The Knowledge Engineering Review, Vol. 20:3, 293–298. © 2006, Cambridge University Press
doi:10.1017/S0269888906000701 Printed in the United Kingdom

Case-based reasoning and law

EDWINA L. RISSLAND[1], KEVIN D. ASHLEY[2] and L. KARL BRANTING[3]

[1]*Department of Computer Science, University of Massachusetts, Amherst, MA 01003, USA;*
e-mail: Rissland@cs.umass.edu
[2]*Graduate Program in Intelligent Systems and School of Law, University of Pittsburgh, Pittsburgh, PA 15260, USA;*
e-mail: Ashley@pitt.edu
[3]*BAE Systems, Inc., Columbia, MD 21046, USA;*
e-mail: Karl.branting@baesystems.com

Abstract

A primary research stream that contributed to the birth of case-based reasoning (CBR) was Artificial Intelligence and Law. Since law is largely about cases, it is a particularly interesting domain for CBR researchers. This article surveys some of the historically significant systems and developments in this field.

1 Introduction

A primary research stream that contributed to the birth of case-based reasoning (CBR) was Artificial Intelligence (AI) and Law. The research pursued in the early 1980s by Rissland, Ashley, Branting, and Skalak explored the rich vein of case-based reasoning in the context of legal argument. Some of these seminal projects were presented in a special 1991 pair of issues of the *International Journal of Man-Machine Studies* (e.g., Ashley 1991; Branting, 1991; Rissland & Skalak, 1991). Ideas from these research projects lay the foundation of what is now termed *interpretive* CBR, that is, how to interpret new cases in light of past interpretations. This work has also influenced the community that develops formal models of argumentation and defeasible reasoning, and these models have in turn contributed more formal models to CBR (e.g., Bench-Capon & Sartor, 2003). The AI and law community continues to provide a rich tributary of ideas and techniques about CBR and for integrating it with other reasoning modalities in CBR hybrids, such as rule-based reasoning, heuristic search, and information retrieval. (See, for instance, the articles in this issue on CBR integrations and on textual CBR.) A special issue of the journal *Artificial Intelligence* provides both a more in-depth look at the field of AI and law as well as a general history (Rissland *et al*., 2003).

Law is a particularly interesting domain for CBR researchers given that it is largely about cases. Anglo-American law is precedent-based; its judicial standard, *stare decisis*, mandates that similar cases should be decided similarly. While the standard focuses on reasoning from case to case, it does not say how to determine 'similarity'. In fact, similarity is not static. It can depend on one's viewpoint and desired outcome. This makes law an excellent domain for exploring reasoning with cases and for studying issues about indexing and retrieval, similarity assessment, and case comparison. Other characteristics make it inviting for CBR, as well: (1) the open-textured nature of legal concepts; (2) a variety of task orientations (e.g., advocacy, adjudication, advising); (3) diverse categories of knowledge (e.g., cases, statutes, regulations) and modalities of reasoning (e.g., case-based, rule-based, hybrid); (4) significant repositories of richly linked case knowledge; and (5) its use of hypotheticals in reasoning (Rissland, 1990; Rissland *et al*., 2003). As well as actual

https://doi.org/10.1017/S0269888906000701 Published online by Cambridge University Press

cases, legal experts (e.g., law professors, Supreme Court justices) employ hypotheticals or *hypos*—made-up cases that are often telling modifications of real cases—to shed light on issues and emphasize points (Rissland, 1983, 1989). Hypos, including streamlined prototypes, are used in Socratic law school dialogues, appellate oral arguments, and legal planning.

2 Examples, hypos, and HYPO

Significantly for CBR, reasoning with hypotheticals in law is related to reasoning with examples in other domains like mathematics and to adapting exemplars to solve problems. In her doctoral research, Rissland created a model of mathematical knowledge and understanding that gave a prominent role to concrete and prototype examples and their interconnections—so-called *Examples-space*—in the network of mathematical knowledge (Rissland, 1977, 1978). In the early 1980s, she studied how one generates examples that satisfy certain desiderata for use, for instance, in teaching and doing mathematics. Constrained Example Generation (CEG) worked by retrieving a good example (e.g., one maximizing the number of constraints satisfied) and then modifying it to try to satisfy additional constraints (Rissland, 1980; Rissland & Soloway, 1980). CEG's 'retrieval-plus-modification' architecture made it an early forerunner of adaptive CBR systems.

Initially concerned with how legal experts create hypos, Rissland and Ashley (then her student) embarked on a research program that led to the landmark HYPO system (Rissland *et al*, 1984; Rissland & Ashley, 1987; Ashley & Rissland, 1988). HYPO was brought to fruition in Ashley's doctoral research (Ashley, 1988, 1990). HYPO was the first true CBR system in the legal domain. While there had been earlier legal reasoning systems (e.g., Gardner, 1987) that used examples, HYPO was the first to use cases, that is, actual precedents. HYPO's initial focus was on the creation of legal hypotheticals—hence its name—by generating a hypothetical case based on an existing case and modifying it according to a particular way of looking at the problem, for instance to make it more extreme (Rissland & Ashley, 1986). Modifications were guided by knowledge captured in *dimensions*, one of the key innovations of HYPO, an idea first introduced in a 1984 AAAI paper (Rissland *et al*., 1984; Ashley, 1988, 1990; Rissland & Ashley, 2002).

3 HYPO and its progeny

HYPO, a paradigm interpretive CBR system, performs the entire sequence of steps of CBR from analysis to argumentation: it analyzes a new case; retrieves relevant cases from its case base; sorts them according to how on-point they are (in a 'claim lattice'); selects best cases for each side of the issue; generates '3-ply' point–counterpoint–rebuttal style arguments; and explores strengths and weaknesses of each side's arguments using hypotheticals (Ashley, 1988, 1989, 1990, 1991). Among HYPO's novel contributions are its definitions of relevance and on-pointness and its methods to compare and contrast cases in terms of dimensions. Dimensions capture the knowledge that certain frequently occurring fact patterns enable one to approach a legal dispute in a particular way and that certain changes in these facts tend to strengthen or weaken a side's claim. For instance, in trade secrets law (HYPO's domain), one way to approach a misappropriation claim is to focus on the number of disclosures the plaintiff made of its putative trade secret. *Ceteris paribus*, the more disclosures there are, the weaker the argument for a misappropriation, or equivalently, the stronger the argument that the defendant is blameless. The plaintiff's position is strongest where it made no disclosures, far weaker where there were many disclosures. This is the basis of HYPO's 'Secrets-Voluntarily-Disclosed' dimension (Ashley, 1988, 1990).

HYPO has had many progeny, including CABARET, BankXX, SPIRE, CATO, SMILE, and IBP. Central to each of these is a HYPO-style analysis. CATO harnessed HYPO's model of legal argument to teach law students how to make arguments with cases (Aleven, 1997, 2003; Aleven & Ashley, 1997). For instance, if CATO noticed that a student has not distinguished an opponent's case in the most effective way possible or has cited a case that can be trumped by an opponent's

https://doi.org/10.1017/S0269888906000701 Published online by Cambridge University Press

more on-point case, CATO generated the opponent's arguments to explain the student's mistakes. In the CATO system, Ashley and his student Aleven addressed the fact that dimensions are often related to each other and to higher-level legal reasons. They introduced *factors*, a construct closely related to dimensions, and a Factor Hierarchy to represent the underlying legal reasons (Aleven, 1997; Ashley & Aleven, 1997; Ashley, 2002). (For a description of the differences between dimensions and factors, see Rissland & Ashley (2002).) Since CATO provided legal reasons why shared factors favor a similar result or why distinctions are or are not legally significant, it could abstractly characterize cases and interpret them differently depending on the side for which it argues (Ashley & Aleven, 1997). In a controlled experiment, Aleven and Ashley demonstrated the effectiveness of CATO's case-based teaching strategies; CATO taught basic skills of case-based legal argument as effectively as an accomplished human instructor. Aleven also defined new ways for selecting the most relevant cases in terms of CATO's and HYPO's argumentation criteria and demonstrated their value in predicting case outcomes (Aleven, 2003). Ashley's student Stefanie Brüninghaus developed an alternative prediction approach, 'issue-based prediction' (IBP), that breaks problems into issues, generates hypotheses about which side should win each issue, tests the hypotheses by attempting to explain away counterexamples, and combines the issue-based predictions into an overall prediction and explanation (Brüninghaus & Ashley, 2003).

4 CBR and legal information retrieval, gathering, and extraction

Several systems have addressed the problem of finding relevant information in cases and the corpus of highly interconnected legal knowledge. SPIRE, developed by Rissland and her student Jody Daniels, used information gleaned from a claim lattice to drive a standard full text retrieval engine to generate queries and then retrieve documents (case opinions) and highlight passages within them (Rissland & Daniels, 1996; Daniels & Rissland, 1997). Brüninghaus and Ashley developed SMILE, a system that employed learning techniques to extract information about factors from textually described cases using 'Propositional Patterns' (ProPs), a representation that captures information about actors, actions, and objects (Brüninghaus & Ashley, 2001, 2005). The combined SMILE and IBP programs analyzed textually described problems, predicted their outcomes, and explained the predictions. They demonstrated the utility of the ProP representation.

The goal of the BankXX system was to harvest balanced collections of information (e.g., on-point cases both for and against one's side, legal theories) to use in argument (Rissland *et al.*, 1996). BankXX discovered information by intelligently crawling through a highly interconnected network of legal information using the mechanism of best-first heuristic search. The evaluation functions used by BankXX captured heuristic assessments of the benefits of including particular kinds of information. For example, if BankXX had already collected a surfeit of on-point cases for one side, collecting more of them—particularly at the possible exclusion of other types of information—was not deemed to be beneficial for preparing a good argument.

5 Reasoning with legal cases and rules

CABARET, a hybrid system, integrated HYPO-style CBR and classic rule-based reasoning (RBR) (Rissland & Skalak, 1991). A pioneer in multi-modal reasoning, CABARET used a dynamic agenda-based control architecture to interleave CBR and RBR so that these two reasoning modes could complement and supplement each other. CABARET produced memos outlining statutory arguments, including supporting and contrary cases, in a statutory area of United States tax law concerning home office deductions. CABARET's agenda was controlled by heuristic rules that embodied a theory of statutory interpretation, that is, a theory of how to argue about the meaning of legal rules and rule predicates, which often pose troublesome problems such as unstated exceptions and incurable open-texture. For instance, if all but one of a rule's prerequisites were satisfied (a rule-based 'near miss'), CABARET could use cases to argue that the predicate had

https://doi.org/10.1017/S0269888906000701 Published online by Cambridge University Press

actually been satisfied, or alternatively, that it was not a necessary condition. CABARET was a landmark system for CBR in that it was the first hybrid system to give equal status to both case-based reasoning and rule-based reasoning, and to allow them to interleave dynamically. Previous hybrids had used cases primarily in a supporting manner, for instance, when RBR failed or as a sanity check on its results (e.g., Gardner, 1987). As a project in AI and law, CABARET was important because it operationalized a three-tiered theory of interpretive statutory argument using 'argument strategies, moves, and primitives' (Skalak & Rissland, 1992). This theory can be viewed as part of an 'argumentation pyramid' whose base contains case facts, intermediate levels, dimensions and factors, statutory rules, court-made theories and tests, culminating in legal policies and principles. Computational models of case-based legal reasoning implement algorithms for navigating among the levels and bringing them to bear on analyzing and explaining problems (Ashley & Rissland, 2003).

The GREBE system pioneered the adaptation and reuse of case justifications for argument creation (Branting 1991, 2003). GREBE's model of argument structure includes both rule applications and *exemplar-based explanations* (EBEs) that relate portions of case facts to the conclusions they justify. This model of argument structure was the basis of a reduction-graph model of legal precedents later elaborated in Branting (1994, 2000). GREBE used a detailed relational representation of case facts and structure matching for similarity assessment. Like CABARET, GREBE is a hybrid CBR/RBR program that reasons with both rules and cases. For instance, it creates case analogies when the rules run out or otherwise fail to show that a legal term has been satisfied. GREBE used a heuristic measure of argument strength to rank the arguments for and against a given conclusion. Branting also tackled the issue of how to evaluate a CBR system that operates in a domain having nuanced answers. He used a domain expert to grade GREBE's analyses and those of law students in a blind test. The expert's evaluation of GREBE's analyses compared favorably with those of the law students.

CBR research in law has created other hybrids. The IKBALS III program combined CBR with induced rules (Zeleznikow *et al*., 1994); it augmented induced rule-based explanations with case examples illustrating alternative conclusions. Later work on the Split-Up project combined CBR with structured connectionist networks to predict judicial allocations of marital property in divorce cases (Zeleznikow *et al*., 1996).

CBR methodologies have also been extended to other legal tasks and to related domains. The CHIRON project explored the use of cases to guide tax planning. Precedents helped evaluate prototypical transaction plans and suggest possible plan adaptations (Sanders, 2001). The Truth-Teller and SIROCCO systems worked in the domain of ethics, a law-like normative system employing case-based techniques such as line drawing with prototypical examples. Truth-Teller took arbitrarily selected pairs of truth-telling dilemmas and explained how they were similar or different in a manner comparable to undergraduates' explanations (Ashley & McLaren, 1995). The SIROCCO program accepted engineering ethics problems and retrieved relevant past cases and ethical principles to help the user evaluate the new case (McLaren, 2003). By turning case knowledge on and off, SIROCCO empirically demonstrated the contribution past case decisions make in fleshing out the meaning of general normative principles.

Lastly, CBR research in law has impacted others in the AI and law community who focus not primarily on reasoning with cases but on natural argument. Responding to the critique that case-based interpretation must involve reasoning with the purposes and policies of law (Hafner & Berman, 2002), Bench-Capon and Sartor have investigated reasoning about values and goals. They have studied how cases give rise to preference rules among competing values; new cases are decided so as to optimize consistency with value preferences (Bench-Capon & Sartor, 2003).

6 Conclusion

Research on case-based approaches to law continues to make a vital contribution to CBR. Not only has it contributed to the birth of CBR and, in particular, interpretive CBR, but it has also spurred

https://doi.org/10.1017/S0269888906000701 Published online by Cambridge University Press

development of hybrid systems that combine case-based and other reasoning modes. It currently drives considerable work in the field of textual CBR. Its focus on comparing and contrasting cases in terms of the reasons why relevant similarities and differences are important has also been 'rediscovered' in the CBR field's recent focus on explanation-guided retrieval (Nugent *et al*., 2005).

Acknowledgements

Edwina Rissland's contribution to this material is based in part upon work supported as independent research while serving at the National Science Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

Aleven, V, 1997, Teaching case-based argumentation through a model and examples, PhD dissertation, University of Pittsburgh Graduate Program in Intelligent Systems, unnumbered technical report, Learning Research and Development Center.

Aleven, V, 2003, Using background knowledge in case-based legal reasoning: a computational model and an intelligent learning environment. *Artificial Intelligence* **150**(1–2), 183–237.

Aleven, V and Ashley, KD, 1997, Teaching case-based argumentation through a model and examples, empirical evaluation of an intelligent learning environment. In Du Boulay, B and Mizoguchi, R (eds) *Proceedings of the Eighth World Conference of the Artificial Intelligence in Education Society*, Kobe, Japan. Amsterdam: IOS Press, pp. 87–94.

Ashley, KD, 1988, Modelling legal argument: reasoning with cases and hypotheticals, PhD dissertation, COINS Technical Report No. 88:01, Department of Computer and Information Science, University of Massachusetts, Amherst.

Ashley, KD, 1989, Defining salience in case-based arguments. In Sridharan, NS (ed.) *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*. San Mateo: Morgan Kaufmann, pp. 537–542.

Ashley, KD, 1990, *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*. Cambridge, MA: MIT Press.

Ashley, KD, 1991, Reasoning with cases and hypotheticals in HYPO. *International Journal of Man-Machine Studies* **34**(6), 753–796.

Ashley, KD, 2002, An AI model of case-based legal argument from a jurisprudential viewpoint. *Artificial Intelligence and Law* **10**(1–3), 163–218.

Ashley, KD and Aleven, V, 1997, Reasoning symbolically about partially matched cases. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, Nagoya, Japan*. San Francisco: Morgan Kaufmann, pp. 335–341.

Ashley, KD and McLaren, B, 1995, Reasoning with reasons in case-based comparisons. In Veloso, M and Aamodt, A (eds) *Case-Based Reasoning Research and Development, Proceedings of the First International Conference on Case-Based Reasoning (Lecture Notes in Artificial Intelligence, 1010)*. Berlin: Springer, pp. 133–144.

Ashley, KD and Rissland, EL, 1988, A case-based approach to modeling legal expertise. *IEEE Expert* **3**(3), 70–77.

Ashley, KD and Rissland, EL, 2003, Law, learning and representation. *Artificial Intelligence* **150**(1–2), 17–58.

Bench-Capon, T and Sartor, G, 2003, A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence* **150**(1–2), 97–143.

Branting, LK, 1991, Building explanations from rules and structured cases. *International Journal of Man-Machine Studies* **34**(6), 797–837.

Branting, LK, 1994, A computational model of ratio decidendi. *Artificial Intelligence and Law* **2**(1), 1–31.

Branting, LK, 2000, *Reasoning with Rules and Precedents: A Computational Model of Legal Analysis*. Dordrecht: Kluwer Academic Publishers.

Branting, LK, 2003, A reduction-graph model of precedent in legal analysis. *Artificial Intelligence* **150**(1–2), 59–95.

Brüninghaus, S and Ashley, KD, 2001, The role of information extraction for textual CBR. In Aha, DW and Watson, I (eds) *Case-Based Reasoning Research and Development, Proceedings of the Fourth International Conference on Case-Based Reasoning (ICCBR-01) (Lecture Notes in Artificial Intelligence, 2080)*. Heidelberg: Springer, pp. 74–89.

https://doi.org/10.1017/S0269888906000701 Published online by Cambridge University Press

Brüninghaus, S and Ashley, KD, 2003, Combining case-based and model-based reasoning for predicting the outcome of legal cases. In Ashley, K and Bridge, D (eds) *Case-Based Reasoning Research and Development, Proceedings of the Fifth International Conference on Case-Based Reasoning (ICCBR-03) (Lecture Notes in Artificial Intelligence, 2689)*. Berlin: Springer, pp. 65–79.

Brüninghaus, S and Ashley, KD, 2005, Reasoning with textual cases. In Munoz-Avila, H and Ricci, F (eds) *Case-Based Reasoning Research and Development, Proceedings of the Sixth International Conference on Case-Based Reasoning (ICCBR-05) (Lecture Notes in Artificial Intelligence, 3620)*. Berlin: Springer, pp. 137–151.

Daniels, JJ and Rissland, EL, 1997, What you saw is what you want: using cases to seed information retrieval. In *Case-Based Reasoning Research and Development, Proceedings Second International Conference on Case-Based Reasoning (ICCBR-97) (Lecture Notes in Artificial Intelligence, 1266)*. Berlin: Springer, pp. 325–337.

Gardner, A vd L, 1987, *An Artificial Intelligence Approach to Legal Reasoning*. Cambridge, MA: MIT Press.

Hafner, CD and Berman, DH, 2002, The role of context in case-based legal reasoning: teleological, temporal, and procedural. *Artificial Intelligence and Law* **10**(1–3), 19–64.

McLaren, BM, 2003, Extensionally defining principles and cases in ethics: an AI model. *Artificial Intelligence* **150**(1–2), 145–181.

Nugent, C, Cunningham, P and Dónal, D, 2005, The best way to instill confidence is by being right. In *Case-Based Reasoning Research and Development, Proceedings of the Sixth International Conference on Case-Based Reasoning (ICCBR-05) (Lecture Notes in Artificial Intelligence, 3620)*. Berlin: Springer, pp. 368–381.

Rissland, E (Michener), 1977, Epistemology, representation, understanding and interactive exploration of mathematical theories, PhD dissertation, Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA.

Rissland, E (Michener), 1978, Understanding understanding mathematics. *Cognitive Science* **2**(4), 361–383.

Rissland, EL, 1980, Example generation. In *Proceedings Third National Conference of the Canadian Society for Computational Studies of Intelligence*, Victoria, BC, pp. 280–288.

Rissland, EL, 1983, Examples in legal reasoning: legal hypotheticals. In Bundy, A (ed.) *Proceedings Eighth International Joint Conference on Artificial Intelligence*. San Mateo, CA: Morgan Kaufmann, pp. 90–93.

Rissland, EL, 1989, Dimension-based analysis of hypotheticals from supreme court oral argument. In *Proceedings Second International Conference on AI and Law (ICAIL-89)*, Vancouver, BC. Association For Computing Machinery, pp. 111–120.

Rissland, EL, 1990, Artificial intelligence and law: stepping stones to a model of legal reasoning. *Yale Law Journal* 99(8), 1957–1982.

Rissland, EL and Ashley, KD, 1986, Hypotheticals as heuristic device. In *Proceedings Fifth International Conference on Artificial Intelligence*, San Mateo, CA: Morgan Kaufmann, pp. 289–297.

Rissland, EL and Ashley, KD, 1987, A case-based system for trade secrets law. In *Proceedings First International Conference on AI and Law*, Northeastern University, Boston, MA. ACM Press, pp. 60–66.

Rissland, EL and Ashley, KD, 2002, A note on dimensions and factors. *Artificial Intelligence and Law* **10**, 65–77.

Rissland, EL, Ashley, KD and Loui, RP, 2003, AI and Law: a fruitful synergy. *Artificial Intelligence* **150**(1–2), 1–15.

Rissland, EL and Daniels, JJ, 1996, The synergistic application of CBR to IR. *Artificial Intelligence Review: Special Issue on the use of AI in Information Retrieval* **10**(5–6), 441–475.

Rissland, EL and Skalak, DB, 1991, CABARET: statutory interpretation in a hybrid architecture. *International Journal of Man-Machine Studies* **34**(6), 839–887.

Rissland, EL, Skalak, DB and Friedman, MT, 1996, BankXX: supporting legal arguments through heuristic retrieval. *Artificial Intelligence and Law* **4**(1), 1–71.

Rissland, EL and Soloway, EM, 1980, Overview of an example generation system. In *Proceedings First National Conference on Artificial Intelligence*. AAAI Press, pp. 256–258.

Rissland, EL, Valcarce, EM and Ashley, KD, 1984, Explaining and arguing with examples. In *Proceedings Fourth National on Artificial Intelligence*. AAAI Press, pp. 288–294.

Sanders, KE, 2001, CHIRON: planning in an open-texture domain. *Artificial Intelligence and Law* **9**(4), 225–269.

Skalak, DB and Rissland, EL, 1992, Arguments and cases: an inevitable intertwining. *Artificial Intelligence and Law* **1**(1), 3–44.

Zeleznikow, J, Stranieri, A and Gawler, M, 1996, Split-Up: a legal expert system which determines property division upon divorce. *Artificial Intelligence and Law* **3**, 267–275.

Zeleznikow, J, Vossos, G and Hunter, D, 1994, The IKBALS project: multimodal reasoning in legal knowledge-based systems. *Artificial Intelligence and Law* **2**(3), 169–203.

https://doi.org/10.1017/S0269888906000701 Published online by Cambridge University Press