International Journal of Law and Information Technology © *Oxford University Press 2006;*
doi:10.1093/ijlit/eal023

A Knowledge Representation Model for the Intelligent Retrieval of Legal Cases

YIMING ZENG[1], RUILI WANG[2], JOHN ZELEZNIKOW[3], AND ELIZABETH KEMP[4]

Abstract

In this paper, we develop a knowledge representation model for the innovative intelligent retrieval of legal cases, which provides effective legal case management. Examples are taken from the domain of accident compensation. A new set of sub-elements for legal case representation (sub-issues, pro-claimant, pro-respondent and contextual features) has been developed to extend the traditional representation elements of issues and factors. In our representation model, an issue may need to be further decomposed into sub-issues; factors are categorised into pro-claimant and pro-respondent factors; and contextual features are also introduced to help retrieval. These extensions can effectively reveal the factual relevance between legal cases. Based on the knowledge representation model, we propose the IPF scheme for intelligent legal case retrieval. Experiment and statistical analysis have been conducted to demonstrate the effectiveness of the proposed representation model and retrieval scheme.

Key words: legal case retrieval, case representation elements, legal knowledge representation, accident compensation

[1] Institute of Information Sciences and Technology, Massey University, New Zealand E-mail: y.zeng@massey.ac.nz

[2] Institute of Information Sciences and Technology, Massey University, New Zealand E-mail: r.wang@massey.ac.nz (corresponding author).

[3] School of Information Systems, Victoria University, Australia E-mail: john.zeleznikow@vu.edu.au

[4] Institute of Information Sciences and Technology, Massey University, New Zealand E-mail: e.kemp@massey.ac.nz

1 Introduction

In legal proceedings under the common law system, past legal cases (precedents) are frequently used to support arguments and judicial opinions. In such a process, precedents can be followed, analogised, distinguished or overruled[5]. The reason for using precedents is because of the open textured nature of Law. Some commentators view open texture as the uncertainty about whether given legal terms match specific facts[6]. Lawyers and judges often use precedents in open textured domains, in order to see how legal terms will apply to the current case facts[7].

Past legal cases can be seen as sources of legal knowledge[8], because they contain the tacit knowledge about how open-textured legal terms were once applied. The process of using precedents for making arguments and giving judicial opinions employs the tacit knowledge contained in the precedents. Developing a knowledge representation model for the intelligent retrieval of legal cases can greatly facilitate the users' ability to correctly access tacit legal knowledge, which is an important goal of legal knowledge management[9].

The need to manage legal knowledge effectively for lawyers and judges to locate knowledge and information becomes urgent, due to the rapidly growing volume of the reported landmark cases[8]. Also, an effective retrieval mechanism is needed especially for case-based legal arguments, reasoning and decisions. The traditional retrieval and management methods for legal precedents include alphabetical indexing and keyword-based searching[10]. Other methods include merely searching on surface features like title, location, publication date, document number, and so on[11]. None of these provides a satisfactory solution because they return a lot of irrelevant cases or missing the relevant cases.

Case-based retrieval, an application of case-based reasoning[12], could be an appropriate solution, because of the nature of use of legal precedents

---

[5] Prakken and Sartor, 'Modelling Reasoning with Precedents in a Formal Dialogue Game' (1998) Artificial Intelligence and Law 6 231–287.

[6] Gardner, *An Artificial Intelligence Approach to Legal Reasoning* (Bradford Books/MIT Press: Cambridge 1987).

[7] Porter, Bareiss and Holte, 'Concept Learning and Heuristic Classification in Weak-Theory Domains' (1990) Artificial Intelligence 45 1–2. Also see Prakken, *Logical Tools for Modelling Legal Argument* (Kluwer Academic: Dordrecht 1997).

[8] Oskamp, Tragter and Lodder, 'Mutual Benefits for AI & Law and Knowledge Management' (1999) 7th Intl Conf on Artificial Intelligence and Law: 126–127.

[9] Zeleznikow, 'Using an Argumentation Based Approach to Manage Legal Knowledge' in Schwartz (ed), *Encyclopedia of Knowledge Management* (Idea Group Inc: Hershey 2005).

[10] Moens, 'Automatic Indexing and Abstracting of Document Texts' (2000) Kluwer International Series on Information Retrieval 6.

[11] Bueno, Gresse von Wangenheim, Mattos, Hoeschl, and Barcia, 'JurisConsulto: Retrieval in Jurisprudential Text Databases Using Juridical Terminology' (1999) Intl Conf on AI in Law.

[12] Kolodner, *Case-Based Reasoning* (Morgan Kaufmann Publishers: San Mateo 1993).

in arguments and judicial opinions[13]. For legal tasks, which generally have a weak or no domain model, decision-makers are known to rely heavily on their knowledge of past cases, since verdicts must be consistent and transparent.

In Ashley's version of legal case-based reasoning[14], issues and factors are the basic elements to represent legal cases. Lawyers frequently make arguments by analysing and interpreting the similarities and differences between cases[15]. Successful arguments depend on whether the cited precedents are factually relevant to the current case, i.e., whether the cited precedents and the current case share issues and factors.

We shall extend Ashley's work by developing a set of sub-elements (sub-issues, pro-claimant factors, pro-respondent factors and contextual features) that better reveal factual relevance between cases. In this paper, we focus upon the domain of New Zealand accident compensation (Accidents Compensation Act 1982 & Accident Rehabilitation and Compensation Insurance Act 2001). We believe this representation model has the potential to be used in other domains.

This paper is organised as follows. In Section 2, we will introduce similarity measuring and existing legal case-based retrieval models. In Section 3, we will introduce the new representation sub-elements. In Section 4, we will illustrate a new knowledge representation model and IPF scheme for retrieving legal precedents. In Section 5, we will test the effectiveness of the proposed knowledge representation model and IPF scheme, and give abundant discussions. Finally, a summary of our research is given in Section 6.

2 Case-based reasoning and retrieval

Case-based reasoning (CBR), which has been used for case/information retrieval, is an AI approach to solve a new problem by remembering and adopting a previous similar situation[16]. A typical CBR cycle comprises four steps[17]: (i) retrieve the most similar cases to the target problem; (ii) reuse the cases to try to solve the target problem; (iii) revise the solution to fit the target problem if necessary; and (iv) retain the confirmed new solution as part of the case base.

---

[13] Zeleznikow, Stranieri and Hunter, 'Beyond Rule Based Reasoning – the Meaning and Use of Cases' (1995) 11th Conf on Artificial Intelligence for Applications 292–298.

[14] Ashley, 'Case-based Reasoning and Its Implications for Legal Expert Systems' (1992) Artificial Intelligence and Law 1(2) 113–208.

[15] Ashley and Rissland, 'Waiting on Weighting: A Symbolic Least Commitment Approach' (1988) AAAI-88 239–244.

[16] Riesbeck and Schank, *Inside Case-Based Reasoning* (Lawrence Erlbaum Associates: Cambridge 1989). Also see Watson and Marir, 'Case-Based Reasoning: A Review' (1994) The Knowledge Engineering Review 9(4).

[17] Aamodt and Plaza, 'Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approach' (1994) AI Communications 7(1) 39–59.

In CBR systems, the previous cases are stored in a case base. Each case normally consists of three parts[18]: (i) description of the problem or situation; (ii) solution to that problem or situation; and (iii) outcome of the solution.

The description of the problem or situation is characterised by a set of predefined dimensions (also called features or attributes). Given values (either numeric or fuzzy values) to these dimensions, a problem or situation can be described definitely, and an index of the case base can be built. When a new problem or situation (also called target case) occurs, case retrieval algorithms will be used to retrieve the similar previous cases by comparing the values of these dimensions between the new problem and the previous cases. Generally, no previous cases will exactly match the new situation. Therefore, the case retrieval algorithms normally perform the partial matching on the aggregate match[19].

Kolodner[20] indicates two steps for computing a similarity measure. The first step is a dimensional match, performed for a dimension by comparing quantitative or qualitative values. The second step is an aggregate match, usually performed by computing all the dimensional matching scores and then measuring the overall matching score.

A commonly used aggregate match algorithm for similarity measure is the nearest-neighbour matching algorithm including the following steps[20]. First, all dimensional scores between a target and a previous case are calculated. Second, an aggregate matching score is computed based on weights and the pre-calculated dimensional scores. As different dimensions have different degrees of importance to a solution, the aggregate match takes into account the weights of different dimensions. This underlying notion is built in the nearest-neighbour matching function. Typically, an aggregate matching score, SIM (T, P), between a target case T and a previous case P is calculated by the following function: $SIM(T,P) = \sum_{i=1}^{n}\left[w_i \times sim(f_i^T, f_i^P)\right]$, where w_i is the weight for the i^{th} dimension, and $sim(f_i^T, f_i^T)$ is the i^{th} dimensional matching score. Then, the previous cases will be ranked based on their aggregate matching scores.

Based on similarity measurement, CBR has been used in case/information retrieval[21]. Other well known methods for case retrieval include induction algorithms[22] and case retrieval nets (CRN)[23]. The induction

---

[18] Aamodt and Plaza, 'Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approach' (1994) AI Communications 7(1) 39–59.

[19] Watson and Marir, 'Case-Based Reasoning: A Review' (1994) The Knowledge Engineering Review 9(4).

[20] Kolodner, *Case-Based Reasoning* (Morgan Kaufmann Publishers: San Mateo 1993).

[21] Ashley and Rissland, 'Law, Learning and Representation' (2003) Artificial Intelligence 150 17–58.

[22] Quinlan, 'Induction Over Large Databases' (1979) Rep. No. HPP-79–14, Heuristic Programming Project, Computer Science Dept., Stanford University, US.

[23] Lenz and Burkhard, 'Case Retrieval Nets: Basic Ideas and Extensions' (1996) Advances in Artificial Intelligence.

algorithms decide which features can best distinguish cases, and creates a decision tree to organise the cases. This method is useful when a single feature dependent on other features is required as the solution. The knowledge-guided induction applies the domain knowledge to manually identify case features. The idea of CRN is to represent the cases and their attributes as a semantic network of interconnected information entities (IE). Starting with the query's IE activated, CRN uses a spreading activation algorithm to retrieve the best matching cases. CRN is regarded as a fast retrieval algorithm for large case bases.

Early legal reasoning was primarily modelled using rule-based reasoning[24]. However, rules are fundamentally inadequate to represent legal cases, because legal rules are not black and white with clear boundaries[25]. Case-based reasoning is very suitable for the legal domain that has open textures and lacks a strong domain model[26], and has been widely and successfully applied in the legal domain.

In legal proceedings under the common law system, in order to deal with the open textured concepts in legal domains, legal precedents are often used to see how legal terms once applied to the past cases, and therefore how these legal terms will apply to the current cases[27]. Two views on how case similarity can be determined have been sketched[28]: interpretation by direct comparison and interpretation by generalisation. In interpretation by direct comparison, the raw facts and circumstances of the precedent and current case are directly compared. While in interpretation by generalisation, the generalised facts and circumstances of the precedent and current case are compared.

HYPO[29] is believed to be the first true case-based legal system that operated in the domain of US trade secrets. Dimensions were introduced in HYPO to represent legal cases. HYPO-style dimensions capture the knowledge that certain facts support one side (i.e., plaintiff or defendant) in a legal dispute and certain changes in these facts tend to strengthen or weaken the claim of a side. Dimensions have a range of values, along which the supportive strength may shift from one side to the other[30].

---

[24] Zeleznikow and Hunter, Building Intelligent Legal Information Systems: Knowledge Representation and Reasoning in Law (Kluwer Computer/Law Series: Deventer-Boston 1994).

[25] Rissland, Ashley and Loui, 'AI and Law: A fruitful synergy' (2003) Artificial Intelligence 150 1–15.

[26] Zeleznikow, Stranieri and Hunter, 'Beyond Rule Based Reasoning – the Meaning and Use of Cases' (1995) 11th Conf on Artificial Intelligence for Applications 292–298.

[27] Porter, Bareiss and Holte, 'Concept Learning and Heuristic Classification in Weak-Theory Domains' (1990) Artificial Intelligence 45 1–2.

[28] Bing, 'The Problem of Finding a Precedent' in Nagel (ed.) *Law, Decision-making and Microcomputers* (Quorum Books: New York 1991).

[29] Ashley, 'Case-based Reasoning and Its Implications for Legal Expert Systems' (1992) Artificial Intelligence and Law 1(2) 113–208.

[30] Ashley and Rissland, 'Law, Learning and Representation' (2003) Artificial Intelligence 150 17–58.

Under HYPO's case retrieval criterion, the most relevant cases are called Best Untrumped Cases (BUCs). After retrieving the legal cases, HYPO constructs the 3-ply legal arguments for a new case. First, when Party A starts arguing, a point for A is produced making a claim that A should win. Second, a response for Party B is produced countering the claim. Last, a rebuttal for Party A is produced to answer the response. To create the initial point for A, HYPO identifies the most on point case whose result favoured A, and makes an argument by analogy. Responses and rebuttals use the legal argument techniques of distinguishing by focusing on differences rather than similarities, and trumping by presenting more on-point counter-examples.

CATO[31] is a tutoring program designed to teach law students how to perform argument tasks. Factors were introduced in CATO to represent legal cases. CATO-style factor is a kind of abstract knowledge related to a stereotypical fact pattern of the case. The factors are categorised into two groups to represent, analogise and distinguish legal cases: pro-plaintiff (pro-p) and pro-defendant (pro-d), favourable to plaintiff and defendant respectively. The pro-p and pro-d factors support one side and potentially influence the outcome of the case, although a single factor is neither necessary nor sufficient to decide a legal claim. A case is represented simply as a set of applicable factors[32].

CATO's case retrieval criterion narrows HYPO's BUC criterion. Only the BUC cases that have no significant distinctions will be selected. CATO's case retrieval criterion is therefore referred to as BUC/NoSignDist. After retrieving the most relevant cases, CATO supports eight basic argument moves, which is much more complicated than HYPO's 3-ply legal arguments.

The IBP model (Issue-Based Prediction)[33] has adapted the HYPO/CATO framework, which was originally developed for argumentation, to make predictions based on argumentation concepts. Most AI and Law models of case-based legal reasoning systems focus on making arguments and skip predicting the results of new cases. By contrast, the IBP model is designed to predict case outcomes. IBP still uses a set of CATO-style factors to represent legal cases. When IBP predicts the outcome of a new case, it combines case-based and model-based reasoning. It uses the Domain Model to identify the issues related to the case facts. It then determines for each of these issues, which side is favoured. In this step, IBP applies CBR

[31] Aleven, 'Using Background Knowledge in Case-based Legal Reasoning: A Computational Model and an Intelligent Learning Environment' (2003) Artificial Intelligence 150 183–237.

[32] Rissland and Ashley, 'A Note on Dimensions and Factors' (2002) Artificial Intelligence and Law 10 65–77.

[33] Ashley and Brüninghaus, 'A Predictive Role for Intermediate Legal Concepts' (2003) 16th Annual Conf on Legal Knowledge and Information Systems 153–162. Also see Brüninghaus and Ashley, 'Predicting Outcomes of Case-based Legal Arguments' (2003) 9th Intl Conf on Artificial Intelligence and Law 234–242.

if necessary. The algorithm makes its final prediction by combining the analysis of the issues following the logical structure in the Domain Model. Unlike the machine learning or statistical techniques, IBP's predictions are accompanied by explanations. Experiments confirm that legal case outcome predictions based on the case-based reasoning method are not only possible but also efficient.

3 Sub-elements for case representation

Traditionally, issues and factors are the basic elements to represent legal cases. During legal proceedings, lawyers frequently make arguments by analysing, assessing, abstracting and interpreting the significance of similarities and differences between cases[34]. Successful arguments depend critically on whether the cited precedents can convince the judge or court. The convincing precedents are those cases that have the same issues and that are more factually relevant to the current case. An important aspect of these arguments is the strategy to abstract points to describe and represent cases, i.e., issues and factors[35]. In this section, we propose the concepts of sub-elements for case representation.

3.1 Issue and sub-issues

In order to introduce the concept of sub-issues, we need to discuss the concept of issues first, including where the concept of issue comes from and what the implications of the concept are.

Legal issues rather than the factual issues are often the focus of the disputes between parties. Legal issues originate from the open texture of the law. Open texture means the uncertainty whether the terms in the legal predicates match specific facts[36]. Even simple legal rules generally contain some open texture[37]. For example, 'personal injury by accident in New Zealand' is a short legal term in the *Accident Compensation Act 1982, s26(2)(a)*. Although simple, this term may potentially generate many conflicts, such as whether the mental consequences of an incident without physical injury belong to the class 'personal injury', whether a failure in medical treatment within the reasonable failure rate belongs to the class 'by accident', whether a doctor in New Zealand improperly advising a New Zealander to receive an operation in a foreign country belongs to

[34] Ashley and Rissland, 'Waiting on Weighting: A Symbolic Least Commitment Approach' (1988) AAAI-88 239–244.

[35] Aleven, 'Using Background Knowledge in Case-based Legal Reasoning: A Computational Model and an Intelligent Learning Environment' (2003) Artificial Intelligence 150 183–237.

[36] Gardner, *An Artificial Intelligence Approach to Legal Reasoning* (Bradford Books/MIT Press: Cambridge 1987).

[37] Zeleznikow, Stranieri and Hunter, 'Beyond Rule Based Reasoning – the Meaning and Use of Cases' (1995) 11th Conf on Artificial Intelligence for Applications 292–298.

the class 'in New Zealand', and even whether a fetus belongs to the class 'personal'.

Legal precedents, to which the open-textured legal terms once applied, can be used to bridge the 'generality gap' between abstract legal terms and specific case facts[38]. Thus, the problem of matching specific case facts with open-textured legal terms is reduced to the simpler problem of matching two sets of specific case facts[39].

Many issues may arise in many different perspectives, such as what the relevant law is, how the relevant law should be interpreted in the context, how the law should be applied, what the facts of the case are, etc. Lawyers identify the issues in dispute, identify the relevant precedents for these issues, and organise their arguments according to these issues[40]. Finally, courts explain their decisions according to the disputed issues[40]. If a precedent contains irrelevant issues other than the disputed issues, normally these irrelevant issues will not be taken into consideration[41].

Many issues are logically compound and focus on the quite broad open textured nature of law. For example, 'whether the claimant suffered from the personal injury by accident' is a familiar issue in the New Zealand accident compensation proceedings. This logically compound issue consists of several subordinate issues. These sub-issues are related to the *Accident Compensation Act 1982, s2, 11, 26*, such as (i) whether the injury was 'personal', whether the injury was 'caused by accident', (ii) whether the accident was 'unexpected and unintentional' and (iii) whether the cause of the accident was 'identifiable or identifiable at the particular time'.

To precisely retrieve the factually relevant precedents, this kind of compound issue can be necessarily decomposed into more refined sub-issues. Figure 1 shows the sub-issues of the compound issue: 'whether the claimant suffered from the personal injury by accident', referred as Personal_Injury_By_Accident. The sub-issues focus on more detailed open textures and disputes, and facilitate the retrieval of the relevant precedents. Given a current case C that has the sub-issue Personal_Injury, if two precedents A and B are very relevant to C except that A has the same sub-issue Personal_Injury but B has a different sub-issue e.g. Injury_By_Accident, then A is potentially more relevant to C than B. Only when we cannot find such a case A, we will study the relevance between B and C, because after all B and C share the same issue.

---

[38] Porter, Bareiss and Holte, 'Concept Learning and Heuristic Classification in Weak-Theory Domains' (1990) Artificial Intelligence 45 1–2.

[39] Branting and Porter, 'Rules and Precedents as Complementary Warrants' (1991) 9th National Conf on Artificial Intelligence 3–9.

[40] Aleven, 'Using Background Knowledge in Case-based Legal Reasoning: A Computational Model and an Intelligent Learning Environment' (2003) Artificial Intelligence 150 183–237.

[41] Prakken and Sartor, 'Modelling Reasoning with Precedents in a Formal Dialogue Game' (1998) Artificial Intelligence and Law 6 231–287.

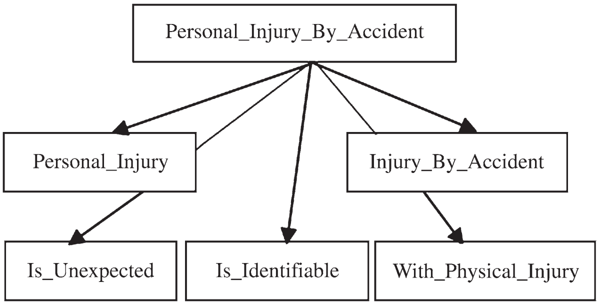

Figure 1. Sub-issues of 'Personal_Injury_By_Accident'

3.2 Factor categories and contextual features

A factor is a kind of abstract knowledge related to a stereotypical fact pattern of the case[42], and a case can be represented simply as a set of applicable factors[43]. In the HYPO, CATO and IBP models, the factors are categorised into two groups to represent, analogise and distinguish legal cases: *pro-plaintiff* (pro-p) and *pro-defendant* (pro-d), favourable to plaintiff and defendant respectively. The pro-p and pro-d factors support one side and potentially influence the outcome of the case, although a single factor is neither necessary nor sufficient to decide a legal claim[44].

The classification of factors in a case into pro-p and pro-d is argument-oriented. It facilitates the organisation of arguments for the new case, while it may present problems of knowledge representation by ignoring how strongly a factor can support one side[45]. Also, it ignores that in practice, the roles of one side as plaintiff or defendant are not fixed. For example, a plaintiff on a claim might be a defendant in a lawsuit who brings a counterclaim. Therefore, an applicable factor with the same underlying fact can be pro-p in one case and pro-d in another one. The lack of a uniform predefined standard to classify factors in precedents makes case retrieval more difficult.

We propose to classify all facts into three categories to represent and retrieve cases: *pro-claimant* (pro-c), *pro-respondent* (pro-r), and *contextual features*, where the *claimant* is the side who makes a claim for damages and the *respondent* is the side who is possibly liable for the accident or

[42] Aleven, 'Using Background Knowledge in Case-based Legal Reasoning: A Computational Model and an Intelligent Learning Environment' (2003) Artificial Intelligence 150 183–237.

[43] Rissland and Ashley, 'A Note on Dimensions and Factors' (2002) Artificial Intelligence and Law 10 65–77.

[44] Aleven, 'Using Background Knowledge in Case-based Legal Reasoning: A Computational Model and an Intelligent Learning Environment' (2003) Artificial Intelligence 150 183–237.

[45] Rissland and Ashley, 'A Note on Dimensions and Factors' (2002) Artificial Intelligence and Law 10 65–77.

compensation damages. This classification gives a uniform predefined standard to label the factors and at the same time keeps the argument-oriented characters of the pro-p and pro-d factors.

The *pro-c* and *pro-r* factors are those important facts favourable to the claimant and respondent respectively. Unlike the pro-p and pro-d factors, whether an applicable factor is pro-c or pro-r can be pre-decided. For example, the factor *claimant has been warned of the potential risk* in accident compensation domain will support the respondent, no matter the role of respondent in cases as plaintiff or defendant.

Contextual features are those important facts that are neutral to both parties. It may seem surprising that a feature can be viewed as neutral, when judicial decision-makers do not have the luxury of making a neutral decision. However Lexis-Nexis explicitly cites the existence of neutral factors[46]. Neutral factors are ones that whilst important do not affect the outcome of a particular case.

The contextual features are different from the pro-c or pro-r factors that are favourable to one side or the other. From the viewpoint of deciding who wins a legal argument, contextual features can be ignored. Thus, traditionally all factors are classified positive or negative to one side and there is no such examination of contextual features.

Contextual features are not surface features, but factual and contextual knowledge related to the evidence. Although contextual features do not affect the outcome of a particular case, the important contextual features still strongly characterise particular kind of cases and potentially imply the presence of particular legal issues. Therefore, contextual features are useful for case retrieval.

For example, the contextual feature *Role of Respondent* (RR) evidentially supports neither party. However, for those accident compensation cases in which the roles of the respondents are 'police', we usually find out in these case facts that the claimants had suffered personal injury during a warranted police action. Therefore, these cases are very much likely to have the same argument points and even refer to the same clauses of a concerned act. For another example, the contextual feature *Compensation Type* (CT) refers to the kind of compensation that the claimants claim. Those accident compensation cases whose claim types are to claim the 'backdated payments' are very much likely to have a common issue, i.e., whether a claim lodged late is still time effective.

Detailed examples about sub-issues, pro-c/pro-r factors, and contextual features are given in Section 5. Figure 2 concludes the extension from the traditional representation elements to the proposed set of sub-elements.

---

[46] See for example Lex K. Larson (2006) Larson's Employment Discrimination, Matthew Bender. In this manual Neutral Factors Peculiar to Sex Discrimination and Neutral Factors Peculiar to Origin are discussed.

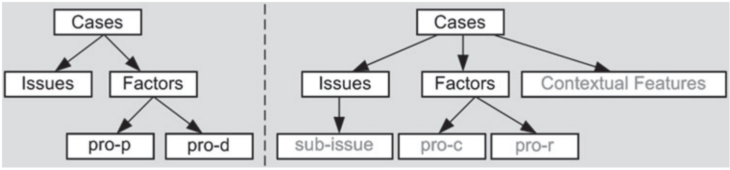

Figure 2. Extension of the traditional representation elements

4 Representation model and retrieval scheme

4.1 Representation model

The purpose of our model for legal precedent retrieval is to effectively represent and precisely position knowledge and information contained in the precedents, based on the concept of introduced sub-elements.

Figure 3 shows the four levels of the representation model. At the top is the *source level*, where a collection of raw legal precedents exists as plain text. A new case also belongs to this level. The second level is the *representation level*, where the representation elements, such as issues, sub-issues, pro-c, pro-r and contextual features, are drawn from each case. The third level is the *repository level*, where the legal precedents are stored in a case base according to the abstract representation. On the bottom is the *output level*, where the most factually relevant precedents are captured through a retrieval scheme.

4.2 Retrieval scheme

Based on the proposed representation model, we have developed a retrieval scheme for the intelligent retrieval system. When a new case arises, users always want to find the factually relevant precedents that are similar in all or most representation elements. However, it is impossible to find a 100% similar precedent with the new case due to the factual diversity of actual cases. In that case, those precedents sharing one or more representation elements with the new case are desired as they are potentially useful for making legal arguments. Given a new case and a large precedents database, we develop the following scheme to identify those relevant precedents step by step.

Step 1: narrow the range based on issues and sub-issues. First, narrow the range to those precedents sharing the same issues with the new case. Obviously, those precedents sharing issues with the new case are considered to be relevant. Those precedents sharing no issues with the new case are irrelevant and can be safely excluded. Second, if the issue has sub-issues, narrow the range further to those precedents sharing the same sub-issues with the new case. Obviously, the more sub-issues they share, the more factually relevant they are. We use I_Counter to record

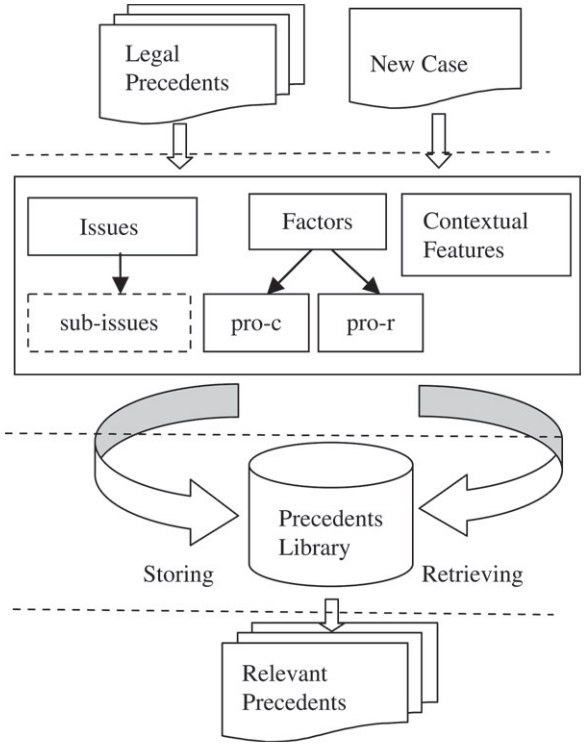

Figure 3. The knowledge representation model for legal precedent retrieval

how many issues and sub-issues a precedent shares with the new case. When the precedent shares an issue that has no sub-issues, I_Counter increases 1 point. When the precedent shares an issue that has sub-issues, I_Counter increases 1 point for each shared sub-issue; otherwise if the precedent shares no sub-issues with the new case, I_Counter only increases 0.5 point because although they share the same issue, they do not share a same sub-issue.

Step 2: narrow the range further to those precedents sharing the same pro-c or pro-r factors with the new case. While under the same issue or sub-issue, the more factors they share, the more factually relevant they are. We use P_Counter to record how many pro-c and pro-r factors a precedent shares with the new case. Simply, when the precedent shares a factor with the new case, P_Counter increases 1 point.

Step 3: narrow the range further to those precedents sharing the same contextual features with the new case. As already mentioned, sharing the same contextual features will demonstrate additional conviction that a precedent has the similar context with the new case. We use F_Counter to record how many contextual features a precedent shares with the new case. Simply, when the precedent shares a contextual feature with the new case, F_Counter increases 1 point.

The three counters (I_Counter, P_Counter and F_Counter) will be used to sort the identified precedents that are more or less factually relevant to the new case. According to the importance decreasing from I_Counter, P_Counter to F_Counter, the identified precedents will be sorted in the order of I_Counter, P_Counter and F_Counter. For any two identified precedents, the precedent with the higher value in I_Counter deserves more attention than the other precedent with the lower value in I_Counter. Further, if the two precedents have the same value in I_Counter, the precedent with the higher value in P_Counter deserves more attention than the other precedent with the lower value in P_Counter. Again, if the two precedents have the same I_Counter and P_Counter, the precedent with the higher value in F_Counter deserves more attention than the other precedent with the lower value in F_Counter. For simplicity, we refer to this scheme as the *IPF* scheme.

Note that the scheme proposed intends only to locate and sort the most factually relevant cases rather than to select a case in legal arguments. Whether a precedent will be cited for the arguments and judicial opinions also depends on other considerations like outcomes, potential relevance, possible transformations of the cases, etc.[47] However, the IPF retrieval scheme provides the abstracted representation elements which can provide additional knowledge for users to explain why a precedent is considered factually relevant. The predefined factors (i.e., pro-c and pro-r factors) are convenient for users to organise arguments. The experiment in next section will demonstrate the effectiveness of the IPF scheme.

5 Experiment and discussion

5.1 The precedents database

We have built a legal case base, which currently contains 152 landmark accident compensation cases, reported by Court of Appeal or High Court of New Zealand from 1958 to 2003. These cases have been represented by a set of elements: issues, sub-issues, factors and contextual features. Totally, there are 36 issues. Five of them have sub-issues. For example, the sub-issues of the issue 'whether the claimant suffered personal injury by accident' are: 'whether the accident must be unexpected', 'whether the accident must be identifiable or identifiable at the particular time', 'whether the external event is required for the injury', 'whether mere mental consequence belongs to personal

[47] Ashley and Rissland, 'A Case-Based Approach to Modelling Legal Expertise' (1988)IEEE Expert: Intelligent Systems and Their Applications 03(3) 70–77.

injury', and 'whether a pregnant mother suffers personal injury due to the injury of the fetus'. For another example, the sub-issues of the issue 'whether the claimant is entitled to claim the compensatory damages' are: 'whether the claimant is entitled to claim the backdated compensatory damages', 'whether the personal representative of the deceased is entitled to claim compensatory damages', 'whether the patient suffered from medical misadventure is entitled to claim compensatory damages', 'whether the next of kin of the deceased is entitled to claim the compensatory damages for mental injury', and 'whether person independent on the deceased is entitled to claim the compensatory damages'.

Besides the issues and sub-issues, there are 16 pro-c factors and 29 pro-r factors. The pro-c factors are those key facts that support the claimant's claim for the compensatory damages. Examples of the pro-c factors include: the claimant had not been warned of the risk, the claimant lives dependent on the deceased, the claimant suffered injury in course of employment, the accident was caused by unlawful office action, etc. The pro-r factors are those key facts that help the respondent to avoid or moderate his liability for the claimed compensatory damages. Examples of the pro-r factors include: the respondent's action is lawful, the claimant suffered the injury in course of committing criminal offence, the deceased committed suicide, the claimant has not lodged on claim timely, the claimant suffered no direct physical injury, the accident was caused by the claimant's own fault, etc.

Last, there are four categories of contextual features. They are *Role of Claimant or Victim* (RC), *Role of Respondent* (RR), *Injury Reason* (IR) and *Compensation Type* (CT). These contextual features support neither side and have no decisive effect to the case outcome, but these contextual features strongly indicate the context and background of a particular kind of cases. RC and RR indicate the special roles of the two parties, such as patient, pregnant mother, spouse, representative of the deceased, next kin, police, etc. IR indicates the reason for the personal injury, such as occupational syndrome, medical misadventure, mental injury, etc. CT indicates the kind of compensation that the claimant claims for, such as backdated payments, rehabilitation, earning-related compensation, etc.

These issues, sub-issues, factors and contextual features have been identified manually after studying the 152 accident compensation cases.

5.2 Experiment

When a new case comes, users (personal injury lawyers or judges) first analyse the case facts and identify its issues, sub-issues, factors and contextual features. After that case analysis stage, the proposed retrieval scheme can be used to help the users to effectively locate the most relevant precedents. Here we first give one example as follows to demonstrate how the IPF scheme works.

We take the latest case in the database as a 'new' case. Since the cases in the database are sorted by ascending date, the latest case is Case 152 *Harrild v Director of Proceedings*[48]. The details of *Harrild* are abstracted in Figure 4.

After identifying the representative elements of *Harrild*, the system goes through the rest 151 cases with the IPF scheme. After the three steps, finally there are 12 precedents that are identified to be more or less relevant to *Harrild*. These 12 precedents are sorted by their values of the three counters (I_Counter, P_Counter and F_Counter) (see Table 1). From the table, we can clearly see the most relevant precedents may be Cases 114 and 105. Actually, Case 114 *Accident Compensation Corporation v E*[49] was cited in the judicial opinion in *Harrild*. The details of Case 114 are abstracted in Figure 5.

The main dispute of both Case 114 and *Harrild* is how to interpret the clause on 'personal injury by accident'. They share the same issues 'whether the claimant suffered personal injury by accident'. They also share one same sub-issue 'whether mere mental consequences belong to personal injury'. Moreover, in these two cases, both accidents occurred without direct physical injury. This is a contextual feature that strongly characterises certain kind of cases, whose potential issue is whether the case arises within the meaning of 'personal injury by accident'.

The reason why another very relevant Case 105 is not cited in the judicial opinion in *Harrild* is that Case 114 is more recent and Case 105 is already cited in the judicial opinion of Case 114.

Besides Case 114, Case 115 at the fifth place in Table 1 is also mentioned in the judicial opinion in *Harrild*.

Since all the cited precedents (Case 114 and 115) have been retrieved by the IPF scheme, we say the cover rate is 100% for Case 152. Also, one precedent (Case 114) is ranked by IPF as at the first place and the other precedent (Case 115) is ranked by IPF as at the fifth place. Both are in the top of the ranking list. From this experiment, we can clearly see that the IPF scheme can effectively identify the most relevant precedents for Case 152.

5.3 Statistical analysis

We have applied the retrieval scheme to the 152 precedents in the same way introduced above, in order to test whether for each 'new' case the precedents retrieved with the scheme were actually cited in the judicial opinion.

We first examine the overall cover rate. There are totally *109* precedents are cited in the judicial opinions of all the 152 cases (e.g. there are two

---

[48] Court of Appeal Wellington, [2003] 3 NZLR 289.
[49] Court of Appeal Wellington, [1992] 2 NZLR 426.

Harrild v Director of Proceedings

The claimants are parents of a stillborn child. Dr Harrild, as one of the respondents, is a medical specialist. Dr Harrild had provided an <u>inadequate medical service</u> [**feature: IR**] to <u>the pregnant mother</u> [**feature: RC**], resulting in the death of <u>the fetus</u> [**feature: RC**]. Obviously the parents including the mother have <u>suffered the mental injury</u> [**feature: IR**] due to this unfortunate. <u>No direct physical injury occurred to the mother</u> [**pro-r**].

The principal issue raised by the proceedings was <u>whether the mother had suffered personal injury with only mental injury</u> [**Sub-Issue**] on account of the death of the fetus <u>without clear physical cause to the mother</u> [**pro-r**]. Also, <u>whether the ordinary meaning of the word 'person' can be applied to an unborn child</u> [**Sub-Issue**], because a fetus could be part of the mother, even though an unborn child was not a legal person, and a mother and fetus were biologically distinct.

Figure 4. Harrild v Director of Proceedings

precedents are cited in the judicial opinion of Case 152). The IPF scheme has successfully retrieved *105* of them (e.g. the IPF scheme has successfully retrieved two for Case 152). The overall cover rate is *(105/109 =)* 96.3%. This means almost all factually relevant precedents can be retrieved by IPF. Reasons for omission of the remaining four cases are: two cases are just mentioned as exceptions that should not be followed; and the other two are mentioned in *obiter dicta* (other things that judges said by the way in the course of deciding a case[50]). Unlike the *ratio decidendi* part of a decision, it is too flexible to predict which precedents would be used for the *obiter dicta* part.

Only a high cover rate is not enough to demonstrate the effectiveness of the scheme. We also examine how precise IPF is by examining the places of the retrieved 105 precedents in the ranking list. We have mentioned that for Case 152, one cited precedent is ranked by IPF at the first place and the other at the fifth place. Obviously, with an effective representation mechanism and retrieval scheme, those cited precedents should be

[50] *Study of Law* (Carswell/Thomson Professional Publishing: Ontario 1992).

Table 1. The relevant precedents for *Harrild*

Case	I Score	P Score	F Score
114	***1.5***	***1***	***1***
105	***1.5***	***0***	***1***
134	*1*	*1*	*0*
110	*1*	*0*	*1*
115	*1*	*0*	*1*
140	*1*	*0*	*0*
119	*1*	*0*	*0*
113	*1*	*0*	*0*
109	*0.5*	*1*	*0*
111	*0.5*	*0*	*0*
104	*0.5*	*0*	*0*
42	*0.5*	*0*	*0*

ranked as top as possible. For all the 105 cited and retrieved precedents, 46 precedents are ranked at the first place *(43.8%)*, 32 precedents are ranked at the second place *(30.5%)*, 12 precedents are ranked at the third place *(11.4%)*, and 11 precedents are ranked at the fourth place *(10.5%)*. Thus, totally 101 precedents are ranked within the top four *(43.8% + 30.5% + 11.4%+ 10.5% = 96.2%)*. This means generally the actually cited precedents are included within the top four places. Therefore, we believe the retrieval scheme can greatly free users from the burdensome works of searching precedents.

5.4 Discussion

5.4.1 Legal ontologies

Breuker et al.[51] claim that unlike engineering, medicine or psychology, law is not ontologically founded. They claim law is concerned with constraining and controlling social activities using documented norms. They have developed a core upper level ontology LRI-core. This ontology has over 200 concepts and has definitions for most of the anchors that connect the major categories used in law – person, role, action, process, procedure, time, space, document, information,

[51] Breuker, Elhag, Petkov and Winkels, 'Ontologies for Legal Information Serving and Knowledge Management' (2002) Jurix 2002: 15th Annual Conf Legal Knowledge and Information Systems 73–82.

Accident Compensation Corporation v E

E was employed in a large organisation where she was a senior staff member holding a position of responsibility. In 1979, E was sent by her employer to a management course. She was the first female employee of the organisation to attend that course. Her fellow employees warned her of the rigors of the course [**pro-r**]. After E had been on the course for four days she suffered a psychiatric breakdown [**feature: IR**] and was admitted to the psychiatric unit of Wellington Hospital. There was no particular physical incident or event [**pro-r**] during the course that could be identified as having triggered the psychotic episode. E had since been in receipt of psychiatric attention for depressive symptoms and had terminated her employment with the organisation for health reasons. E lodged a claim for compensation under the Accident Compensation Act. The Accident Compensation Corporation rejected the claim. That decision was overturned in the High Court. The corporation appealed by way of case stated. The questions posed for determination in the case stated were:

(a) Whether the particular injury suffered by the claimant in the present case properly arises within the meaning of the words 'personal injury by accident' [**Sub-Issue**]....

(d) Whether mental consequences or disturbance not accompanied by direct physical injury [**pro-r**] to the claimant can come within the meaning of the words 'personal injury by accident' [**Sub-Issue**].

Figure 5. Accident Compensation Corporation v E

intention, and so on. The main intended use is supporting knowledge acquisition for legal domains, but a real test of its semantics is whether it enables natural language understanding of common sense descriptions of simple events, as in the description of events in a legal case documentation. This is of course the core principle of the Semantic Web initiative of WC3.

An ontology is defined as an explicit conceptualisation of a domain[52]. Bench-Capon and Visser[53] has examined the development of legal ontologies. Ontologies have benefits for:

- Knowledge sharing;
- Verification of a knowledge base;
- Knowledge acquisition; and
- Knowledge reuse.

At a future stage in this project we will be constructing an ontology for modelling Accident Compensation in New Zealand. Currently we are focusing upon intelligent legal retrieval rather than carefully modelling domain knowledge.

5.4.2 Related conceptions

The sub-issues in our model focus the subordinate and finer open texture or dispute of an issue. From the viewpoint of retrieval, the links between an issue and its sub-issues are that an issue consists of the sub-issues.

The term 'sub-issue' is also used in the legal negotiation and decision-making systems like AdjustedWinner[54] and Family_Winner[55]. In their systems, however, the purpose to decompose an issue into sub-issues is to seek possible agreement within the issue when the issue cannot be resolved in its current form.

In the Factor Hierarchy of CATO, the terms 'high-level factors' and 'abstract factors' are used to refer to the Legal Issues and Intermediate Legal Concerns, where the latter are subordinate to the former. The links between Intermediate Legal Concerns and Legal Issue is an evidential support relation, from which the system will organise the argument.

In IBP's Domain Model for legal outcome prediction, an issue is not logically connected with CATO's factors. Instead, the knowledge about an intermediate legal concept is captured. However, the links between issues and intermediate legal concepts is logic AND and OR.

Contextual features are different from surface features. Contextual features are factual and contextual knowledge related to a fact of a

---

[52] Gruber, 'Towards Principles for the Design of Ontologies used for Knowledge Sharing' (1995) Intl J Human-Computer Studies 43 907–928.

[53] Bench-Capon and Visser, 'Ontologies in Legal Information Systems: the Need for Explicit Specifications of Domain Conceptualizations' (1997) 6th Intl Conf Artificial Intelligence and Law 132–141.

[54] Bellucci and Zeleznikow, 'A Comparative Study of Negotiation Decision Support Systems' (1998) 31st Hawaii Intl Conf System Sciences 254–262.

[55] Bellucci and Zeleznikow, 'Representations of Decision-making Support in Negotiation' (2001) J Decision Systems 10(3–4) 449–479.

case. Contextual features strongly characterise certain kind of cases and potentially imply the presence of certain legal issues, although not necessarily. Therefore, contextual features are useful for case retrieval.

In IBP's Domain Model, CATO's Factors are further distinguished into two groups, knock-out factors and weak factors, depending on the factor's predictive strength for the outcome. Those weak factors cannot provide sufficient evidence for issue-based prediction and therefore are ignored by the program for the particular purpose of determining whether there is conflicting evidence on a particular issue. Clearly, even a weak factor in IBP model is different from the contextual features in our model as the weak factor is not neutral.

5.4.3 Working with traditional retrieval technologies

By restricting retrieval as the filtering of cases based on requested issues, factors and contextual features, one may restrict severally the search possibilities of the users. Sometimes users may still want to use traditional retrieval methods, such as alphabetical indexing, keyword-based searching and searching by some surface features (e.g., title, location, publication date and document number).

This problem can be solved by integrating this proposed filter with traditional retrieval technologies. For example, users can first search by surface features, then search within the returned cases based on issues, factors and contextual features. Alternatively, users can give the system keywords, factors and issues at the same time, the system will return the relevant cases and notify if these cases contain the keywords.

5.4.4 Managing the representation elements

As new cases are continuously stored into the legal case base and the legal statutes can be changed, new issues, factors and contextual features are likely to arise. A contrary situation is that some dated representation elements may not appear any more. These kinds of changes reflect the changes of social policies, interests and values. For example, in some legal proceedings in 1950s and 1960s, it is not uncommon to see the arguments about whether the illegitimate children can have the same right with those children within marriage to claim for damages. And now in many countries, this is not a question any more. This reflects the change of social view on the equal right of all the children.

For these new or dated factors and issues, a retrieval system may provide the administrator with the following functions: (i) to add new factors and issues; (ii) to add new sub-issues and link them to their super-issues; (iii) to note dated factors and issues as 'deleted'; and (iv) to physically delete the dated factors and issues.

5.4.5 Perspectives for the retrieval systems

As mentioned above, most of the traditional commercial retrieval systems, like LexisNexis, provide alphabetical indexing, keyword-based searching and searching by surface features. The retrieval method developed in this paper employs AI technologies to mimic the lawyers' real case-based searching logic. Therefore, it conforms to lawyers' working custom more than the traditional searching and browsing methods. Therefore, there is potential to integrate our technology in commercialised legal retrieval systems.

Besides the commercialised potential, we believe this system can be used further develop and integrate in other research in the area of AI and Law. Currently, we are integrating this retrieval system into our intelligent tutoring system for legal argument.

6 Summary

A knowledge representation model for effective and precise legal case retrieval is needed to aid the everyday work of the lawyers and judges. Based on existing representation elements, we have proposed a set of sub-elements to better reveal factual relevance between cases. First, we have introduced the concept of sub-issues. Second, we have adapted the HYPO-style factors to our pro-claimant and pro-respondent factors. Also, we have introduced a new category of contextual features. Based on this set of sub-elements, we have proposed a new model for legal precedent retrieval and the IPF retrieval scheme. Experiments and statistical analysis have shown that the model and IPF scheme can achieve satisfactory results.

Future research might focus upon using the model to improve legal argument and reasoning systems.