The Fun Begins with Retrieval: Explanation and CBR

Edwina L. Rissland

Department of Computer Science, University of Massachusetts
Amherst, Massachusetts, U.S.A. 01003

Abstract. This paper discusses the importance of the post-retrieval steps of CBR, that is, the steps that occur after relevant cases have been retrieved. Explanations and arguments, for instance, require much to be done post-retrieval. I also discuss both the importance of explanation to CBR and the use of CBR to generate explanations.

1 Introduction

Some of the most interesting aspects of CBR occur after relevant cases have been retrieved. Explanations—and here I include argument—are some of the most important, and they play a central role in CBR. They are needed to elucidate the results of the case-based reasoning—why a case was interpreted or classified in a particular way, how a new design or plan works, why a particular diagnosis is most compelling, etc.—and explanations can themselves be created using CBR. For CBR to create arguments, designs, plans, etc., much work must be done, and most of it begins after relevant cases have been retrieved [18], [23]. That is, a good part of the core of case-based reasoning occurs post-retrieval.

Since some systems like Branting's GREBE [5] and Koton's CASEY [19] create their explanations using adaptive mechanisms, it is not clear how to draw a line between so-called interpretive and adaptive CBR systems. However, it is abundantly clear that in both types the lion's share of the work is done post-retrieval. While explanation is not the focus of other adaptive CBR systems like Hammond's CHEF [16] or Cheetham's FORM TOOL [8], they do indeed accomplish their tasks post-retrieval. That is, retrieval is only an initial step in case-based problem-solving, and the fun—and most of the hard work—occurs post-retrieval.

The ability to explain one's reasoning is a hallmark of intelligence, and is—or should be—one of the keystones of CBR systems. This is so whether CBR is being used to interpret or classify a new case, or to adapt an old solution in order to solve a new problem. Too often our CBR systems—particularly those used to classify new cases—de-emphasize or even forget about the post-retrieval "R's" in CBR, like "reuse, revise, retain" [1]. Retrieval is, of course, an absolutely crucial step in CBR, but it is only one of several: it is one of the six R's in Göker & Roth-Berghofer's formulation [14] and one of the eleven in Derek Bridge's [7].

Explanation is really a kind of teaching, and can be viewed as the other side of the coin of learning. Both explanation and learning are inextricably intertwined with

T.R. Roth-Berghofer et al. (Eds.): ECCBR 2006, LNAI 4106, pp. 1 – 8, 2006.

concepts, conceptual emergence, and concept change. We really thus have a longterm cycle in which cases play an integral role. Although I won't really consider the closely related problems of similarity assessment and credit assignment in this presentation, they are indeed very important to both this overarching cycle and to the inner workings of CBR, including retrieval.

Most of us know how critical the choices of similarity metric and case space structure are in CBR. Both choices are motivated by what we want to bring to the fore in the reasoning. They also dictate what will be possible to accomplish in it or explain about it. That is, there is another inescapable intertwining in CBR between notions of similarity and explanation. One can thus say that the fun also begins before retrieval.

This is especially true in systems that stop at retrieval or a slight bit beyond—what we might call CB-little-r systems—for instance, those that use retrieved examples to classify a new case (e.g., with nearest neighbor methods), or that use the results of the early steps of CBR to initiate other types of processing, like information retrieval. For instance, the SPIRE system stopped short of argument creation, but used retrieval and similarity assessment (e.g., HYPO-style claim lattices) to generate queries for a full-text IR engine [9], [44], [45]. In CB-r systems there is perhaps a more critical dependence on getting the space and metric "right" than in CBR systems that keep on processing or that can explain themselves.

In fact, explanations can help lessen the burdens of CBR systems since they make their reasoning less opaque, a requirement, I believe, for intelligent systems. Explaining the behavior of CBR systems to users is receiving new attention in recent work, with goals such as enabling systems to explain their questions [31] or to explain the space of retrieval possibilities [37]. Leake & McSherry's [24] collection on CBR and explanation demonstrates new activity in a number of directions, but current work just scratches the surface of possibilities. Even with regard to similarity and retrieval, we don't, in my opinion, have enough variety in our ideas. So, in addition to pressing for more consideration of the post-retrieval R's, I would also press for more research on the first R: retrieval.

2 Cases as Both Drivers and Aids

Cases (called exemplars or examples in other contexts) not only are drivers of the inter-looped processes of explanation and concept evolution, but they can also serve as central elements in the representation of concepts and the teaching of the art of explanation. For instance, examples can be used by themselves to produce a totally extensional representation; that is, a concept is simply considered to be the set of its positive exemplars. They can participate in hybrid representations in concert with other mechanisms like rules or prototypes or statistical models. Examples can serve as extensional annotations on rules; these can serve to help resolve ambiguities in rules or terms and to keep them up to date with new interpretations and exceptions. Concrete examples can be used to capture some of the information that statistics summarize but cannot explicitly represent. Cases—like atypical borderline examples, anomalies, penumbral cases—are particularly useful in the tails of distributions where data can be sparse.

Hybrid approaches, both in representation and reasoning, have been used in a variety of systems from the earliest days of CBR to the present: CABARET, GREBE, ANAPRON, CAMPER, CARMA, and FORM TOOL, for instance. (For concise overviews of such hybrids, see [27], [28].) Cases in many of these systems serve to complement and supplement other forms of reasoning and representation. For example, ANAPRON used cases to capture exceptions to rules [15]. An early landmark system in AI and Law by Anne Gardner to model the issue-spotting task issue on problems of the kind found on law school and bar exams used examples as sanity checks on rule-based reasoning and when rule-based reasoning failed, or when as Gardner puts it, "the rules run out" [12]. CABARET used cases in these ways as well [46]. In addition CABARET used cases to help carry out a repertoire of strategies and tactics for performing statutory interpretation, that is, determination of the scope and meaning of legal rules and their ingredient predicates [51]. CAMPER used cases and rules to generate plans for nutritional menus [26]. CARMA used cases and rules together with models [17].

Related to our interests are two paradigms from psychology concerning reasoning with and representing concepts and categories: the *prototype* and *exemplar* views [34]. (Murphy's *The Big Book of Concepts* provides an extensive overview.) Pioneered by Rosch, Medin and others about thirty years ago, the prototype paradigm focuses on the "typicality" of examples [49], [32], [33], [52]. Sometimes a prototype is taken to be a maximally typical actual example; other times it is more of a summary or a model like a frame in AI. Prototypes have been extensively investigated in psychology. In the exemplar view, a concept is represented by the set of its positive examples. It has not been as thoroughly considered, and it is nowhere as sophisticated as our own work in CBR to which it is obviously closely related. For instance, we have many highly elaborated and computationally well-defined mechanisms for case retrieval and comparison. Hybrid representations—of prototypes and examples, say—have not been used much at all in psychology. On the other hand, hybrid approaches have been extensively explored in CBR and closely related fields like AI and Law. For example, McCarty and Sridharan early on proposed a hybrid "prototype-plus-deformation" approach [29]. (For an overview of AI & Law, see [43].)

3 The Centrality of Explanation

Explanation is central to all types of CBR. In interpreting a new case using past interpretations from the case base, many CBR systems reason with relevant similarities and differences. Such interpretive CBR can involve analogically mapping over explanations from existing precedents, for instance by structure mapping, or by constructing a completely new rationale, for instance, by HYPO-style dimensional analysis.

Many of the earliest CBR systems focused on explanations. For instance, HYPO used highly relevant previously interpreted cases—that is, precedents—to interpret the case at hand and generate arguments both for and against a particular conclusion [4]. HYPO elucidated interpretations with explanatorily-relevant hypotheticals. Branting's GREBE re-used and mapped over past explanations to new situations [5], [6]. It employed the structure mapping model of analogical reasoning developed by Gentner and others [10], [11], [13]. Koton's CASEY used a causal model of heart disease and a calculus of explanatory differences to adapt a previous explanation to diagnose a new patient's symptoms [19]. Kass, Leake and Owen's SWALE directly addressed the problem of explaining a phenomenon—particularly an unexpected one like the collapse of a racehorse in the prime of its life—by recalling and adapting relevant past explanations [50]. Leake's work also illustrated the centrality of explanations by showing how they can serve many different goals, for system reasoning as well as external performance tasks [22].

More recently, CBR has been used to foster learning of how to perform specialized types of explanation like appellate argument. Aleven and Ashley's CATO tackled the task of teaching law students how to make good precedent-based arguments [2], [3]. McLaren's SIROCCO used examples to help explain ethics principles [30]. In his research, Aleven demonstrated that CATO-trained law students do as well as those trained in the traditional ways involving written and oral exercises [2]. This comports well with what has been found experimentally in psychology.

Psychologists have shown that explicit comparison of past exemplars with a new instance can promote more nuanced and better learning. This is true across a whole range of learners from toddlers to business school students. For instance, Gentner showed that exploring explicit analogical mappings in a concept categorization task can lead to better categorization in the sense that the children focused more on deep properties (like functionality) rather than on shallow ones (like appearance) [20], [35]. Business school students were better able to choose the appropriate negotiation strategy for a new problem case when they had already practiced making analogical comparisons [25].

The examples and cases so vital to CBR and explanation can themselves be constructed using CBR. Example generation is the twin task to example interpretation. In it one creates examples that meet specified criteria; these typically serve the needs of other processes like explanation, argument, teaching, supervised learning, etc. CEG can be viewed as a design task. In my lab, we developed a "retrieval-plus-modify" architecture for creating examples satisfying such prescribed constraints, and called it *Constrained Example Generation* or CEG [39], [47]. In CEG, a person or machine tries to satisfy as many of the desiderata as possible by retrieval—that is, finding examples that possess as many of them as possible, and then trying to satisfy the remaining properties through modification. This is essentially "adaptive" CBR. However, given the nasty way that constraints can interact, this is not easy. A fuller model of example generation should integrate techniques from constraint satisfaction problem (CSP) solving into CEG. There has been important work on CSP and on integrating CBR and CSP since CEG was developed (See [27], [28]).

CEG was initially directed at generating counter-examples in mathematics and grew into a larger effort to explore the use of examples in other types of explanations (e.g., on-line help systems) and arguments (e.g., appellate-style legal argument) [48]. Counter-examples are like designs having specified properties. For instance, if one wants to show that not all continuous functions are necessarily differentiable, that not all quadrilaterals are necessarily convex, that not all primes are odd, one needs examples that possess precise mathematical properties. Such closely crafted examples are pivotal in the dialectic of proofs and refutations—they can annihilate conjectures and force concept change, for instance [21], [36], [41]. Examples of various types—start-up, reference, counter, etc.—play a very important role in developing understanding in mathematics [38]. Where such interesting examples come from is an intriguing question. One answer is through an adaptive process like CEG.

Generating hypotheticals can also be viewed as a kind of CEG task, and thus amenable to adaptive methods. For instance, hypos can be created by taking a seed case and making it more extreme, or by combining two hypos to create a conflict hypo [42]. Hypos can be used with surgical precision in oral argument and law school dialogues to show the limits of a line of reasoning or to uncover fatal flows in its logic [40]. The reasoning context provides the desiderata for the hypos and a CBR process can be used to produce them. The HYPO system actually grew out of our work on examples and hypos using the CEG approach [48].

The study of legal explanations, including argument, is a rich area for study and it can provide us both with interesting data to "explain" and learn from, and with interesting techniques to borrow and apply in other domains. For instance, if we want to build CBR systems that can analogize and distinguish cases as a way of explaining why a particular outcome should obtain, there is a plethora of examples from law that we can examine. There are many kinds of legal argument moves and strategies—slippery slope, strawman, chicken-turkey-fish, reduction loop—that can profitably be used in non-legal domains as well [51].

While there have indeed been many insightful landmark systems on explanation and argument, they have by no means exhausted the topic. It's time to push the envelope further.

4 Conclusion

In summary, in my talk I focus on explanations and CBR, and the larger issue of what can be accomplished in the post-retrieval stages of CBR. This is not to diminsh the importance of similarity assessment and of retrieval in CBR, but rather to suggest that we cannot ignore what happens once relevant cases have been retrieved. Some of the most interesting work for CBR—adapting old solutions to solve new problems, using existing precedents to interpret new facts—is done post-retrieval. Of late, we have shied away from these stages of CBR, some of which I grant can be quite difficult. But for us to miss out on all the post-retrieval fun in CBR would indeed be a shame.

Acknowledgments

I acknowledge the support of the US National Science Foundation in preparing this article. Any opinions, findings, conclusions, or recommendations expressed in this material are my own and do not necessarily reflect the views of the NSF. I also thank David Leake for his generous feedback.

References

1. Aamodt, A. & Plaza, E. "Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approaches." *Artificial Intelligence Communications*, IOS Press, 7(1): 39-59, 1994.
2. Aleven, V. "Using background knowledge in case-based legal reasoning: A computational model and an intelligent learning environment." *Artificial Intelligence*. 150(1-2): 183-237, 2003.

3. Aleven, V. & Ashley, K.D. "Automated Generation of Examples for a Tutorial in Case-Based Argumentation." *Second International Conference on Intelligent Tutoring Systems* (ITS-92), pp. 575-584, 1992.
4. Ashley, K.D. *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*. MIT Press, 1990.
5. Branting, L.K. "Building Explanations from rules and structured cases." *International Journal of Man-Machine Studies,* (34)6: 797-837, June, 1991.
6. Branting, L.K. *Reasoning with Rules and Precedents*. Kluwer, 2000.
7. Bridge, D. "The Virtue of Reward: Performance, Reinforcement, and Discovery in Case-Based Reasoning." Invited talk, *Proceedings 6th International Conference on Case-Based Reasoning (ICCBR 2005)*. Chicago, August, 2005.
8. Cheetham, W. "Tenth Anniversary of the Plastics Color Formulation Tool." *AI Magazine*, 26(3): 51-61, 2005.
9. Daniels, J.J., & Rissland, E.L. "What You Saw Is What You Want: Using Cases to Seed Information Retrieval." *Case-Based Reasoning Research and Development* (Springer Lecture Notes in Artificial Intelligence No. 1266) the *Proceedings Second International Conference on Case-Based Reasoning, ICCBR-97*, Providence, July 1997, pp. 325-337.
10. Falkenhainer, B, Forbus, K.D. & Gentner, D. "The structure-mapping engine: algorithm and examples." *Artificial Intelligence*. 41(1): 1-63, 1989.
11. Forbus, K., Gentner, D., & Law, K. "MAC/FAC: A model of similarity-based retrieval." *Cognitive Science*, 19, pp.141-205, 1994.
12. Gardner, A vdL. *An Artificial Intelligence Approach to Legal Reasoning*. MIT Press, 1987.
13. Gentner, D. "Structure-mapping: A theoretical framework for analogy." *Cognitive Science*, 7(2): 155-170, 1983.
14. Göker, M., & Roth-Berghofer, T. "Development and utilization of a Case-Based Help desk Support Sytem in a Corporate Environment." *Proceedings Third International Conference on CBR*, pp. 132-146, 1999.
15. Golding, A.R. & Rosenbloom, P.S. "Improving rule-based systems through case-based reasoning." *Proceedings of the Ninth National Conference on Artificial Intelligence* (AAAI-91), pp. 22-27, 1991.
16. Hammond, K. J. *Case-based planning*. Academic Press, 1989.
17. Hastings, J., Branting, L.K., & Lockwood, J. "CARMA: A case-based rangeland management advisor." *AI Magazine*, 23(2): 49-62, 2002.
18. Kolodner, J.L. *Case-Based Reasoning*. Morgan-Kaufman, 1993.
19. Koton, P. "Reasoning about evidence in causal explanation." *Proceedings Seventh National on Artificial Intelligence (AAAI-88),* pp. 256-261.
20. Kotovsky, L., & Gentner, D. "Comparison and categorization in the development of relational similarity." *Child Development*, 67, 2797-2822, 1986.
21. Lakatos, I. *Proofs and Refutations: The Logic of Mathematical Discovery*. Cambridge University Press, 1976.
22. Leake, D. B. *Evaluating Explanations: A Content Theory*. Lawrence Erlbaum, 1992.
23. Leake, D. B. (ed). *Case-Based Reasoning: Experiences, Lessons, & Future Directions*. AAAI Press/MIT Press, 1996.
24. Leake, D.B., & McSherry, D. (eds). *Artificial Intelligence Review*: Special Issue on Case-Based Reasoning and Explanation, 24(2), 2005.
25. Loewenstein, J., Thompson, L. & Gentner, D. "Analogical encoding facilitates knowledge transfer in negotiation." *Psychonomic Bulletin & Review*, 6, pp. 586-597, 1999.
26. Marling, C.R., Petot, G.J., & Sterling, L.S. "Integrating case-based and rule-based reasoning to meet multiple design constraints." *Computational Intelligence* 15(3): 308-332, 1999.

27. Marling, C., Muñoz-Avila, H., Rissland, E.L., Sqalli, M., & Aha, D. "Case-Based Reasoning Integrations." *AI Magazine*, 23(1): 69-86, Spring 2002.
28. Marling, C., Rissland, E.L., & Aamodt, A. "Integrations with case-based reasoning." *Knowledge Engineering Review*, 20(3), Special issue on Case-Based Reasoning, 2006.
29. McCarty, L.T. & Sridharan, N.S. "The Representation of an Evolving System of Legal Concepts: I. Logical Templates." *Proceedings Third National Conference of the Canadian Society for Computational Studies of Intelligence*, Victoria May 1980, pp. 304-311.
30. McLaren, B.M. "Extensionally defining principles and cases in ethics: An AI model." *Artificial Intelligence*. 150(1-2): 145-181, 2003.
31. McSherry, D. "Explanation in Recommender Systems." In Leake, D.B., and McSherry, D. (eds). *Artificial Intelligence Review*: Special Issue on Case-Based Reasoning and Explanation, 24(2), 2005.
32. Medin, D. "Concepts and conceptual structure." *American Psychologist*, 44, pp.1469-1481, 1989.
33. Medin, D.L. & Shaffer, M.M. "Context theory of classification learning." *Psychological Review*, 85, pp.207-238, 1978.
34. Murphy, G.L. *The Big Book of Concepts*. MIT Press, 2002.
35. Namy, L.L., & Gentner, D. "Making a silk purse out of two sows' ears: Young children's use of comparison in category learning." *Journal of Experimental Psychology: General, 131*, pp. 5-15, 2002.
36. Polya, G. *Mathematical Discovery*. Volume 2. John Wiley & Sons, 1965.
37. Reilly, J., McCarthy, K., McGinty, L., & Smyth, B. "Explaining Compound Critiques". In Leake, D.B., and McSherry, D. (eds). *Artificial Intelligence Review*: Special Issue on Case-Based Reasoning and Explanation, 24(2), 2005.
38. Rissland, E. (Michener) "Understanding Understanding Mathematics." *Cognitive Science*, 2(4): 361-383, 1978.
39. Rissland, E.L. "Example Generation." *Proceedings Third National Conference of the Canadian Society for Computational Studies of Intelligence*, Victoria, May 1980, pp. 280-288.
40. Rissland, E.L. "Examples in Legal Reasoning: Legal Hypotheticals." *Proceedings Eighth International Joint Conference on Artificial Intelligence (IJCAI-83).* Karlsruhe, August 1983, pp. 90-93.
41. Rissland, E.L. "The Ubiquitous Dialectic." *Proceedings Sixth European Conference on Artificial Intelligence* (ECAI-84), Pisa, pp. 367-372, North Holland, 1984.
42. Rissland, E.L., & Ashley, K.D. "Hypotheticals as Heuristic Device." *Proceedings Fifth National Conference on Artificial Intelligence (AAAI-86),* Philadelphia,, pp. 289-297, 1986.
43. Rissland, E.L. & Ashley, K.D., & Loui. R.P. "AI and Law: A Fruitful Synergy." *Artificial Intelligence*. 150(1-2): 1-15, 2003.
44. Rissland, E.L. & Daniels, J.J. "Using CBR to Drive IR." *Proceedings Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95)*, Montreal, pp. 400-407, 1995.
45. Rissland, E.L. & Daniels, J.J. "The Synergistic Application of CBR to IR". *Artificial Intelligence Review*: Special Issue on the use of AI in Information Retrieval, 10, pp. 441-475, 1996.
46. Rissland, E.L., & Skalak, D.B. "CABARET: Statutory Interpretation in a Hybrid Architecture." *International Journal of Man-Machine Studies*, 34(6): 839-887, June 1991.
47. Rissland, E.L., & Soloway, E.M. "Overview of an Example Generation System." *Proceedings First National Conference on Artificial Intelligence (AAAI-80),* Stanford, pp. 256-258, 1980

48. Rissland, E.L., Valcarce, E.M., & Ashley, K.D. "Explaining and Arguing with Examples." *Proceedings Fourth National on Artificial Intelligence (AAAI-84),* Austin, pp. 288-294, 1984
49. Rosch, E, & Mervais, C.B. "Family resemblance: Studies in the internal structure of categories." *Cognitive Psychology* 7: 573-605, 1975.
50. Schank, R.C., & Leake, D.B., "Creativity and Learning in a Case-Based Explainer." Artificial Intelligence, 40(1-3): 353-385, 1989.
51. Skalak, D.B. & Rissland, E.L. "Arguments and Cases: An Inevitable Intertwining." *Artificial Intelligence and Law,* 1(1): 3-44, 1992.
52. Smith, E.E., & Medin, D.L. *Categories and Concepts*. Harvard University Press, 1981.