Case Retrieval through Multiple Indexing and Heuristic Search

Edwina L. Rissland, David B. Skalak and M. Timur Friedman

Department of Computer Science
University of Massachusetts
Amherst, MA 01003, USA

rissland@cs.umass.edu

Abstract

We discuss the indexing of cases for use in precedent-based argument. Our focus is on how multiple, related indices into a case base of legal precedents are exploited by an argument-generation program called BankXX. This system's architecture and control scheme are rooted in a conceptualization of legal argument as heuristic search. Although our framing argument as search is not discussed in detail in this paper, we describe the main features of this view to provide context for a discussion of an indexing scheme that facilitates argument creation. We describe five inter-related index types—citation, prototypical story, factor, family resemblance, and legal theory indices—and show how they can be used to access, view, widen, or filter a set of cases. The application domain is a U.S. Federal statute that governs the approval of bankruptcy plans.

1 Introduction

The thesis of this paper is that different access paths to information can be found by applying related types of indices in complementary ways and that composite indexing strategies can lead to improved case retrieval. We explore these hypotheses in the context of legal argument using a test set of five related index types into a case base of legal precedents.

Legal experts working in Anglo-American jurisdictions index cases in a variety of ways, using:

(1) Traditional citation linkages between cases. A case cites the precedent cases to which it refers for support of the various legal propositions advanced. Indexing services such as Shepard's Citations [Shepard's, 1992] track the citations to many published cases. Terms such as *see, but see,* and *cf.* are reserved to introduce precedents in precise ways that are specified in legal style manuals such as *The Blue Book* [BlueBook, 1986]. *See* [Ashley and Rissland, 1987].

(2) Recurring prototypical fact situations. Generic cases or recurring fact patterns have been used in legal reasoning [Gardner, 1987] as well as in other domains [Kolodner, 1983]. If, for example, the current problem involves a former student with large educational debts who files for bankruptcy immediately after graduating, a bankruptcy expert may recall previous student loan cases. In particular, knowing that a legal theory has successfully been applied to cases fitting a particular fact pattern provides a basis for creating analogies to justify applying that theory.

(3) HYPO-style dimensions. In domains where cases can be compared with respect to a stable set of discernible factors, the factors can be conceptualized and implemented as "dimensions" [Rissland, Valcarce and Ashley, 1984; Ashley, 1990]. Dimensions may be used to index and retrieve cases from a case base and to order precedents by their relevance to a problem situation as in the HYPO and CABARET systems [Ashley, 1990; Rissland and Skalak, 1991].

(4) Family resemblance or prototypicality indices. Measures of family resemblance and prototypicality originating in part from the psychological research of Rosch can be used to create a conceptual, graded landscape of cases in which the highest peaks or most centrally located instances represent the cases with the greatest family resemblance rating [Rosch and Mervis, 1975; see also McCarty, 1983; Bareiss, 1989]. Given a case family, one can find member cases of varying degrees of prototypicality. Given an individual case, one can assess the strength of its membership within a family.

(5) Legal theories. Courts and advocates usually strive to provide a legal theory as to why a case should be decided a certain way. Legal theories often explicitly specify factors or other features to be considered. Knowing what cases have been argued under a theory is a means to access other cases, such as cases in which a theory was clearly held to control a decision. In addition, relations between legal theories themselves, such as refinement, permit "nearby" theories to be retrieved, along with the cases that apply them.

(6) Rules. There is no shortage of rules in the law [Twining and Miers, 1982]: statutes; agency regulations; "blackletter" rules, which are generalizations of case law found in restatements of the law; and the rules of a case, which state the holding of the case. Each type of rule provides a means

to access cases: for example, the cases leading to the rule, the cases elaborating the rule, and those following the rule.

(7) Domain taxonomies. Commercial publishers have also developed indexing schemes, such as the key number system used in WestLaw [West, 1992], in which legal topics are assigned key numbers. Such schemes provide a useful taxonomy of the law and index legal opinions by the topics they address.

(8) Terms of art. Through dictionary, digest, and encyclopedia entries legal practitioners find cases that define, interpret, elaborate and refine the meaning of legal terms whose scope is often the source of litigation (*cf.* "dictionary-based reminding" [Schank, 1982]).

Each of these index types imparts its own emphasis or "view" on the case base and displays its own strengths and weaknesses as an indexing medium. In the next sections we discuss five of these indices in further detail — citation linkages, recurring prototypical fact patterns, factors, family resemblance, and legal theories — and show how they may be used in conjunction with each other to improve case retrieval[1].

In Section 2 we provide a context for our discussion of indexing by describing how argument creation can be couched as heuristic search. We also provide a brief introduction to the domain in which we have instantiated our system, personal bankruptcy plans under Chapter 13 of the U.S. Bankruptcy Code. In Section 3, we discuss the particular indices used by our system, called BankXX. We show in Section 4 how multiple, related indices can be used to overcome some practical problems in case retrieval — such as too many or too few cases retrieved — and provide examples of retrieval using BankXX's network of indices. We show how the composite use of related indices can outperform a single index. Section 5 discusses related research and summarizes contributions.

2 Background

2.1 Argument Formation as Search

One can view the generation of argument as heuristic search. At each stage in developing an argument, choices need to be made. Should one seek a broad set of supporting cases, anticipate the best cases for the opposing side, or create a telling hypothetical? Each choice takes the emerging argument to a new state of development. Limited resources force the arguer to make choices about which avenues to pursue.

In one implementation of argument as search, the search space would be the space of all arguments, the start state would be an empty argument, and the search operators would represent ways to advance the argument. However, our current system models argument as the emerging by-product of the search and research that an expert might perform in a space of domain knowledge. We perform search in domain space rather than in argument space in part because we are interested in modeling the legal research activities of attorneys and partly because the indexing fabric of the domain space is better understood. In our approach, domain space search identifies nodes that contain domain knowledge that can support an argument and collects the support in an argument data structure.

In particular, this project uses classical best-first search [Nilsson, 1980] in a domain space consisting of legal theories and of cases viewed in various ways. The heuristic evaluation function driving the search is based on a checklist of desired ingredients in an argument. The 12 argument *desiderata* implemented in BankXX include supporting cases, supporting legal theories, supporting citations, most on-point and best cases according to a factor analysis, classification of the case as a prototypical story, and favorable cases with the highest family resemblance rating. How supporting, most on-point and best cases are used in case-based legal argument has been investigated in HYPO [Ashley, 1990]. The argument *desiderata* are analogous to the higher level features often identified in applications such as game playing (e.g., [Samuel, 1967]'s "center control" or "piece advantage"). In the current BankXX implementation, each argument piece contains a functional predicate that determines if a case node can supply that useful piece of an argument.

Thus, BankXX's implicit control paradigm is search in the space of arguments. Its explicit, implemented control scheme uses heuristic search in a domain space to create arguments through the exploitation of a variety of indices.

2.2 The Chapter 13 Bankruptcy Domain

Our system searches a domain network of cases and other legal knowledge to support arguments in favor of or against granting an individual relief from debts via a partial repayment plan through Chapter 13 of U.S. Bankruptcy law (11 U.S.C. §§ 1301-1330). Chapter 13 was enacted to permit overwhelmed debtors with regular income to rehabilitate themselves by repaying part of their debt and having the balance discharged, thus allowing debtors a fresh start.

3 Description of BankXX

The purpose of BankXX is to experiment with the connection between (a) argument generation realized as heuristic search and (b) case retrieval accomplished with a variety of indexing strategies. In this section we present details of the case representation and indexing schemes used in BankXX. In particular we describe how various case perspectives can be reflected in a partition of case memory into spaces of interconnected cases.

[1] We have previously addressed the role of rules as indexes to legal cases in [Rissland and Skalak, 1991] and [Skalak and Rissland, 1992].

3.1 Case Spaces

The case base in BankXX consists of a semantic network whose nodes represent cases and legal theories, and whose labeled links represent connections between the nodes. We refer to this network as the *case graph*, which consists of *case-graph nodes* together with labeled link edges.

Case-graph nodes are legal cases represented: (1) as factual situations, (2) as bundles of citations, (3) as stereotypical stories or scripts, (4) in terms of various legal factors, and (5) by the measure of their prototypicality. Legal theories are also represented as case-graph nodes.

The case graph is partitioned into spaces, similar to the partition of a blackboard application's working memory into spaces. Each space contains case-graph nodes that represent cases (or legal theories) according to a particular perspective that has proven useful to human legal reasoners (Figure 1). Nodes in the case graph are highly interconnected; in-space links connect objects within a space and cross-space links connect objects in different spaces. During search of the case graph, links are traversed by *neighbor methods*, operators that expand nodes in the graph by following either in-space links, cross-space links or a sequence of links. Traversing a link is tantamount to using the link label as an index.

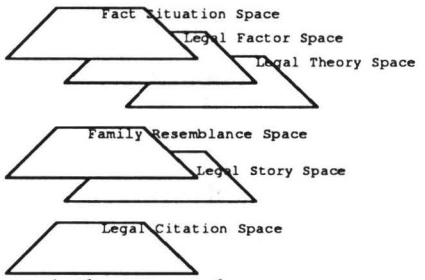

Figure 1. Spaces in the case graph.

We now describe each space of case-graph nodes, including its intra-space indexing links. Cross-space links are discussed in Section 3.2.

Fact Situation Space. Fact situation case nodes encode legal cases as sets of facts. Each situation is represented as a tree of frames implemented as CLOS instances. Examples of frames at this level describe the proposed plan and payments, the debt, the debtor's income, and generic information about the case. This level, which is the level at which cases are input, is the surface level of factual description. Cases at this level of representation are linked to each other through case citations (Figure 2).

Legal Citation Space. Citation case nodes encode cases as sets of citations: those that are cited by a case and those (later cases) that cite it. Citations include a citation signal that specifies the sense in which a case is cited. Citation instances also include the West key numbers indexing the topics addressed in the case. Citation instances are not linked to each other but provide links to the citing case and the cited case in Fact Situation Space (Figure 2).

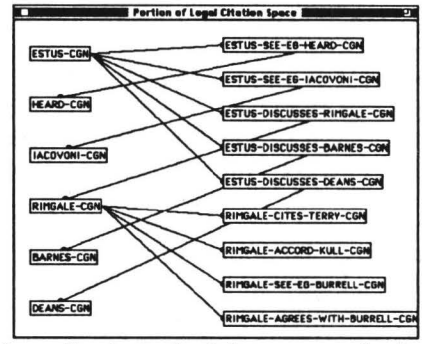

Figure 2. A small subgraph of the case graph, showing cases and inter-case citation links.

Legal Factor Space. Legal cases can be represented in terms of their values on domain dependent factors [Rissland, Valcarce and Ashley, 1984; [Ashley, 1990]. Factors are derived features recognized by domain experts as strongly influencing a case's outcome. A factor compares cases as stronger or weaker with respect to the factor's perspective. In Legal Factor Space, a case is represented by a vector composed of the magnitude of a case on each dimension.

Examples of factors in BankXX are the percent of disposable income that is allocated to payments under the plan, the length of proposed plan, and the portion of the debt that is attributable to educational loans. Using a HYPO-style analysis [Ashley, 1990], the system creates links between factors and the cases to which they apply (Figure 3).

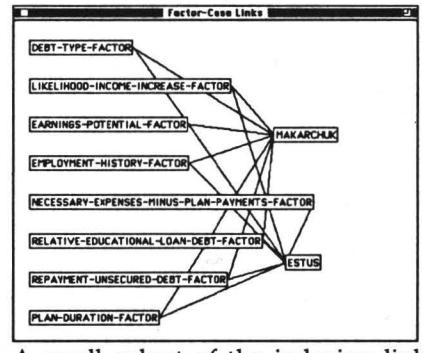

Figure 3. A small subset of the indexing links between domain factors and cases.

Legal Story Space. Two of the bankruptcy story prototypes used by BankXX are: (1) *the student loan story* — student incurs educational debts and soon after graduating files for bankruptcy protection from his educational loan creditors, and (2) *the dishonest debtor* — debtor commits fraud or some other offense, a judgment is entered against debtor, debtor files for bankruptcy.

BankXX does not link story prototypes to each other. Exploiting such links would require an understanding of how stories can be related and,

ideally, an automated means to recognize them (e.g., plot units [Lehnert, 1981]).

Family Resemblance Space. We have begun to incorporate some of the research of Rosch, who proposed a model of the internal structure of categories that is captured in the family resemblance hypothesis: "the most prototypical members of categories are those with most attributes in common with other members of that category and are those with least attributes in common with other categories" [Rosch and Mervis, 1975, p.576]. While Rosch proposed the family resemblance hypothesis as a cognitive structural model, we are experimenting with family resemblance as an indexing and processing model. BankXX can calculate the degree of family resemblance of a case to a given set of cases and select the cases within that family having the greatest family resemblance. For instance, the system can calculate the family resemblance of all student loan cases, and find the most prototypical.

Legal Theory Space. Legal theories are represented as a list of factors (see the discussion of Legal Factor Space) that are necessary to determine how a theory applies to a case. Legal theory nodes are linked by pointers that describe the relationships between them (Figure 4), such as "overlaps with," "rejects," and "agrees with." In the next implementation, a legal theory node will specify how to combine the factors — for instance, via weighting — to apply the theory. Legal theories have been culled from opinions by hand.

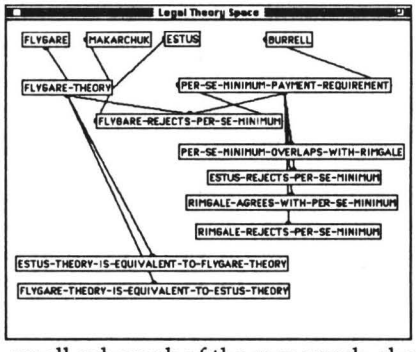

Figure 4. A small subgraph of the case graph, showing inter-theory links and links from theories to cases.

3.2 Cross-space Case Links

In addition to the in-space links we have described, a variety of bi-directional, cross-space links exist. For instance, links exist between factors and legal theories that use those factors, and between story prototypes and cases instantiating them.

We also have identified sequences of links that may be traversed serially to yield cases that satisfy a set of criteria. These multiple link pathways function as domain-dependent macro-operators [Fikes, et al., 1972]. Examples are given in Section 4.

BankXX currently contains 54 bankruptcy cases (and the same number of factor analyses in Legal Factor Space), 70 inter-case citations using 6 citation signals, 27 domain factors, 18 legal theories related by 21 inter-theory links of 8 types, 4 prototypical stories, and 2 measures of family resemblance. A mouse-sensitive graphical interface displays the case graph from user-specified indexing perspectives.

3.3 Control Flow and System Features

The control flow of BankXX is straightforward in that it follows from our framing argument as classic heuristic search, as described in Section 2.1.

The user inputs a frame-based description of a debtor's situation, including a description of the repayment plan and the debt. Unless the user specifies a start node in the case graph (e.g., a legal theory to be tested), the system begins by analyzing the problem in terms of its domain factors. After this analysis, which provides a starting node in Legal Factor Space, BankXX begins search of the case graph. At each iteration in the search, a node is expanded by the neighbor methods and the generated nodes are placed on the list of open nodes; the case-graph node with the highest evaluation score is removed from the open list. (BankXX's evaluation functions are discussed in detail in [Rissland, Skalak and Friedman, 1993].) This best node is passed to the argument pieces, which attempt to use that node for support. Search iterations continue until a user-provided space or time bound is exceeded (the time bound is measured in "billable seconds" in keeping with law firm time accounting) or until the open list maintained by the search algorithm is empty. The argument pieces are output in their final state (Figure 5).

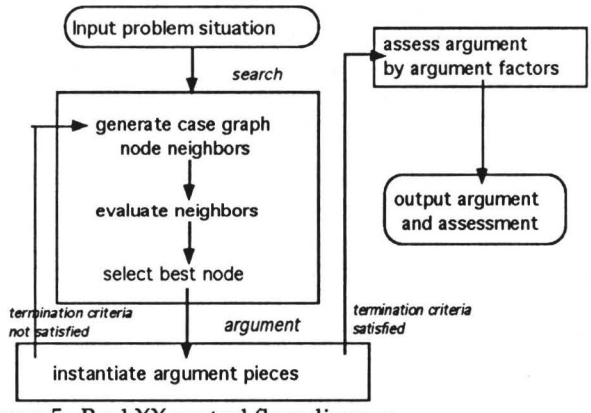

Figure 5. BankXX control flow diagram.

4 Examples and Discussion of Retrieval using Multiple Indexing

4.1 Benefits of Multiple Indexing of Cases

Because cases are embedded in a case graph rather than a discrimination tree, several paths to each case in BankXX's memory may be traversed (see also, e.g., [Kolodner, 1983; Turner, 1988]), and this flexibility aids case retrieval in several ways. Multiple paths to cases, found through the sequential application of distinct types of indices, can be coupled with case representations at different abstraction levels, and can

yield a finer retrieval granularity. The use of multiple types of indices also increases the robustness of case retrieval in "real world" domains in which noisy cases can be indexed incorrectly. Mis-indexing a case by one index does not render it inaccessible when other indices still provide a path. Finally, from a cognitive vantage point, in a richly connected domain like the law people use a variety of indices for reminding or to access information [Schank, 1982].

4.2 Examples

Example 1: Retrieval of Too Many Cases

One *desideratum* of a legal argument is a legal theory that can be applied to support the desired conclusion. One use of multiple indices is to prune a set of theories (and the cases applying them) if initial retrieval probes return too many. This first example shows how to filter potential theories by using two distinct retrieval paths, which assures that the remaining theories will be relevant from at least two perspectives.

Consider as a new problem the facts of *Matter of Akin*, 54 B.R. 700 (Bkrtcy. 1985): a graduate of the National College of Business in Rapid City, South Dakota, filed a Chapter 13 plan to discharge student loan and other debts. *Akin* is in the system's case base, but we have removed it and its linkages so it can be used *de novo* as a problem for the system.

To begin processing BankXX performs factor analysis on the facts of *Akin*, constructs a claim lattice, and extracts the most on-point cases: *Baez, Ashton, Canda, Hawkins, Chura, Ali* and *Gibson*. A claim lattice is a data structure that represents a partial order of cases by similarity, developed for the HYPO program [Ashley, 1990]. These seven most on-point cases are too many to base an argument on; it is desirable to reduce this group to a smaller set that are relevant from several indexing perspectives.

Next, one neighbor method of the system traces the links from each of these most on-point cases to the legal theories that are used by them. The *Gibson, Canda* and *Hawkins* cases did not apply any recognizable theory, but the following theories were applied by the remaining cases: *Kitchens-Kull-theory, Old-Bankruptcy-Act-Good-Faith-Definition, Flygare-theory,* and *Estus-theory*.

In this first access route, BankXX retrieved these theories by visiting in turn the following subspaces in case memory: **Fact Situation Space — Legal Factor Space — Fact Situation Space — Legal Theory Space.** To get to Legal Factor Space, the current fact situation was analyzed for its applicable domain factors. In Legal Factor Space, applying the factor representation of cases, the most on-point cases were extracted from a claim lattice. These most on-point cases exist in Fact Situation Space; links from these cases are followed to the legal theories in Legal Theory Space.

BankXX's neighbor methods reveal another route to Legal Theory Space. The story prototype of *Akin* is *student-loan*: the unsecured debt is primarily student loan debt and *Akin's* fact pattern follows the usual script for debtors attempting to discharge student loans. The case graph includes links from the *student-loan* story prototype to the cases that have followed that script, yielding a list of student loan cases: *Severs, Ponanski, Marsch, Gunn, Estus, Canda, Ali, Dos-Passos, Gibson, Hawkins, Makarchuk, Owens* and *Sanabria*.

From those cases BankXX retrieves the legal theories applied by them. Most of the student loan cases do not espouse a particular theory, but the theories that have been applied are *Estus-theory Flygare-theory, Principal-Purpose-Student-Loan-Discharge* and *Owens-3-factors*

To recap, on this second route through the case graph, BankXX's neighbor methods provided the path: **Fact Situation Space — Legal Story Space — Fact Situation Space — Legal Theory Space.** In Legal Story Space, the cases that follow the student loan script were identified, which brought the search back to Fact Situation Space. Links to the theories applied to those student loan cases were traced to arrive in Legal Theory Space.

There are now two paths to *Estus-theory* and to *Flygare-theory* (Figure 6). That these two theories have been found by two different lines of reasoning reinforces their importance: BankXX's neighbor methods are used to reduce the list of potentially useful theories to these two.

The utility of composite indexing is borne out by comparing the actual judicial opinion in the *Akin* case with BankXX's retrieved findings. In the *Akin* opinion, the judge did apply *Estus-theory*, one of the two closely related theories identified by BankXX. This example of sequential application of multiple indices also has identified a starting point for further legal research. The *Ali* case is similar to *Akin* from three perspectives: it is a student loan case, it is dimensionally most on-point, and it applies the *Estus-theory*.

In this example, multiple indexing paths permitted partial resolution of the "too many" retrieval problem. Too many most on-point cases and too many legal theories were retrieved by the first probes of the case base. Following indices along useful paths filters the list of most on-point cases and reduces the list of previously applied legal theories to those that are relevant from several perspectives.

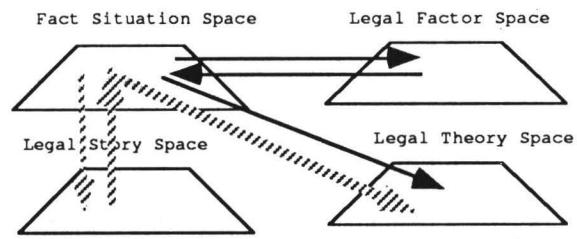

Figure 6. Two paths (solid arrows and dashed arrows) from Fact Situation Space to Legal Theory Space.

Example 2: Retrieval of Too Few Cases

We consider in less detail a second example to illustrate that multiple indexing can provide recourse when an initial probe of the case base retrieves insufficiently many cases. *In re Chura*, 33 B.R. 558 (Bkrtcy. 1983), involved a debtor who, as a bank teller, had handled a checking account on behalf of her disabled cousin Debra. Chura applied Debra's welfare and social security payments to her own personal use. After a lawsuit, a judgment of $27,000 was entered against Chura. Chura filed for Chapter 13 bankruptcy, and proposed a 9-month plan[2] under which she would pay $1 in total to Debra to discharge the $27,000 judgment. Finding the support to argue for confirmation of such a plan on behalf of such a debtor presents a challenge.

To see how BankXX handles this challenge, as in the previous example we remove the *Chura* case and links from memory and take the facts of *Chura* as a fresh problem situation. Again, to get a toehold into case memory, BankXX begins with a factor analysis, and extracts the most on-point cases from a claim lattice: *Hawkins, Rasmussen, Owens, Gunn, Easley, Gibson* and *Akin*.

All but the *Owens* case were decided for the creditors, but *Owens* does not involve outrageous or illegal debtor conduct and is therefore clearly distinguishable from *Chura*. Thus the debtor has no most on-point cases for its side: insufficiently many cases have been retrieved. However, the obvious designation of this case as of the dishonest debtor factual prototype indexes into 8 dishonest debtor cases: *Rimgale, Rasmussen, Okoreeh-Baah, Brown, Boyd, Baez, Schaitz* and *Girdaukas*.

One neighbor method of BankXX then computes the family resemblance ratings of the dishonest debtor cases to find the most prototypical dishonest debtor case. In order to situate the current problem within Family Resemblance Space, Chura itself is included in the dishonest debtor category. The family resemblance measure of prototypicality reveals in fact that *Chura* is the most prototypical dishonest debtor case. This high prototypicality measure supports analogizing *Chura* with the dishonest debtor cases that were decided for the debtor, *Rimgale, Okoreeh-Baah* and *Baez*.

Finally, each of these three cases is linked to the legal theory applied in it: *Rimgale-theory, Okoreeh-Baah-theory,* and *All-the-facts-and-circumstances*.

Examining each of these theories reveals that the *Okoreeh-Baah-theory* holds that the all the facts and circumstances should be taken into account, not just the debtor's unsavory pre-plan conduct. Thus BankXX has found a theory and a case to support the debtor's side in *Chura*, notwithstanding the difficulty of that task occasioned by the debtor's conduct.

[2]A plan typically lasts at least 36 months.

5 Related Research and Summary

Like the BankXX system itself, our line of inquiry relies on diverse precedents. BankXX uses a case base structure similar to conceptual legal retrieval designs in which a graph of cases and concepts with labeled edges captures influences and taxonomic information (e.g., [Hafner, 1987], [Bing, 1987]; see generally [McCarty, 1983]).

Other case retrieval systems have organized case memory as a graph and permit multiple paths to a case (e.g., Kolodner's CYRUS [1983]). Turner's MEDIC [1988] uses a case memory of linked discrimination nets, which allows multiple paths to diagnostic schemata. PROTOS [Bareiss, 1989] uses a fixed strategy for classification that takes advantage of three kinds of indexing knowledge. Rose and Belew's [1991] SCALIR uses a variety of inter-case links including Shepard's citation links and West's key taxonomy, but applies spreading numerical activation to search the network, rather than heuristic search. Direct Memory Access Parsing (DMAP, [Martin, 1990]) is a case-based architecture that uses a semantic network of case frames that is searched via a marker-passing algorithm to instantiate frames that are expected in the problem context. BankXX is similar to [Owens, 1989], in that repeated probes to the case base take account of previous memory probes. [Bhatnagar, 1989] and [Branting, 1991] have applied A* search to the respective tasks of creating arguments and of matching fact situations and justifications to a graph that represents the current case.

However, our focus differs from these diverse systems. We emphasize how related, multiple index types can rectify retrieval failures, and more generally, how the constraints of argument formation can be incorporated in a heuristic evaluation function that drives case base search .

In summary, BankXX retrieves case and theory support for legal arguments. The indexing scheme used to find the bases for argument takes account of the different kinds of knowledge required for argument and the inter-connections between them. In our experiments, improved case retrieval results (1) when cases that are similar to the current problem from one vantage point are filtered to yield cases that are similar according to several perspectives, and (2) when a complementary indexing scheme is applied if other indexing methods have failed to retrieve sufficiently many relevant cases. Our experience with the system is that complementary indexing methods can yield several paths to a case to improve case retrieval. Future research will test the appropriateness of best-first search for controlling search of the case graph.

Acknowledgments

This research was supported in part by the National Science Foundation, contract IRI-890841, and the Air Force Office of Sponsored Research, contract 90-0359. We thank Matt Cornell, Kate Sanders, Elaine Whitlock and Barbara Fell-Johnson.

References

[Ashley, 1990] K. D. Ashley. *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*. M. I. T. Press, Cambridge, Massachusetts, 1990.

[Ashley & Rissland, 1987] K. D. Ashley and E. L. Rissland. But, See, Accord: Generating Blue Book Citations in HYPO. *Proceedings, First International Conference on Artificial Intelligence and Law*, 67-74. Boston, Massachusetts, 1987. ACM.

[Bareiss, 1989] E. R. Bareiss. *Exemplar-Based Knowledge Acquisition*. Academic Press, Boston, Massachusetts, 1989.

[Bhatnagar, 1989] R. K. Bhatnagar. *Construction of Preferred Causal Hypotheses for Reasoning with Uncertain Knowledge*. Ph.D. Dissertation, University of Maryland, College Park, Maryland, 1989.

[Bing, 1987] J. Bing. Designing Text Retrieval Systems for "Conceptual Searching". *Proceedings, First International Conference on Artificial Intelligence and Law*, 43-51. Boston, Massachusetts, 1987. ACM.

[BlueBook, 1986] *A Uniform System of Citation* (14th ed.) Harvard Law Review Association, Cambridge, Massachusetts, 1986.

[Branting, 1991] L. K. Branting. Building Explanations from Rules and Structured Cases. *International Journal of Man-Machine Studies*, 34, 797-837, 1991.

[Fikes, Hart & Nilsson, 1972] R. E. Fikes, P. Hart and N. J. Nilsson. Learning and executing generalized robot plans. *Artificial Intelligence*, 3, 251-288, 1972.

[Gardner, 1987] A. vdl. Gardner. *An Artificial Intelligence Approach to Legal Reasoning*. M. I. T. Press, Cambridge, Massachusetts, 1987.

[Hafner, 1987] C. D. Hafner. Conceptual Organization of Case Law Knowledge Bases. *Proceedings, First International Conference on Artificial Intelligence and Law*, 35-42. Boston, Massachusetts, 1987. ACM.

[Kolodner, 1983] J. L. Kolodner. Maintaining Organization in a Dynamic Long-Term Memory. *Cognitive Science*, 7(4), 243-280, 1983.

[Lehnert, 1981] W. Lehnert. Plot Units and Narrative Summarization. *Cognitive Science*, 5(4), 1981.

[Martin, 1990] C. E. Martin. *Direct Memory Access Parsing*. Ph.D. Thesis, Yale University, New Haven, Connecticut, 1990.

[McCarty, 1983] L. T. McCarty. Intelligent Legal Information Systems: Problems and Prospects. *Rutgers Computers and Technology Law J.*, 9(2), 265-294, 1983.

[Nilsson, 1980] N. Nilsson. *Principles of Artificial Intelligence*. Tioga Press, Palo Alto, California, 1980.

[Owens, 1989] C. Owens. Integrating Feature Extraction and Memory Search. *Proceedings, The 11th Annual Conference of the Cognitive Science Society*. Ann Arbor, Michigan, 1989. Lawrence Erlbaum.

[Rissland & Skalak, 1991] E. L. Rissland and D. B. Skalak. CABARET: Rule Interpretation in a Hybrid Architecture. *Int. J. Man-Machine Studies*, 34, 839-887, 1991.

[Rissland, Skalak & Friedman, 1993] E. L. Rissland, D. B. Skalak and M. T. Friedman. BankXX: A Program to Generate Argument through Case-Based Search. *To appear, Proceedings, Fourth International Conference on Artificial Intelligence and Law*. Amsterdam, The Netherlands, June, 1993. ACM.

[Rissland, Valcarce & Ashley, 1984] E. L. Rissland, E. M. Valcarce and K. D. Ashley. Explaining and Arguing with Examples. *Proceedings, AAAI-84*, Austin, Texas, 1984. AAAI.

[Rosch & Mervis, 1975] E. Rosch and C. B. Mervis. Family Resemblances: Studies in the Internal Structure of Categories. *Cognitive Psychology*, 7, 573-605, 1975.

[Rose & Belew, 1991] D. E. Rose and R. K. Belew. A Connectionist and Symbolic Hybrid for Improving Legal Research. *Int. J. Man-Machine Studies*, 35, 1-33, 1991.

[Samuel, 1967] A. L. Samuel. Some Studies in Machine Learning using the Game of Checkers II - Recent Progress. *IBM J. Research and Development*, 11, 601-617, 1967.

[Schank, 1982] R. C. Schank. *Dynamic Memory*. Cambridge University Press, Cambridge, England, 1982.

[Shepard's, 1992] *Shepard's Federal Citations*. Shepard's/McGraw-Hill, Colorado Springs, Colorado, 1992.

[Skalak & Rissland, 1992] D. B. Skalak and E. L. Rissland. Arguments and Cases: An Inevitable Intertwining. *Artificial Intelligence and Law: An International Journal*, 1(1), 3-48, 1992.

[Turner, 1988] R. Turner. Organizing and Using Schematic Knowledge for Medical Diagnosis. *Proceedings, Case-Based Reasoning Workshop*, 435-446. Clearwater Beach, Florida, 1988. Morgan Kaufmann.

[Twining & Miers, 1982] W. Twining and D. Miers. *How To Do Things With Rules* (Third ed.). Weidenfeld and Nicolson, London, England, 1982.

[West, 1992] *WestLaw*. West Publishing Company, Saint Paul, Minnesota, 1992.