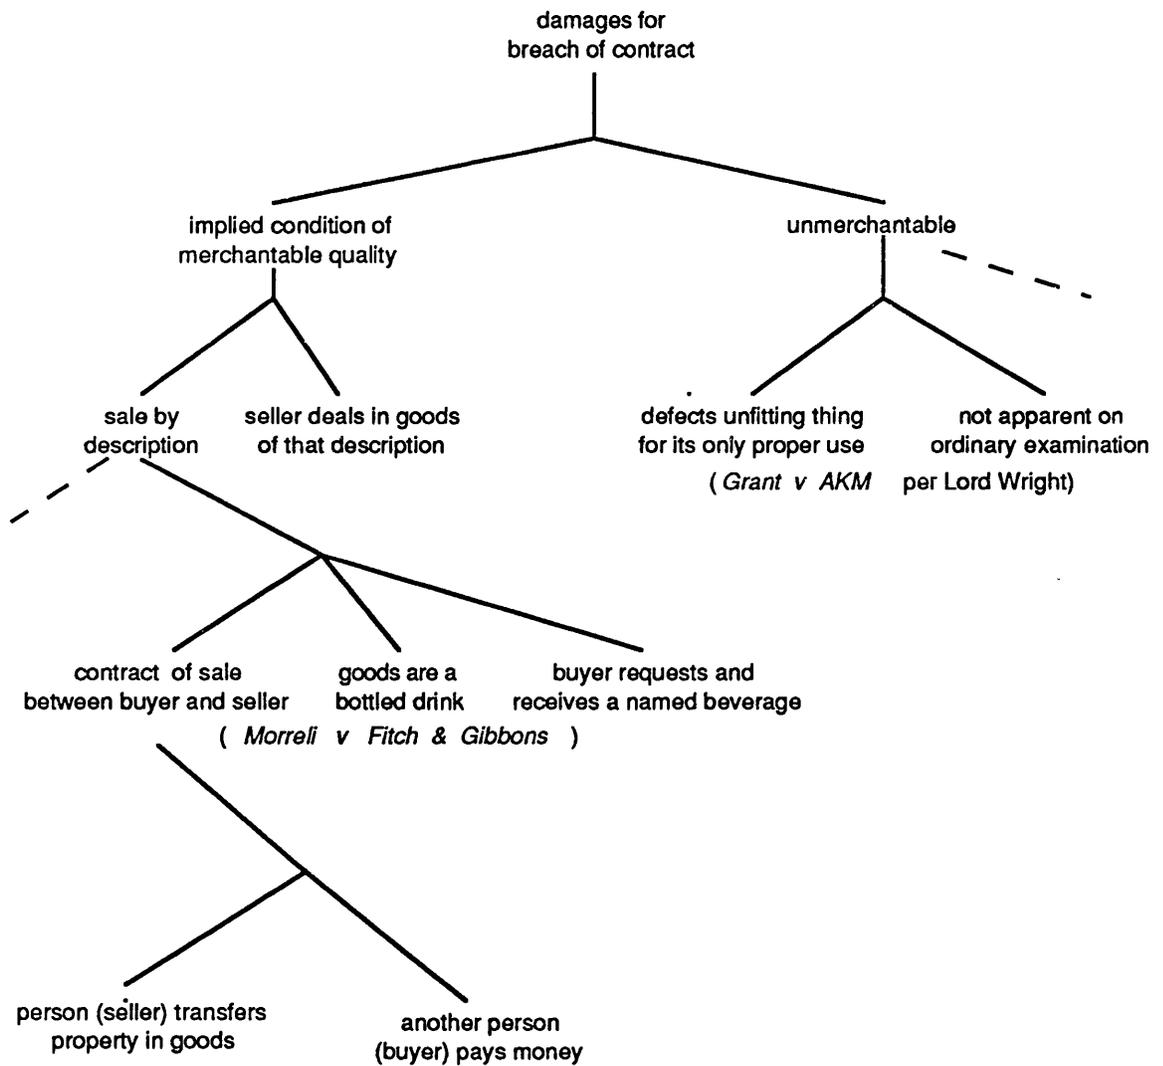


Figure 1. Decision Tree.

However, it is always possible for the authority of a rule derived from a case to be questioned. Thus a conclusion which resulted from inferencing which at any stage employed a case derived rule should be qualified. This qualification could simply be a statement of the values of the factors which a lawyer must consider, or weigh, when assessing the authority of a case (e.g. its age, jurisdiction in which it was decided, hierarchical position of the court which decided it; and maybe also size of the majority (if an appellate case), and comments on the decision in subsequent important cases). The lawyer would then have to assess the strength of the system's findings himself, given such information for each case derived rule the system used.

Additionally, the expert system might incorporate some numerical weighting of the authority of a case rule, given the different values of these factors, in order to compute a certainty for its (the system's) final conclusion. It has been commented that numerical quantification of matters relating to the certainty of an outcome is not something that lawyers are comfortable with. On the other hand others claim that 'Lawyers are accustomed to talking about a cases' chances of success or failure in percentage terms among themselves and to clients'¹⁴. Which view is true may, of course, depend on the lawyer concerned, however, at least one reported research project, the Malicious Prosecution Consultant, has successfully developed a scheme for weighting precedents according to age, hierarchical position, and jurisdiction.¹⁵ Admittedly the purpose was to allow comparative assessment of the strength of cases for the purpose of case based reasoning, rather than to assess objectively the strength of any one case, but the case weights could be converted to certainty factors. Testing of the Malicious Prosecution Consultant on decided cases produced good results, so the weighting scheme would appear to have performed satisfactorily for its intended purpose. Whether a set of certainty factors could be derived from such a weighting scheme, which would produce a useful indication of the certainty of the system's goal being satisfied (particularly when more than one case rule, and therefore more than one certainty factor, was involved, and so some mechanism for combining them became necessary), could only be revealed by comparing the certainty factors produced from test runs with expert assessments of certainty.

Above I cited one case derived rule which could be used to expand on a statutory term. However, it would not be enough for the legally trained knowledge engineer to identify one such clear (explicit or implicit) case rule for a particular term; he would (obviously) have to identify all such rules, because he would not know which would be relevant to a new problem or fact situation (NFS). For example, another case rule might have involved a car sold by make and model. Thus statutory rules would form the core of the decision tree, and each case rule would be encoded as an alternative condition branch from the particular statutory term which it defined, thereby providing specific example clarifications of it. Together these two

¹⁴ Deedman, C. & Smith J. C. The Nervous Shock Advisor: A Legal Expert System in Case-based Law. (*paper*) p. 13

¹⁵ Kowalski, A. Case Based Reasoning and the Deep Structure Approach to Knowledge Representation. *ICAIL-91*, pp. 21-30

types of rule would constitute the essence of the RKB.

However, case derived rules, theoretically at least, merely provide an illustration of some of the ways in which the terms to which they relate may be satisfied, not a complete set of definitions and examples. Thus it might be argued that for each such term a final alternative branch should allow the user to treat the term as satisfied, even though his NFS does not match any of the case rule branches. Any findings of the system relying on this branch would be qualified by the reminder that there was no authority for holding the term satisfied, given the fact situation, and that the user himself would have to find an argument for extending the term to cover his particular facts.

Two points need to be made to qualify this simple idea of a set of alternative case rules offering possible ways to satisfy a particular statutory term, followed by a final alternative permitting the user to treat the term as satisfied irrespective of these rules.

Firstly, if the expert system failed to fire any of the case derived rules, and it contained a case based reasoning module, as discussed later, then that reasoner could be invoked and may retrieve relevant precedents which could be used to support an argument that the NFS satisfied the term. These cases would be identified by looking for close, (as well as just exact), matches between precedents and the NFS. An attempt to satisfy the term in this way would be made before allowing the user to resort to the 'unsupported speculation' final alternative in the RKB, as an authority supported argument that the term is satisfied would obviously be preferable to the user simply making such an unsupported speculation.

Secondly, even if the case derived rules, and case based reasoning (CBR), failed to satisfy, or at least arguably satisfy, the term, the provision of a final alternative triggering rule, which simply allowed the user to speculate that his NFS satisfied the term, is, though theoretically justifiable, practically questionable, for the following reason.

When a term is contained in an old statute, and has already been the focus of many cases attempting to apply it to various factual situations, it is likely that judges will regard its range of application as fairly settled (i.e. inflexible) and so it is unlikely that

they will extend its application to a scenario which has no resemblance whatsoever to decided cases. (I say 'whatsoever' because CBR would already have identified partial resemblances.)

Thus in such situations a final 'unsupported speculation' alternative would be almost bound to be redundant, so from a practical point of view it might be argued that it should be omitted, or at least the user strongly advised against its use, as it would allow (and by the fact of its presence perhaps even tempt), the user to make a speculation which, in reality, a judge would almost never allow. On the other hand, where the term involved was from a recent statute, on which there was little case law, the inclusion of the alternative would seem less contentious.

2.3 Generic Factual Rules

It was demonstrated in the previous section, that it is possible to derive rules from cases which give examples of, or definitions of, a statutory term. Given that actual case examples are often indicative of types of situations or facts that judges regard as satisfying particular terms, it is submitted that efforts should be made to knowledge engineer generic factual rules (GFRs) from particular case rules/examples. Such a GFR, for a particular case rule, would employ broader factual terms than those present in the case rule, and these might subsume the facts of the NFS when it failed to match the facts present in the case rule.

To illustrate this idea I return to the *Morelli* case. I have already described the specificity of the case rule example of 'sale by description' which MacCormick found implicit in *Morelli*. The generation of a GFR from that case rule might simply involve generalising from the terms 'bottled drink' and 'beverage' to 'sealed item of sustenance', or, more generally still, simply 'consumable'. Then, if the user's facts failed to match the actual fact 'bottled drink', but did satisfy the more general term 'sealed item of sustenance' (e.g. suppose a tin of soup was involved), and assuming all other conditions in the case rule were satisfied, a qualified satisfaction of the conclusion of the rule would result.

It might be qualified in two senses. Firstly, in the sense that the user would be reminded that satisfaction of the conclusion of this rule

was based, partly at least, on a knowledge engineered projection from an actual case fact to a generic category. This reminding would take the form of a message presented when the satisfaction of the conclusion in this manner actually occurred, and also in the qualifications attached to the final output. Thus, if the output was an explanation of how the system reached its final conclusion, and the user decided to use the inferences it made when presenting his case, he would know from the qualification of the *Morelli* rule conclusion that he would have to argue that 'bottled drink', and 'tin of soup', are of the same genus, and so the rule should apply.

Secondly, in the sense that a certainty or generality factor (GF) might be associated with the term (or category) 'sealed item of sustenance', which would affect the case rule's conclusion in the following way. The GF would modify the certainty of the condition containing the case fact to which the generic term related, in order that the certainty of the conclusion of the rule was lessened (even if only marginally), the more general the projected term was compared to the actual fact present in the case. More than one of the conditions in a rule might only be satisfiable by employing a generic term, so a suitable method for combining the certainty values of the conditions, to produce the value for the certainty of the rule's conclusion, would need to be identified. (Propagation of such values throughout the RKB would of course determine the certainty of the rule based reasoner's overall conclusion.) However, before that method for determining the certainty of a rule based reasoner's advice could be employed, the values of the GFs themselves would need to be determined. On this matter it is worth noting that Susskind, in his proposal of 'case law-generalisations', which parallel my idea of GFRs, suggests that 'their certainty factors may be drawn from secondary sources or from human experts'¹⁶.

The following discussion rests on the assumption that it would be possible to obtain GFs (either from an expert, or perhaps from a knowledge engineer's examination of the secondary sources) in spite of the claim made by some that legal experts are unwilling to quantify matters of uncertainty numerically.

The simplest version of the GFR idea would require the domain expert to identify those general factual terms, or categories, which he thinks a particular fact in a case rule condition could fall under (e.g.

¹⁶ *ESL*, p. 105

'house' is definitely a type of 'building'), and which he thinks the judiciary might be willing to regard the case rule as applying to (i.e. would they want the rule to apply to buildings in general, or just to a house). A GF would thus represent the likelihood of a case fact being accepted as merely indicative of a particular category, as far as the rule containing that fact was concerned. In other words GFs would be context (i.e. rule) sensitive, as the GF for a particular fact-category combination might differ from one rule to another. The task of determining whether a particular fact is subsumable under a particular category would hopefully not involve any semantic uncertainty (e.g. a 'house' is clearly a 'building'), but if it did (e.g. is a 'scooter' a 'vehicle'), this would have to be an additional, preliminary, consideration of the expert when determining the GF.

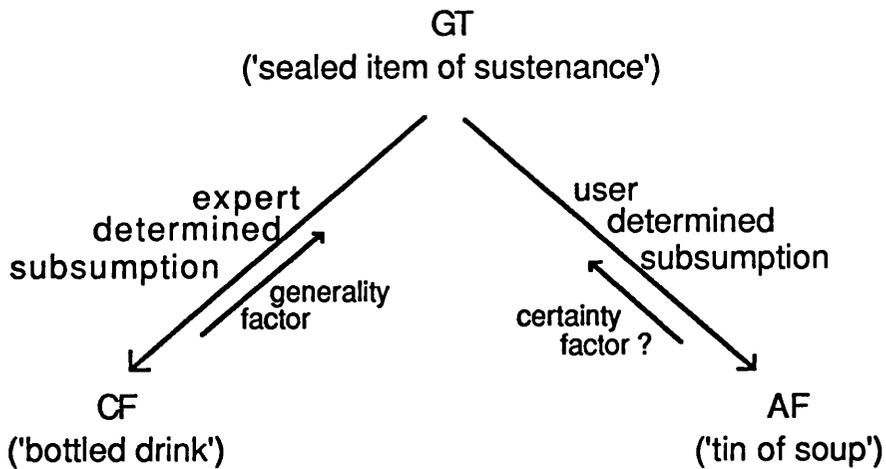
If the user's actual facts failed to match the case fact contained in a case rule condition, then he would be offered the opportunity to try to satisfy that condition by selecting one of the generic terms (or categories) associated with the case fact. As noted, the certainty of the condition would be modified by the GF associated with the term selected. However, as the user himself might not be completely certain that the court would be willing to accept that his actual facts (semantically) satisfied a particular category/term offered (e.g. Is a 'three walled shed' a 'building?'), he would perhaps be allowed to attach a certainty factor to his response. This would further qualify the certainty of the particular condition involved (and therefore that of the rule's conclusion).

If the user was permitted to express his certainty, then, as the expert's GF values would probably decrease the greater the distance between a case fact and its associated generic terms, while the user's certainty would probably increase, and as this would probably be true of all the conditions in a case rule, the system would ideally allow the user to see resultant conclusion certainty values for a rule based on selecting various combinations of generic terms for its different (case rule) conditions before deciding on particular ones.

To summarise the main points above. The user would not be asked, or allowed, to assess the probability of the judge accepting a fact-category generalisation for a particular rule; this would be exclusively an expert (or perhaps knowledge engineer) task. However, he might have to consider the issue of whether a NFS fact is (semantically) subsumable under a particular generic term, and

semantic uncertainty might also be one of the expert's considerations.

Figure 2. Relationship between the fact in a case rule condition (CF), a generic term (GT), and the actual fact (AF) in a NFS.



The expert's generality factor, and user supplied certainty factor, would combine to produce a certainty value for the 'bottled drink' condition, which the GT would now replace.

The tasks of determining generic terms (or categories), and GFs, for a fact in a case rule condition, would ideally be done for each fact, in each case rule condition. Although these have been discussed mainly as expert tasks, it is arguable that even if an expert was available the first of these could satisfactorily be achieved, to a large extent, by the knowledge engineer. Consultation of the secondary sources, and examination of related domains, could provide the knowledge engineer with likely generalisation, including, perhaps, guidance on the semantic issue. The expert would then simply approve, or disapprove, the suggestions, and supply GFs. Each fact in a case rule could have a number of generic terms associated with it, and by offering these to the user the system would allow him to determine that generalisation of the case rule (i.e. that GFR), which was closest to his NFS and the case rule itself.

2.4 Generic Factual Hierarchy

There are aspects of the GFR idea which one might question; here I will consider two. Firstly, there is a knowledge representation issue. Under the simple outline described, a fact in a particular rule condition might be associated with a number of generic terms. Furthermore, the same generic term could be related to the same fact in different rules, or several different facts in the same or different rules. Although these term-fact relationships could differ (e.g. a particular generic term might be the closest of two to one fact, and the most remote of two from another) and so could the associated GFs, would it be necessary to duplicate the generic terms themselves in the course of associating them with facts? Secondly, there is the issue of the user making certainty assessments. It would be better if the possibility of his feeling the need to do so could be reduced.

As an attempt to deal with these issues I propose that a hierarchy of domain facts, implemented as a semantic network, could be constructed. Facts further up the hierarchy would be generic in relation to those they subsumed, i.e., the facts would form a generic factual hierarchy (GFH). This idea can be found in various researchers' work,¹⁷ but here I will discuss it in the context of Karl Branting's recent work on his CBR system GREBE.

Branting uses a hierarchy to express semantic relationships between facts in his target domain of Texas worker's compensation law. For example, 'passenger' has the generalisation 'traveller', and 'traveller' has the specialisation 'driver'.¹⁸ Indeed GREBE's semantic network is an extremely ambitious taxonomy of objects, from the most general 'thing', through, for example, 'tangible entity', down to 'food', and so on; or through 'legal concept', down to 'sub category' or 'legal relation' or 'rule' or 'case', with, for example, 'legal relation' subsuming 'employment', and 'worker's compensation liability'. Each object in the network is represented by a frame, in which (causal, temporal and intentional) relationships between it and other objects in the network can be explicitly represented.¹⁹ This hierarchy enables GREBE to perform its sophisticated form of CBR, which is discussed later. However, for my current purposes, a hierarchy simply relating the basic tangible entities, or facts, of the domain would suffice. In

¹⁷ e.g. That of L T McCarty, and that of Anne von der Lieth Gardner

¹⁸ Branting, L. K. Building Explanations from Rules and Structured Cases. *IJMMS-1*, p. 808

¹⁹ ref. 18, pp. 808-9

the context of *Morelli*, the taxonomic relations, or, as Branting also calls them, semantic rules,²⁰ which such a structure would support would be of the nature 'bottled drink', has generalisation 'sealed item of sustenance', which has specialisation 'tin of food'.

Branting states that 'differing case terms can be matched if both are manifestations of the same abstract term, e.g., if both are symptoms caused by the same underlying physiological state or both have the same generalisation'.²¹ This statement describes a practice of the court, whereby it will sometimes hold that the facts of the case it is considering satisfy a particular open textured phrase, and justify its holding by referring to a precedent in which the phrase was held satisfied by facts which were different from, but of the same type as, those at issue. For example, the court justified the assertion that food was 'reasonably essential' to ensure the proper performance of work duties in an instant case, by observing that iced water had been held to be reasonably essential in a precedent.²² Branting refers to this practice as 'case elaboration', and employs it (or at least instances of it), in GREBE's form of CBR.

I have already proposed that the expert should suggest what generalisation might be permitted for a fact in the context of its rule. That proposal can be viewed as the first part of an attempt to anticipate this practice of the court. I now propose that the expert, or knowledge engineer, should also suggest hypothetical, or novel, specialisations of generalisations which might arise. Such a hypothetical specialisation would be 'tin of food' in the simple example derived from *Morelli*. This would complete the attempt to anticipate this court practice. Again a parallel with this idea can be found in Susskind's work, where he suggests the use of 'case law-derivations' to supplement the RKB.²³ Such specialisations, or derivations, would hopefully lessen the need for the user to attach certainty factors to his responses, as if his facts failed to match the fact contained in a case rule condition, the rule based reasoner would identify generic terms subsuming that fact, and offer him specific alternative specialisations of those terms, suggested by the expert, one of which would hopefully match his facts. Reducing the user's

²⁰ Branting, L. K. Reasoning with Portions of Precedents. *ICAIL-91*, p. 150

²¹ Branting, L. K. & Porter, B. W. Rules and Precedents as Complementary Warrants. (paper for AAAI-91) p. 3

²² ref. 21, pp. 3-4

²³ *ESL*, p. 104

desire to make certainty assessments was the second issue raised at the start of this section. The first was avoiding duplication of generic terms. A semantic network would achieve this in the following way.

Each case rule condition would reference that part of the semantic network which contained the most specific generic term for its fact (that term would of course have a GF for that fact). This would allow the rule based reasoner to offer the user that term, and then each successive subsuming term which had a GF for the fact (plus those terms' specialisations). Thus each generic term, and any associated information (e.g. authority for, or source of, the generalisation) would only need to be represented once, as references to the semantic network, and links within it, would provide the means by which the generic terms for each fact were identified.

Chapter 3

Reasoning by Analogy

3.1 Introduction

I have suggested that GFRs, and a GFH, could be used to infer that a condition of a case derived rule is satisfied when the facts of a NFS do not match the case fact contained in that condition. This idea can be viewed as simulating one way of reasoning by analogy, i.e., drawing attention to the similarity of specific facts in the instant case, and a precedent, in order to support a claim. I now wish to examine how CBR could be used to reason by analogy when there is doubt over whether the NFS satisfies an open textured legal phrase, or term, contained in a node of the basic decision tree.

The examination reveals a CBR approach to the question of what can be argued when the facts in a precedent, and those in the NFS, do not match exactly. The first part of the examination considers one particular way in which the facts may only partially match. The problem this part addresses is similar to that which was being addressed when discussing GFRs, but the exact nature of the problem is different. Here the issue is not simply whether the mismatching facts can be subsumed under a particular term, but rather that they represent different degrees of that term.

3.2 Approximate Factual Match

Looking at the decision tree in Figure 1, one can see that two of the terminal nodes constitute the conditions of a general rule articulated by Lord Wright in *Grant v. A.K.M.* (i.e. If a thing 'has defects unfitting it for its only proper use' and they are 'not apparent on ordinary examination' then it is 'not merchantable'.) Although given the facts of *Daniels* this dictum is clearly satisfied, both these terminal node antecedents are potentially open textured, and so actual case examples might assist the user in deciding if his NFS satisfied them. (Of course it should still be open to the user to decide that his NFS satisfied an antecedent phrase irrespective of any support from case

examples.)

Regarding the first antecedent, a fact situation which satisfied it could be found in *Grant* itself, where the defective article was sulphite impregnated underpants, which caused dermatitis to the wearer. The legally trained knowledge engineer could also look to *Morelli*, in which unmerchantability was held, but which was decided before Lord Wright's dictum, as an example of the sort of scenario which he intended it to cover. If we assume the new problem occurred after *Daniels* then it too could provide an example of what Lord Wright meant.

The expert or knowledge engineer would have to identify the critical facts in each of these cases. (This is, of course, a matter of interpretation; thus even if it seemed obvious that the facts identified were an example of the dictum, it would be possible for a court to deny it.) e.g. A foreign substance, in a bottle of drink, which causes sickness when consumed, is a defect unfitting the drink for its only proper use - *Daniels*. If a corked bottle breaks when an attempt is made to remove the cork, with a corkscrew, in the ordinary way, then it has a defect unfitting it for its only proper use - *Morelli*.

Obviously all three of these fact situations could be formed into case rules with conjunctive conditions. However, suppose the NFS only satisfied some of the condition facts of such a rule exactly; a lawyer might argue that the rule still applied by claiming that the other condition facts were merely instances of the matter to which their conditions related, and that the NFS also satisfied those conditions, although to a different degree. For example, suppose the NFS, like *Daniels*, involved a foreign substance in a bottled drink, but the foreign substance it contained was harmless, and its only affect was to alter the taste. The lawyer might claim that the plaintiff's enjoyment of the drink was diminished, as it was in *Daniels*, and so the first part of the dictum was satisfied, as enjoyment of a drink is part of its 'only proper use'.

On the other hand, what if the altered taste was not unpleasant, just (slightly) different from what was expected. For example, if an energy giving drink was involved, an opposing lawyer could argue that its only proper use was to be drunk and give energy, not to be drunk and have its taste appreciated, its taste being irrelevant, as long as

not unpleasant. He might concede that what *Daniels* was really about was a consumer being adversely affected by a defective drink, but say that it only covered adverse effects in the sense of a consumer being harmed by the drink. He would claim that *Daniels* did not go so far as to establish that a defect in a drink makes it unfit for its only proper use if that defect merely alters the drink's taste, particularly if the altered taste was not unpleasant, and therefore argue that the first part of the dictum was not satisfied.

However, even if the altered taste was not unpleasant, the first lawyer might argue that the NFS facts were still sufficiently close to *Daniels* to satisfy the claim 'unfit for only proper use', albeit not as strongly as they would if sickness had resulted from consuming the drink. His argument would be to reassert that appreciation of the taste of a drink is part of its only proper use, even if not the main reason for consuming it. (If one were to accept this view then what if the defect did not alter the drink in any way. Would the knowledge of, or even mere presence of, a defect in a drink be enough to unfit that drink for its only proper use?)

3.3 'Factors' and HYPO like Case Based Reasoning

The above hypothetical debate reflects the fact that lawyers may, in effect, accept that the fact of a particular case rule condition is merely an example of the matter which the condition is intended to cover, but argue over whether it justifies the assertion that the NFS also satisfies that condition. In order to model the arguments which may be advanced in such a debate it is submitted that, as well as deriving case rules, an expert, or knowledge engineer, should try to identify, from commentary and cases on such an open textured phrase as 'unfit for only proper use', a rule, or rules, which purport to define, or be relevant to the satisfaction of, that phrase, and whose conditions could contain terms which were more general than the exact facts of particular cases.

This idea is borrowed from, and its feasibility (at least for certain domains), has been demonstrated in, the HYPO,²⁴ and CABARET,²⁵

²⁴ a) *MLA*

b) Ashley, K. D. Reasoning with cases and hypotheticals in HYPO. *IJMMS-1*, pp. 753-796

²⁵ Rissland, E. L. & Skalak, D. B. CABARET: rule interpretation in a hybrid architecture. *IJMMS-1*, pp. 839-887

projects, although the term 'factor' is used rather than 'rule'. I will adopt the term factor, and talk about its 'elements', rather than conditions, as these terms better convey the idea of a group of facts whose presence in a case can merely strengthen, or weaken, (i.e. be a plus factor for, or a minus factor against), the claim that a phrase is satisfied,²⁶ rather than alone being determinative of the satisfaction, or non satisfaction, of that phrase. (Obviously this idea, of extracting from the domain sources, groups of elements which purportedly contribute to the satisfaction of, or non satisfaction of, an open textured term or phrase, parallels the construction of GFRs, discussed in the previous chapter, but there are differences, which will be introduced.)

Much of the rest of this chapter discusses how factors can be used to identify the arguments that a precedent can support in a NFS, and as the discussion borrows from, or derives from, Kevin Ashley's observations when discussing his CBR system HYPO, I should first of all consider his comments on the task of actually identifying factors in the domain of trade secret law in which HYPO operates, and any implications the comments may have for the task of identifying factors in the domain I am considering, as unless factors can be identified the ensuing discussion will be purely academic.

Ashley observes that although the legal claims about whose satisfaction HYPO reasons contain open textured phrases, e.g. 'trade secret', they are phrases on which there are many purported definitions in cases and treatises.²⁷ From these purported definitions factors are extracted. However he does also discuss the possibility of extending HYPO to deal with more abstract legal phrases e.g. 'reasonable time', or 'ordinary care', for which there might be less purported definitions, and the problems that might face. e.g. It might not be clear how a phrase is satisfied in a particular case (though it definitely is satisfied) so identifying the relevant facts will be a more subjective matter of interpretation.²⁸ Factors would be identifiable for these more abstract phrases, though the task of identifying them might be harder, and they would be more controversial.

Thus, although the identification of factors might not be as

²⁶ Ashley, K. D. Reasoning with cases and hypotheticals in HYPO. *IJMMS-1*, p. 757

²⁷ *MLA*, p. 239

²⁸ ref. 27, pp. 238-241

straightforward as it was for the open textured phrases HYPO reasons about in the domain of trade secret law, it is feasible to suggest that factors could be identified for 'unfit for only proper use' (which is not as abstract as some of the phrases Ashley comments on) like the one I will now suggest. (Regarding this factor I should note that it is purely my own speculation, my purpose being to discuss the usefulness of CBR as an approach to dealing with open texture, not to identify what factors are identifiable for any particular phrase.)

A factor which might have been extracted from *Daniels* and other cases could have contained two elements which were actual facts of *Daniels*, i.e. 'bottled drink', and 'contains a foreign substance', and a third, denoting a more general category than the third actual fact of *Daniels*, could have been 'adversely affects consumer'. In accordance with the HYPO model, the third element, unlike the first two, would not simply be a binary predicate, taking the value 'yes' or 'no' (and succeeding or failing totally), but would take a value from a range varying from, for example, 'causes alteration of taste', then 'causes an unpleasant taste', through 'causes severe illness', to 'causes death'.

To satisfy such a factor, a case would have to match one of the values for the variable element (as well as, of course, satisfying the other elements in the factor). If the NFS satisfied the elements of such a factor, then any precedent which also satisfied them could provide an argument that the open texture phrase at issue was satisfied, or was not satisfied, by the NFS, the strength of that argument depending on the (relative) values of the variable element that the precedent and the NFS contained.

If we consider the case of *Daniels*, which obviously satisfied the factor's elements, and a NFS which also satisfied them, then we can see what arguments *Daniels* could support in the NFS, regarding the satisfaction of the open textured phrase 'only proper use', given the different ways in which the NFS might satisfy the variable element.

The first point to make is as follows. If both the facts of the NFS and those of the precedent (*Daniels*) satisfied the elements of the factor suggested, i.e. if it was present in both cases, then, irrespective of whether their values matched exactly, this would indicate a relevant similarity between the two cases. On the basis that like cases be treated alike, and before any other matters were taken into account,

this basic similarity would constitute an initial argument that, regarding the phrase the factor concerned, the decision in the NFS should be the same as the precedent. Thus if, as suggested, the only difference between the NFS and *Daniels* was the way in which the foreign substance in the drink affected the consumer's enjoyment, the plaintiff could argue that the drink in the NFS was unfit for its only proper use.

However, if the value of the third element of the factor was lower on the range in the NFS than in the precedent (e.g. if there was 'alteration of taste' rather than 'severe illness'), then the defender could claim that *Daniels* was distinguishable and should not be followed. Indeed he could claim that *Daniels* represented the minimum requirements needed for the phrase to be satisfied, and thus cite it in favour of the opposite outcome. This is an example of a discretionary issue, in which a judge must decide what effect the weaker value of an element of a factor in a NFS, compared to its value in a precedent, has on the support that factor lends to each side in the NFS. i.e. Is the element so weak in the NFS that the factor containing it no longer favours the side it did in the precedent? A case based reasoning system, such as HYPO, can not resolve this question, but it can identify, as a lawyer should, that because of the variable element's differing values, the same case constitutes a potential argument for both sides.

If the third element had the same range value in the precedent and the NFS, then neither the distinguishing argument nor the stronger argument noted would be open to the defender. Moreover, if the range value was higher in the NFS than in *Daniels* (e.g. 'death') then the plaintiff could argue that the precedent's result re the open textured phrase should be given *a fortiori* in the NFS, because the reason for that result (i.e. the presence of the factor) was even stronger in the NFS than in the precedent. A HYPO like case based reasoner would identify this, and the aforementioned arguments for each side, simply by comparing the range values of the factor in the two cases.

3.4 Partial Matching of Cases and Combinations of Factors

Ashley comments, that for each kind of legal claim domain law sources identify factors which are relevant when assessing the strength of each side's position; some of which tend to favour one

side, and some the other.²⁹ Typically a case will contain a number of such competing factors, none of which is individually determinative of the result, each being just one consideration of the judge when making his decision. Consequently, even if the NFS facts satisfy the elements of a factor with the same, or stronger values than those of any precedent which contains the factor, the outcome which that factor favours may still not follow, because it is just 'an argument' that the phrase it considers is (or is not) satisfied.

CBR addresses the question of what can be said about an open textured phrase when it occurs in a NFS containing a combination of competing factors. CBR is employed because in law it is not possible to articulate a rule which will always correctly determine what result will follow given a novel combination of factors (in a new NFS) even if it is clear what outcome each factor supports. It is not possible simply to sum the factors for each side, as some will be more important than others. Nor is it possible to allow for their relative importance simply by weighting them, as the importance of a factor depends on its context, i.e. its relationship with other factors in the case in which it occurs.³⁰ Objective weighting of factors is not possible in law, and even if it was CBR would still be necessary as it justifies legal arguments, not just in terms of factors, which alone are insufficient when presenting a case, but by reference to precedents.

Reasoning with cases cannot determine the outcome in a new case, but it can suggest the best arguments each side can advance in support of the outcome that side is advocating. In CBR precedents are used as judicially sanctioned arguments, or warrants, that when a particular combination of factors occurs a certain result should follow. In other words, because lawyers are expected to justify their claims by reference to legislation, or cases, and because there is a legal tenet that like cases should be treated alike, if the NFS has factors in common with a precedent then a case based reasoner may cite that precedent as an argument that the result given in the precedent should be given in the NFS.³¹

Obviously the more similar a precedent is to the NFS (in terms of factors in common) the stronger the argument for its outcome.

²⁹ Ashley, K. D. Reasoning with cases and hypotheticals in HYPO. *IJMMS-1*, p. 757

³⁰ ref. 21, p. 758

³¹ ref. 21, p. 758

However the issue of which, of two precedents, is most similar, or most on point, is not simply a matter of summing the factors each case has in common with the NFS and comparing the totals. As already noted, some factors will be more important than others, and their relative importance can be different in different combinations. Additionally, when trying to assess the relative similarity to the NFS of two precedents, there is the unknown influence of any factor present in one precedent but not present in the NFS or the other precedent. Consequently a case based reasoner cannot be too ambitious in its attempts to identify the most on point precedents.

Therefore, in HYPO, one case is regarded as more on point than another (with respect to the NFS) only if the latter's relevant factors, i.e. factors shared with the NFS, constitute a proper subset of the former's.³² Additionally, if a particular precedent's set of relevant factors is not subsumed by that of another precedent, then, with respect to those precedents whose factors it subsumes, it is the most on point (MOP) precedent. In HYPO no attempt is made to assign priority between two precedents, each of which has a factor in common with the NFS that the other lacks. Rather they are said to be equally on point. Consequently HYPO may identify more than one MOP case,³³ but it will attempt to suggest which is the best case to use of several MOP cases favouring one side³⁴.

This limited extent to which HYPO ranks the on pointness of precedents to the NFS is merely a reflection of Ashley's recognition that it is not possible to objectively assess the relative importance of the factors they contain. Only if a case has all the factors in common with the NFS that another has, and some more besides, can it be argued that it is more on point than that other, and even then it may be distinguished because of the presence of additional factors not present in the NFS, or because of a difference in the value of a variable element of a factor.

The net result of all this is that, regarding the satisfaction of a particular open textured phrase, HYPO may return several MOP cases, and if these favour different sides then both sides have strong arguments, and no attempt is made to decide between them. According to MacCormick's theory, the way such deadlock would be

³² ref. 21, see diagram p. 759

³³ ref. 21, p. 780, also see diagram p. 781

³⁴ ref. 21, p. 760

resolved is by appealing to principles; this is discussed in the next chapter.

It will be noted that I have only suggested one factor which might be relevant to the open textured phrase 'unfit for only proper use', whereas the discussion in this section has been about what can be argued when there are several, competing, factors relevant to the satisfaction of a particular open textured phrase. However the important point is that, as an example of an open textured phrase, 'unfit for only proper use' represents the sort of situation where a CBR facility, such as HYPO, should be available. i.e. A situation where, given a NFS, the law may be inconclusive as to whether the phrase is satisfied, but where the best case supported arguments available to each side can be established by identifying the factors present in the NFS, and comparing it with those precedents which contain some or all of the same factors.

3.5 Representation (and Utilisation) of Knowledge in HYPO

Ashley has designed his own frame-based case representation language (CRL) to represent the facts of a case at two levels of abstraction.

At the first level of abstraction, legal case frames (LCFs) are used to represent detailed factual objects, relationships, and events that are typically present in disputes of the subject domain. Any particular case, whether a precedent or a new problem, is described in terms of a number of nested LCFs.

At the second, more abstract level, factual predicates (FPs) are used to summarise the lower level case facts represented in LCFs. FPs are generalised factual statements that indicate whether certain legally significant relationships are true in the case. HYPO infers whether a FP is satisfied for a particular case by testing information about the case contained in the LCFs.

The other main knowledge structures in HYPO are dimensions. Each dimension represents a factor and contains a number of prerequisites. These prerequisites correspond to what I have termed the elements of a factor. In order for a dimension to apply to (i.e. for the corresponding factor to be regarded as present in) a particular case,

each of that dimension's prerequisites must be satisfied by the case's FPs.

In the simple scenario I have outlined in my domain the second level of abstraction, FPs, would not be necessary. The question of whether the factor (or dimension) I have suggested is satisfied by a particular case would simply require an examination of the basic facts of that case which would be described by LCFs. (However this is not to deny that there might be other factors identifiable in the domain whose elements are more general factual statements, and for which statements examination of the basic facts of the case would determine satisfaction.)

HYPO's dimensions each have a focal slot prerequisite representing a factual category for which a range of values exists. This corresponds to the variable element I discussed in the context of factors. It allows for closer matching between the NFS and a precedent, for example, strengthening or weakening the force of the argument the precedent supports, in the manner I have discussed, if the range values of the element differ in the NFS and the precedent.

All precedents relevant to the claim (or phrase) at issue are described using the CRL, and are stored in a cases knowledge base (CKB). They are indexed by a library of dimensions. In order to present his problem to the system the user describes the NFS by filling in LCFs for it via a menu driven case editor. Thereafter the task of identifying the MOP precedents proceeds in the manner outlined when discussing factors.

This is a simplification of what HYPO does, but it is sufficient to indicate the potential utility of a HYPO type module for advising on open texture.

3.6 GFRs and Factors

In the previous chapter I outlined how a simple rule based expert system could be built, based on statute and case law, which would be capable of advising on a case like *Daniels*, in the deductive manner MacCormick describes. I then suggested that, in order for the simple rule based system to be able to advise on a wider range of problems, in which the facts may be similar to, but not the same as, those in

the case derived rules, those rules should be supplemented with GFRs, ideally contained in a GFH. In this chapter I examined how CBR can be used to identify the strength of precedential argument for, and against, the satisfaction of a particular open textured term, which may occur either in a statutory, or case-derived rule. The examination was conducted in the context of one particular form of CBR, an alternative to which will be discussed later.

As the task of identifying factors, necessary for the HYPO form of CBR just discussed, and that of identifying GFRs might seem similar, it might be useful to emphasize the different purposes of CBR and reasoning with GFRs, in order to affirm the distinct nature of GFRs and factors, and to illustrate the potential of CBR and reasoning with GFRs to complement one another.

A GFR would be used to augment a clearly defined (case) rule. i.e. When a particular set of case facts can clearly be identified as constituting the necessary and sufficient conditions to satisfy a particular term (those facts constitute the conditions of a case rule), a GFR would express the expert's assessment of one, or more, similar facts being allowed to satisfy one, or more, particular conditions of that case rule. CBR is used when one (or more), precedents satisfy a particular term or phrase, but it is not clear which groups of facts (i.e. factors) are necessary, or sufficient, for that satisfaction. In CBR what is at issue is whether a particular (novel) combination of factors can satisfy a particular term or phrase (which constitutes a condition in a rule, whether statute based or case based). With GFRs, the issue would be whether a particular fact can replace another particular fact (which constitutes a condition in a case rule).

A GFR would be generalised from the fact(s) of one case, with the generic terms it contains not being found in any case, whereas a factor is identified as being common to several cases, but each possibly having different values for one particular generic element, which is clearly discernible as common to all. Hence the idea of a GFR is to allow one to express the likelihood that a particular case rule will be broadened in order to allow an argument that the case is analogous to a NFS. Whereas factors offer the potential for ranking several cases which already contain those factors according to how analogous they are to the NFS, without any expression of certainty.

Consequently, these two approaches to the inadequacies of a rule

based model are not intended to be alternatives, but rather complements, and one might even suggest that they be integrated. For example, some of the elements in a case based reasoner's factors might contain (or contain references to) progressively more generic terms, plus GFs. This would allow the authority of a precedent, which has a factor in common with the NFS, to be modified by the generality, or certainty, of that factor, relative to the NFS, as determined by the GFs of its component elements. For example, the first element of the factor I have suggested above, 'bottled drink', might be generalised to 'sealed item of sustenance', or even simply 'item of sustenance'. Such a generalisation would allow the 'tin of soup' scenario I suggested earlier to be covered by the factor. However this generalisation would affect the certainty of the factor's acceptance as a valid indicator of similarity between the NFS and the precedents the factor indexed, as one could not be sure that the court would accept the analogy between a tin of food, present in the NFS, and a bottled drink, present in those precedents, and this would in turn undermine the authority of the precedents. However, before attempting to assess the utility of integrating GFRs and CBR, the basic idea of GFRs would need to be implemented and tested so that it could be evaluated. Much work has already been done on CBR, and test results seem to demonstrate its ability to model legal argumentation, and return useful advice.

GFRs constitute an attempt to anticipate the exercise of one type of judicial discretion. i.e. Discretion as to whether a fact, similar to one present in a condition of a relevant case derived rule, should be allowed to satisfy that condition. This sort of discretion parallels that which is involved in MacCormick's discussion of how judges resolve disputes when hard cases arise; this will hopefully become apparent in the ensuing chapters.

Chapter 4

Hard Cases and Principles

4.1 Introduction

I will now examine how MacCormick suggests judges reason in hard cases, in which discretion is clearly available, and the extent to which his theory on this matter can be modelled.

First of all one should establish what prevents a case from being deductively resolved. What makes it a 'hard case', in which judges are free to exercise discretion, as opposed to a 'clear case', such as *Daniels*, in which the availability of discretion is all but non-existent, and for whose resolution simple rule based expert systems are, as was demonstrated, adequate. MacCormick suggests ways in which a case may be 'hard', rather than 'clear', although he warns that the majority of cases do not fall neatly into either category.

A case may be hard because one of the relevant rules at issue is ambiguous or unclear, due to an open textured, or vague, phrase or concept.³⁵ The question of how a choice between the competing versions of the rule which are available is made, is what MacCormick refers to as the 'problem of interpretation'.³⁶ Obviously CBR, or 'reasoning by analogy', provides one approach to this problem, as has been explained. i.e. Find that case which has interpreted the phrase which is most similar to the NFS, and use that similarity to justify interpreting the phrase in the same way in the NFS. However, what if the cases are not very analogous to the NFS, or not persuasive or decisive either way? Such situations would seem to transcend the 'problem of interpretation', and fall within the second type of hard case MacCormick identifies, which is discussed next.

Sometimes a decision needs to be given regarding a matter for which there is no directly relevant rule of law. i.e. Justification of a decision by reference to a particular statutory provision, or case, is not possible. (Alternatively, as just noted, the relevant rules may be

³⁵ *LRLT*, p. 65-6

³⁶ ref. 35, p. 68

of little use.) Indeed the very matter at issue may be 'whether there is or is not any legal ground at all for some claim or decision at law'.³⁷ What the judges then have to decide is "Does the law in any way justify a decision in favour of this party against that party in this context?".³⁸ 'The problem is, in effect, whether it is justifiable in law to assert, or to negate, some proposition ... which covers the facts of the instant case'.³⁹ This MacCormick refers to as the 'problem of relevancy',⁴⁰ and it is in the context of one particular case, which epitomises this problem, that MacCormick presents his account of how judges exercise their discretion in hard cases. This case is *Donoghue v. Stevenson* [1932] A.C. 562; 1932 S.C. (H.L.) 31, which concerned the liability of the manufacturer of a product to those harmed as a result of a defect in that product which resulted from negligence in the production of the product.

Basically what judges (and counsel) do, according to MacCormick, when, in the way just noted they 'run out of rules', is appeal to legal principles. So, in order to be able to model, in an expert system, MacCormick's theory of how judges resolve hard cases, the first task must be to identify general principles relevant to the domain in question.

4.2 Acquisition of Principles

The question of whether it is possible to identify the general principles which underlie the surface discourse of the law, and influence the resolution of hard cases, is something which has been disputed. Sometimes principles are clearly expounded in case law, obviously identifying these is not a problem. However, sometimes principles are relied on by judges which have not previously been expressly propounded. Can such principles be identified and incorporated into an expert system? Such a principle, the 'neighbour principle', was articulated by Lord Atkin, who delivered the leading judgment of the majority in the case of *Donoghue v. Stevenson*. MacCormick claims that Lord Atkin's reasoning when framing this principle exemplifies how all judges frame 'new' principles, so for the purpose of modelling MacCormick's theory we should examine to what extent this process can be imitated.

³⁷ ref. 35, p. 100

³⁸ ref. 35, p. 69

³⁹ ref. 35, p. 70

⁴⁰ ref. 35, p. 70

Susskind would appear to regard the exercise as pointless from the outset. He advises against trying to build expert systems to resolve hard cases, and his comments on the prospect of building a model of the legal concepts in a domain suggest he would argue against trying to identify and model principles, as concepts often appear in principles. Susskind observes that he has indicated the written materials which may be consulted to find the units with which a rule based model, capable of dealing with clear cases, may be built. He then comments that 'There are no direct counterparts to those repositories in respect of the constituents of conceptual modelling.'⁴¹

I submit that text books and articles, by academics and practitioners, do articulate existing and inchoate principles, and policies, be they contentious or uncontroversial. Such sources must exist in order that the principles may be taught to students, and debated by academics and practitioners. This view, which of course directly challenges Susskind's assertion, finds support in the work of J C Smith⁴². Smith advocates the need for a conceptual model in order to represent principles, and, furthermore, he notes that it is possible to extract definitions of legal concepts (for example 'estate'), from the secondary sources.⁴³ These definitions might initially be in terms of other legal concepts, but can be progressively refined to simple terms for which decided cases give examples of satisfying facts. By examining the user's facts it can be established whether the initial concept is satisfied or not. (It is worth noting that this 'concept refined to simple fact(s)' relationship seems to parallel MacCormick's theory of 'law as institutional fact'⁴⁴, alluded to on p. 191 of his book.) Smith and others at the Faculty of Law and Artificial Intelligence Research, at the University of British Columbia (UBC), have put his theory into practice in the Nervous Shock Advisor, the Hearsay Rule Advisor, and other projects, and obtained encouraging test results.

The task of identifying all the relevant principles in a domain would probably be beyond the ability of the legally trained knowledge engineer, so here a departure must be made from Susskind's expert

⁴¹ *ESL*, p 155

⁴² Professor of Law at the University of British Columbia

⁴³ Smith, J. C. A Conceptual Model for the Representation of Legal Knowledge. (*paper*)
p. 1

⁴⁴ MacCormick, D. N. Law as Institutional Fact. *Law Quarterly Review*, 1974, Vol. 90, p. 102

system development plan in which the expert does not become involved in a major way until tuning the model. The expert would have prime responsibility for identifying the principles (for most of the UBC projects the expert was Smith himself). It will no doubt be commented that this would compound the existing 'acquisition of knowledge from the expert' bottleneck, already a major problem in the development of expert systems in law (as Susskind notes in the preface to the paperback edition of his book) and this cannot be denied.

The attraction of suitably motivated academic experts is perhaps the best hope of alleviating this knowledge acquisition bottleneck, and in this context probably the best approach they could take to the task of identifying legal principles is to be found in the 'deep structure approach' to knowledge representation, developed by Smith and Cal Deedman, which has been used in the UBC projects.

Deep structure analysis involves searching for "deep structures or fact patterns underlying a legal doctrine which account for and explain the decisions in the cases. Once these patterns have been identified a broadly stated meta-rule or principle may be formulated which explains the general direction of the case law in the domain independently of the 'surface discourse' of law at the doctrinal level".⁴⁵

Deedman makes the following comment about deep structure analysis. "In order to make sense of a body of case law, some general unifying principle is needed. A solution which has been proposed in some detail elsewhere [i.e. in the proceedings of the First International Conference on Artificial Intelligence and Law] has been called the deep structure approach. In essence it calls for the formulation of a broadly stated underlying rule which explains the general direction that all the cases in a domain of law have taken independently of the surface discourse. Without such a rule to act as a compass, one finds oneself drifting aimlessly in a sea of a cases. There is no hard and fast way of discovering such a rule. It must be discerned by a careful examination of the facts of the cases themselves."⁴⁶

⁴⁵ Kowalski, A. Case Based Reasoning and the Deep Structure Approach to Knowledge Representation. *ICAIL-91*, p.22

⁴⁶ Deedman, C. Exposing some Myths about Law and Dispelling some Myths about Computers. (*paper*) p. 5

The significance of these comments on deep structure analysis is that what Smith and Deedman are saying one should do, and do do themselves, when building an expert system, is exactly what MacCormick identifies Lord Atkin as doing when he formulated the 'neighbour principle' on manufacturer's liability. i.e. Lord Atkin looked at the relevant cases and argued that in them a general trend was perceptible, and that therefore from them a general principle could be inferred. So in effect what Smith and Deedman claim, and through their projects demonstrate, is that it is possible (for an expert) to extract principles from a domain, (whether they be regarded as not existing, or merely unarticulated, before hand), and their theory and achievements mirror this process which, according to MacCormick, all judges undertake when faced with hard cases.

Indeed the parallel between Smith's theory, and that of MacCormick, can be seen in the following comment on deep structure analysis. "This theory postulates that judges and lawyers use deep structure fact patterns to decide cases and analyse problems ... whether ... at a subconscious level ... or in ... [a] more overt goal-oriented manner ... where solving a client's problem or arriving at a correct legal decision."⁴⁷

The task of framing principles through identifying trends by deep structure analysis, might be formidable, even for a luminary in the domain field, but the UBC projects demonstrate that it is feasible. On the other hand, the task of identifying a particular trend, in order to articulate a principle, might be quite straightforward. MacCormick himself gives an example of identifying a trend, from case law, concerning the 'neighbour principle'; that trend being from regarding it as a permissible to a mandatory principle.⁴⁸ Another recent example of a clearly discernible trend in an area of law, which all lawyers should have been aware of, would be the trend towards making marital rape a crime in England.

Obviously the principles identified through deep structure analysis would not be uncontentious, even if the work of more than one expert, but that is not a valid objection to the exercise as the opinions of counsel are not indisputable (if they were then there would be no point in going to court), and the most that an expert system can do,

⁴⁷ Kowalski, A. Case Based Reasoning and the Deep Structure Approach to Knowledge Representation. *ICAIL-91*, p.22

⁴⁸ *LRLT*, p. 159

after all, is offer an opinion.

4.3 Deep Structure Analysis and HYPO

It is perhaps instructive at this point to make some comparisons between one of the UBC projects, the Nervous Shock Advisor,⁴⁹ and the HYPO project discussed in the previous chapter, to suggest why Ashley did not propose a deep structure analysis for his system, as both deal with case based law.

Deep structure analysis involves an attempt to identify the main facts which must be established to succeed in a claim of nervous shock, and this seems to parallel the process of identifying factors in the domain of trade secret law, which denote the main elements relevant to success in particular trade secret claims, e.g. trade secret misappropriation.

The reason why Ashley does not propose a deep structure analysis, whereas Smith and Deedman do, is because nervous shock is a vaguely defined area of case law (deliberately chosen by Smith and Deedman for that reason), in which the relevant facts are not already clearly identified in the secondary sources, whereas in the domain of trade secret law they are. The nervous shock domain, being a vaguely defined area of tort (or delict) case law, parallels the *Donoghue v. Stevenson* area of manufacturers' liability more closely than Ashley's technical domain of trade secrets. Hence my earlier proposal, when discussing CBR, for the domain expert to extract factors from the primary and secondary sources, might require what Smith and Deedman would describe as a deep structure analysis.

4.4 Framing Principles at various Levels of Generality

It might be noted that a court only considers what general principle it can formulate in the context of a specific case (i.e. the NFS). Therefore the level of generality of that general principle is probably not going to be much higher than what is necessary to cover the NFS, because it should be consistent with existing rules and principles, and the more general it is the greater the chance of inconsistency.

⁴⁹ Deedman, C. & Smith, J. C. The Nervous Shock Advisor: A Legal Expert System in Case-based Law. (*paper*)

(MacCormick observes that the requirement of consistency takes priority over other considerations, such as the result of employing consequentialist argument, on which see later, because no matter how desirable a particular ruling may be on consequentialist grounds, it must not contradict existing valid rules.⁵⁰)

However an expert system would require the formulation of a general principle in a vacuum; i.e. out of the context of a specific NFS. Thus, from the relevant precedents, a general principle would be formulatable at various levels of generality. MacCormick himself observes that out of the context of a specific case, a principle capable of covering an area of law can be framed at various levels of generality,⁵¹ and he illustrates this by borrowing an example of the different levels of generality at which a principle covering manufacturers' liability to consumers could have been phrased at the time of *Donoghue v. Stevenson*.⁵²

Therefore, as the expert system is constructed in advance of a NFS arising, it is arguable that the expert should be free to generalise the facts of the relevant cases to provide more than one version of a potential general principle so the system (or user) could select the one closest to his NFS.

4.5 The need for a Certainty Factor (and MacCormick's contribution to its assessment)

Having projected a general principle at one or more levels of generality, the expert should then attach a certainty factor to each version of the principle, to indicate to the user his (the expert's) certainty that the court would accept that version if advanced. This certainty factor would reflect two things. Firstly, the expert's certainty that the court would want to accept the principle, which was previously unarticulated, secondly, that the court would want to accept the version of the principle at the particular level of generality concerned. MacCormick's theory provides guidance on the assessment of the first of these matters. Basically it would require the expert to speculate as to the court's view of the desirability of the consequences which would ensue if that principle existed, using

⁵⁰ *LRLT*, p. 106

⁵¹ ref. 50, p. 118

⁵² ref. 50, p. 117-8

certain criteria which MacCormick claims judges evaluate when making this desirability assessment. (A fuller explanation of this will be given shortly when discussing consequentialist argument.) Additionally, it would be necessary for the expert to reflect in this first assessment his certainty that he had correctly identified the underlying trends and developments in the domain when framing his proposed principle.

The second of these matters parallels the expert's task of determining the GFs to be associated with generic terms, in regard to their relationship with particular relevant facts in a particular case rule, when formulating a GFR. (See section 2.3 in chapter 2.) The difference here would be that the expert would probably have generalised from various different case facts, in many precedents (or case rules). As with GFRs, the resultant certainty values, or GFs, for the different conditions in the principle, would have to be combined to produce a composite certainty value for the principle's likely acceptance, on the grounds of generality, at the particular level of generality at which it was pitched. This in turn would have to be combined with the certainty value produced for the first matter, in order to produce an overall certainty factor for the principle at this particular level of generality. This process would be repeated for each version of the principle.

4.6 Consequentialist Argument

Although it might be possible for an expert to identify principles in a domain of law, there are other problems which judges (and counsel) face when attempting to apply principles. The first problem that MacCormick notes is that more than one principle may seem relevant, and they may conflict. Alternatively there may be dispute as to whether or not what is claimed to be a relevant principle covering the instant case is in fact a valid principle of law (this was in fact the situation in *Donoghue v. Stevenson*). Thus if a judge (or expert) claims to have identified a previously unarticulated principle in the domain law, others may challenge its validity. In both these situations, says MacCormick, judges justify their decisions, whether it concerns which of two conflicting principles to support, or whether to support or deny a proposed principle, by applying consequentialist argument.

The following discussion of consequentialist argument looks firstly

at its use in the context of the latter type of situation, as this is the main context in which MacCormick discusses it; following that examples of its application in the former type of situation will be given.

Basically a judge decides, when there is dispute over the validity of a principle, whether the consequences which would ensue if the principle were adopted, are desirable (or acceptable) or not. He is faced with a specific context in which to make that assessment (i.e. the instant case before him), however he must bear in mind that the principle will apply to future analogous scenarios, and therefore he must speculate on such scenarios, and make his evaluation in the light of them as well.

MacCormick notes four potentially competing criteria to be evaluated when making this decision,⁵³ although he says there may be others (i.e. these constitute a minimum) which are as follows. The concept of 'corrective justice', i.e. (paraphrasing MacCormick) that someone who suffers wrong because of the fault of another should be compensated by that other.⁵⁴ The question of what is in the 'public interest', e.g., in the context of *Donoghue v. Stevenson*, should the manufacturer of a product which may cause harm always be liable if it does, or would the price increase which would result if he had to insure himself against any possible action be contrary to the public interest in cheapness of goods? The criterion of 'legal expediency', or 'convenience', which is sometimes advanced to show that a claim which might be meritorious on the ground of 'corrective justice' should not be allowed as it would result in the courts being overwhelmed by a flood of similar claims. 'Common sense', i.e. what would people in general think was the right result. This last criterion reflects the view that, according to MacCormick, judges try to assess and reflect contemporary positive morality (e.g. the current social climate) when making their decision, and as such it lends support to my view that there is validity in someone other than a judge attempting to anticipate a judge's evaluation of the principle's acceptability. A recent example of law making in England (referred to earlier), seems to emphasise the importance of the contemporary positive morality ground of evaluation (i.e. the importance of reflecting the public view). The Lords decided that the importance of reflecting current social attitudes should take priority over the

⁵³ ref. 50, pp. 110-111, 114

⁵⁴ ref. 50, p. 111

established law on marital rape, even though there was no controversy over what the law said. (The English cases were consistent and contrary to the decision the Lords made, though they (the Lords) did find support in the Scottish cases.)

Examples of the sort of analogous scenarios which an expert would have to consider, when deciding on consequentialist grounds whether to support or deny the existence of a principle, can be found in *Donoghue v. Stevenson*. The facts of the case were as follows. Mrs. Donoghue, the pursuer, drank some of the contents of an opaque bottle of Stevenson's ginger beer bought by a friend. She then discovered in the remainder of the contents the remnants of a decomposing snail, and this caused her to suffer gastroenteritis and nervous shock. She averred the presence of the snail in the ginger beer was due to the failure of the manufacturer to take reasonable care in the preparation and bottling of the ginger beer. The principle at issue was whether the manufacturer of a product owed a duty, independent of contract, to the ultimate consumer of the product, to take reasonable care in its manufacture.

As well as considering whether (on the basis of at least the four criteria noted), the consequences of the principle existing given the instant facts were desirable, it was also necessary to contemplate other, possible future scenarios, which it would cover. For example, Lord Tomlin suggested that everyone injured in a train accident which was caused by a defective axle, which was the result of the manufacturer's negligence, would have an action against him. Although, on the one hand, it would seem just that the manufacturer should be accountable to all those harmed due to his negligence, other consequences, it was suggested, were undesirable. For example, the manufacturer might go bankrupt, because of the possibility of having to meet so many unforeseeable (because non contractual) claims, alternatively the cost of insuring against many claims might make the business itself not worth conducting. Additionally there is the above noted inconvenience argument of the court being overburdened by a flood of similar claims.

This scenario was quite different from the actual situation existing in the case, and other scenarios suggested, such as arsenic in bread. Consequently it illustrates that the expert would need to consider scenarios apparently unrelated to each other, which the proposed principle would cover, when considering the desirability of that

principle, but this is something any lawyer should be able to do (and the expert would of course be a lawyer).

4.7 Conflicting Principles and Consequentialist Argument

An example of the application of consequentialist argument when two principles conflict can be seen in *Read v. J. Lyons & Co. Ltd.* 1947 A.C. 156. Two principles, of the law of tort (i.e. delict), were potentially applicable, i.e. 'a man acts at his peril', and 'a man is not liable unless he is to blame'. The appellant claimed the respondent was liable in damages for injury without proof of fault (i.e. negligence). In other words she alleged strict liability, founding her claim on *Rylands v. Fletcher* (1868) L.R. 3 H.L. 330, which applied the first of the above principles. The Lords denied this. e.g. Lord Porter said that the (strict) liability (of *Rylands*) alleged 'is an extension of the general rule, and, in my view, it is undesirable to extend it further.' MacCormick says of this comment that 'The passage ... explicitly acknowledges that in such a case, where there are competing principles which could provide the necessary legal support for a decision either way, the final choice between them must be based upon an evaluation of which general principle it is desirable to follow in the type of case in question.'

Lord MacMillan, also in *Read*, said that the principle of no fault liability (i.e. strict liability) for personal injury should not be extended to this case. He referred to Holdsworth's 'History of English Law' as tracing the evolution of the relevant law 'from the principle that every man acts at his peril and is liable for all the consequences of his acts, to the principle that a man's freedom of action is subject only to the obligation not to infringe on any duty of care which he owes to others.' From which he inferred that 'as the law now stands an allegation of negligence is in general essential to the relevancy of an action of reparation for personal injuries.' This passage further supports the earlier stated view that legal trends, and therefore principles themselves, can be identified from the secondary sources, and that in doing so one can indeed expect to anticipate the practice of the courts.

Another example is provided by the Scottish case of *White & Carter (Councils) Ltd. v. McGregor* [1962] A.C. 413; 1962 S.C. (H.L.) 1. A contract was made between two parties whereby the appellants were

to advertise for the respondent. The respondent repudiated the contract by letter on the same day. The appellants went ahead with the advertising and sued for the full price due under the contract. The respondent claimed that the appellants were only entitled to damages for breach of contract based on loss of profit, not the full price. The court held for the appellants on the principle that a contract cannot be revoked by the unilateral repudiation of one party. As against that there was the equally well established general rule, or 'principle', that a party who has sustained damage as a result of another's breach of contract must take reasonable steps to minimise his own loss, and thus minimise the compensation recoverable. There were two equally applicable principles available, the justification of a choice between them being, according to MacCormick, a matter of evaluative argument. To show how close the decision was in this case MacCormick notes some of the arguments of counsel for the respondent which were adopted by the dissentient Lords Morton and Keith, and gave some embarrassment to the majority. It is worth reciting one of the arguments here as it illustrates the sort of hypothetical situation which an expert would have to consider when contemplating the desirability of the adoption of a particular principle, and because it illustrates that this desirability evaluation is, as MacCormick states, irreducibly subjective, because although the scenarios seemed very convincing arguments for the respondent, they were in fact rejected by the majority.

For example, it was argued that it would be contrary to public policy if an expert was employed by a company to go abroad and prepare a report, and the company repudiated the contract before anything had been done, if then he was still entitled to go ahead and prepare the report, because in such a case he could extort in settlement far more than reasonable compensation. (i.e. He could threaten to go ahead and prepare the unwanted report unless the company 'paid him off'.)

4.8 Consequentialist Argument and Artificial Intelligence

The need to understand, for the purpose of expert systems development, how judges reach their decisions in areas of uncertain law is acknowledged by two long term researchers in AI and law, Donald Berman and Carole Hafner. They observe that "lawyers do not make decisions based on what the rules 'say'. Rather, they base their decisions on how decision makers will apply those rules to specific

factual situation.”⁵⁵ What MacCormick offers, in his theory of consequentialist argument, is an explanation of how decision makers apply the rules, in the sense of how they decide which ‘rules’ to apply.

Unfortunately current AI and law techniques cannot assist in the task of anticipating a judge’s assessment of the desirability of the consequences of a particular principle. The best that can be done is, as indicated in the section before last (section 4.6), for the expert to make the assessment, and to use a certainty factor to reflect his confidence in both this desirability (or likelihood of adoption on consequentialist grounds) assessment, and his appreciation of the general legal, or judicial, trend regarding law making in the domain. (For example, based on the desirability of the two principles at issue in *Read*, and the legal trend identifiable in Holdsworth, the expert would (hopefully) attach a higher certainty factor to the ‘no liability without fault’ principle than to the ‘strict liability’ principle.) Ideally the expert should, in addition, take into account such things as the social and political background of the judges who will hear the cases, and reflect this also in the certainty factor. However as he would not know who will hear a particular case, this ‘personal biases’ factor would have to be left to the user to assess. As already discussed, the above noted certainty factor would constitute part of an overall certainty value attached to each version of a particular principle, the other component reflecting the expert’s certainty that the level of generality of the version would be acceptable. (Of course this other component would not be necessary for well established principles, such as those at issue in *Read*, but it would for ‘new’ principles the expert had framed.)

Current AI and law techniques might, however, be of assistance in another task which is an integral part of MacCormick’s theory. i.e. Establishing how the domain cases which are arguably covered by a proposed principle can be reconciled with it. This possibility is discussed next.

⁵⁵ Berman, D. H. & Hafner, C. D. The Potential of Artificial Intelligence to Help Solve the Crisis in Our Legal System. *Commun. ACM*. Aug. 1989, Vol.32, No.8, p. 3

Chapter 5

Principles and Relevant Cases

5.1 Introduction

So far I have stated that in order to build an expert system which accords with MacCormick's theory of hard case resolution it will be necessary for an expert to identify or formulate the principles which a court may hold applicable, and for him to forecast the probability of them being adopted, on the basis of an evaluation of the desirability of their consequences, knowledge of trends in the domain law, and the levels of generality at which they are formed. I have found support for the assertion that the expert can formulate potential principles, by identifying trends, in the work of J C Smith and his colleagues at UBC.

I have stated that the expert's probability assessments would take the form of certainty factors attached to the principles, and on this matter it might be noted that the Nervous Shock Advisor attaches percentage certainty factors to its opinions,⁵⁶ but it is not clear what particular criteria are evaluated in determining these. MacCormick's theory provides guidance on this determination. I now wish to examine the extent to which computational techniques can be used to establish that the principles that the expert proposes are consistent with the domain cases, in a manner which mirrors the way MacCormick says judges perform this task.

5.2 Reconciling Principles with Case Law

According to MacCormick's theory of how judges formulate a 'new' principle, one of the major tasks they undertake is that of reconciling the principle with (or, one might say, justifying the principle in terms of), the relevant case law. It is not enough simply for the principle to be desirable on consequentialist grounds, it is also necessary to find support for it in the underlying case law, and any relevant cases which are unfavourable to it must be explained or

⁵⁶ ref. 55, p. 2

distinguished. Therefore I will suggest how an expert, assisted by AI techniques, might reconcile and refine a 'rough' principle on manufacturer's liability (possibly framed as a result of identifying trends through deep structure analysis, or simply speculated through his, i.e. the expert's, familiarity with the domain law), in a way that mirrors the reconciliation performed by Lord Atkin who formulated the 'neighbour principle' in *Donoghue v. Stevenson*. I use this example because the formulation of this principle, according to MacCormick, exemplifies the way judges formulate new, or articulate unexpressed, principles, in a hard case in any domain.

This reconciliation task would obviously precede the actual building of the expert system in which the principle was to be employed. Therefore for later convenience, and to emphasise the distinction between the use of AI techniques when building the system, and when using it, one might refer to the program, or programs collectively, which would implement the techniques discussed in the following pages, as constituting a 'principle-cases reconciliation system'; although some of the programs might also be used in the expert system itself.

5.2.1 Identifying Relevant Cases

As a starting point, as suggested, the expert would have framed a rough principle, based on his knowledge of the cases. For example, from his knowledge of the cases on manufacturers' liability, he might surmise the following. That as each case considered whether manufacturers of a particular type of defective product were liable to the ultimate user if harmed by it, it might be possible to generalise and try to establish a general proposition of law that "any manufacturer of a defective product is liable to the ultimate user if harmed by it". He would know that the likelihood of the court accepting this general principle would increase the more the individual cases found for the product user (i.e. the plaintiff), and if he could distinguish and explain those that did not.

It would be necessary for the facts of the domain cases to be represented in a formalism suitable for CBR to be performed on them. More detail on that will be given in due course. However, it would also be necessary for each case represented to contain an account of the main legal points at issue in that case, e.g. "manufacturers'

liability to consumer". This information could be manually extracted from the headnotes contained in the case reports of each decided case. Alternatively, rather than using the actual headnotes contained in case reports, significant legal phrases, such as the above, might be automatically identified from the full case texts using the text analysis mechanism employed by Gelbart and Smith in their FLEXICON system.⁵⁷ This mechanism recognises complex legal phrases, based on approximate matches of words and word ordering,⁵⁸ and incorporates them into structured case profiles.⁵⁹ (More on what FLEXICON does will be noted shortly.)

Using the accounts, or summaries, of legal points at issue a program could identify potentially relevant cases, either through simple keyword search, if the summaries had been manually constructed, or by employing Gelbart and Smith's more sophisticated legal phrase, or concept, searching, if FLEXICON profiles, (plus a concept index) had been generated. Gelbart and Smith themselves specifically note one of the purposes of FLEXICON as being to allow a user to search for cases that share common legal issues.⁶⁰ Some of the cases on manufacturers' liability which would have been identifiable in this way at the time of *Donoghue v. Stevenson* will be discussed shortly.

5.2.2 Distinguishing Factually Dissimilar, Unfavourable Cases

The next task would be to identify which of the potentially relevant cases held for the user, and thus could be advanced as supporting the principle. This should be a fairly straightforward text analysis task, as it would be noted in whose favour the case was decided in the case headnote.

It would then be necessary to explain and distinguish the unfavourable relevant cases (i.e. those which did not hold for the user). The first way in which this might be attempted would require the facts of each case to have been represented in a CBR formalism, such as GREBE's (discussed in section 2.4 and illustrated shortly),

⁵⁷ Gelbart, D. & Smith, J. C. Beyond Boolean Search: FLEXICON, A Legal Text-Based Intelligent System *ICAIL-91*, pp. 225-234

⁵⁸ ref. 57, p. 226

⁵⁹ ref. 57, p. 225

⁶⁰ ref. 57, p. 226

which would allow a case based reasoner to identify which unfavourable cases were factually similar, and which dissimilar, to the favourable cases. Those which were dissimilar could obviously be argued as (factually) distinguishable from the favourable ones, and their identification would prompt the expert to refine his principle so that it did not embrace them. For example, moving away, temporarily, from the consideration of a principle covering the liability of a manufacturer of a defective product, suppose the principle of 'strict liability of a manufacturer of a dangerous substance to those harmed by it' was being considered after the cases of *Rylands* and *Read* had occurred. The case of *Rylands* would support the principle, but the case of *Read*, which considered *Rylands*, but in which the principle of 'no liability without fault' was given priority, would undermine it. However, the reason why a different decision could be given in *Read* was because of a factual difference, i.e. the harm suffered occurred on the manufacturer's premises in *Read*, but outwith it in *Rylands*. So, assuming a case based reasoner would detect such a difference (and this would of course depend on the level of factual detail at which the cases were represented, a matter which is considered later, see chapter 7), the expert would add the rider "where the harm occurs outwith the manufacturer's premises" to the 'strict liability' principle, so that *Read* fell outwith its scope.

At this point some comments should be made about the actual process of casting cases into a CBR formalism, and the promise that the FLEXICON project holds. FLEXICON incorporates an attempt to address one of the fundamental bottlenecks concerning the computer manipulation of cases, i.e., the fact that the representation of case information in the required formalism is a time consuming, manual task. Gelbart and Smith note that case reports can now easily be provided in electronic form, either directly from the courts, or by scanning hard copy documents, and they state that their system can automatically build up a profile of a case in terms of the relationships between four parameters: the legal concepts involved, cases and legislation cited, and the critical facts.⁶¹

The ability of FLEXICON to identify the critical factual elements of a case suggests that it could be used to automatically obtain the components required for a GREBE like representation of a case. (As noted in section 2.4, for the purpose of CBR in GREBE a structured

⁶¹ ref. 57, p. 225

representation of cases is used,⁶² i.e., each case is represented in terms of its objects, e.g. 'tangible entities', and 'legal relations'; and the relationships among them. Examples are given shortly.) However, comparison with manual analyses of cases would need to be conducted to determine how successful FLEXICON was in identifying all the relevant facts, and it is not clear whether FLEXICON actually identifies the relationships among facts themselves.

It should be noted that the authors themselves recognise the potential of FLEXICON's automatic text analysis capability to provide a front end to a case based reasoner. '... we wish to explore to what extent case-based reasoning can also be automated and incorporated into the existing system to produce case retrieval as well as expert predictive capability without the tremendous manual effort required to construct traditional advisory systems.'⁶³

5.2.3 Reconciling Factually Indistinguishable, Unfavourable Cases

Returning to the issue of reconciling the principle with the case law, some method other than distinguishing on factual grounds would be necessary to reconcile cases identified as factually similar to those supporting the principle, but unfavourable in their outcome.

5.2.3.1 Wrongly Decided

One way in which a case which is factually similar to one supporting the principle, but unfavourable, might be explained, is by claiming it was wrongly decided. To do this it must be shown to be (arguably) inconsistent with the previous case law. For example, the following two cases would be identified as similar, one favourable, the other unfavourable. (Two more cases will be introduced a bit later. These four cases are just a sample of the many cases that were relevant to the issue of manufacturers' liability at the time of *Donoghue v. Stevenson*, and were considered in that case.) I will outline the facts of these cases to show how I believe they might partly have been expressed, and identified as similar, using a GREBE like formalism.

⁶² Branting, L. K. Building Explanations from Rules and Structured Cases. *IJMMS-1*, p. 808

⁶³ ref. 57, p. 232

Mullen v. Barr & Co. 1929 S.C. 461. The facts were as follows. A bottle of ginger beer was bought by Mrs. Mullen from a retailer. Unknown to her it contained a dead mouse. She gave some of it to her children to drink. They were poisoned by drinking the contents. Held, by majority, no duty was owed by the manufacturers to a consumer who had not contracted with them, and that in any event negligence had not been proved.

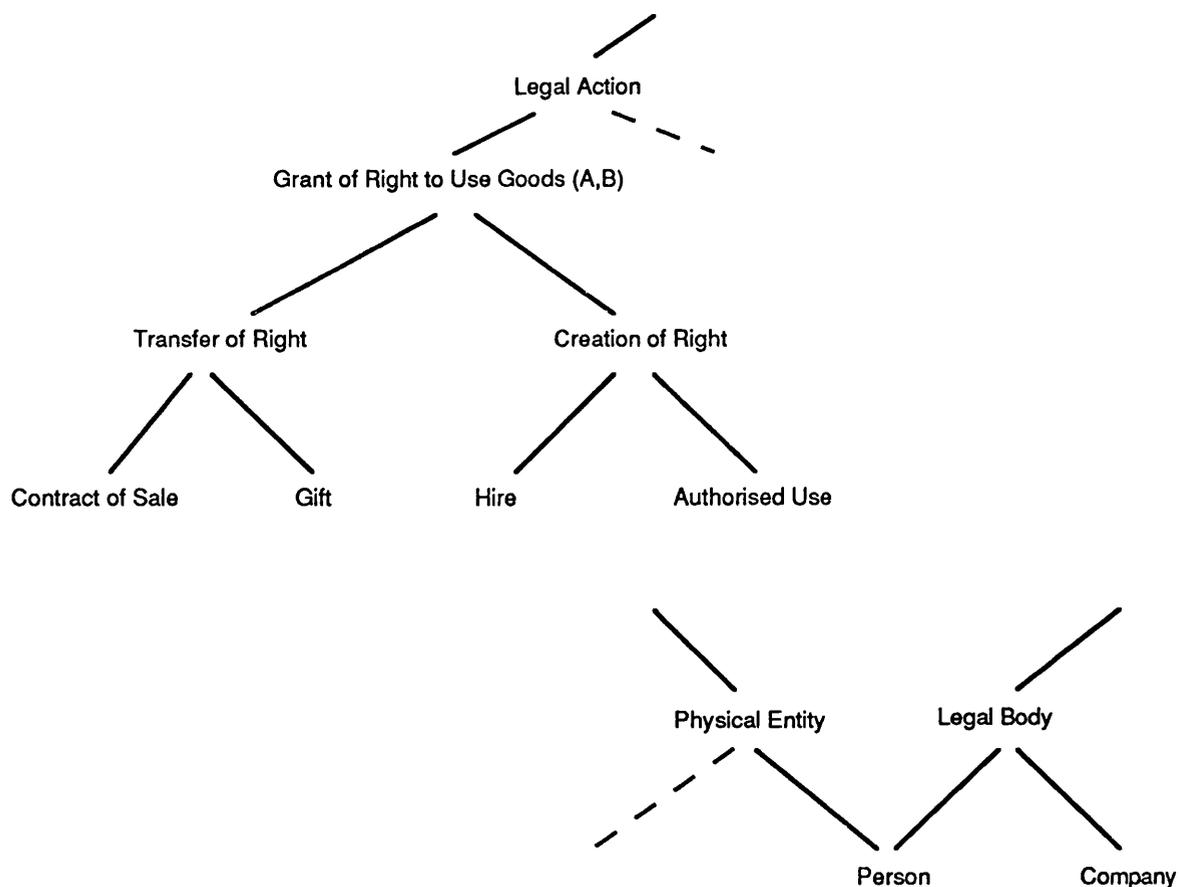
This was the most recent case on manufacturers' liability at the time of *Donoghue v. Stevenson*.

George v. Skivington (1869) L.R. 5 Exch. 1. The facts were as follows. A husband bought shampoo from a chemist (the manufacturer) for his wife. His wife used it and was injured. Held, the manufacturer/chemist was liable for unskillfulness and negligence in manufacture to those for whom the product was purchased.

If the cases had been represented in a GREBE like format, all objects in the cases, and relationships among them, would be expressly represented.

In order to produce these representations the system would contain a (GREBE like) hierarchy or taxonomy of domain objects. The taxonomy/hierarchy would be implemented as a semantic network; the objects would be represented as frames in the network, and relationships among them would be represented as slots in the frames.⁶⁴ Such a structure would include within it the simple GFH (discussed earlier, in section 2.4), and it might be termed the domain hierarchy.

⁶⁴ ref.62, pp. 808-9



A and B are Legal Bodies

Figure 3. Example of possible parts of the domain hierarchy.

The 'Grant of Right to Use Goods (A,B)' object would subsume particular Legal Actions, in which Legal Body A bestows the right to use an item, or items, (not represented in the diagram), on Legal Body B. A and B would be slots in a frame representing a particular Grant of Right, thereby expressing the relationship between that Grant, and the Legal Bodies which are parties to it.

Mullen (unfavourable), and *George* (favourable), would be represented in terms of these objects as follows. (Each 'T' is an instance of the particular Legal Action 'Transfer of Right').

e.g.1 *Mullen v. Barr & Co*

T1(A,B) T2(B,C) T3(C,D)

A = manufacturer/defender i.e. Barr
B = retailer
C = purchaser/pursuer's wife i.e. Mrs. Mullen
D = injured parties i.e. the children
T1 = Sale i.e. T1 is a specific type of Transfer (and Grant) of Right;
a Sale.
T2 = Sale
T3 = Gift i.e. T3 is a specific type of Transfer (and Grant) of Right;
a Gift.

e.g.2 *George v. Skivington*

T1(A,B) T2(B,C)

A = manufacturer/defender i.e. Skivington
B = purchaser/co-pursuer i.e. husband
C = injured party/co-pursuer i.e. wife
T1 = Sale
T2 = Gift

When attempting to match cases, GREBE attempts to find the closest match, both factually (i.e. in terms of what objects are involved), and structurally (i.e. in terms of how the objects are related to each other). This allows GREBE to attempt to improve a match, once it has established a basic degree of similarity by identifying shared facts, by investigating different variations of the structure of the NFS, and thus identifying the best structural mapping.⁶⁵ Presumably this is intended to simulate the practice of a lawyer whereby, when faced with a particular unfavourable precedent, he attempts to identify an alternative, plausible reading, or construction, of the established facts or events of the NFS, so that he can argue that what happened in the NFS was not the same as what happened in the precedent, even though the facts involved in the two cases were the same. (This is what some researchers might refer to as a 'legal heuristic'.) If it is possible to improve the match between cases in this way, it is presumably also possible to do the opposite and attempt to worsen the match between factually similar cases, by trying to identify incompatible structures. This would be something to attempt at the end of the stage just discussed, (i.e. that of trying to distinguish relevant unfavourable cases by comparing their facts), once an

⁶⁵ ref. 62, p. 817

attempt to distinguish them simply on the basis of factual dissimilarity had failed. However, this technique would be of little use for the cases we are considering, due to the simplicity of the structures they involve.

Structurally both these cases constitute a simple chain of Transfers, transferring the right to use a product from a manufacturer to an ultimate consumer. Factually both contain a non-contractual transfer in their chain of Transfers. (The matching algorithm might check for the presence of such a transfer in a chain, as the presence of a non-contractual transfer in one case and not in another could constitute a significant difference between the two. For example, it might be that certain contractual rights can be transferred through a series of contracts.) The only actual difference between the two cases would be minor, i.e. an extra Transfer, and an extra Person, involved in *Mullen*.

CBR would thus identify these cases as very closely matched, factually and structurally, and therefore equally relevant to the desired principle. So how could the principle-cases reconciliation system explain away *Mullen*, while keeping *George*? Especially as, on the basis of other considerations, *Mullen* might be regarded as the most compelling authority. For example, it was decided more recently than *George*, and, as regards a NFS arising in Scotland (which was the country in which the facts of *Donoghue v. Stevenson* took place), it was decided in the same jurisdiction, and at a high level of court.

To do this it would be necessary to establish that the majority of cases prior to *Mullen* were reconcilable with, or favourable to, the principle. Achieving this would support the argument that the court in *Mullen* had failed to comprehensively consider all the relevant case law, or that it had drawn the wrong conclusions from its consideration, and that consequently the case was wrongly decided.

Thus, when conducting its examination of the relevant domain case law, the principle-cases reconciliation system would note if a recent unfavourable case, which, on the grounds of age, jurisdiction, and level of court, could be argued as the most compelling authority for a possible NFS, was in conflict with the majority of previous factually similar precedents.

For example, another case which would be identified at the outset as

considering the legal issues relevant to the proposed principle on manufacturers' liability, and then as being supportive of that principle, and also as being factually similar to *George v. Skivington*, and *Mullen*, would be *Hawkins v. Smith* (1896) 12 T.L.R. 532. Here the defendant, the manufacturer of a defective sack, was held liable for injury sustained by a dock worker using the sack, he (the defendant) knowing the use to which the sack was intended to be put, on the ground of negligence. A very close factual and structural match could be identified between this and the other two cases in spite of some differences; these being that rights were Created, rather than Transferred, in *Hawkins*.

These differences would not have been significant for the principle proposed because the important thing, in this, as in the other cases, was that the injured party was a legitimate user of the defective product. This the system would establish as a result of inferring, from the specific types of Grant of Right to Use present in the precedent, a sequence of Grants from the manufacturer to the injured party.

However, in order to be able to make such inferences, (and as a consequence to be able to treat a case involving a Hire as as on point as one involving a Sale), the system would need to know that, for the principle being considered, none of the specific terms subsumed by the generic term Grant was to be regarded as remote from it, i.e. they and Grant were to be treated as synonyms. This in turn would require that when the expert described his proposed 'rough' principle to the principle-cases reconciliation system he did so in some degree of detail, so that the system could establish exactly what groups of terms in the hierarchy it could regard as interchangeable when attempting to identify relevant cases. Thus when the expert described the principle under consideration (noted at the start of this section) to the system, it would question whether the 'user' concerned was a legitimate user, i.e. one who had a right to use the product. On establishing this to be the case, the system would know that it was necessary to establish an unbroken chain of Grants of the right to use the product, starting with the manufacturer and ending with the injured user, for a precedent to be factually relevant to (i.e. subsumable under) the principle. Furthermore it would know that such a precedent would be relevant no matter how many individual Grants there were in the chain (hence the unimportance of the extra Grant and Legal Body in *Mullen* compared with *George*), or the type of

Grants involved.

The system would also question the expert about the type of 'user' involved in the principle, establish it to be a Legal Body, and thus know that the Grants in the precedents could involve either a Person or a Company. Thus cases differing only as regards type of Grants or Legal Bodies involved, or number of Grants involved in a chain, would be treated as equally on point with regard to the principle. *Hawkins* would therefore be matched to *George* and *Mullen*.

e.g. 3 *Hawkins v. Smith*

C1(A,B) C2(B,C) C3(C,D)

A = manufacturer/defendant i.e. Smith

B = consignees

C = dock company

D = dock worker/injured party/pursuer i.e. Hawkins

C1 = Hire i.e. C1 is a specific type of Creation (and Grant) of Right; a Hire.

C2 = Authorised Use i.e. C2 is a specific type of Creation (and Grant) of Right; an Authorised Use.

C3 = Authorised Use

Given the bases of the judgments in the two similar, favourable cases, *George v. Skivington* and *Hawkins v. Smith*, the expert would note that in both it was necessary for the plaintiff to establish negligence to succeed, and also that in *Mullen* negligence was not proved, so he could refine his general principle to produce a version which asserted that the manufacturer would be liable if negligent in his production. This would have a higher certainty factor than the original version. However the argument this refinement would sustain would be in addition to, not instead of, the argument that *Mullen* was simply wrongly decided given the contrary weight of the previous case law.

It is important to note that what the reconciliation system would be doing here would be attempting to construct an argument which, depending on the NFS presented to the expert system, it might not be possible to use. For example, as regards an attempt to explain away *Mullen*, if the court in which the product user was bringing the action was in the same jurisdiction, and at a lower or equal position in the

court hierarchy, than *Mullen*, then the proposed principle would not be allowed, because the court would (normally) regard itself as being bound by the doctrine of *stare decisis*, and not free to assert any 'new' principle.

This situation would be accounted for in the expert system itself by not making the principle available if such was the procedural context of the user's case. (It would of course be a simple task for the expert system to compare the jurisdiction, and court position, of decided cases with that of the NFS, such information, plus the date of the decision, being part of the ancillary information attached to the factual representation of each recorded case.) If the possibility of asserting the principle was not foreclosed by legal doctrine, then the opportunity would exist to undermine *Mullen*, and the expert system would make the principle available.

5.2.3.2 Basis of Decision

Another way in which a legally relevant, factually similar, but unfavourable case, might be reconciled with a principle, is by arguing that the ground of the decision in the case is (or refining the principle so that it can be argued that it is), in fact compatible with the principle. For example, in *Winterbottom v. Wright* (1842) 10 M. & W. 109, a remedy was denied to the defendant on the ground that the claim being made was flawed in law. In this case the defendant contracted with the Postmaster-General to provide mail-coaches, and to keep the mails in safe condition. Atkinson contracted with the Postmaster-General to convey the road mail-coach from Hartford to Holyhead. The plaintiff was employed by Atkinson to drive the mail-coach. The mail-coach was defective, gave way, and the plaintiff was injured. The only right to recovery from the manufacturer alleged was one arising out of contract (to put the coach in good repair). i.e. It was claimed that because the manufacturer had been negligent in his duties under the contract of repair, as a result of which the plaintiff was injured, he (the plaintiff) was entitled to recover from the manufacturer. It was held that only a party to the contract had rights arising from it.

e.g. 4 *Winterbottom v. Wright*

C1(A,B) C2(B,C) C3(C,D)

A = manufacturer/defendant i.e. Wright
B = Postmaster-General
C = Atkinson
D = plaintiff/injured party i.e. Winterbottom
C1 = Hire
C2 = Authorised Use
C3 = Authorised Use

Part of the headnote information of each case is a statement of the reason for the decision. e.g. Judgment for the defender because the pursuer failed to prove what was required (e.g. negligence) in order to establish his case, or (irrespective of whether he could or could not prove negligence) because the legal basis of the claim itself was invalid. This information, along with (as already discussed) the legal issues considered, and a CBR representation of the facts involved, would be incorporated into case profiles of each domain case. Thus when the system identified *Winterbottom v. Wright* as relevant, the expert would note that the reason for the decision was that in *Winterbottom* the claim was founded, erroneously, on the assertion that a right arose out of a contract to which the pursuer was not a party, compare it with *George v. Skivington*, in which it was alleged simply that a duty of care is owed by a manufacturer to the person for whom the product was intended, and reconcile *Winterbottom* with the principle by stating the principle to be one of delict (or tort) rather than contract.

The foregoing analysis of how these four cases could be reconciled with the principle which, on consequentialist grounds, might be considered desirable, reflects what in fact happened in the case of *Donoghue v. Stevenson*. The facts in *Donoghue v. Stevenson* were virtually identical to *Mullen*, and when the Court of Session considered *Donoghue v. Stevenson* it followed the decision in *Mullen*, because of the doctrine of *stare decisis*. The case went to the House of Lords, which was not bound by *Mullen*, but the majority, who wanted to find for Mrs. Donoghue, still had to justify a departure from it if they wanted to adopt a 'new' principle. This they did by noting its inconsistency with the underlying case law. Thus when asserting the neighbour principle the Lords discounted *Mullen*, explained *Winterbottom v. Wright*, and advanced *George v. Skivington* and *Hawkins v. Smith* in support of it (i.e. the principle).

5.3 Using the Same Cases to Deny a Proposed Principle

However, although those cases were used by the majority to support the principle, they could just as easily have been used to deny it, by noting that *Donoghue v. Stevenson* was a Scottish case, and that *Mullen* was the most on point Scottish authority, and that the other authorities, being English cases, were distinguishable. This fact serves to remind us that once deductive reasoning has failed to produce a definite answer, and reasoning by analogy has not returned compelling arguments, the court can usually manipulate the relevant law to support any principle it wishes to assert, or to support its denial. (This principle's desirability having been determined, according to MacCormick's theory, through the application of consequentialist argument.) Hence the importance, for the purpose of building an expert system, of the expert's initial task of anticipating the principle which the courts will want to establish, as the case law relevant to a purported principle may often be capable of both supporting and denying it.

5.4 Obiter Dicta

When considering a case, judges may also rule on related points of law which are not essential to a decision on the point at issue. It is accepted that the distinction between the part of the judgment essential to the decision on the matter at issue (the *ratio decidendi*), and the ancillary comments (*obiter dicta*), is often not clear. Hence, if the case headnotes were unclear on this matter, or not relied on, the expertise of the person (or sophistication of the system), interpreting the cases for the purpose of computer representation would be vital, and the resulting decisions, as to what is *ratio*, and what *obiter*, would, for some cases, necessarily be contentious.

In the case of *Winterbottom v. Wright* it seems clear that the ground on which the pursuer sought to establish a right of recovery against the defendant was his (the defendant's) breach of the contract that he (the defendant) had with the Postmaster-General. The majority held that only those who are parties to a contract have rights under it, so the pursuer could not recover. They also commented that he would have no right under tort (delict), i.e. independently of a contract, either. These comments could be included in the case profile for

Winterbottom, but noted as *obiter dicta*.

As already noted, the system could identify *Hawkins v. Smith* as being similar to *George v. Skivington*, and supporting the principle, and *Winterbottom v. Wright* as being similar to *George v. Skivington*, and unfavourable to the principle. The expert could then reconcile *Winterbottom* with the principle by refining the principle, i.e. noting that it should be advanced as a right based arising from delict, not contract. However, comments made *obiter* do carry weight if made by respected judges, so if, as just suggested, comments made *obiter* were included in the case profiles of cases, the expert would have to (further) modify the certainty of his principle being adopted, even if based specifically in delict (tort), according to how much authority he thought would, in later cases, be attributed to the comments in *Winterbottom*.

On the matter of allocating authority to *obiter dicta* it should be noted that if comments which are apparently *obiter* are supportive of the ruling a judge wishes to make, then he may claim that they are in fact binding, (implying that he does not regard them as *obiter*). This in fact happened in *Donoghue v. Stevenson* where one of the minority, Lord Buckmaster, specifically adopted the comments in *Winterbottom* as a ground for denying the neighbour principle. Some may say that this illustrates the reality of the uncertainty which may sometimes exist over what is *ratio*, and what is *obiter*, and consequentially the necessary contentiousness (or even futility), of any attempt to specifically delimit them. Others may say that it just illustrates that a judge may try to cloud the distinction between comments made *obiter*, and the *ratio decidendi*, if it suits his purpose. Whichever view is true, it certainly does emphasise the importance, and difficulty, of the expert's task of correctly anticipating the essentially subjective consequentialist arguments of the judges, concerning whether to support or deny a particular principle. This is emphasised because the only reason that Lord Buckmaster advanced the comments made *obiter* in *Winterbottom* was because his consequentialist evaluation of the proposed principle had led him to wish to deny it.

Therefore given the proposal of a principle on manufacturers' liability the expert would have to foresee that there were consequentialist arguments against adopting the principle which might be advanced, and also that the denial of the principle could be

supported by the case law. He would then have to weigh these considerations against those which supported the principle, and reflect his deliberations in the certainty factor attached to it.

Of course remarks made *obiter* in cases such as *Winterbottom* would have to be weighed in the context of the relevant cases as a whole. Thus even if made by an important judge their authority would be diminished the more similar cases were decided on the opposite view of the law, particularly if such decisions were made at the highest level, and in recent years. e.g. *Oliver v. Sadler & Co.* 1929 S.C. (H.L.) 94; [1929] A.C. 584. A firm of stevedores and a portorage company were employed to unload a cargo of bags of maize. The stevedores used their own rope slings to raise the cargo to the ship's deck, and permitted the portorage company to use those slings in transporting the cargo to the quay. A sling broke while bags were being transported to the quay and the bags it contained fell and killed an employee of the portorage company. Held, the stevedore company owed a duty to the porters to see that the slings were fit for use, which they failed to discharge.

5.5 Unreconcilable Cases

In any domain there may be cases which simply cannot be reconciled with a proposed principle. Smith discovered this when building the Nervous Shock Advisor. The system would identify such contrary cases (along with the other legally relevant, factually similar cases, which were reconcilable), and the expert would have to modify the certainty factor attached to the principle according to how much importance he thought would be ascribed to them. For example, he might decide they were simply wrongly decided (which seems to be Smith's view on the contrary cases in the domain of nervous shock), and attach very little weight to them, on the basis that the court would do all it could to avoid being bound by them.

5.6 Procedural Context

A further refinement to the task of reconciling a principle with the case law which may be proposed here is that the expert should consider the procedural, or judicial, contexts in which he contemplates his principle being advanced, and determine the

'precedential pedigree' of the relevant unreconcilable contrary precedents (and perhaps also that of the favourable ones), as part of the task of deciding how much authority he thought would be accorded to these cases, and its effect on the principle's chances of succeeding. The precedential pedigree would be determined by weighting the formal parameters of jurisdiction, level of court, and age. (On weighting these parameters see earlier discussion of the Malicious Prosecution Consultant in section 2.2. As evidence of the success of the Malicious Prosecution Consultant's weighting scheme, it might be noted that Deedman and Smith are intending to incorporate it into the Nervous Shock Advisor.⁶⁶) The consideration would result in (for example), lessening the certainty of a principle in relation to a particular procedural context (i.e. court and jurisdiction) in which that principle might be advanced, the greater the authority of a particular unfavourable, unreconcilable precedent relative to that procedural context.

The consequence of the expert undertaking such a consideration is that he would have an extra factor to take into account when determining the certainty factors for each version of the principle. (The versions reflecting different levels of generality.) Moreover, there would no longer be just one certainty factor associated with a particular version, rather it would vary depending on the procedural context of the the NFS presented to the expert system. For example, if the NFS concerned an action being raised in the House of Lords, on appeal from the Court of Session, and the only unreconcilable case occurred in the Court of Exchequer in England, then the certainty of the principle would be higher than if the action was being raised in the House of Lords, on appeal from the English Court of Appeal, and the only unreconcilable case was another Lords case, heard on appeal from the Court of Session.

If the certainty of the principle which was determined, when the expert system was informed of the procedural context of the NFS, was low enough, then it might be wise to regard the principle as completely untenable, and not make it available for that NFS. This would cover the aforementioned suggestion that it would have been pointless to have tried to advance a principle on manufacturers' liability, which was contrary to the decision in *Mullen*, in the Court of Session, because the court would have been bound by the doctrine

⁶⁶ Deedman, C. & Smith, J. C. The Nervous Shock Advisor: A Legal Expert System in Case-based Law. (*paper*) p. 19

of *stare decisis*.

However, as regards a case arising in the House of Lords, and there existing contrary Lords decisions, it might still be advisable to make the principle available, as the Lords do not regard themselves as bound by their own decisions, though they believe there must be good reason for a departure. Thus if the expert's consequentialist evaluation, and knowledge of the general legal trend, suggested a very high certainty factor for a principle, he should not modify it too much in contemplation of a NFS arising in the House of Lords, even though there existed contrary Lords decisions, on the basis that they would regard the contrary decisions as being just one of their considerations (but not a decisive one), when reaching a judgment. Again support for this view can be seen in the recent House of Lords English decision on marital rape.

5.7 Parallel between Framing a Principle, and Generating a GFR

Obviously the expert's initial speculation of a 'rough' general principle, from his knowledge of the domain case law (noted earlier, at the beginning of this chapter, as the starting point in any attempt to incorporate unarticulated principles into an expert system), possibly at various levels of generality, entails a generalisation from the facts of the specific relevant cases he is aware of to more general terms. As alluded to earlier (in section 4.5), this parallels the proposal (in section 2.3), of generalising from the fact(s) of a specific precedent, to a GFR, for the purpose of arguing that a new case and the precedent are sufficiently similar for it to be necessary for the precedent's result to be followed. The difference for a principle is that the generalisation would be made from the facts of more than one precedent, and that the principle (or general rule) itself would provide the reason, or 'legal warrant', that a particular result should be given.

5.8 Lack of Distinction between 'Argument by Analogy', and 'Argument from Principle'

The similarity, between the two generalisation processes required to build into an expert system the ability to argue by analogy, using a

GFR, and the ability to argue from principle, seems to tie in with MacCormick's observation,⁶⁷ that there is no clear line between 'argument by analogy', and 'argument from principle'. For example, suppose an expert was trying to formulate a principle to cover the right of recovery of a party injured while trying to prevent the damage which would result from another party's wrongful act (or that of someone for whom the latter was responsible). In the light of the settled law, that a person injured while saving another person endangered by the wrongful act of a third party is entitled to recompense from the latter (e.g. *Woods v. Caledonian Rly.* (1886) 13 R. 1118, *Wilkinson v. Kinneil Cannel & Coking Co.* (1897) 24 R. 1001), and his knowledge of the trend in the law of reparations, the expert might suggest a principle by generalising from 'saving a person' to 'saving a person or his property'. Alternatively he might generalise from just one of the cases to produce a GFR to the same effect with a GF of course attached to the resultant generalisation.

MacCormick illustrates the similarity between 'argument by analogy' and 'argument from principle' with the case of *Steel v. Glasgow Iron and Steel Co. Ltd.* 1944 S.C. 237, in which the court 'found' the principle just suggested in the case law, but talked about analogy (which is what a GFR would support). In the *Steel* case the guard of a shunting train was killed while trying to save the property of the defender, endangered due to the carelessness of an employee of the defender. There was no directly on point authority, but the majority decided there was sufficient analogy between saving life, the situation in the precedents, and saving property, the situation in the instant case, (they said the two were different in degree, not in kind),⁶⁸ to find a principle covering both, and thus give the pursuer a remedy.

(If the expert had proposed a principle this 'difference in degree' would be reflected in the 'level of generality acceptable' part of the certainty factor attached to the principle. He would know that the level of risk to which the injured party had voluntarily exposed himself, in attempting to prevent the damage, which would be acceptable to the court, would be greater if life was involved, rather than merely property. If he had produced a GFR this difference of degree would be reflected in the GF.)

⁶⁷ *LRLT*, p. 161

⁶⁸ ref. 67, p. 162

5.9 GFRs and Principles

This similarity of the expert system construction tasks of generating a GFR, for the purpose of reasoning by analogy, and forming a principle, for reasoning from principles, plus this example from MacCormick of the potential for the court to cloud the distinction between the two processes, raises the question of whether there is any real distinction between a GFR and a principle. To answer this question it is worth noting firstly that the GFR was conceived as a way to implement one aspect of reasoning by analogy, (it can be viewed as a supplement to the HYPO type of CBR, see chapter 3, or as one aspect of GREBE like CBR, termed 'case elaboration'), and then to examine how, or whether, one can remain faithful to MacCormick's comment that there is no clear dividing line between the two reasoning process in a case like *Steel*, and justify the need to maintain the GFR, and the principle, as two distinct entities in an expert system.

If a 'right of recovery' principle had been formulated by the expert and the principle-cases reconciliation system, in the way I have outlined in this chapter, then the relevant, favourable cases would have been identified, and the possibility of arguing that they were analogous to the NFS would simply be a by-product of the fact that the principle already (factually) subsumed them, and must also (factually) subsume the NFS, in order to be advancable. Thus the reconciliation system would reflect the fact that if one is 'arguing from principle', it is always going to be possible to 'argue by analogy', as a principle is based on previous cases, and to be usable must subsume the current case. So the expert system would deal with a scenario like *Steel* by reasoning with principles, although the argument produced could also be interpreted as an argument by analogy, and that argument could be substantiated by referencing the precedents which supported the principle. The final certainty factor computed for the principle (resulting from a combination of the expert's assessments) would indicate the likelihood of the principle being accepted in a NFS, but could also be viewed as the likelihood that the relevant precedents would be considered sufficiently analogous to the NFS for the court to be bound by their outcomes.

This might suggest that by formulating principles one is forestalling the need to generate GFRs, as by identifying a principle one probably

embraces the analogical arguments which a GFR would support. However, a GFR is perceived as involving a smaller degree of generalisation than a principle, the method of generating it is simpler than what is required for a principle, its purpose is more limited, and as a consequence of these factors the chances of it providing useful advice are probably greater. As already noted, it would be generated from a single case (rule), and not require the same process of case law reconciliation as a principle, because it would not involve the same degree of generalisation. Its purpose is primarily intended to be to suggest the likelihood of a match between a particular precedent case rule, and a few possible NFSs, being accepted, (by suggesting an argument by analogy to advance where deductive reasoning has faltered), and not to provide a more tentative, broad argument, capable of embracing many possible scenarios, which is the purpose of a principle.

Consequently I would say that, as far as the type of expert system I am proposing is concerned, 'argument by analogy' can be distinct from 'argument from principle', because argument by analogy can involve the use of a GFR, and a GFR is a distinct entity from a principle, although the two are related. On the other hand, an argument from principle can always be interpreted as an argument by analogy. Finally, just to reinforce the answer, to the question posed at the start of the section, I would say that GFRs and principles are clearly distinguishable, both in the procedure required to form them, and in the use to which they would be put.

Chapter 6

Representation and Utilisation of Principles and Cases

6.1 Representing a Principle

It is proposed that the expert's refined and reconciled domain principles would be represented in the expert system simply as individual, unrelated rules, each with a certainty factor (reflecting the expert's evaluation of the criteria previously discussed), and references to profiles (see later) of the relevant cases.

6.2 Satisfying the Components of a Principle

A principle would contain a combination of any, or all, of simple facts, open textured legal phrases, and legal concepts defined (by the expert), in terms of a necessary and sufficient set of facts. For example, the principle on manufacturers' liability might contain facts and an open textured phrase, viz. 'If a manufacturer fails to take reasonable care in the production of his product and the ultimate user is harmed as a result then he (the manufacturer) is liable in damages to that user.' Principles in other domains might contain definable legal concepts, such as 'estate', or 'theft', which Smith considers. In the expert system a 'principle based reasoner', a program distinct from both the rule based reasoner and the case based reasoner, would be used to determine whether any of the principles, contained in a principles knowledge base (PKB), were satisfied by the NFS.

6.2.1 Facts

The principle based reasoner's first step in trying to satisfy a principle would be to examine the user's NFS to see if the necessary facts were present. As the user would have already described his NFS to the system for the purpose of RBR, and CBR, (see section 7.1 on NFS entry), all the relevant information would probably be available

by the time the principle base reasoner was invoked. This information would be used either for direct matching, or for inferred matching using the domain hierarchy discussed earlier, which would contain the factual terms used in the principles. e.g. The principle based reasoner would need to establish that there was a manufacturer of a product. In the NFS the relationship of 'manufactures', might already have been established between the objects 'manufacturer', and 'car'. A simple and certain inference would then be made from 'car' to 'product' to establish that there is a manufacturer who manufactures a product. (Note that the inference here would be certain because there is no semantic uncertainty involved, e.g. a 'car' is definitely a 'product'; and because the purpose of the inference is not to try to establish the likelihood that a case based rule which concerned car would be broadened to include products in general, i.e. the purpose is not to generate a GFR, but merely to use the common sense knowledge that a car is a type of product. (Although it is arguable that the certainty should not be quite one hundred percent as the court may choose to define product restrictively, i.e. give it a particular legal definition for this principle which would exclude vehicles.) The user would then be expressly questioned about any of the principle's facts whose presence or absence was not ascertainable from the NFS description.

6.2.2 Legal Concepts

Definable legal concepts would be included in the domain hierarchy, and therein refined to the level of simple, or in MacCormick's terminology 'brute',⁶⁹ facts, in accordance with both MacCormick's and Smith's theories (see section 4.2). The principle based reasoner's initial approach to the task of trying to satisfy concepts would again be to examine the NFS for satisfaction of the relevant facts, and if that was not sufficient, to question the user. The inclusion of legal concepts (or 'institutional' facts as MacCormick calls them), as well as simple facts, in the domain hierarchy, would result in a structure which would parallel closely the sophistication of GREBE's own hierarchy of objects, as can be seen from the speculative examples given in section 5.2.3.1 (see figure 3). The utility of such a domain hierarchy, both for the construction and the functioning of the proposed expert system, should hopefully now be apparent. It would be used for GREBE like CBR when trying to reconcile cases with

⁶⁹ MacCormick, D. N. Law as Institutional Fact. *Law Quarterly Review*, 1974, Vol. 90, p. 102

proposed principles before incorporating them into the system, for the generation of GFRs when doing RBR in the system, possibly for CBR when using the system if GREBE's form of CBR was found to be more suitable for the domain in question than HYPO's, and for trying to match principles to the NFS when doing PBR. This general utility of the structure suggests that it should be carefully created right at the start of building the system, before any other task was undertaken.

6.2.3 Open Textured Phrases

If the principle incorporated vague legal phrases, such as 'reasonable care', or 'negligence', (in the production of goods), then, as when such issues arose in a statute based, or case derived rule of law, more than one approach could be employed. The simplest, and arguably the least satisfactory approach is simply to ask the user if he thinks the phrase is satisfied in his NFS. However, although some researchers have criticised this approach as being unsatisfactory, it is worth noting that questioning the user about the satisfaction of such a term as 'negligence' is not apparently regarded as unsatisfactory by Deedman and Smith; this they do in the Nervous Shock Advisor. Presumably they would justify this approach by observing that the intended user of their system is a lawyer, and that as lawyers themselves they are in the best position to judge whether or not a lawyer is well enough equipped to make such an assessment. However, another open textured term, which they perhaps regard as more domain/context sensitive, they do not leave to the user to determine satisfaction. Consequently there is another approach employed in the Nervous Shock Advisor, whereby Deedman and Smith, due to their deep structure analysis, have identified a group of facts which satisfy the term 'reasonably foreseeable', as it relates to nervous shock, and the system simply questions the user about them.⁷⁰ (I would suggest that these facts could be expanded to generalisations, and coupled with GFs, to cover analogous situations in the manner suggested earlier when discussing GFRs.) This approach of Deedman and Smith seems akin to Branting's notion of reasoning with portions of precedents, in which the specific facts in a relevant case (or cases) which satisfy an open textured phrase have been identified, and a case based reasoner attempts to establish the

⁷⁰ Deedman, C. & Smith, J. C. The Nervous Shock Advisor: A Legal Expert System in Case-based Law. (*paper*) pp. 4-5

phrase as satisfied in the NFS by matching only those facts.⁷¹ This CBR technique of Branting's constitutes another approach to the task of satisfying an open textured term in a principle. (Again, I would claim that the possibility of generalising the facts in those precedential portions exists.)

To illustrate the idea of reasoning with portions of precedents one might consider the simple rule based system suggested in chapter 2 (see figure 1) for dealing with straightforward cases like *Daniels*. The case of *Grant* contains two open textured phrases intended by Lord Wright to clarify what is meant by 'unmerchantable', viz. 'defects unfitting a thing for its only proper use' which are 'not apparent on ordinary examination'. From the cases two different sets of facts might be identified as satisfying these terms. e.g. A foreign substance, in a bottle of drink, which causes sickness to the consumer, is a 'defect unfitting a thing for its only proper use'; a non discolouring foreign substance, in a sealed bottle of drink, is a defect which is 'not apparent on ordinary examination'. Both of these precedent constituents were in fact elements of *Daniels*, but the point is that in order for a NFS to satisfy either of the terms, it would not be necessary for it to match all the facts of *Daniels*, but merely those that had been identified as satisfying the term in that case. Moreover, it might be the case that the NFS only matched one of those term definitions in *Daniels*, but as regards the other term, it matched an alternative definition in another precedent. In such a situation satisfaction of the two open textured terms (and consequently the subsuming term 'unmerchantable'), would be achieved by matching with portions of different precedents, which is what Branting illustrates.⁷²

I propose that when the problem of vague concepts arises in principles, then, as with statutory or case derived rules, the best idea is to use a combined approach. Thus, for example, *Donoghue v. Stevenson* could provide a factual description, or precedent constituent example, of the non satisfaction of the phrase 'reasonable care in manufacturer', which a NFS like *Daniels* would fail to match. If it failed to match any of the other precedent constituent examples of the satisfaction, or non satisfaction, of the

⁷¹ a) Branting, L. K. Reasoning with Portions of Precedents. *ICAIL-91*, p. 147

b) Branting, L. K. & Porter, B. W. Rules and Precedents as Complementary Warrants. (paper for AAAI-91) p. 8

⁷² ref. 71, p. 147

phrase, (and possibly even if it did), the user would then be allowed to decide for himself what he thought a judge would decide. (One might also suggest that where the authority on a particular term is sparse, an expert should form hypothetical examples of its satisfaction, and non satisfaction, with certainty factors to indicate his confidence in them, to assist the user in his assessment.)

6.3 Identifying all the Principles applicable to a NFS

As previously discussed (in section 4.4), the domain principles the expert would articulate might be formed at different levels of generality, and the likelihood of a principle succeeding would be greater the less general it was in relation to the NFS. Therefore, when invoked, the principle based reasoner would examine all the principles in the PKB, and identify the most specific version of any principle which the NFS satisfied. For this task the domain hierarchy would be used. For example, based on the cases favourable to the principle on manufacturers' liability which I noted at the start of the chapter (in section 6.2), the most specific version of that principle which the expert could form might have contained the phrase 'product intended for use by an individual'. (i.e. Given the cases of *George v. Skivington*, where defective shampoo was involved, *Hawkins v. Smith*, where a defective sack was involved, and *Oliver v. Sadler & Co.*, involving a defective rope sling, this would be the most certain version of the principle the expert could generate.) A more general version, which he might have speculated, could have employed simply 'product', as I originally suggested. A specialisation of 'product' (not found in the cases) could be 'product for public use'. Thus if the NFS had involved an accident on public transport, the user would have satisfied this latter phrase when questioned about his problem. The NFS would not have matched the case facts (hence CBR would not produce a strong argument), but the principle based reasoner would infer satisfaction of the more general version of the principle, speculated by the expert, and this would be presented to the user as the most specific principle which he may argue as embracing his case (it being a generalisation of the most specific version discernible).

This version of the principle would of course have a certainty factor attached to it, indicating the expert's certainty that the court would want to adopt the principle, and would want to adopt it at this level of generality. However, if the specialisation 'product for public use'

had not been added to the hierarchy by the expert, then the user would have been questioned simply about the generic term 'product' when inputting his NFS. In such a situation he may be allowed, as discussed under GFRs, to attach a certainty factor to his response to reflect (semantic) uncertainty, e.g. over whether the court would want 'product' to subsume 'public transport'. (As with GFRs, this would affect the certainty of the satisfaction of the condition of the principle which the phrases concerned. This would in turn affect the certainty of the principle's conclusion. The final certainty value would indicate the certainty of its application (to the NFS), not the certainty of its acceptance by the court as a valid principle. This is what the expert's certainty assessments alone would indicate.)

Thus the most specific versions of any principles which the user's NFS satisfied would be determined, plus certainty factors indicating the likelihood of those versions succeeding. When presenting those (versions of) principles, and certainty factors, to the user, the principle based reasoner would also present the expert's reasons for his assessments, stored with the principles, plus, if required, the cases relevant to the principles, and accounts of how they were reconciled with the principles, which would be stored in the CKB. The user would of course be free to view less specific versions, but the certainty factors would be less. Potentially conflicting principles would be identified, but hopefully their certainty factors would differ thus suggesting which were most likely to succeed.

6.4 Case Profiles Construction and Utilisation

It is proposed that each potentially relevant domain case would be represented by a case profile, in which would be stored its critical facts, in a formalism suitable for CBR, plus other necessary information such as the legal issues involved, verdict, reason for the decision, and 'precedential pedigree' details. The method of obtaining, and use of, each of these items has already been discussed in the course of the preceding chapters. Here I will relate the items to the construction and use of the case profiles themselves.

The aforementioned information, necessary for each case profile, would be extracted from the raw texts either in the usual manual way, or hopefully, by employing a text analysis facility like that of FLEXICON to automatically identify it. The factual information would

be cast into a GREBE like representation, and this would be a manual task unless the analysis which identified the critical facts also identified their relationships *inter se*. The legal phrases used to describe legal issues, and employed in the reasons for the decisions, would be linked to a synonym index to enable comparison of these matters between the cases (and with the proposed principles). e.g. 'Failure to take reasonable care' in production, might be the expression used in one case, whereas 'negligence' in production could be used in another.

Using these case profiles the principle-cases reconciliation system would enable the expert to attempt to reconcile the cases with a proposed principle, and if necessary refine the principle, in the manner discussed in the preceding chapter. Then, for each case, an account of the way in which the reconciliation was achieved would be appended to its profile. (e.g. 'This case supports the principle because it is subsumable under it, and has the required outcome', or, 'Though unfavourable to the principle this case is arguably reconcilable with it because the following facts make it distinguishable from any of the favourable cases, and the wording/ambit of the principle does not subsume it.' Or, for a contrary, unreconcilable case, 'this authority contradicts the principle, but, in the opinion of the expert, is likely to be viewed as wrongly decided.'). The resulting profiles would constitute the CKB to be used by the expert system. Thus one CKB could serve both the principle-cases reconciliation system, and the expert system itself; in the expert system the case based reasoner would use the representations of the facts of the cases, and the principle based reasoner would use the reconciliation accounts produced by the principle-cases reconciliation system.

6.5 Use of the Principle Based Reasoner

The user would be encouraged to invoke the principle based reasoner, to ascertain what principles might be argued as applicable to his case, after the rule based reasoner, and then the case based reasoner, had been invoked. This would accord with the way in which a lawyer works. First see if a straightforward deductive (fairly incontestable) outcome can be predicted, as determined by unambiguous statutory rules, and clear, authoritative case based rules (see the section on rule based reasoning). Then, if deductive reasoning proves inadequate, see what support reasoning with cases (by analogy; GFRs and CBR) can

lend to each side's argument. Finally, if there is still the slightest element of doubt over what outcome the primary sources support, (and ideally, even if there appears to be no doubt), see if appeal to principle might support either side. This exhortation, that one should always see what arguments appeal to principle can produce, even if the case seems clearly determinable by deductive reasoning, or analogical reasoning, is merely a way of paying heed to MacCormick's observation that the majority of cases do not fall neatly into either the 'clear', or 'hard', case categories.

6.6 User Projection of Principles

I have proposed that the expert should identify potential domain principles (and build them into the system). As noted (in chapter 4), unlike a court, he would be performing the formulation task out of the context of a specific case. Therefore it might be the case that the principles he produced failed to cover a particular NFS. A more contentious idea would be to allow the user himself to speculate as to possible principles when actually running the expert system. It is contentious because unless he was himself (or had access to), a domain expert, he would be unlikely to make a good assessment of the likely acceptance of his projections. That is the disadvantage of this idea. The advantage of it is that he would be trying to construct a principle knowing the factual scenario (i.e. his NFS), which it was to cover, which is of course what judges do.

In outline this idea would work as follows. Having described his NFS to the system, and it finding no relevant statutory or case based rules, and CBR failing to identify very on point cases, the user would be allowed direct access to the domain hierarchy. The (most specific) terms in the hierarchy which his NFS satisfied would already have been established, and from them the user would be able to undertake generalisation inferencing in the hierarchy, identifying what he believed the court would perceive as generic terms (which could be simple facts, or legal concepts), which could constitute a principle. These generic terms would subsume other specific terms which, if present in other precedents, would identify cases which could be argued as analogous to the NFS, and subsumed by the principle, if advancing the principle. These arguably factually similar precedents would be retrieved by a (GREBE like) case based reasoner, which would use the specific terms subsumed by the generic terms the user

had selected as a sort of 'new' (or analogous) NFS, for which matches were to be sought among the domain cases.

To illustrate this idea, using the example parts of a domain hierarchy proposed in figure 3, suppose that part of the user's NFS involved a Person; he might generalise to a principle involving Legal Bodies, and if so the case based reasoner would look for precedents involving Companies, as well as Persons. It would ultimately be up to the user to assess the likelihood of the court regarding as analogous any case identified as arguably similar to the NFS, through inferencing in the hierarchy and CBR, when trying to form a principle, (although the GFs attached to links in the GFH, for identifying GFRs, might be of some assistance to him in this task, see chapter 2).

Having identified cases which were arguably analogous to the NFS, and subsumed by his proposed principle, the user would then have to reconcile them with the principle. At this point the principle-cases reconciliation system would be used to assist him in explaining away unfavourable cases (in the same way as it is proposed it would assist the expert), although obviously distinguishing would not be an option, unless his generalisations had only been made from some of the critical facts in his NFS. The precedential weight of the relevant cases would be determined (using their precedential pedigree details, and comparing them to the procedural context of the NFS), and would be an important consideration, as not being an expert the user could not afford to be as confident as an expert could about the court simply overriding authoritative, unfavourable, unreconcilable precedents.

Perhaps the most important point about this idea of user constructed principles is that it would allow him to identify the generalisations of his facts which subsumed those cases which would best support his claim. Thus if one set of generalisations (i.e. a 'principle'), resulted in too many unfavourable, unreconcilable cases being identified as potentially analogous, he could try a variation of those generalisations. However he would constantly have to reevaluate the chances of his proposed principle being accepted by the court, even if reconcilable with the case law, basing his assessment on the same criteria as the expert would use. i.e. Knowledge of trends in the domain law, and appraisal of judicial consequentialist evaluation, and of the level of generality at which the principle would be accepted.

As noted, this idea of the user creating principles would have the advantage that it would allow a principle to be formed knowing the NFS it was to cover, this being a consequence of the fact that the principle was being formed while using the expert system. To put it another way, the best time to form a principle would be when using the system. It has also been noted that the best person to form a principle is an expert. Therefore, one might conclude, that the best person to use the system would be an expert, (which some might say defeats its purpose).

6.7 Other Domain Principles

The above discussion on principles in this chapter has mainly concerned, as the whole previous discussion has, those speculative 'new' principles formulated by the expert, which he believes will be accepted by the court. However the PKB would also contain those principles he has identified as potentially conflicting, and regarding which he has assessed the court's likely preference, plus relevant principles which are clearly established, and uncontentious, and the principles based reasoner would attempt to apply them to the NFS in the same way as the others.

Chapter 7

Factual Descriptions of Cases

7.1 New Fact Situation entry

The domain expert would decide on a basic set of factual terms, and it would provide the basic level of representational detail at which each domain case, stored in the CKB, could be described. To obtain a description of the the NFS a menu-driven case editor would be invoked at the start of the expert system consultation, primarily to obtain all the information required for RBR, and it could be reinvoked if required if additional information was needed by the case based reasoner or principle based reasoner. The determination of a complete factual description of the NFS prior to processing is of course necessary for CBR, however most rule based reasoners simply acquire factual information as required in the course of inferencing. The reason I suggest that a full account be obtained right at the start of the consultation is so that all possible rule based arguments would be identified, not just the first one the facts satisfy, which is what conventional rule based processing achieves. This is desirable because a lawyer should use all possible lines of attack, and must be aware of all possible counter arguments.

The actual means of entry of individual items of information would follow the conventional practice of requiring the user to respond to system queries so that it could establish his knowledge about certain matters, or obtain factual information. Where it was possible to avoid a number of detailed questions by asking a more general one this would of course be done, however, whenever the user was uncertain about the answer, the subsumed, detailed questions would need to be presented. Additionally, even if a particular detailed set of factual questions was avoided when RBR, it might be necessary to ask them later for the task of CBR.

7.2 Problem of Grain Size

At this point some observations should be made about the level of

detail at which cases would be described to, and represented in, the expert system. Ideally, the specific facts of every precedent, and the NFS, would be recorded. In practice there would be (or a least there is in present systems), a limit to how specifically any case could be described (i.e. as noted above, there would be a basic level of terms which would constitute the greatest level of detail at which any case could be described), and this could cause problems.

For example, suppose a system was operating in the area of law in which the *Rylands* and *Read* cases arose, i.e. that concerning the liability of a person in charge of a dangerous substance to those harmed by it (see section 4.7). Suppose the *Read* scenario was in fact the NFS presented to the system. Unless the expert had decided that the fact that the harm in *Rylands* occurred outside the defendant's property was a relevant factor, the location of the harm might not be included in a case rule derived from the case, or a CBR representation of its facts, and the question of where the harm occurred in the NFS (*Read* scenario) would not be asked of the user. (The harm actually occurred on the defendant's property in *Read*.) The consequence of this is that the rule based reasoner might establish that the NFS satisfied the case rule, or the case based reasoner might find that *Rylands* and the NFS were very analogous. Thus, when the system noted that 'strict liability' was established in *Rylands*, it would predict 'strict liability' as the most likely outcome in the NFS (the *Read* scenario), i.e. it would advise judgment for the plaintiff, whereas the verdict actually given in *Read* was for the defendant.

However, it is possible that if the system user then invoked the principle based reasoner, it would identify the 'strict liability', and 'no liability without fault' principles as potentially relevant, and that the later might have the highest certainty factor (assuming the expert had correctly assessed the legal trend, and consequentialist arguments). The user would then have to examine the cases himself, to identify a factual disparity which the expert system had failed to reveal due to its level of representational detail, and which the courts might highlight (as in fact they did in *Read*), in order to achieve the end they desired.

This hypothetical example is intended to illustrate the problem facing the expert, and knowledge engineer, when deciding to what level of detail to go when developing a case description language for the domain. It also hopefully illustrates the utility of modelling

principles, and why the user should always invoke the principle based reasoner when using the system, as the advice based on principles may conflict with that produced from applying RBR or CBR to the NFS, and that advice may alert the user to the possible inadequacies of these forms of reasoning.

Ideally an expert system would be capable of representing every legally significant fact, but as the law is an ever evolving entity the importance of the critical ones is often not apparent until the judges decide to accord them legal significance, (e.g. the question of where the harm was suffered was not considered important before the *Read* case arose), and one cannot expect an expert to be able to anticipate every such 'yet to become important' fact. Thus I would question the utility of producing a complete factual representation of every domain case, even if it was achievable, as that would not achieve an identification of the critical facts. It is much more useful to identify the applicable principles, which reflect the relevant legal trends.

Chapter 8

A Principle-Cases Reconciliation System and Expert System based on MacCormick's Theory

8.1 Modelling MacCormick's ideas

To bring together the observations made in the proceeding chapters, I now propose that an attempt to model, computationally, the ideas Neil MacCormick expresses in 'Legal Reasoning and Legal Theory', using the currently available AI techniques, might involve the following three stages.

8.2 Stage One

Stage one would involve the following. The construction of a semantic network domain hierarchy of facts and legal concepts (which would subsume a hierarchy of facts alone, i.e. the GFH). The construction of case profiles, containing GREBE like CBR representations of the facts, plus additional case information, which would constitute a CKB. The construction of a synonym index of legal phrases present in the domain. The construction of a RKB, comprised of legal rules, extracted from the statutes and cases. (Although the construction of the RKB could be delayed until the end of stage two, as it would not be used by the principle-cases reconciliation system.) These structures would serve the systems used in the following two stages.

8.3 Stage Two

The construction of a principle-cases reconciliation system, which would do the following. Retrieve domain cases which considered the legal issues relevant to the proposed principles, employing the synonym index while carrying out this task. Use GREBE's form of CBR in order to identify, and also to try to actively create, factually similarity (or dissimilarity), between cases. Employ a weighting scheme (such as that developed at UBC), to assess the 'precedential

pedigree' of relevant precedents, relative to possible procedural contexts. Supplement the case profiles (i.e. the CKB) with the precedential weightings, plus accounts of how proposed principles had been reconciled with the cases. Store the refined (versions of) principles, plus their certainty factors, plus the reasons for those certainty factors (i.e., for each principle, reference to any relevant, contrary, unreconcilable cases whose precedential weight made the principle untenable, or extremely unlikely to succeed, in a particular procedural context; and the expert's reasons for his certainty assessments), in a PKB. Link each principle to the case profiles relevant to it.

8.4 Stage Three

The development of an expert system, which would be composed of a rule based reasoner, a case based reasoner, and a principle based reasoner. The rule based reasoner would reason deductively by performing RBR on the RKB of statutory and case derived rules, and perform generalisation and specialisation inferencing in the domain hierarchy. (Although when discussing deductive reasoning I illustrated generalising and specialising in the context of a simple GFR, it is more realistic to envisage the rule based reasoner using the entire domain hierarchy, as a legal rule can contain legal concepts.) Inferencing in the domain hierarchy, to identify generic rules, would be the technique the rule based reasoner would use when the NFS failed to match specific condition terms of a rule in the RKB for which the expert had proposed generalisations. When there was uncertainty over whether a rule condition was satisfied, due to the open texture of a phrase it contained (whether the rule was one of those contained in the RKB, or a generic rule), then the rule based reasoner would pass control to the case based reasoner.

The case based reasoner would identify the best arguments for and against an open textured phrase being held satisfied in a NFS, by attempting to match to, and distinguish from, the NFS, the domain cases (contained in the CKB) which ruled on that phrase. The type of CBR it would employ to produce these 'arguments by analogy' would depend on the nature of the domain involved, and the degree of sophistication decided on for the reasoner. If the nature of the domain was such that there were quite readily identifiable, and maybe fairly fixed in number, groups of facts, whose *intra* group

relationships were also fairly rigid, and if the presence or absence of those groups unequivocally strengthened, or weakened, the claims the system was to advise on, in other words if HYPO 'factors' were identifiable, then the simpler, less flexible, HYPO form of CBR might be adequate. If this was not the case, and the domain was one in which it was less clear what particular facts, or structured groupings of them, were critical to the outcome of domain cases, then the overall similarity of facts (and relationships among them) in different cases would be more important, and a system like GREBE which could look for new similarities (or try to create differences) would be more appropriate. This sophistication of GREBE's form of pattern matching is of course an advantage, when compared to the more schematic form of case matching performed by HYPO; its disadvantage is that it is potentially very computationally expensive, as Branting himself notes.⁷³

It is probable, if the HYPO form of CBR was to be used in the expert system, that the factual information contained in the case profiles, which would be cast in the GREBE formalism at the outset, could be automatically examined and used to identify what HYPO factors were present in the CKB cases. This assertion is based on the fact that HYPO itself automatically infers what factual predicates (which are used to satisfy the prerequisites/elements of its dimensions/factors, see section 3.5) are satisfied by a NFS, from the user's description of the detailed facts of that case. Thus it can be assumed that, even if the factual details of each case needed to be reexamined in order to enable HYPO like CBR, the same set of case profiles (and thus the same CKB), could be used for this purpose as was used by the principle-cases reconciliation system, and that consequently it would not be necessary to scrutinise the raw materials again.

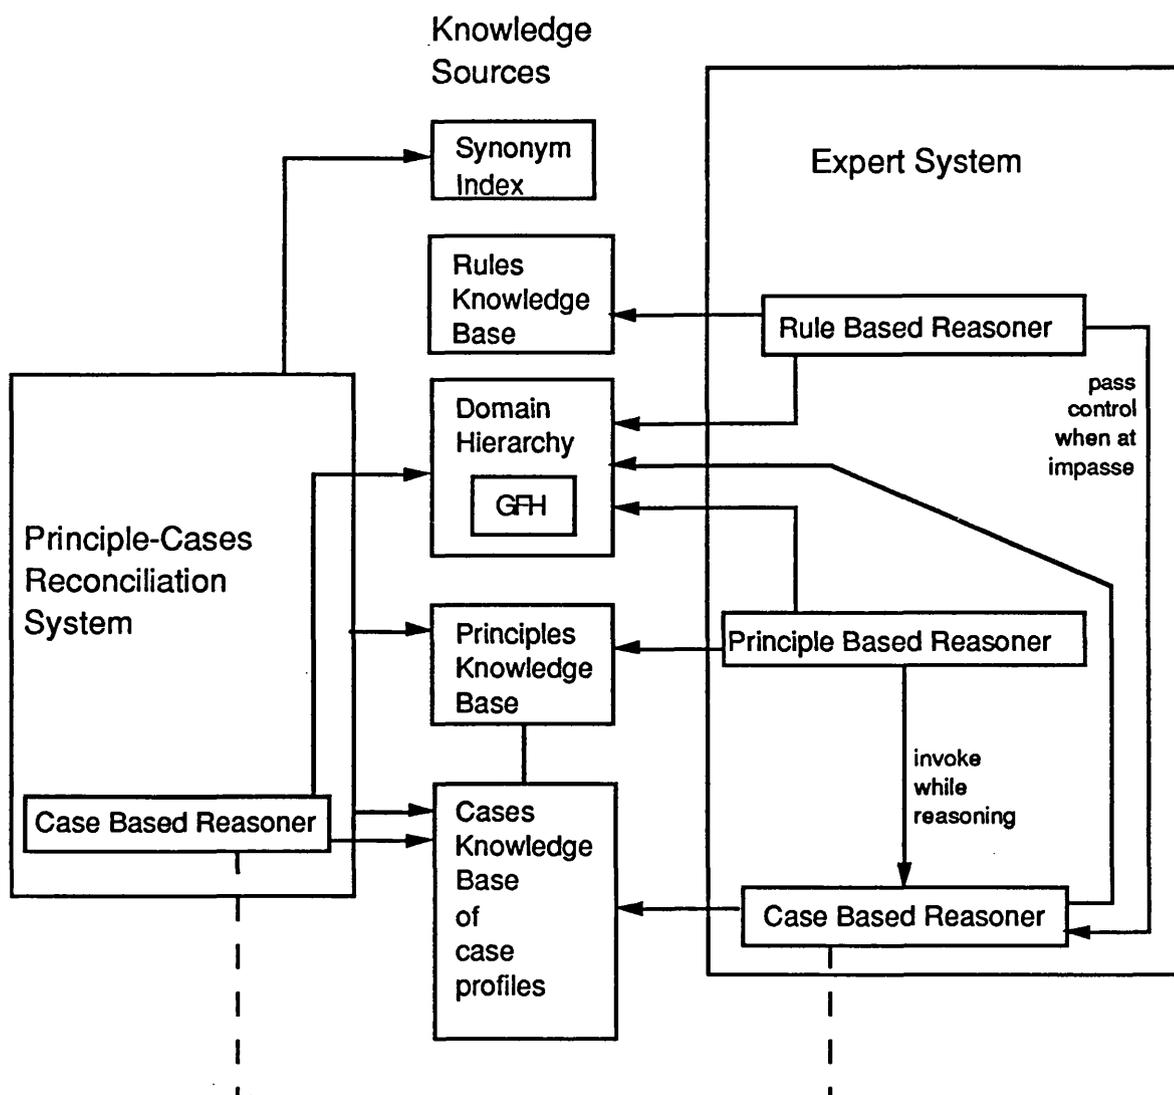
The case based reasoner would also note the precedential pedigree of the cases identified as factually analogous to the the NFS. This would enable the user to rank the cases it retrieved either according to their degree of similarity with the NFS, or according to their precedential authority relevant to it, but it would be left to him to offset one against the other.

The principle based reasoner would note the procedural context of the NFS, in order to exclude from its reasoning any proposed principle for

⁷³ Branting, L. K. Building Explanations from Rules and Structured Cases. *IJMMS-1*, p. 817

which there was an unfavourable case, which could not be reconciled with the principle unless it was argued as having been wrongly decided, and which, due to the doctrine of *stare decisis*, the court hearing the NFS would not be able to overrule, even if it considered that the case was wrongly decided. The reasoner would then attempt to produce 'arguments from principle', by identifying the most specific version of any usable principle contained in the PKB which could be advanced as subsuming the NFS. To do this it would generalise, in the domain hierarchy, from the facts of the NFS, to the legal concepts and facts contained in (versions of) the principles; and it would use CBR to identify the arguments for and against the NFS satisfying any open textured terms they contained. The versions (of principles) returned would have certainty factors to indicate their chances of success, and to reconcile conflicts, plus an explanation of the reasons for the individual certainty evaluations involved, plus the accounts of how the principles were reconciled with the domain case law.

Figure 4. Diagram of the relationships between the proposed systems and knowledge sources.



The principle-cases reconciliation system would have its own dedicated case based reasoner, for performing GREBE's form of CBR, if it was decided that the expert system should use HYPO's form; otherwise the same reasoner would be used by both systems.

Chapter 9

Conclusions

9.1 Recapitulation

As stated at the start of this work, its main purpose has been to examine the extent to which one might model, computationally, part of a specific theory of judicial legal reasoning; as part of that examination I have also suggested how one might attempt to build a legal expert system which is rooted in that theory. From these considerations one might draw the the following conclusions.

9.2 The extent to which MacCormick's Theory can be modelled

Deductive justification can be predicted by using RBR, reasoning by analogy can be simulated, to an extent, by using generic rules and CBR. However the use of principle in legal reasoning can only be computationally expressed to a limited extent at present. Assessing the desirability of the consequences of proposed principles would remain a human task, and even though the person undertaking this (the expert) should be aware of all the constraints to which the person whose assessment he was trying to anticipate (the judge) was subject, there would remain an unavoidable element of guess work as, at the end of the day this assessment is, as MacCormick says,⁷⁴ at least in part irreducibly subjective.

Expert involvement would be required even more in a project to implement MacCormick's ideas than it is for the development of current expert systems, and getting that participation is, as previously stated, already a major problem. However, for the task of actually forming principles, and reconciling them with the case law, AI techniques can be of assistance, as was explained. In fact they can provide the tools with which the user himself can speculate as to new principles, as was suggested, and so if expert assistance was not available they could provide useful service by enabling the user to

⁷⁴ LRLT, p.112

formulate additional arguments to consider, and put before the judge (if favourable), or decide how to counter (if unfavourable).

9.3 The need for Principles in any Expert System

I have suggested that three types of reasoning would need to be modelled in an expert system based on MacCormick's theory. Reasoning with legal rules i.e. rules derived from statutes or cases, reasoning by analogy (using cases as descriptions of factual situations), and reasoning with principles (formulated as rules by the expert).

The two legal claims which have provided the main context in which the discussion of these types of reasoning has been undertaken (i.e., right of recovery against the seller of a defective product for harm suffered, and right of recovery against the manufacturer of a defective product for harm suffered), each rely predominantly on a different type of reasoning (one on reasoning with statutory and case based rules, one on reasoning with principles). Together they constitute (at least part of) a specific area of law with which one can imagine a lawyer having to deal on a daily basis, and in which an expert system could provide useful service. i.e. Consumer (or user) rights of recovery for harm suffered from a defective product. A system offering advice on just one of these claims would be of limited use, as the lawyer's task is to identify and pursue all possible courses of action open to his client, not just those available against a specific party. Therefore, as a lawyer should offer comprehensive advice, in order to be of practical use an expert system operating in an area of law such as this should be able to reason with principles, as well as rules and by analogy.

This area of law basically describes the legal context in which MacCormick conducts his discussion of deductive justification, and the use of principles when deductive justification is not possible (i.e. in hard cases). He notes that appeal to principle is made in many areas of law when hard cases arise. However, from a practising lawyer's point of view, one might suggest that investigating the possibility of appealing to principles should not be confined merely to what appear to be hard cases. Even if a deductively justifiable argument does appear possible (i.e. the case appears clear), one should still attempt to identify alternative, principle based

arguments, as the deductive one might fail for some reason. Consequently it is submitted that, the above noted view, that in order to be of practical use an expert system in law should offer comprehensive advice, and therefore should be able to reason with general principles of law, as well as with case or statutory rules, and by analogy, is extendible, and can be argued as being applicable to expert systems in law operating in practically any area of law.

Although the likes of Smith have recognised the need for a legal expert system to be able to reason with principles, and through the UBC projects he and his researchers have proved that it is possible to meet that need, I do not believe any one has yet constructed a system capable of all three types of reasoning. However such a system cannot be far away. So called hybrid systems, which use the techniques of RBR and CBR, to implement reasoning deductively with legal rules, and reasoning by analogy, have already been built, e.g. GREBE, and CABARET; and the UBC projects use the same techniques to reason with principles. The attraction of combining the two types of system cannot have escapes notice. I believe that the development of such a system is the next step towards the widespread acceptance and use of expert systems in law.

References

Books

- ASHLEY, K. D. *Modelling Legal Argument: Reasoning with Cases and Hypotheticals*. The MIT Press, Cambridge, Massachusetts, 1990
- MACCORMICK, D. N. *Legal Reasoning and Legal Theory*. Clarendon Press, Oxford, 1978
- SUSSKIND, R. E. *Expert Systems in Law: A Jurisprudential Enquiry*. Clarendon Press, Oxford, 1987

Articles

- ASHLEY, K. D. Reasoning with Cases and Hypotheticals in HYPO. *IJMMS-1*, pp. 753-796
- BERMAN, D. H. & HAFNER, C. D. The Potential of Artificial Intelligence to Help Solve the Crisis in Our Legal System. *Communications of the ACM*, August 1989, Volume 32, Number 8, pp. 928-938
- BRANTING, L. K. Building Explanations from Rules and Structured Cases. *IJMMS-1*, pp. 797-837
- BRANTING, L. K. Reasoning with Portions of Precedents. *ICAIL-91*, pp. 145-154
- BRANTING, L. K. & PORTER, B. W. Rules and Precedents as Complementary Warrants. (*paper received from Branting, to appear in AAAI-91*)
- DEEDMAN, C. Exposing some Myths about Law and Dispelling some Myths about Computers. (*paper received from Professor J. C. Smith*)
- DEEDMAN, C. & SMITH, J. C. The Nervous Shock Advisor: A Legal Expert System in Case-based Law. (*paper received from Smith*)
- GELBART, D. & SMITH, J. C. Beyond Boolean Search: FLEXICON, A Legal Text Based Intelligent System. *ICAIL-91*, pp. 225-234
- GORDON, T. F. An abductive theory of legal issues. *IJMMS-2*, pp. 95-118
- KOWALSKI, A. Case Based Reasoning and the Deep Structure Approach to Knowledge Representation. *ICAIL-91*, PP. 21-30
- MACCORMICK, D. N. Law as Institutional Fact. *Law Quarterly Review*, Volume 90, 1974, pp. 102-129
- RISSLAND, E. R. & SKALAK, D. B. CABARET: rule interpretation in a

hybrid architecture. *IJMMS-1*, pp. 839-887

SERGOT, M. The representation of law in computer programs: A Survey and Comparison, in: *Knowledge Based Systems and Legal Applications*, Bench-Capon, T. J. M., Ed. Academic Press, 1990

SMITH, J. C. A Conceptual Model for the Representation of Legal Knowledge. (*paper received from Smith*)

Cases

Daniels & Daniels v. R. White & Sons & Tarbard [1938] 4 All E.R. 258

Donoghue v. Stevenson [1932] A.C. 562; 1932 S.C. (H.L.) 31

George v. Skivington (1869) L.R. 5 Exch. 1

Grant v. Australian Knitting Mills [1936] A.C. 85

Hawkins v. Smith (1896) 12 T.L.R. 532

Morelli v. Fitch & Gibbons [1928] 2 K.B. 636

Mullen v. Barr & Co. 1929 S.C. 461

Oliver v. Sadler & Co. 1929 S.C. (H.L.) 94; [1929] A.C. 584

Read v. J. Lyons & Co. Ltd. [1947] A.C. 156

Rylands v. Fletcher (1868) L.R. 3 H.L. 330

Steel v. Glasgow Iron & Steel Co. Ltd. 1944 S.C. 237

White & Carter (Councils) Ltd. v. McGregor [1962] A.C. 413; 1962 S.C. (H.L.) 1

Wilkinson v. Kinneil Cannel & Coking Co. Ltd. (1897) 24 R. 1001

Winterbottom v. Wright (1842) 10 M. & W. 109

Woods v. Caledonian Rly. (1886) 13 R. 1118

Other Material

Books

- ALTY, J. L. & COOMBS, M. J. *Expert Systems: Concepts and Examples*. Manchester, NCC, 1984
- GARDNER, von der Lieth, A. *An Artificial Intelligence Approach to Legal Reasoning*. The MIT Press, Cambridge, Massachusetts, 1987
- KOERS, A. W. et al. *Knowledge Based Systems in Law: In search of methodologies and tools*. Kluwer, Deventer, 1989
- Proceedings of the Third International Conference on Artificial Intelligence and Law*, June 26-28, 1991. ACM, New York

Articles

- BAINBRIDGE, D. I. A Computer System to assist with the Sentencing of Offenders. *Computer Applications in Business and in the Social Sciences*, Volume 1, Number 1, January 1990, pp. 31-45
- BAINBRIDGE, D. I. Expert Systems and the Law. *The Law teacher*, 1989, Volume 23, Number 3, pp. 279-292
- CUTHILL, B. A Dynamic Approach to Cross-Context Case Retrieval. (*paper received from Smith*)
- McCARTY, L. T. Artificial Intelligence and Law: How to get There from Here. *Ratio Juris*, July 1990, Volume 3, Number 2, pp. 189-200
- REED, C. Expert Systems and Legal Expertise. *Computer Law and Practice*, Volume 5, 1988-89, pp. 122-125
- SMITH, J. C., GELBART, D., GRAHAM, D. Building Expert Systems in Case-Based Law. (*paper received from Smith*)
- WALKER, R. F., OSKAMP, A., et al. PROLEXS: Creating law and order in a heterogeneous domain. *IJMMS-2*, pp. 35-67
- WALKER, R. F., ZEINSTRA, P. G. M., van den BERG, P. H. A Model to Model Knowledge about Knowledge, or, Implementing Meta-Knowledge in PROLEXS. in *Advanced topics of information technology and law*, Vandenberghe, G. P. V. Ed., Kluwer, Deventer, 1989

