

Evaluating CBR Explanation Capabilities: Survey and Next Steps

Lawrence Gates and David Leake

Luddy School of Informatics, Computing, and Engineering, Indiana University
Bloomington IN 47408, USA
gatesla@indiana.edu, leake@indiana.edu

Abstract. The promise of case-based reasoning (CBR) methods for explaining intelligent systems has prompted much interest, including interest in explaining the CBR process itself. However, there is an impediment to understanding and advancing explanations of the CBR process: There has been little assessment of how well they perform. This paper develops strategies for filling this gap. We first highlight selected work on evaluating systems in explainable artificial intelligence, and then present a selective survey of work on explainable CBR systems to identify those that have been assessed with human subjects studies. From this we develop a proposed set of dimensions for characterizing explanation components of CBR systems. These dimensions provide a framework that can be used to guide evaluation, as well as a vehicle for identifying new research questions. We close with proposed next steps for research and for community initiatives.

Keywords: Explainable Case-based Reasoning, Explanation, Evaluation, Human Subjects, Survey

1 Introduction

Explanation of intelligent systems is a major area of AI research. Progress has been driven both by AI research initiatives, most notably DARPA’s Explainable Artificial Intelligence (XAI) Program [17], and by practical needs arising from the European Union directive on the “right to explanation” for conclusions from decision-making systems [16]. The CBR community has long seen interpretability as a key advantage of CBR (e.g., [25, 27]). There is substantial evidence for human use of case-based reasoning [26], and experience with CBR applications suggests that people find prior examples a natural vehicle for explaining decisions. Numerous approaches have been proposed for using CBR to generate explanations for black box systems (e.g., [20, 21, 36]), including “twinning” CBR and black box systems and using cases to explain the black box system results (see Keane and Kenny [20] for a survey). There are also strong arguments for focusing directly on the construction of interpretable systems rather than attempting to augment noninterpretable ones with explanatory capabilities [44]—another opportunity for CBR. Workshops on explainable CBR (XCBR) [2, 6]

have highlighted opportunities both for applying CBR to explain other systems and for explaining the behavior of CBR systems.

This paper focuses on needs for explaining CBR system behavior. Researchers have observed that simply presenting the most similar case in a case base may not be sufficient for optimal explanation (e.g., [11, 31, 32, 43]). A challenge for CBR is assessing the explanations that are generated. The obvious approach—and the only definitive method—for assessing the usefulness of explanations to end users is human subjects studies. However, such studies require substantial resources and few have been done for CBR explanation systems. Keane and Kenny refer to this gap as “[t]he embarrassment of user testing.” In their extensive survey of twinned CBR-ANN systems, they “found less than a handful (i.e., < 5) that performed any adequate user testing of the proposal that cases improved the interpretability of models” [20]. Belén Díaz Agudo’s ICCBR-19 keynote talk on explanation listed human subjects studies of human-oriented explanations as one of four key explanation challenges for the CBR community.¹

This paper aims to solidify understanding of the state of the art in assessing explanation of CBR systems. First, it presents general background on the nature of explanation and human subjects evaluation in XAI. Second, it presents a survey of human-subjects evaluations of explanation capabilities in CBR, to provide a reference for the methods and capabilities that have been assessed. Third, it categorizes this work, to identify the space to be explored and current gaps. Finally, based on the categorization and observations of prior research, it proposes five steps—both short and longer term—for advancing the evaluation of explanation of CBR systems and their explanation capabilities.

2 Aims of Explanation

Numerous studies have examined what makes a good explanation and how to explain in XCBR. Kenny et al. [21] propose a “tricorn” user model for XAI, with explanations relating to a user’s domain model, model of the system, and model of how particular explanation strategies (e.g., graphical presentation of information) relate evidence to a conclusion. Schoenborn et al. [45] survey explainable CBR and Sormo, Cassens and Aamodt [47] survey explanation in CBR with an emphasis on criteria for good explanations. In general, explanations can serve a rich range of needs, depending on the user’s knowledge state, goals, and tasks [24]. Here we focus on two fundamental goals. One is *transparency*, understanding how the system produced its answer. This may be useful both to foster trust (e.g., [14]) and to diagnose problems to guide system repair or refinement of system knowledge or processes. The other is *justification*, understanding why the system’s answer is a good answer [47].

¹ <https://icbr2019.com/wp-content/uploads/2019/09/Keynote-BelenDiazAgudo.pdf>

3 Assessing Explanation Systems

Given the goal of serving a human user, the gold standard for system assessments must be the results of human subjects studies. Human subjects studies can be designed directly to assess and compare explanation quality of alternative systems (e.g., [10, 30, 37, 38]). However, such assessments are challenging to design and may be time-consuming and expensive, especially given that many aspects of system processing (e.g., situation assessment, similarity assessment, or adaptation) might need to be explained. As a result, few XCBR efforts use human subject evaluations and the space of possible factors affecting explanation effectiveness is sparsely covered. Some studies rely on informal evaluations, seeking reactions of a small set of subjects, and many system designs simply rely on the system-builder’s intuitions. A goal of this paper is to provide a foundation for choosing aspects of a system to evaluate and to help identify the aspects that are most and least understood, to illuminate areas for future study.

4 Categorizing Explanations of CBR Systems

This section presents a categorization of research on explaining CBR systems. We begin by introducing the dimensions used and then place research according to those dimensions. The dimensions primarily consider the aspects of the CBR process being explained and the knowledge (derived from the CBR knowledge containers) that is the focus of those explanations. Of particular interest is identifying systems that have been evaluated with human subjects studies and opportunities for future human subjects studies. Schoenborn et al. [45] present a complementary survey, categorizing XCBR systems according to four types of aspects: definitional (explanation goal and kinds such as why-explanations or how-explanations), model-based or model-agnostic explanation, and (overlapping with this survey) the medium of presentation. We note that XCBR is not limited to explanation of CBR systems, including, for example, much interesting work on CBR to support explanation in non-CBR systems (such as [20, 41]). However, that work is outside the scope of this paper.

Explanations of the conclusions of a CBR system may focus on the solution provided by the system or the results of any of the intermediate steps in the CBR cycle—retrieve, reuse, revise, or retain [1]—with an account of how those results contribute to the solution. For example, CBR system behavior could be explained by the retrieved case, or a retrieved case plus how it was adapted. In principle, the retrieval could be explained both by similarity criteria and by how previous retention decisions affected the choice of cases to retain.

For any of the steps, the generation of a particular result depends on the knowledge brought to bear. Knowledge from any of the CBR knowledge containers [42] could be relevant to an explanation. In addition, given the well known tight coupling of knowledge containers, there may be a choice in determining which type of knowledge to highlight for a user (e.g., indexing or similarity knowledge accounting for retrieval of a case, if case relevance may not be obvious to the user, or case adaptation knowledge, if case relevance is intuitive but

the adapted solution may not be). Each of these choices provides a dimension for categorizing explanations of CBR systems.

4.1 Categorization Table

The table in Figure 1 categorizes the explanation approaches taken by sixteen projects. Research evaluated with human subjects studies is highlighted in bold. Rows are divided into four sections. The first section categorizes the research by *Explanation Medium*, the presentation of the explanation. The medium may be in graphical/visual or textual (explanations in internal system symbolic representations are considered to be in textual form). The second section categorizes the research by the *Target of Explanation*. Targets of the explanations may be the system’s non-adapted final result, the reason for the retrieval of the case used, the result when the retrieved case was adapted, or aspects of the ongoing system process.

The third section of rows categorizes the research by the primary *Knowledge Focus* of the explanation. A *holistic* focus provides an entire retrieved case to the user; the user determines how the case relates to explanatory needs. A *key features* focus provides information about the case features salient for retrieval. An *index knowledge* focus explains retrievals in terms of memory organization, the knowledge used to generate indices from cases, and/or knowledge underlying how those indices may be revised or transformed during the retrieval process. An *adaptation knowledge* focus explains the generation of a solution in terms of the application of knowledge such as adaptation rules, adaptation cases, or other knowledge from the adaptation knowledge container. A *similarity* focus explains retrieval in terms of the similarity measure and/or analysis of case features in terms of similarity criteria. In principle, a system explaining CBR system behavior could have multiple knowledge focuses. However, in practice, most research has focused on single focus types. The final two sections of rows categorize the research by *Explanation Recipient* and the *Process or Result Focus*. The *recipient* of the explanation may be a human or an AI system (as in systems that explain internally to guide processing). The *focus* of explanation is the final the *system process* used to obtain it (for transparency) or *system result* (for justification).

4.2 The Systems Categorized

Armengol, Ontanón and Plaza [4] present a method for explaining similarity to justify solutions in classification tasks. The approach was applied to provide justification to other agents in a multi-agent system. This system focuses on explaining the *similarity* of the *non-adapted final result in text* (symbolic) format, for *AI systems*.

Burke and Kass [7] present a system that explains the *retrieval result* for video clips with “bridging text” accounting for the relevant indices, focusing on *index knowledge*, for presentation to a *human user*.

Cunningham, Doyle and Loughrey [10] performed, to our knowledge, the first human subjects study on prior cases as explanations. Their study, using

		Ong et al., 97	Nugent, Cunningham & Doyle, 05	Muhammad et al., 15	McSherry, 01	McSherry, 04	Maximini, Frelman & Schaaf, 04	Massie, Crow & Wiratunga, 04	Lieber et al., 08	Leake, 91	Lamy et al., 19	Kolodner, 83	Kass, 90	Doyle, Cunningham & Walsh, 06	Cunningham, Doyle & Loughrey, 03	Burke & Kass, 96	Armengol, Ontanon & Plaza, 04
Explanation Medium	Visualization																
	Text	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Target for Explanation	Non-Adapted Final Result	✓		✓							✓			✓			
	Retrieval Result		✓									✓					
	Adaptation Result						✓						✓				
	Ongoing Process																✓
Knowledge Focus	Case Only (Holistic)			✓	✓			✓		✓							
	Key Features																✓
	Index Knowledge		✓									✓					
	Adaptation Knowledge							✓									
	Similarity	✓															
Explanation Recipient	Human User		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	AI System	✓								✓							
Process or Result Focus	System Process							✓	✓								
	System Result	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Fig. 1. A Categorization of Explanations of Case-Based Reasoning.

37 subjects, compared the convincingness of no explanation, presentation of the case on which a CBR system's conclusion was based, and rule-based explanation for predictions of blood alcohol levels. The system's target is the *non-adapted final result* in *text* format, with *holistic* knowledge focus, to a *human user*.

Doyle, Cunningham and Walsh [12] study explanation via a selected prior case (not necessarily the nearest case), justification text and a confidence level, in the domain of hospital admission and discharge decisions. The justification text highlighted features that support or undermine relevance of the explanation case. The study had two evaluation stages, a formative evaluation and a second stage of two studies, one with 106 domain experts judging confidence in the system and the appropriateness of the explanation case, and the second with 14 domain experts judging whether the nearest neighbor or explanation case better

explained the system decision. The system's target is the *non-adapted final result* in *text* format, with *holistic* knowledge focus, to a *human user*.

Kass [19] presents a case adaptation system that provides textual descriptions of adaptation paths that could be presented to explain adaptations to a user. The system's target is the *adaptation result* in *text* format, with knowledge focus of *adaptation knowledge*, for a *human user*.

Kolodner [22] presents explanation of retrieval by a question transformation trace leading to the *retrieval result* in a *text* format, with the knowledge focus *index knowledge* to a *human user*.

Lamy et al. [23] focus on visualization of differences and similarities between a query case and similar cases, in the context of a CBR system to diagnose breast cancer. Their system was evaluated by 11 users who selected a treatment based on the data presented by the system and indicated their confidence level in the choice. The target is the *non-adapted final result* in the form of a *visualization*, with a *holistic* knowledge focus, to a *human user*.

Leake [24] presents a set of criteria for evaluating CBR-generated explanations of anomalous events in stories for various goals (e.g., prediction of an event or prevention of an event). The system's target is the *non-adapted final result* in *text* format, with a knowledge focus of *holistic*, for an *AI system*.

Lieber et al. [29] present an approach to explaining adaptations, in the context of breast cancer treatment, by presenting adaptation paths with annotations of each step. The system's target is the *adaptation result* in *text* format, with knowledge focus on *adaptation knowledge*, to a *human user*.

Massie, Crow and Wiratunga [30] present a visualization tool aimed at presenting feature similarity between a new problem and retrieved case, in the domain of drug design. The work is assessed by asking two domain experts to assess similarity and express their confidence in the answer the solution provided. The target for explanation is the *retrieval result* in the form of a *visualization*, with a *holistic* knowledge focus and explanation to a *human user*.

Maximini, Freßman, and Schaaf [31] present a system including a range of explanation capabilities for an intellectual property domain. The system includes two explanation components. The first explains the set of cases retrieved in terms of attributes and supports user navigation through the candidates, as well as presenting the most decisive attributes to narrow the retrieval set. When no candidates are retrieved, it can also help the user understand which query features could not be matched, to aid query refinement. The other supports examination of the result, with color coding of features by similarity as well as a graphical visualization and textual explanation of feature connections. The system's target is the *non-adapted final result* in *text* format, with a knowledge focus of the *similarity* for a *human user*.

McSherry [32] presents an explanation method that provides the features of the nearest cases that support or oppose a prediction. This system focuses on explaining the *non-adapted final result* in form of case *text*, with knowledge focus on *key features* to a *human user*.

McSherry [33] presents an explanation method for conversational CBR for diagnosis of computer faults. The system provides explanations of the relevance of questions. This system focuses on explaining the *ongoing process* in the form of *text*, with knowledge focus on *index generation*, to a *human user*.

Muhammad et al. [34] present an approach to explaining recommendations based on positive and negative features compared to user preferences, in a hotel recommendation domain. The recommender personalizes explanations by highlighting the features the given user seeks. This system focuses on explaining the *non-adapted final result* in form of *text*, with a knowledge focus on *key features* to a *human user*.

Nugent, Cunningham and Doyle [37] extend their original study [10] for an explanation component providing the retrieved case, text explaining the difference between the query and explanation case, and a confidence value. If confidence is below a threshold, an additional counter case is provided, with a description of how feature values in both cases impact classification. The study used 12 human subjects. The system target is the *non-adapted final result* in text format, with a *holistic* knowledge focus and explanation for a *human user*.

Ong et al. [39] present a system that explains similarity by presenting similarity scores for matching features, in a cancer diagnosis domain. This system focuses on explaining the *non-adapted final result* in form of *text*, with knowledge focus on the case as a whole, so *holistic*, to a *human user*.

5 Next Steps

Examination of research in the survey suggests five steps for advancing evaluation of explanation of CBR systems. The first, building on the dimensions of the categorization, is to fill gaps in coverage of evaluation dimensions in the categorizations. The second and third involve community effort to make evaluations more practical to perform and results more comparable. This requires both advancing automated evaluations, to enable gathering initial evaluation data more rapidly, and establishing standard datasets and standardized evaluation designs, in order to make studies more useful for understanding the relationships between different explanation focuses and how they contribute to the benefit of explanation. The fourth and fifth are to move beyond the current space of explanations, focused on local explanation of a single solution, to build explanations to help users better understand how systems will perform in the future and to explain system failures. Understanding the full scope of explanation of CBR systems will require experiments that assess the full system, including effects of component interactions.

5.1 Using the Categorization Scheme to Suggest Evaluation Targets

The categorization table of the previous section shows the distribution of human subjects studies for explaining CBR. We observe that existing evaluations have primarily focused on case-oriented explanations—explaining by presenting

cases and highlighted features to the user. This has been important for assessing the fundamental CBR premise that cases are a natural mode of explanation. However, in domains unfamiliar to the user, explanatory support for assessing similarity may be crucial. This received some recognition [4, 30], but so far, to our knowledge, only one method has been tested with human subjects, and in a small-scale assessment (by two experts) [30]. Adaptation-focused explanation has also received little attention.

The survey also suggests the potential for studies of how different modes of information presentation affect efficacy of explanations. Systematic evaluation will require understanding:

- Alternative methods and tradeoffs for conveying to the user information about system knowledge in each of the CBR knowledge containers [42]: the vocabulary, similarity measure, case base, and adaptation knowledge.
- How best to explain the processes underlying each step of the CBR cycle: retrieve, reuse, revise, and retain [1].
- The influence of information needs [24, 47]: how to select parts of that process to highlight to particular users.

The ideal result would be a toolkit of best practices for explaining classes of CBR systems based on the tasks they perform. For example, explaining a “compare and contrast” task, as in American legal reasoning [5], will necessarily differ from explaining anomalies [24]. Expert systems research on *generic tasks* [8] suggests the potential to identify a comparatively limited set of task classes. This in turn would increase the practicality and usefulness of human subjects studies by encouraging studies with wide task applicability. A community effort to develop such a collection could leverage both XCBR research and the application of XCBR in domains for which explanation is important.

5.2 Advancing Automatic Evaluation

There is a longstanding tradition of testing conversational CBR systems with simulated users (e.g., [3]). Research on case-based explanation generation has examined criteria for evaluating the goodness of explanations generated [18, 24, 46]. This prior work raises an interesting question: Would it be possible to develop automated evaluation schemes to test aspects of the explanations generated by CBR systems? Even if such methods did not have complete fidelity to human subjects studies, automatic testing could be a valuable intermediate step.

Work on case-based explanation for story understanding systems has developed criteria for evaluating explanations according to their ability to resolve anomalies in stories, address system questions, and satisfy other explanation goals [24, 40]. Applying such methods to a broader class of tasks would require analysis of the information needs associated with those tasks and how they can be satisfied. Providing a general solution for this would be an enormous task. However, it could be more feasible in testbed microdomains—were the community to embrace developing such resources.

5.3 Developing Standardized Evaluation Resources

The ability to compare and replicate human subjects studies is impeded by the lack of agreed-upon testbeds for such evaluations. Different research groups use different data sets, making it more difficult to assess the relative contributions of explanations of different parts of the CBR process and how those components interact. From the five human subject studies reported in this paper’s categorization, two (done by the same research group) used the blood-alcohol domain [10, 37], one used the tablet formulation domain [30], one used Bronchiolitis medical data [12], and one used breast cancer data [23]. Of the 4 unique domains, only the breast cancer data [13] is publicly available. None of these data sets involve CBR systems including case adaptation.

The development of the needed data sets—and agreement to use them—must be a community endeavor. The CBR community has a history of developing shared resources and advancing the state of the art with competitions, as in the Computer Cooking Competition series (e.g., [35]). This series included various forms of human evaluations, both by human judges and by the conference attendees as a whole. An explanation challenge, based on standard tasks and data sets and informally evaluated by humans as part of the competition, could “jump start” the development of standard resources. In addition, the design of the competition could help establish evaluation protocols to form the basis for more formal human subjects studies.

5.4 Explaining the prior cases themselves

Explanations based on presenting prior cases assume the quality of the retrieved case can be trusted—that no justification is needed for why the case should be believed. However, especially if cases have been automatically generated by adaptation processes, understanding and accepting a case as an explanation depends on understanding the presented case and its relevance. That may depend on understanding how it was derived—its provenance [28]. The use of provenance for explaining CBR is an interesting area for future research.

5.5 Broadening CBR Explanations to Explaining Failures

Existing studies of explaining CBR focus primarily on accounting for why the system reasoned as it did. Users may use that information to determine the applicability of a proposed solution. This is sufficient to determine trust in the current solution. However, when users determine that the system went astray, the explanation components generally have no capability to explain why the CBR process failed. In the absence of such an explanation it is harder to determine trust in the system’s future processing.

Some CBR research has applied metareasoning to explain and repair reasoning failures within the CBR process for internal system learning (e.g., [9, 15]). This points to another topic for CBR system explanation: accounting for expectation failures in how the system operates, to help users refine their predictive model of the CBR system overall.

6 Conclusion

This paper has presented an initial categorization of research on explaining CBR systems, highlighting dimensions for which human subjects evaluations have been done—and for which they are still needed. Based on this categorization it has proposed XCBR research opportunities including increased attention to explaining similarity and adaptation. It has also advocated new initiatives on automatic evaluation, building community evaluation resources, explaining case provenance, and explaining system failures.

7 Acknowledgments

We acknowledge support from the Department of the Navy, Office of Naval Research (Award N00014-19-1-2655), and the US Department of Defense (Contract W52P1J2093009).

References

1. Aamodt, A., Plaza, E.: Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications* **7**(1), 39–52 (1994)
2. Aha, D., Agudo, B.D., Garcia, J.R. (eds.): Proceedings of XCBR-2018: First workshop on case-based reasoning for the explanation of intelligent systems. ICCBR (2018), <http://iccb18.com/wp-content/uploads/ICCBR-2018-V3.pdf>
3. Aha, D., Breslow, L.: Refining conversational case libraries. In: Proceedings of the Second International Conference on Case-Based Reasoning. pp. 267–278. Springer Verlag, Berlin (1997)
4. Armengol, E., Ontanón, S., Plaza, E.: Explaining similarity in CBR. In: ECCBR 2004 Workshop Proceedings. pp. 155–164 (2004)
5. Ashley, K., Rissland, E.: Compare and contrast, a test of expertise. In: Proceedings of the Sixth Annual National Conference on Artificial Intelligence. pp. 273–284. AAAI, Morgan Kaufmann, San Mateo, CA (1987)
6. B. Diaz Agudo, J.R.G., Watson, I. (eds.): Proceedings of XCBR-2019: Second workshop on case-based reasoning for the explanation of intelligent systems. ICCBR (2018)
7. Burke, R., Kass, A.: Retrieving stories for case-based teaching. In: Leake, D. (ed.) *Case-Based Reasoning: Experiences, Lessons, and Future Directions*, pp. 93–109. AAAI Press, Menlo Park, CA (1996)
8. Chandrasekaran, B.: Generic tasks in knowledge-based reasoning: High level building blocks for expert system design. *IEEE Expert* **1**(3), 23–30 (1986)
9. Cox, M., Ram, A.: Introspective multistrategy learning: On the construction of learning strategies. *Artificial Intelligence* **112**(1-2), 1–55 (1999)
10. Cunningham, P., Doyle, D., Loughrey, J.: An evaluation of the usefulness of case-based explanation. In: *Case-Based Reasoning Research and Development: Proceedings of the Fifth International Conference on Case-Based Reasoning, ICCBR-03*. pp. 122–130. Springer-Verlag, Berlin (2003)
11. Doyle, D., Cunningham, P., Bridge, D., Rahman, Y.: Explanation oriented retrieval. In: Funk, P., González Calero, P.A. (eds.) *Advances in Case-Based Reasoning*. pp. 157–168. Springer, Berlin, Heidelberg (2004)

12. Doyle, D., Cunningham, P., Walsh, P.: An evaluation of the usefulness of explanation in a case-based reasoning system for decision support in bronchiolitis treatment. *Computational Intelligence* **22**(3-4), 269–281 (2006)
13. Dua, D., Graff, C.: UCI machine learning repository (2017), <http://archive.ics.uci.edu/ml>
14. Floyd, M., Aha, D.: Special issue on case-based reasoning. *AI Communications* **30**(3-4), 281–294 (2017)
15. Fox, S., Leake, D.: Introspective reasoning for index refinement in case-based reasoning. *The Journal of Experimental and Theoretical Artificial Intelligence* **13**(1), 63–88 (2001)
16. Goodman, B., Flaxman, S.: European Union regulations on algorithmic decision-making and a “right to explanation”. *AI Magazine* **38**(3), 50–57 (2017)
17. Gunning, D., Aha, D.W.: DARPA’s explainable artificial intelligence program. *AI Magazine* **40**(2), 44–58 (2019)
18. Johs, A.J., Lutts, M., Weber, R.O.: Measuring explanation quality in XCBR. In: Minor, M. (ed.) *Workshop proceedings ICCBR 2018*. pp. 75–83 (2018)
19. Kass, A.: *Developing Creative Hypotheses by Adapting Explanations*. Ph.D. thesis, Yale University (1990), northwestern University Institute for the Learning Sciences, Technical Report 6
20. Keane, M.T., Kenny, E.M.: How case-based reasoning explains neural networks: A theoretical analysis of XAI using post-hoc explanation-by-example from a survey of ANN-CBR twin-systems. In: *Case-Based Reasoning Research and Development: ICCBR-19*. pp. 155–171. Springer, Berlin (2019)
21. Kenny, E.M., Ford, C., Quinn, M., Keane, M.T.: Explaining black-box classifiers using post-hoc explanations-by-example: The effect of explanations and error-rates in XAI user studies. *Artificial Intelligence* **294**, 103459 (2021)
22. Kolodner, J.L.: *Reconstructive memory: A computer model*. *Cognitive Science* **7**(4), 281–328 (1983)
23. Lamy, J.B., Sekar, B., Guezennec, G., Bouaud, J., Séroussi, B.: Explainable artificial intelligence for breast cancer: A visual case-based reasoning approach. *Artificial intelligence in medicine* **94**, 42–53 (2019)
24. Leake, D.: Goal-based explanation evaluation. *Cognitive Science* **15**(4), 509–545 (1991)
25. Leake, D.: CBR in context: The present and future. In: Leake, D. (ed.) *Case-Based Reasoning: Experiences, Lessons, and Future Directions*, pp. 3–30. AAAI Press, Menlo Park, CA (1996), <http://www.cs.indiana.edu/~leake/papers/a-96-01.html>
26. Leake, D.: Cognition as case-based reasoning. In: Bechtel, W., Graham, G. (eds.) *A Companion to Cognitive Science*, pp. 465–476. Blackwell, Oxford (1998)
27. Leake, D., McSherry, D.: Introduction to the special issue on explanation in case-based reasoning. *Artificial Intelligence Review* **24**(2), 103–108 (2005)
28. Leake, D., Whitehead, M.: Case provenance: The value of remembering case sources. In: *Case-Based Reasoning Research and Development: ICCBR-07*. pp. 194–208. Springer-Verlag, Berlin (2007)
29. Lieber, J., d’Aquin, M., Badra, A., Napoli, A.: Modeling adaptation of breast cancer treatment decision protocols in the KASIMIR project. *Applied Intelligence* **28**(3), 261–274 (2008)
30. Massie, S., Craw, S., Wiratunga, N.: A visualisation tool to explain case-base reasoning solutions for tablet formulation. In: *Proceedings of the 24th SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence*. Springer-Verlag, Berlin (2004)

31. Maximini, R., Freßmann, A., Schaaf, M.: Explanation service for complex CBR applications. In: *Advances in Case-Based Reasoning, ECCBR*. pp. 302–316. Springer, Berlin (2004)
32. McSherry, D.: Explaining the pros and cons of conclusions in CBR. In: *Advances in Case-Based Reasoning*. pp. 317–330. Springer, Berlin (2004)
33. McSherry, D.: Interactive case-based reasoning in sequential diagnosis. *Applied Intelligence* **14**(1), 65–76 (2001)
34. Muhammad, K., Lawlor, A., Rafter, R., Smyth, B.: Great explanations: Opinionated explanations for recommendations. In: *Case-Based Reasoning Research and Development*. pp. 244–258. Springer, Berlin (2015)
35. Najjar, N., Wilson, D.: Computer cooking contest workshop at the twenty-fifth international conference on case-based reasoning (ICCBR 2017). In: *ICCBR 2017 Workshops, Doctoral Consortium, and Competitions (2017)*, http://ceur-ws.org/Vol-2028/XXCCC17_preface.pdf
36. Nugent, C., Cunningham, P.: A case-based recommender for black-box systems. *Artificial Intelligence Review* **24**(2), 163–178 (2005)
37. Nugent, C., Cunningham, P., Doyle, D.: The best way to instil confidence is by being right. In: *International Conference on Case-Based Reasoning*. pp. 368–381. Springer (2005)
38. Nugent, C., Doyle, D., Cunningham, P.: Gaining insight through case-based explanation. *Journal of Intelligent Information Systems* **32**(3), 267–295 (Jun 2009)
39. Ong, L.S., Shepherd, B., Tong, L.C., Seow-Choen, F., Ho, Y.H., Tang, C.L., Ho, Y.S., Tan, K.: The colorectal cancer recurrence support (CARES) system. *Artificial Intelligence in Medicine* **11**(3), 175–188 (1997)
40. Ram, A.: AQUA: Asking questions and understanding answers. In: *Proceedings of the Sixth Annual National Conference on Artificial Intelligence*. pp. 312–316. Morgan Kaufmann, Seattle, WA (July 1987)
41. Recio-García, J.A., Díaz-Agudo, B., Pino-Castilla, V.: CBR-LIME: A case-based reasoning approach to provide specific local interpretable model-agnostic explanations. In: *Case-Based Reasoning Research and Development: ICCBR-20*. pp. 179–194. Springer, Online (2020)
42. Richter, M.: Introduction. In: Lenz, M., Bartsch-Spörl, B., Burkhard, H.D., Wess, S. (eds.) *CBR Technology: From Foundations to Applications*, chap. 1, pp. 1–15. Springer, Berlin (1998)
43. Roth-Berghofer, T.: Explanations and case-based reasoning: Foundational issues. In: *Advances in Case-Based Reasoning*. pp. 389–403. Springer, Berlin (2004)
44. Rudin, C.: Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* **1**, 206–215 (2019)
45. Schoenborn, J., Weber, R., Aha, D., Cassens, J., Althoff, K.: Explainable case-based reasoning: A survey. In: *Proceedings of the AAAI Explainable Agency in Artificial Intelligence Workshop (2020)*
46. Sizov, G., Öztürk, P., Bach, K.: Evaluation of explanations extracted from textual reports. In: *Proceedings of the Twenty-Ninth International Florida Artificial Intelligence Research Society Conference, FLAIRS-16*. pp. 425–429. AAAI Press (2016)
47. Sormo, F., Cassens, J., Aamodt, A.: Explanation in case-based reasoning—perspectives and goals. *Artificial Intelligence Review* **24**(2), 109–143 (2005)