

AD-A279 897 		ON PAGE- Form Approved OMB No. 0704-0188	
1. AG		3. REPORT TYPE AND DATES COVERED FINAL/15 SEP 90 TO 28 FEB 94	
4. TITLE AND SUBTITLE CAED-BASED REASONING IN MIXED PARADIGM SETTINGS AND WITH LEARNING		5. FUNDING NUMBERS 7518/05 AFOSR-90-0359	
6. AUTHOR(S) DR. RISSLAND		8. PERFORMING ORGANIZATION REPORT NUMBER AFOSR-TR- 94 0326	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) UNIVERSITY OF MASSACHUSETTS DEPT OF COMPUTER SCIENCES AMHERST, MA 01003		10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFOSR-90-0359	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NM 110 DUNCAN AVE, SUTE B115 BOLLING AFB DC 20332-0001		11. SUPPLEMENTARY NOTES	
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE: DISTRIBUTION IS UNLIMITED		12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) In this project we investigated: (1) CBR in mixed paradigm settings, in particular in a blackboard-based system, called FRANK, that generated various types of explanations and arguments where supporting tasks, such as case-and rule-based reasoning, were dynamically configured to reflect the user's intended purposes for the report: (2) "pure" CBR, particularly issues concerning the use of multiple indices and types of case representations, in a system called BankXX, that used classic heuristic best-first search to retrieve information needed for case-based argument: and (3) the application of machine learning techniques to core issues in CBR, such as the problems of learning indices and prototype based and estimating concept theory drift. DTIC QUALITY INSPECTED 2			
14. SUBJECT TERMS		15. NUMBER OF PAGES	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED		16. PRICE CODE	
18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED		19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	
20. LIMITATION OF ABSTRACT SAR(SAME AS REPORT)		21. LIMITATION OF ABSTRACT	

94 6 3 055

94-16610

Edwina L. Rissland, AFOSR-90-0359, YEAR III FINAL REPORT

ORGANIZATION: University of Massachusetts

SUBCONTRACTORS: none

PRINCIPAL INVESTIGATOR: Edwina L. Rissland,

TITLE: Case-Based Reasoning in Mixed Paradigm Settings and with Learning

Approved for public release;
distribution unlimited.

OBJECTIVE: In this project we investigated: (1) CBR in mixed paradigm settings, in particular in a blackboard-based system, called FRANK, that generated various types of explanations and arguments where supporting tasks, such as case- and rule-based reasoning, were dynamically configured to reflect the user's intended purposes for the report; (2) "pure" CBR, particularly issues concerning the use of multiple indices and types of case representations, in a system called BankXX, that used classic heuristic best-first search to retrieve information needed for case-based argument; and (3) the application of machine learning techniques to core issues in CBR, such as the problems of learning indices and prototype cases and estimating concept and theory drift.

APPROACH: We designed, built, and tested several mixed paradigm, pure, and learning-enhanced CBR systems in domains where reasoning from past (precedent) cases is a key or necessary tool. Target application domains included legal reasoning, decision support in medicine, and classic classification domains from the UC/Irvine datasets.

DETAILED SUMMARY: Here we provide a brief summary of the results of our research. Fuller exposition can be found in the published research articles that are included in the appendix.

(1) CBR in a Mixed Paradigm Setting: We explored issues concerning the use of CBR in a mixed paradigm setting in several blackboard-based systems: GBB-CABARET, ABISS, and particularly, FRANK. Our goal was to explore control issues and the use of CBR in a mixed paradigm context, particularly a blackboard-based architecture. In FRANK, we focused on the influence of high level reasoning goals—such as advocacy or advice-giving—on the details of CBR and vice versa, the affect of CBR results on higher level tasks.

GBB-CABARET and ABISS

During the first year of this project, we built two precursors to FRANK. Both were built in the Blackboard Development (GBB) environment. GBB-CABARET was a variant of our previous CABARET system [Rissland & Skalak, 1991]. (Work on CABARET was supported under previous grants.) It used GBB to implement CABARET's agenda-driven control strategy. We also explored various CBR and related control issues in ABISS (A Blackboard Integer Sequence Solver) [Rissland et al. 1991]. We built upon our preliminary experience with these two blackboard-based systems in building FRANK [Rissland et al., 1993].

The FRANK Project

FRANK (Flexible Report and ANalysis System) demonstrates dynamic context-sensitivity and use of internal feedback in performing key tasks such as case retrieval. FRANK is implemented in the Generic Blackboard Development (GBB) environment and instantiated in the application domain of medical diagnosis and treatment of back complaints and injuries. FRANK uses both HYPO-style and nearest neighbor CBR for case-based reasoning; it uses OPS5 for rule-based reasoning.

As well as using a blackboard-based mixed paradigm architecture, FRANK is novel in how it dynamically config .es supporting tasks—such as, case-based and rule-based inference—to reflect the

user's goals, and, vice versa, provides feedback on these goals through its on-going internal assessment of supporting tasks. It dynamically configures tasks according to the user's high-level expository purposes for the requested report. It then monitors their performance, and, if necessary, reconfigures them and provides feedback to the user on the efficacy of generating the requested report. In addition to the usual domain knowledge, like cases and heuristic diagnostic rules, FRANK uses taxonomies of report types and reasoning strategies for generating them. These were also developed during this research; they are completely novel.

FRANK produces various types of reports (e.g., balanced pro-con decision brief, biased pro-position advocacy brief) by using case-based, rule-based and procedural reasoning in ways that reflect the user's purposes. For instance, if the user specifies that the information is for a balanced decision briefing, alternative courses of action are presented in a neutral manner. In this context, FRANK would specify that both pro and con supporting cases and analyses must be retrieved, analyzed, and incorporated into the report. FRANK dynamically configures the supporting tasks needed to produce its reports, for instance by specifying the similarity metric and definition of "best" case to be used in the CBR modules incorporated into FRANK. FRANK evaluates the executed tasks, and, if necessary, can reconfigure the reasoning mechanisms and cause them to try to accomplish their tasks again; if all attempts fail, the user is told that the report will be hard to generate and that another report type, selected from FRANK's taxonomy, might be more appropriate.

An obvious benefit of a system like FRANK, which combines different types of reasoning (case-based, rule-based, etc.) in one generic architecture, is leverage and robustness. The different reasoning mechanisms can complement and supplement each other. Another is the ability to make supporting tasks, like case retrieval, sensitive to the use of the information: what you retrieve and how you reason with it depend on what you are going to do with it. FRANK offers a promising step in the direction of intelligent decision support systems using component technologies, such as case-based reasoning and information retrieval, by making them task-sensitive. For instance, FRANK's report and report-generation taxonomies can be used to guide information retrieval, where knowledge of the intended use of information is key to the actual retrieval of it.

Notable results of our work on the FRANK Project are:

- dynamic configuration of lower level, supporting tasks (e.g., CBR) according to high-level tasks (e.g., report needs).
- internal monitoring and possible reconfiguration of tasks (e.g., CBR) during processing.
- selection of strategies and plans affects the type of CBR used.
- interactive assistance to user through suggestion of changes in report type or strategy in response to previous cases.
- demonstration of the utility of a blackboard environment for integrating CBR and procedural knowledge.
- identification of taxonomies of justification strategies and report types typical of a diagnostic domain.

(2) Heuristically guided Retrieval for Case-Based Reasoning. We explored various fundamental CBR issues, such as similarity assessment and case retrieval, in two systems: BankXX and BROADWAY. BROADWAY is based on a blackboard (GBB) architecture [Skalak, 1992]. BankXX is based on the classic architecture of heuristic best-first search [Rissland et al., 1994].

The BANKXX Project

BankXX is a CBR system whose purpose is to search out, harvest, and analyze information needed for case-based argument. BankXX generates case-based argument by gathering and analyzing various types of multiply indexed and multiply represented domain knowledge. BankXX demonstrates the use of classic heuristic search as a control regime for gathering and analyzing information for case-based argument. BankXX is unique in that it uses classic heuristic search with resource bounds to guide and

constrain its reasoning and information gathering in a highly connected directed graph of domain information of various types (e.g., actual cases, prototype cases, case citations, typical factual scenarios, domain theories). BankXX is implemented in Common Lisp. It is instantiated in the legal domain of personal, Chapter 13, bankruptcy.

Completely novel about BankXX is its characterization of case-based argument as heuristic search. BankXX's process of "information-grazing" in an extensive body of domain knowledge is modeled on the way a research assistant might rummage and gather information in a large library or on-line information service where information is highly interconnected through standard indexing schemes, case citation systems, and secondary sources. BankXX is notable in its use of multiply indexed information to handle classic problems, such as too few or too many retrieved cases, and its ability to view cases in light of other knowledge to assist retrieval.

The "bottom up" process modeled by BankXX complements the first research focus of this grant, the FRANK Project, as well as our earlier DARPA-sponsored projects, like CABARET, which take a "top down" perspective. Since BankXX uses some of the commercial indexing methods (e.g., case citation schemes), it could eventually be linked with existing on-line data bases both to extend its own knowledge store and to enhance the type of information gathering they provide. Using CBR with IR would help remediate problems of both current CBR systems, such as the knowledge acquisition bottleneck (e.g., small case bases), as well as well-known shortcomings of current IR systems, such as domain- and task-ignorance.

BankXX models argument in terms of a collection of a dozen or so "argument pieces" which form the building blocks of any precedent-based argument. These include supporting cases, best supporting cases, contrary cases, best contrary cases, leading cases, factor analyses, supporting citations, applicable and nearly applicable legal theories, prototype stories, etc. [Rissland et al., 1994]. BankXX views argument creation as the filling out of these argument pieces. While this is a very simplified representation of an argument, it is, we feel, a prerequisite to more sophisticated representations, such as those modeling logical or rhetorical structure.

We used three different evaluation functions in BankXX. These encode knowledge about case-based argument at different degrees of sophistication. The evaluation function that is least knowledge-intensive uses only general knowledge about information types (e.g., cases, domain theories). The next most knowledge-rich evaluation function is based on knowledge of the argument pieces (e.g., most on-point cases both pro and con, a winning theory, a persuasive prototype, a stereotypical domain story) that go into a good argument. The third evaluation function encodes knowledge about how to evaluate an argument (e.g., according to the centrality of cases cited, win record of a chosen theory on similar cases).

During the last funding period, we performed an extensive series of empirical experiments to assess BankXX's performance. These experiments fell into three classes: (1) evaluation of BankXX performance through comparison against domain expert performance, that is, by comparing BankXX output with court opinions; (2) evaluation of BankXX as a CBR program by variation of several of BankXX's internal mechanisms (e.g., evaluation functions); and (3) comparison of BankXX against previous CBR programs, namely HYPO, which was also developed at UMass. This involved running BankXX on already litigated cases and comparing its output with the published court opinions. Thus far our assessment is that BankXX performs well given that it is a heuristic bottom up approach to argument.

BankXX is notable in that it made the following contributions to our knowledge about CBR:

- The process of gathering information to support an argument can be modeled as heuristic search.
- Case retrieval is effectively performed through a combination of search and knowledge-based indexing.

- The search for cases and other supporting domain knowledge, such as theories or re-occurring scenarios, can be performed within a single retrieval framework.

The BROADWAY Project

In addition to BankXX, we explored fundamental CBR issues in the BROADWAY project [Skalak, 1992]. The BROADWAY system, which was implemented in GBB, represents cases as knowledge sources in a blackboard-based architecture. It explored an alternative model of case representation and retrieval. BROADWAY was instantiated in the domain of BROADWAY was instantiated in the common-sense domain of the selection of suitable automobile models for purchase, given a set of user requirements.

Rather than relying on global metrics to assess case similarity, BROADWAY applies local similarity metrics that apply only in particular sub-regions of the case space, for instance, in neighborhoods of certain noteworthy, prototypical cases. This approach is motivated by the intuitive idea that the way to determine which cases are similar will depend on the types and particulars of cases being considered. (This theme was also explored in BankXX.) Another emphasis of BROADWAY was the representation of cases as dynamic knowledge sources that can recognize when they are relevant to a problem situation. Such an approach may be more amenable to parallel implementations than classic retrieval schemes with more centralized or uniform control.

(3) Machine Learning and CBR. We explored issues of machine learning and CBR in three separate projects:

- (1) CABOT
- (2) learning prototypical cases
- (3) a study of concept drift.

In the CABOT system [Callan et al., 1991], we explored the use of an inductive learning algorithm to adjust case retrieval and adaptation mechanisms. CABOT was instantiated in the classic ML domain of Othello. CABOT demonstrated a 50% reduction in the number of stored cases needed and an increased level of game-playing performance.

Learning Prototypical Cases

OFF-BROADWAY, MC1 and RMHC-* are three CBR-ML systems that learn case prototypes. We feel that methods that enable the learning of prototypes would represent a significant advance since they clearly can play such a large role in both CBR and ML. Since these programs use a variety of so-called weak methods, such as Monte Carlo sampling, random mutation hill-climbing and genetic algorithms, they are applicable to a variety of domains.

All three programs are fully implemented in Common Lisp and have been tested on benchmark databases from the University of California at Irvine Machine Learning Repository, including heart disease and breast cancer databases.

OFF-BROADWAY, MC1 and RMHC-* made the following notable contributions:

- Simple search techniques with a random component can be used to decrease substantially the on-line storage costs of nearest neighbor retrieval algorithms.
- Classification accuracy can be improved in some domains through the identification of prototypes to be used in nearest neighbor classifiers.
- Feature selection can be performed using random search techniques to further reduce classification costs.

Concept Drift

We have developed an approach using standard CBR and ML techniques (e.g., ID5R) to recognize concept drift, through examination of case exemplars and the decision trees that they can induce. In preliminary work [Brodley & Rissland, 1993], we developed a metric to measure "instability" in decision trees (or any partially ordered set). By monitoring the stability of a decision tree as it is incrementally generated, we can detect episodes of potential concept drift. These hypothesized episodes of drift can then be examined more fully—by statistical analyses, by generation of pre- and post-episode trees, etc.—to determine if the concept has indeed changed. If there has been drift, the concept can be altered accordingly.

The ability to detect and respond to concept drift is critical to any system which deals with the real world. This is especially true for systems with some degree of autonomy (e.g., those not constantly re-engineered by human experts). Our approach of combining CBR, whose forte is knowledge of cases, with ML, whose is the induction of concept representations from examples (i.e., cases), offers great promise in attacking this fundamental problem. This research thrust, developed in the last period of this research grant, is currently being pursued with other funding support and is the subject a new proposal.

TECHNOLOGY TRANSFER: We have disseminated the results of this funding through a large variety of conferences and publications. For instance, results of this research have been published in the National Annual Conferences of the AAAI, the biannual international IJCAI conferences, DARPA-sponsored conferences on CBR, Machine Learning Conferences, the Annual Cognitive Science Conference, International Conferences on AI and Law, and AAAI-sponsored workshops and symposia on CBR, IR, and ML. In addition, we have ported our ideas and methods to a large NSF-sponsored project on intelligent information retrieval at University of Massachusetts at Amherst. This is a long term (3-6 year) project, on which Rissland is a co-PI, and it is sponsored by NSF's State-University-Industry Initiative. Preliminary results indicate that a combined CBR-IR approach can lead to striking results.

In addition, ideas from this and previous landmark research, also sponsored through DARPA, are being used by other researchers throughout the world (e.g., several Australian AI and Law projects). Ideas about case-based argument previously developed by our group are now being formalized in various logics by several AI researchers interested in abduction and truth-maintenance (e.g., Prakken, Loui, Gordon).

SYSTEMS DEVELOPED DURING GRANT FUNDING

CABARET — the first mixed paradigm (CBR-RBR) hybrid. It combines HYPO-style CBR with standard (forward and backward) rule-based reasoning through the use of heuristic control rules. The control rules embody *theory of statutory argument* composed of *argument stances, moves* and *primitives*. CABARET-style arguments address key issues in rule interpretation, including *open-textured predicates* and rules with emergent exceptions and unstated prerequisites.

- Rissland, E. L. & Skalak, D. B. (1991). CABARET: Rule Interpretation in a Hybrid Architecture. *International Journal of Man-Machine Studies*, 34(6), 839-887.
- Skalak, D. B. & Rissland, E. L. (1992). Arguments and Cases: An Inevitable Intertwining. *Artificial Intelligence and Law: An International Journal*, 1(1), 3-48.

BankXX — the first CBR system using the framework of heuristic best-first search to guide retrieval of cases and other pertinent knowledge for adversarial argument. It searches through a *case-domain-graph* whose nodes represent key items of legal knowledge, such as cases, legal theories, stereotypical factual scenarios. Search is guided by one of three heuristic evaluation functions—called *node-type, argument-piece, argument-dimensions*. Each takes into consideration how the information in an (open) node can contribute to the evolving argument being built up as BankXX searches.

- Rissland, E. L., Skalak, D. B. & Friedman, M. T. (1993). BankXX: A Program to Generate Argument through Case-Based Search. *Proceedings, Fourth International Conference on Artificial Intelligence and Law*. Amsterdam, The Netherlands. ACM.
- Rissland, E. L., Skalak, D. B. & Friedman, M. T. (1993). Case Retrieval through Multiple Indexing and Heuristic Search. *Proceedings of IJCAI-93*. Chambéry, Savoie, France. International Joint Conferences on Artificial Intelligence.

FRANK — a blackboard-based architecture to create diagnostic reports tailored to the user's prescribed goals and specified report type. This mixed paradigm system, incorporating CBR, rule-based and planning components, dynamically modified its retrieval strategies and queries with feedback from the system's previous success or failure in retrieving cases that support a user's rhetorical and pragmatic goals for the report.

- Rissland, E. L., Daniels, J. J., Rubinstein, Z. B. & Skalak, D. B. (1993). Case-Based Diagnostic Analysis in a Blackboard Architecture. *Proceedings of the Eleventh National Conference on Artificial Intelligence*, 66-72. Washington, DC. AAAI Press/MIT Press.

In addition to these CBR systems, Rissland's CBR Lab has developed a number of CBR-ML experiments and systems:

CABOT — This project in case-based search developed an inductive learning algorithm that adjusts its retrieval and adaptation mechanisms. CABOT demonstrated a 50% reduction in the number of stored cases and an increased level of performance in a game-playing application.

- Callan, J. P., Fawcett, T. E. & Rissland, E. L. (1991). CABOT: An Adaptive Approach to Case-Based Search. *Proceedings, 12th International Joint Conference on Artificial Intelligence*, 803-808. Sydney, Australia. International Joint Conferences on Artificial Intelligence, Inc.

BROADWAY — This system represents cases as blackboard knowledge sources whose preconditions invoke local similarity functions that apply only within a closed neighborhood of each in the space of cases. Broadway provided a new dynamic model of case retrieval in which individual cases are assigned their own similarity metrics, so that retrieval varies according to the type of case being retrieved.

- Skalak, D. B. (1992). Representing Cases as Knowledge Sources that Apply Local Similarity Metrics. *The 14th Annual Conference of the Cognitive Science Society*, 325-330. Bloomington, Indiana. Lawrence Erlbaum.

OFF-BROADWAY — A case retrieval and classification system that learns a set of distinguished cases (actual prototypes) that have demonstrated classification power. The system maintains a population of sets of potentially prototypical cases; each set is evaluated by its classification accuracy. A genetic algorithm is used to search the space of prototype sets ("spanning sets" in the language of §I.2.B) for one of superior accuracy.

- Skalak, D. B. (1993). Using a Genetic Algorithm to Learn Prototypes for Case Retrieval and Classification. *Proceedings of the AAAI-93 Case-Based Reasoning Workshop (Technical Report WS-93-01)*. Washington, D.C. American Association for Artificial Intelligence, Menlo Park, CA.
- Skalak, D.B. (1994).

ABISS — a blackboard system that integrated CBR with "from-scratch" first principles problem solving to predict integer sequences. The system was one of the first to use a blackboard architecture to build a hybrid system that incorporated case-based reasoning with other reasoning techniques. The program also investigated the utility of case-based control strategies and investigated

the application of a number of similarity metrics for case retrieval at different points in the course problem solving.

- Rissland, E.L., Basu, C., Daniels, J., McCarthy, J. Rubinstein, Z. & Skalak, D.B. "A Blackboard-based Architecture for CBR: An Initial Report". *Proceedings of the Case-Based Reasoning Workshop*, May 1991, pp. 77-92.

PUBLICATIONS

- Rissland, E.L., Skalak, D.B., and Friedman, M.T. "BankXX: Supporting Legal Arguments through Heuristic Retrieval", in preparation for submission to the *Journal of Artificial Intelligence and Law*.
- Rissland, E.L., Skalak, D.B. and Friedman, M.T. "Evaluating Legal Arguments" in preparation for submission to the *Journal of Artificial Intelligence and Law*.
- Skalak, D.B. "Prototype and Feature Selection by Sampling and Random Mutation Hill Climbing Algorithms". To appear in *Proceedings of the Eleventh International Conference on Machine Learning*, New Brunswick, NJ. 1994. Morgan Kaufmann.
- Rissland, E.L., Skalak, D.B. & Friedman, M.T. "Heuristic Harvesting of Information for Case-Based Argument" *Proceedings of the Twelfth National Conference for Artificial Intelligence (AAAI-94)*, Seattle, WA, (in press)
- Rissland, E.L., Daniels, J., Rubinstein, Z., & Skalak, D.B. "Case-based Diagnostic Analysis in a Blackboard Architecture." *Proceedings of the Eleventh National Conference for Artificial Intelligence (AAAI-93)*, Washington, D.C. pp. 66-72.
- Rissland, E.L., "Legal Reasoning and Artificial Intelligence: Some recent advances with an emphasis on case-based approaches." *Proceedings First NISCALE Workshop on Criminality and Law Enforcement*. Netherlands Institute for the Study of Criminality and Law Enforcement (NISCALE), Den Haag, The Netherlands, October 1993, pp. 187-205.
- Rissland, E.L., Daniels, J.J., Rubinstein, Z.B. and Skalak, D.B. "Diagnostic Case Retrieval Guided by Evaluation and Feedback". *Proceedings of the AAAI-93 Workshop on Case-Based Reasoning*, Washington, DC, July 1993, pp. 135-140.
- Rissland, E.L., Skalak, D.B. and Friedman, M.T. "Using Heuristic Search to Retrieve Cases that Support Arguments. *Proceedings of the AAAI-93 Workshop on Case-Based Reasoning*, Washington, DC, July 1993, pp. 5-11.
- Rissland, E.L., Skalak, D.B., & Friedman, M.T. "BANKXX: A Program to Generate Argument through Case-Based Search." *Proceedings of the Fourth International Conference on AI and Law (ICAIL-93)*, Amsterdam, The Netherlands, June 1993, pp. 117-124.
- Rissland, E.L., Skalak, D.B., & Friedman, M.T. "Case Retrieval Through Multiple Indexing and Heuristic Search." *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93)*, Chambéry, Savoie, France, August 1993, pp. 902-908
- Brodley, C. E. & Rissland, E. L. "Measuring Concept Change." *Working Notes of the AAAI-93 Spring Symposium on Training Issues in Incremental Learning*, 1993. pp. 98-107.

- Skalak, D.B. "Representing Cases as Knowledge Sources that Apply Local Similarity Metrics". *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, pp. 325-330. 1992.
- Skalak, D.B. and E.L. Rissland., "Using Case-Based Reasoning to Extend the Expertise of Expert Systems" in *Expert Systems in Law*, A. Martino (Ed.), Elsevier Science Publishers. 1992. pp. 321-338.
- Skalak, D.B. & Rissland, E.L. "Arguments and Cases: An Inevitable Intertwining." *Artificial Intelligence and Law —An International Journal*, Kluwer , pp. 3-44, Fall 1992.
- Skalak, D.B. & Rissland, E.L. "Argument Moves in a Rule-Guided Domain." *Proceedings of the Third International Conference on AI and Law (ICAIL-91)*, Oxford, June 1991, pp. 1-12.
- Callan, J.P., Fawcett, T.E. & Rissland, E.L. "CABOT: An Adaptive Approach to Case-Based Search". *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, pp. 803-808.
- Rissland, E. L., & Skalak, D. B., "CABARET: Statutory Interpretation in a Hybrid Architecture." *International Journal of Man-Machine Studies (IJMMS)*, June, 1991, (34):839-887.
- Rissland, E.L., Basu, C., Daniels, J., McCarthy, J. Rubinstein, Z. & Skalak, D.B. "A Blackboard-based Architecture for CBR: An Initial Report". *Proceedings of the Case-Based Reasoning Workshop*, May 1991, pp. 77-92.
- Callan, J.P., Fawcett, T.E. & Rissland, E.L. "Adaptive Case-based Reasoning". *Proceedings of the Case-Based Reasoning Workshop*, May 1991, pp. 179-190.

INVITED LECTURES, PANELS, COLLOQUIA AND CONFERENCE PRESENTATIONS

Rissland

October 1993 - Invited Speaker, First NISCALE Conference on Criminality and Law Enforcement (The Netherlands)

August 1993 - International Joint Conf. on Artificial Intelligence, IJCAI-93 (Chambery, France)

July 1993 - AAAI-93 (Washington, DC)

July 1993 - AAAI Workshop on CBR, (Washington, DC)

June 1993 - International Conference on AI and Law, ICAIL-93 (Amsterdam, NL)

May 1993 - Brandeis University

March 1993 - AAAI Spring Symposium on CBR and IR (Stanford)

March 1993 - AAAI Spring Symposium on Machine Learning (Stanford)

March 1993 - Lockheed Symposium Series (Palo Alto)

August 1992 - Series of six lectures at Department of Electronics, Research Labs, New Delhi, India.

August 1992 - IEEE/Delhi, India.

July 1992 - Invited plenary lecture "Recent Progress in AI and Legal Reasoning." Tenth National Conference for Artificial Intelligence (AAAI-92) San Jose, CA.

May 1992 - Daylong tutorial on CBR, Hartford Graduate Center.

May 1992 - CBR Roundtable, Boston, MA

November 1991 - Distinguished Lecture, University of Wisconsin (2 talks)

June 1991 - International Conference on AI and Law, ICAIL-91 (Oxford England)

May 1991 - DARPA Case-based Reasoning Workshop (Washington, DC)

March 1991 - AAAI Spring Symposium (Palo Alto, CA)
February 1991 - Apple Corporate Research Labs (Cupertino, CA)

Skalak

July 1992 - Fourteenth Annual Conference of the Cognitive Science Society, Bloomington, IN

Rubinstein

July 1993 - AAAI-93 - Case-Based Diagnostic Analysis in a Blackboard Architecture. - Washington, D.C.

Friedman

July 1993 - AAAI Workshop on CBR - Washington DC

Daniels

July 1993 - AAAI CBR Workshop - Diagnostic Case Retrieval Guided by Evaluation and Feedback

Conferences Attended:

Rissland

January 1994 - CBR Roundtable, Cambridge MA
August 1993 - International Joint Conf. on Artificial Intelligence, IJCAI-93 (Chambery, France)
July 1993 - AAAI-93 (Washington, DC)
July 1993 - AAAI Workshop on CBR, 2 talks (Washington, DC)
June 1993 - International Conference on AI and Law, ICAIL-93 (Amsterdam, NL)
March 1993 - AAAI Spring Symposium on CBR and IR (Stanford)
March 1993 - AAAI Spring Symposium on Machine Learning (Stanford)
March 1993 - CBR Roundtable, Boston, MA
July 1992 - Invited plenary lecture "Recent Progress in AI and Legal Reasoning."
Tenth National Conference for Artificial Intelligence (AAAI-92) San Jose, CA.
May 1992 - CBR Roundtable, Boston, MA
June 1991 - International Conference on AI and Law, ICAIL-91 (Oxford England)
May 1991 - DARPA Case-based Reasoning Workshop (Washington, DC)
March 1991 - AAAI Spring Symposium (Palo Alto, CA)
October 1991 - CBR Roundtable, Amherst, MA

Skalak

January 1994 - CBR Roundtable, Cambridge MA
July 1993 - Eleventh National Conference on Artificial Intelligence, Washington, DC
July 1993 - AAAI Case-Based Reasoning Workshop, Washington, DC
June 1993 - Fourth International Conference on Artificial Intelligence and Law, Amsterdam, The Netherlands
April 1993 - The Second International Conference on Artificial Intelligence Applications on Wall Street, New York, NY
July 1992 - Fourteenth Annual Conference of the Cognitive Science Society, Bloomington, IN
March 1993 - Tenth International Conference on Machine Learning, Amherst, MA
June 1991 - Third International Conference on Artificial Intelligence and Law, Oxford, England

May 1991 - Third DARPA Case-Based Reasoning Workshop, Washington, DC
March 1993 - CBR Roundtable, Boston, MA
May 1992 - CBR Roundtable, Boston, MA
October 1991 - CBR Roundtable, Amherst MA

Rubinstein

July 1993 - AAAI CBR Workshop - Washington DC.
July 1993 - AAAI-93 (Washington, DC)
May 1991 - DARPA Case-based Reasoning Workshop - Washington, DC

Friedman:


January 1994 - CBR Roundtable, Cambridge MA
June 1993 - International Conference on AI and Law, ICAIL-93 (Amsterdam, NL)
July 1993 - AAAI CBR Workshop - Washington DC.
March 1993 - CBR Roundtable, Boston, MA
May 1992 - CBR Roundtable, Boston, MA
October 1991 - CBR Roundtable, Amherst MA
June 1991 - International Conference on AI and Law, ICAIL-91 (Oxford England)
May 1991 - DARPA Case-based Reasoning Workshop - Washington, DC

Daniels:

July 1993 - AAAI CBR Workshop - Washington DC.
May 1991 - DARPA Case-based Reasoning Workshop - Washington, DC

Off Broadway Quad Chart

Picture



New Ideas

- Prototypes may be identified by a genetic algorithm for use in a nearest neighbor classifier
- Features may be filtered using a genetic algorithm
- Accuracy of sampling and search methods shown to depend on degree of clustering of data set


Current and Future Work

- Monte Carlo sampling used to identify prototypes and features
- Random mutation hill climbing has also been used
- Explore combinations of simple nearest neighbor classifiers

Impact

- Decreases on-line storage costs of nearest-neighbor algorithms without sacrificing accuracy
- In tested UCI domains, approximately 1% of instances retained
- Classification accuracy actually increased in some cases where fewer prototypes used

FRANK Quad Chart



New Ideas

- Case retrieval plan should reflect user's expository goals and perspective
- Case retrieval can be dynamically adjusted on basis of feedback
- Specific suggestions can be made to user to change perspective or goals, based on case findings

Future Work


- Use FRANK to drive IR
- Use FRANK to assist decision-makers (e.g., physicians)
- Integrate FRANK and BankXX technologies

Impact

- Decision support:
Makes CBR responsive to user's goals
- Rationale analysis:
System dynamically adjusts retrieval strategy to locate useful cases
- Report preparation:
Suggests alternative presentation strategies to user when support for current approach is insufficient

BankXX Quad Chart

Picture



Case-Base Search

Impact

- Domains: Applicable to "weak-theory" domains generally (e.g., strategic planning, intelligence analysis)
- Tasks: In such domains, provides case and other support as building blocks for arguments and explanations
- Practicality: Search can scale to large knowledge bases and use existing index structures

New Ideas

- Arguments can be created bottom-up through case-base search.
- Indexing and search are complementary case retrieval methods.
- Diverse knowledge types (e.g., theories, stereotypical situations) are placed in same framework as cases to improve retrieval of all types of support

Future Work

- BankXX evaluation completed
- Prototype-learning
- Theory-learning
- Integration with bottom-up and other control strategies for argument