What you saw is what you want: Using Cases to seed Information Retrieval.*

Jody J. Daniels and Edwina L. Rissland

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

Abstract. This paper presents a hybrid case-based reasoning (CBR) and information retrieval (IR) system, called SPIRE, that both retrieves documents from a full-text document corpus and from within individual documents, and locates passages likely to contain information about important problem-solving features of cases. SPIRE uses two case-bases, one containing past precedents, and one containing excerpts from past case texts. Both are used by SPIRE to automatically generate queries, which are then run by the INQUERY full-text retrieval engine on a large text collection in the case of document retrieval and on individual text documents for passage retrieval.

1 Introduction

A good indication of what to look for in a new problem situation is often given by examples of what has worked in the past. This idea—the fundamental tenet of case-based reasoning—is applicable in information retrieval (IR) as well. We have employed this idea at two levels in a hybrid CBR-IR approach:

1. within a corpus of documents, to find documents relevant to a new problem situation, retrieve documents similar to those that are already known to be relevant;
2. within an individual document, to find passages that address a particular aspect of a situation, retrieve passages that are similar to those that illustrate past discussions of the topic.

We call these two levels of retrieval the *corpus* and *document* levels. At both levels, exemplars of relevant text—documents or excerpts—from past problem solving experiences provide good clues of what to look for in a new situations.

In this paper, we describe our system called SPIRE (Selection of Passages for Information REduction) that performs retrieval at both the document and passage levels. Given a new problem situation input as a case frame, SPIRE retrieves relevant documents from a full-text document corpus, and, within each of these, highlights the passages most likely to discuss important problem-solving

* This research was supported by NSF Grant no. EEC-9209623, State/Industry/-University Cooperative Research on Intelligent Information Retrieval.

features. In this paper, we emphasize retrieval at the passage level; other papers have detailed retrieval at the document level [DR95, RD96].

After problem entry, SPIRE performs all of its processing, including the generation of needed queries, without *any* intervention on the part of the user. Thus, the only representational or processing burden on the user is the specification of the problem case, which is done in a manner that is standard practice in CBR systems (without natural language front-ends) [Kol93]. Thus SPIRE locates relevant textual regions within documents without imposing on the user the burden of reading entire documents.

SPIRE uses two case-bases:

1. A case-base of past, resolved problem situations (precedents) represented as case-frames of features for use by a HYPO-style CBR module.
2. For each case feature in the case-frame, a case-base of actual text excerpts, culled from past cases, that contain useful information about the value of the feature.

Both case-bases are used by SPIRE to generate queries, which are then acted upon by the INQUERY retrieval engine in its usual manner [CCH92].

The first case-base is the standard type of case-base used by many generations of our own HYPO-style CBR systems with their concomitant mechanisms of dimension-based analysis, sorting into a claim lattice, etc. [Ash90]. The second is simply a collection of textual fragments partitioned into sub-case-bases, one for each problem feature of interest. Indexing and selection are minimal in the second case-base at this point; the feature (name) serves as the index and all fragments are selected. We note that in the future, as these collections grow, more highly attenuated indexing and selection will most likely be needed. However, even now, the question of what cases, that is, text fragments, to include is an interesting one. In Section 5.3 we discuss the impact of the composition of the excerpt case-base on the performance of the system.

Although SPIRE does not actually extract the information contained in the passages it retrieves, we believe SPIRE could play a key role in the text extraction process by focusing an extractor's (human or machine) attention on those passages that are worth the effort of careful "reading." Currently, it is not feasible nor reasonable to apply an extraction effort across a long text, especially when there are only a few small portions that are relevant. Thus, we can use the passages highlighted by SPIRE as input to an extraction process. The output of the extraction effort can then be plowed back into a knowledge base, used by our system or some other symbolic reasoner. SPIRE can thus aid in executing the full loop of case-based reasoning by assisting in the acquisition of new, symbolically represented cases. This is particularly important in domains where a large volume of data already exists in textual form.

2 System Description

SPIRE works in two stages (as shown in Figure 1):

1. from a large text collection, SPIRE retrieves documents that are relevant to the presented problem case, and

2. within those retrieved documents, SPIRE highlights passages that contain information relevant to specific case features.

In the first stage, SPIRE is given a new problem situation. It uses its HYPO-style CBR module to analyze it and select a small number of most relevant cases from its own case-base consisting of symbolically represented texts. In the usual CBR fashion, SPIRE determines the similarity of each known case to the new problem and represents the results of this analysis in a standard claim lattice [Ash90].

The most relevant cases from this analysis—typically the cases in the top two layers of the claim lattice—are then used to "prime the pump" of INQUERY's relevance feedback module. This set of cases is called the *relevance feedback case-knowledge-base* or RF-CKB [RD95, RD96]. (These are labeled as "Best Case Texts" in Figure 1.) The original texts of the cases in the RF-CKB are passed to the INQUERY IR engine, which then treats them as though they had been marked relevant by a user. INQUERY automatically generates a query by selecting and weighting terms or pairs of terms from within this set. This query is then run against the larger corpus of texts, with the result that new documents are retrieved and ranked according to INQUERY's belief as to their relevance to the posed query. (A detailed description of this first stage can be found in [DR95, RD96].)

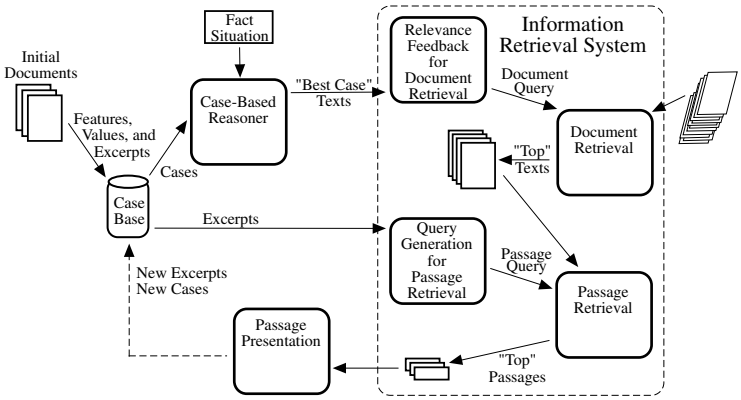

Fig. 1. Overview of SPIRE.

In the second stage, SPIRE locates germane passages within each of the texts retrieved in stage one. In this stage SPIRE locates passages within a single document rather than documents within a collection. Again SPIRE uses a hybrid CBR-IR approach. This was motivated by our belief that past discussions of a topic would provide good clues to the location of new discussions.

To locate passages, SPIRE generates queries by using excerpts from past discussions of a feature. Each excerpt is an actual fragment of text containing relevant information about a case feature and comes from an episode of information location/extraction performed on a past case. Example excerpts are given in Section 3.

There are numerous techniques for transforming the excerpts into passage retrieval queries. (A fuller discussion of this can be found in [Dan97].) SPIRE presents the query along with a specified document to the IR engine, which, in turn, retrieves the top ranked passages for presentation to the user or possibly to an information extraction system. Thus, excerpts are used analogously to the RF-CKB's of stage one: their terms are used to generate queries.

We created these case-bases of excerpts by asking an individual familiar with the representation of the problem domain to read a small number of opinions corresponding to cases in SPIRE's case-base and to highlight any portion of text—whether it be just a few terms, a phrase, or several sentences or more—that was useful for determining the feature's value. It was common for pieces from different locations throughout the text to be highlighted. Normally, this step would be done in conjunction with the creation of the case-base for the domain and the encoding of the first few cases and thus would not require a full review of the textual sources. However, since we were re-using a portion of the bankruptcy case-base used in the BankXX project [RSF96], this highlighting of textual examples was done *post hoc*.

As each document and feature is addressed in stage two, the user (or information extraction program) can examine the presented passages, determine (if possible) the actual value of the feature, and add it to the representation for the text, for instance, as a case. The user may also decide to add one or more of the retrieved passages, or selected portions of them, to the appropriate excerpt case-base along with the feature and value. In this way, SPIRE may aid in the acquisition of additional knowledge about the textual context of each feature.

3 Example

Our example comes from the domain of personal bankruptcy under Chapter 13 of United States personal bankruptcy law (11 U.S.C. §1301-1330). The question presented is whether the plan proposed by the debtor to settle with the creditors has been proposed in "good faith". This question is central to approval of the debtor's plan.

We use the situation as described in the *In re Makarchuk*, 76 B.R. 919 (Bankr.N.D.N.Y.1987) opinion as our problem case. In *Makarchuk*, the debtor proposed a plan whose dominant purpose was to discharge two student loans. The debtor had recently unsuccessfully attempted to discharge these same loans under Chapter 7 of the code.

We submit this case to SPIRE, which compares it to those situations found in its own case-base. The full-text documents associated with the most similar cases—the RF-CKB—are passed to the IR system. The IR system creates a document-level query, poses it, and retrieves a set of documents. The ten top-rated documents for the *Makarchuk* situation are listed in Table 1. We note that *Ali, Akin*, and *Hawkins*, as well as *Makarchuk* were already known to SPIRE (i.e., represented in its case-base of documents,) while only *Ali* has text fragments in the excerpt case-base. Thus, the other six of the top ten cases must be "read"

in order for their facts to be ascertained in preparation for any use in a legal analysis for *Makarchuk*. This completes SPIRE's stage one.

Rank	Doc-Id	Case Name	Belief Score
1	877	In re Stewart	(0.518380)
2	178	In re Ali	(0.508805)
3	565	In re Makarchuk	(0.504300)
4	353	Matter of Akin	(0.503330)
5	427	In re Gathright	(0.500740)
6	177	Matter of Hawkins	(0.498640)
7	693	In re Ellenburg	(0.496048)
8	915	In re Carpico	(0.493452)
9	733	In re Newberry	(0.492010)
10	764	In re Porter	(0.491779)

Table 1. The most highly ranked documents for the *Makarchuk* problem.

We would like to examine specific facts in these newly retrieved cases, such as, finding out whether the court found the debtors to be sincere when proposing their plans. (Other features of bankruptcy cases are discussed in Section 4.1.) To do this, we direct SPIRE in stage two to locate passages within the top case texts that concern the feature called *sincerity*. SPIRE uses excerpts from its case-base of excerpts on *sincerity* to form a query to retrieve passages. Sample excerpts from this case-base are:

- "represents an earnest effort to repay his unsecured creditors"
- "sincerity is tempered by her desire to avoid returning to Maine."
- "The Court believes the debtors' motivation and sincerity are genuine."
- "The Chapter 13 petition was intended to wipe out BNY's claims rather than to repay them."
- "this couple makes a concerted effort to live sensibly and substantially within their means."

To illustrate passage retrieval, we use the *In re Stewart* case, the top rated document for the *Makarchuk* problem. INQUERY, the IR engine, divides the *Stewart* opinion into overlapping windows of 20 words each, approximating the length of a sentence. Since the windows overlap, each word in the opinion will appear in two windows (except for the first 10 words). The retrieved passages are ranked by INQUERY according to its belief that each is relevant to the query.

For this example, we allow SPIRE to use two simple methods to generate queries. The first combines the terms from all the excerpts about a feature into a single "natural language" query. Each word in each excerpt provides a possible match against the words in the window. Regardless of whether two words were in different excerpts, each contributes to the total belief. We refer to this type of query as a *bag of words* query. The second type of query places a restriction so that terms from within an excerpt need to be found co-occurring in the passage. We refer to this type of query as the *sum* query because it is formed by wrapping

an INQUERY #Sum operator around each excerpt. Part of both queries are shown below:

```
#Passage20(
 represents an earnest effort to repay his unsecured creditors
 sincerity is tempered by her desire to avoid returning to Maine. ...)

#Passage20(
 #Sum( represents an earnest effort to repay his unsecured creditors)
 #Sum( sincerity is tempered by her desire to avoid returning to Maine.)
 ...);
```

Posing these two queries over the *Stewart* opinion causes INQUERY to retrieve passages. In this particular example, both the *bag of words* and *sum* queries retrieve the same top five passages, although with differing belief values. Below are the top five passages for the *bag of words* query, annotated with whether or not each is relevant:

Bag of Words			
Rank	Psg	Strt	Belief
1	3390	(0.410120)	Rel
2	3400	(0.409335)	Rel
3	2580	(0.405726)	
4	2570	(0.405726)	
5	2160	(0.404761)	

The following is the text of the 3390 and 3400 passages, top-ranked by both retrievals. We boldface content terms that match those found in the excerpts and show word counts along with the text.

...(9) the frequency with which the debtor
3390 | has **sought** relief under the Bankruptcy Reform Act; (10) the
3400 | **motivation** and **sincerity** of the **debtor** in seeking **Chapter 13**
3410 | relief; *1004 (11) the burden which the **plan's** administration would place upon the trustee.

From these passages we cannot determine whether the court found the debtor to be *sincere*, however the text is highly on-point to the topic. The next few passages are not relevant to *sincerity*, but because most of them discuss the debtor's effort to make an "earnest" or "substantial" effort to repay creditors, they are are highly ranked.

The next cluster of relevant passages are ranked 8 through 14. Passages 4030, 4040, and 4050, which received ranks 11, 10, and 13, respectively, by the *bag of words* query, and 14, 8, and 12, respectively, for the *sum* query, are given below.

...not a requirement for confirmation of every Chapter
4030 | **13 plan**, was one **intended** purpose of **Chapter 13's** enactment.
4040 | Failure to provide **substantial repayment** is certainly evidence that a
4050 | **debtor** is **attempting** to manipulate the statute **rather** than **attempting**
4060 | to honestly **repay** his **debts**....[sic]

In stage two, SPIRE has thus located passages relevant to the *sincerity* feature without requiring a user to pose a query. Unlike other approaches, which

merely retrieve entire documents, SPIRE is able to retrieve documents and then present a significantly reduced amount of text about features contained within the document. This greatly decreases the amount of text a user must inspect.

By comparison, if we had intervened after SPIRE's first stage, and manually generated a query for *sincerity*, we might have posed the following query:

```
#Passage20( motivation sincerity genuine sensible earnest );
```

On the *Stewart* opinion, this query yields:

```
Rank Psg Strt Belief
  1    3400   (0.443848) Rel
  2    3390   (0.443848) Rel
```

The text of both of these passages is given above. While they are relevant to the feature, unfortunately, they are the *only* relevant passages that are retrieved by the manual query. In fact, they are the only passages that this query retrieved at all.

4 Domain Knowledge

We now describe the various types of features we examined, the case-bases of textual excerpts, generation of answer keys, and the evaluation metric.

4.1 Features examined

We selected ten features from a bankruptcy good faith case representation. There were five types of values that these features could have: Boolean, date, category, set, or numeric. For our set of ten features, we included two of each type. They were: *sincerity* (was the debtor sincere in proposing the plan), *special-circumstances* (were there any extenuating conditions affecting the debtor), *loan-due-date*, *plan-filing-date*, *procedural-status* (such as appeal or affirmation), *future-income* (the likelihood that the debtor will experience an increase in income), *debt-type* (such as educational or consumer), *profession*, *monthly-income*, and *duration* (of the proposed plan in months).

4.2 Excerpt case-bases

For the above set of ten features we gathered excerpts from 13 case opinions. Once SPIRE stemmed and removed non-content terms, the average number of remaining unique content terms for the ten features was 46.7, although two of the features only have 18 content terms. (See Table 2.)

In the previous section, we included example excerpts for *sincerity*. The following are examples for the feature of *future income*:

Future income – this is text that discusses whether the debtor's income is projected to increase in the future. The text might be negative or positive on this matter.

- "the Court cannot see any likelihood of future increases"
- "the prospect of a regular job with substantially increased income is not great. "

Feature	Num Excerpts	Total Words	Num Unique Terms	Num Unique Content Terms
Plan Duration	14	212	92	59
Monthly Income	13	110	52	34
Sincerity of the Debtor	9	123	89	52
Special Circumstances	8	188	117	71
Loan Due Date	4	47	32	18
Plan Filed Date	10	145	66	45
Debt Type	10	164	102	63
Profession	3	36	29	18
Future Income	8	88	68	36
Procedural Status	13	194	100	71

Table 2. Number of terms contained in the excerpts.

- "her health brings into question her future ability to work."
- "no evidence that raises are likely."

Examples of excerpts for the feature of *special circumstances*, which include unusual events (e.g. pending divorce, being in prison) that can affect the debtor's ability to repay debts, include:

- "The Court believes the debtors' medical expenses will increase as time goes on and believes this is a 'special circumstance' under factor 8."
- "This debtor has not been the victim of extraordinary 'outside' forces."
- "The debtor is now in treatment for the condition that may have contributed to the debtor's need for Chapter 13 relief."
- "Debtor was incarcerated in the New Mexico State Penitentiary for fraudulent practices"

4.3 Answer keys

In order to evaluate SPIRE's ability to locate relevant passages, we needed to create answer keys specifying where within our test documents there was text discussing each of the features. These answer keys were created by outside readers.

We hired two undergraduates to read case opinions and underline any text that they perceived as being about a given feature. They were given a set of written instructions that described each feature and samples of the sort of text they should mark.

4.4 Evaluation metric

Most retrieval systems are judged on the basis of precision and recall. These measure what percentage of the retrieved items are relevant (coverage) and what percentage of the relevant items are retrieved (accuracy), respectively.

We are not concerned with locating *every* relevant item, so recall is not a concern. If we only look at precision, by examining the passages retrieved at certain cutoff depths, we will lose information about the ordering of the relevant and non-relevant passages. We are concerned with how much effort will be wasted by users as they examine retrieved passages. This can be measured by *expected search length* (esl)[Coo68], which measures the number of non-relevant items encountered before finding a specified number of relevant ones. In this work we use esl_1, esl_3, and esl_5, which are esl values when 1, 3, or 5 passages are specified.

5 Experiment results

We ran SPIRE using three problem cases and collected the top documents for each. Removing duplicates and documents that had been used to derive the excerpt case-bases, we made a test collection from among the top 10 retrievals for each problem to make a test set of 20 documents. Using various methods for passage query generation, we tested SPIRE on these 20 documents with 10 different case features.

5.1 Query types

In the experiments reported here, we are concerned primarily with the two previously mentioned query formation methods: *bag of words* and *sum*. These are the two *base* methods that performed the best. The others in this set were: *bag of words plus phrases, sum plus phrases*, and *set of words*. Formation and results for these queries is discussed in more detail in [Dan97].

We had SPIRE build two other sets of queries. The first is based on a term weighting scheme suggested by Kwok [Kwo96] and the second set is what we called *semi-random*. The latter incorporated only one-half or one-third of the available query terms from the excerpt case-base. Neither of these sets performed better than the two base queries. (See [Dan97] for details.)

To provide another point of comparison, we also had a human expert, familiar with both the domain and INQUERY query operators, create queries. These manual queries are highly refined expert queries and provide a very high baseline. We used the best manual query for each feature as a point of comparison and refer to this set as the *manual* queries.

5.2 Results

Comparison of the *bag of words* and *sum* queries revealed that they performed about equally well as measured by esl scores. Across all 20 documents and 10 features, the *sum* queries performed slightly better when requesting one or five relevant passages, and the *bag of words* queries performed slightly better when requesting three passages. Overall, SPIRE-generate queries performed just about equally to the expert manual queries. (See Table 3, which provides a comparison between the manual and SPIRE-generated queries on half of the test document collection; results on the other half are similar.)

When we look at the results broken down by feature, there are noticeable differences. There were two features where the manual queries did better: *procedural status* and *plan filed date*, and two features where the SPIRE-based queries

Doc-ID	Debt Type	Duration	Future Income	Loan Due	Mthly income	Plan Filed	Proc. Status	Profes-sion	Sincere	Special Circ
001	=	M	=	=	=	SP	M	SP	M	s
180	M	SP	=	M	=	M	M	=	=	s
188	s	M	M	SP	SP	M	M	=	SP	s
204	SP	SP	=	SP	SP	M	SP	M	SP	SP
206	M	M	M	=	SP	SP	=	SP	SP	=
260	M	SP	SP	=	SP	M	M	=	=	b
289	=	M	=	=	=	M	M	M	SP	SP
353	=	M	=	SP	=	=	s	SP	=	=
407	SP	=	=	=	s	M	b	b	=	SP
427	=	M	M	=	SP	SP	M	=	=	s

Table 3. Comparison between the esl_3 of manual and SPIRE-generated queries. An "SP" indicates that both SPIRE queries performed better than the manual. An "M" indicates that the manual query performed better. If the manual fell between the two, the SPIRE query performing the best is given: "b" for *bag of words* and "s" for *sum*. If all three queries performed equally well, an "=" is shown.

did distinctly better: *sincerity* and *loan due date*. With the other features, the results were closer.

For *procedural status* this difference is easily explained: discussion about this feature normally includes at least one of a small set of easily enumerated keywords, such as "confirmation" and "appeal". Not all of these terms were present in SPIRE's excerpt case-base, but all were included in the manual query. For example, "affirmation" and "convert" were never given as the status of any of the cases found in our small corpus. This is an instance where knowledge of a domain-specific vocabulary, particularly a small set of technical terms, is easily enumerated and should be used to form the query.

The difficulty SPIRE had in finding the *plan filed date* is partially due to the way in which the opinions express the date. For example:

- "At the time of filing the Chapter 13 proceeding," [case opinion 289]
- "LeMaire signed a promissory note evidencing a debt to his parents of $12,722 only one day prior to filing his bankruptcy petition. Prior to this filing, LeMaire had ..." [case opinion 860]

In neither is a calendar date given. Additionally, the first text fragment is the only relevant passage within that text. We note that pattern matching techniques or concept recognizers, would also be unable to locate these passages.

5.3 Reexamining the excerpt case-base

In the course of examining the retrieved passages for *plan filed date* we noticed that they often included specific names of debtors in a prior case. In our excerpt case-base, such names had sometimes been included, for instance, "Debtors-Appellants, Mr. and Mrs. Okoreeh-Baah, filed for bankruptcy on November 19, 1985." Since the case name, "Okoreeh-Baah", was included in the excerpt, it caused SPIRE to rate passages that included it very highly, even though the

presence of this specific name does not make a passage relevant to the issue of *plan filed date*.

Based on this realization, we reexamined SPIRE's excerpt case-base. Within the excerpts for several features, proper names were frequently included. Additionally, there were instances where excerpts contained text that had no real bearing on discussion of the feature. Where reasonable (i.e., at the beginning or end of an excerpt), we subsequently removed any proper names or superfluous text from the excerpt case-base to create a second excerpt case-base. A second case-base of excerpts was created for the features of *debt type, duration, future income, monthly income, plan filed date,* and *procedural status*. We then recreated and reran the *bag of words* and *sum* queries.

For all of the features, the queries from the new case-base showed improvement over the original one. *Plan filed date* had the largest improvement. For this feature it was not uncommon for the relevant passages to move up in the ranking by as many as ten to twenty, or even forty positions. Besides proper names, deleted text included several instances of "under Chapter 7;" these have no bearing on this feature. Table 4 shows the average number of non-relevant passages that were no longer required to be read before reaching the requested number of relevant passages.

Average reduction in esl

ESL level	Bag of Words	Sum
1	2.30	3.95
3	9.85	9.70
5	10.53	10.94

Table 4. Difference in esl between the two excerpt sets for *plan filed date*.

The results were similar for the other features. For instance, the *monthly income* results benefited from the deletion of a reference to the *In re Flygare* opinion, and *future income* results improved with the deletion of a fragment: "Mr. Severs testified that".

From this experience with the modified case-bases, we conclude that one must be a bit more careful when creating the excerpt case-base. This is particularly true regarding the inclusion of proper names. On the other hand one cannot simply use a few generic keywords to form a good query, since the excerpts did better than the manual queries for many of the topics.

6 Conclusion

We have presented the SPIRE system, which incorporates a two-stage approach to first, retrieve documents relevant to a given problem situation and second, locate passages within them that discuss particular aspects of the case. SPIRE automatically generates the queries needed for both of these stages in a case-based manner. SPIRE minimizes the amount of effort expended—by human or machine—in locating important pieces of information without sacrifice in

performance. We found that SPIRE does as well or better than manually crafted passage queries for many of the case features we tested.

SPIRE is a hybrid CBR-IR system. Its CBR processing makes use of two case-bases: a traditional HYPO-style case-base of precedents, and a case-base of specific text excerpts. While the question of case indexing in SPIRE's current excerpt case-base is not particularly interesting (at this point), the question of what excerpts to include, or alternatively, the level of generality needed in such excerpts, is indeed interesting. We discussed how the content of the excerpt case-base can affect performance, and noted that overly specific excerpts that contain specific names, dates, and dollar amounts, can hurt performance. In our comparison of SPIRE against manually created queries of generic keywords, we also demonstrated that an overly general approach is not optimal either.

References

[Ash90] Kevin D. Ashley. *Modeling Legal Argument: Reasoning with Cases and Hypotheticals.* M.I.T. Press, Cambridge, MA, 1990.

[CCH92] James P. Callan, W. Bruce Croft, and Stephen M. Harding. The INQUERY Retrieval System. In A. M. Tjoa and I. Ramos, editors, *Database and Expert Systems Applications: Proceedings of the International Conference in Valencia, Spain*, pages 78–83, Valencia, Spain, 1992. Springer Verlag, NY.

[Coo68] William S. Cooper. Expected Search Length: A Single Measure of Retrieval Effectiveness Based on the Weak Ordering Action of Retrieval Systems. *American Documentation*, 19:30–41, 1968.

[Dan97] Jody J. Daniels. *Retreival of Passages for Information Reduction.* PhD thesis, University of Massachusetts, Amherst, Amherst, MA, May 1997.

[DR95] Jody J. Daniels and Edwina L. Rissland. A Case-Based Approach to Intelligent Information Retrieval. In *Proceedings of the 18th Annual International ACM/SIGIR Conference on Research and Development in Information Retrieval*, pages 238–245, Seattle, WA, July 1995. ACM.

[Kol93] Janet L. Kolodner. *Case-Based Reasoning.* Morgan Kaufmann, 1993.

[Kwo96] K. L. Kwok. A New Method of Weighting Query Terms for Ad-Hoc Retrieval. In *Proceedings of the 19th Annual International ACM/SIGIR Conference on Research and Development in Information Retrieval*, pages 187–195, Zurich, Switzerland, August 1996. ACM.

[RD95] Edwina L. Rissland and Jody J. Daniels. Using CBR to Drive IR. In *Proceedings, 14th International Joint Conference on Artificial Intelligence*, pages 400–407, Montreal, Canada, August 1995. AAAI.

[RD96] Edwina L. Rissland and Jody J. Daniels. The Synergistic Application of CBR to IR. *Artificial Intelligence Review*, 10:441–475, 1996.

[RSF96] Edwina L. Rissland, D. B. Skalak, and M. Timur Friedman. *BankXX: Supporting Legal Arguments through Heuristic Retrieval. Artificial Intelligence Review*, 10(1-71), 1996.

This article was processed using the LaTeX macro package with LLNCS style