

Advances in conversational case-based reasoning

DAVID W. AHA¹, DAVID MCSHERRY² and QIANG YANG³

¹*Navy Center for Applied Research in Artificial Intelligence, Naval Research Laboratory, Code 5515,
4555 Overlook Avenue SW, Washington DC 20375-5337, USA;
e-mail: aha@aic.nrl.navy.mil*

²*School of Computing and Information Engineering, University of Ulster, Coleraine BT52 1SA, Northern Ireland, UK;
e-mail: dm.mcsherry@ulster.ac.uk*

³*Department of Computer Science, The Hong Kong University of Science and Technology, Hong Kong;
e-mail: qyang@cs.ust.hk*

Abstract

A considerable amount of research in case-based reasoning (CBR) has recently focused on *conversational* CBR as a means of providing more effective support for interactive problem solving. We review progress made to date and identify challenges that remain to be addressed.

1 Introduction

In many case-based reasoning (CBR) applications, a complete description of the target problem is assumed to be available in advance. This is an unrealistic assumption in domains such as interactive fault diagnosis, where it is natural for users to provide only a brief initial description of the problem. Often the user can also provide the results of simple tests, but cannot be expected to provide a complete description of the problem, as this may require an unreasonable amount of time, or involve difficult or expensive tests that the user is unable or reluctant to perform.

In *conversational* CBR (CCBR), a query describing a target problem is incrementally elicited in an interactive dialogue with the user, often with the aim of minimizing the number of questions the user is asked before a solution is reached (Aha *et al.*, 2001; McSherry, 2002a). An initial query based on the user's brief description of the problem in free text may instead be incrementally *extended* in the problem solving dialogue. In each dialogue cycle, the user is shown any cases with an above-threshold similarity to the current query and invited to select from a list of questions ranked in order of usefulness (e.g., their ability to discriminate between competing cases). Alternatively, the user may be directly asked the most useful question, and if she is unable to answer that question, the dialogue moves on to the next most useful question as determined by the system's *question selection* strategy.

At any stage, the user can terminate the dialogue by selecting the solution of a displayed case as the solution to the target problem. Otherwise, the dialogue continues until some predefined *termination criteria* are satisfied. At this point, the solution for the most similar case is presented as the proposed solution to the target problem. Yang *et al.* (1997) describe a CCBR approach to configuration in which the user can adapt a proposed solution through a constraint satisfaction process. More typically, there is no need for adaptation in CCBR, although a minimum level of similarity may be required of the retrieved case on which a proposed solution is based.

CCBR was pioneered by Inference Corporation, who marketed a series of CCBR shells (e.g., CBR Express, *k-Commerce*) that targeted the help desk and related customer support market. Many papers describe success stories involving these products, often focusing on cost savings due

to decreased call volume and related measures (e.g., Acorn & Walden, 1992; Nguyen *et al.*, 1993). Buoyed by Inference's success, other companies began marketing CCBR-related products. For example, eGain purchased Inference and now markets the eGain Guided Help product, while CaseBank Technologies applies CCBR to the diagnosis of jet engines and other complex equipment, processes, and systems (Gupta, 1998).

Much of the research discussed in this commentary focuses on CCBR as an approach to interactive problem solving in domains such as helpdesk support, fault diagnosis, and product recommendation (e.g., Göker *et al.*, 1998; McSherry, 2001a; Ricci *et al.*, 2002). Other contributions have investigated the integration of CCBR with other problem-solving methodologies and its deployment in multi-agent systems. For example, Muñoz-Avila *et al.*'s (1999, 2001) HICAP system uses CCBR to assist users with decomposing individual tasks in a hierarchical generative planner. This complements the planner's ability to autonomously apply other task decomposition methods. Giampapa & Sycara (2001) demonstrate the ability of Aha *et al.*'s (2001) NaCoDAE, a generic CCBR shell, to proactively dialogue with other agents in a goal directed approach to information gathering. McSherry (2005a) examines the potential role of CCBR in a multi-agent approach to product recommendation in which a local recommendation failure triggers a referral of the user's query to other recommender agents.

In Section 2, we review recent developments in case representation, acquisition, and maintenance in CCBR. In Section 3, we turn our attention to the management of CCBR dialogues and related issues such as question selection, dialogue inferencing, and dialogue termination. Our conclusions are presented in Section 4.

2 Case representation, acquisition, and maintenance

In this section, we focus initially on aspects of case representation that are specific to CCBR and then go on to examine related developments in case acquisition and maintenance.

2.1 Case representation

Previous problems are stored with their solutions in the case library of a CCBR system and retrieved in response to an often incomplete query describing the target problem as a set of (question, answer) pairs. A case's problem description is typically represented in a similar way, although other approaches have been investigated. In *taxonomic* CCBR, for example, relevant domain features (e.g., test results) are arranged in feature taxonomies in which levels of abstraction are explicitly represented by subsumption links (Gupta, 2001). Cases need only reference features at a taxonomy's leaf nodes because, by definition, a parent feature must appear in all the cases in which any of its children appear. This taxonomic representation of cases is designed to address problems often associated with the use of different levels of abstraction to describe case features. It also plays an important role in question selection by focusing conversations on questions that lead downwards in the taxonomies.

In more recent work, Gu & Aamodt (2005) describe a function that maps concepts to questions, where concepts are represented as nodes in a semantic network. This enables relations among case features to be inferred more easily in a CCBR system. In a similar vein, Shimazu (1999) describes a mapping from scripts to the features used to describe cases.

It is common in CCBR case libraries for each solution (or class) to be represented by a single case. Such a case library is said to be *irreducible* in that the deletion of any case means that its solution is no longer represented (McSherry, 2001b). A *heterogeneous* case structure is also typical in domains such as fault diagnosis, where tests may be applicable only in certain types of faults (Aha *et al.*, 2001). McSherry (2002b) investigates the effects of retrieval based on incomplete queries on the *separability* of cases in an irreducible case library.

2.2 Case acquisition and maintenance

Contributions that aim to reduce the overheads traditionally associated with case acquisition and maintenance in CCBR include the use of formal approaches to knowledge modeling such as KADS to validate manually constructed case libraries (Trott & Leng, 1997) and approaches to automate case acquisition and maintenance tasks. Racine & Yang (1997) present automated techniques for detecting inconsistent and redundant cases, while Gupta & Aha (2004) describe a semi-automated approach to acquiring CCBR case indices from text documents. Aha & Breslow (1997) investigate an automated approach to revising CCBR case libraries in line with case authoring guidelines. Their case revision tool, Clire, transforms cases to a decision tree representation and then edits the tree before extracting the revised cases.

Zhang & Yang (1999) describe an algorithm for tuning feature weights based on user feedback on the relative ranking of cases. They report that their algorithm, inspired by the error back-propagation procedure, converges quickly enough to enable it to be trained for use in practical settings. Weber *et al.* (2004, 2005) describe a CCBR system called CBRFlow that assists workflow management by enabling users to capture and reuse cases representing exceptions to workflow rules in reference to process schemas. Exceptions that frequently recur can be used to guide case abstraction.

3 Problem-solving dialogues

We focus in this section on issues related to the management of problem-solving dialogues in CCBR. Examples of CCBR (and related) systems and tools that support some of the features mentioned in our discussion are given in Table 1. While focusing on incremental query elicitation, CCBR does not impose a specific interface modality for interactive problem solving. One alternative to question answering and visual presentation of solutions via a standard GUI is to wrap the CCBR tool within a speech processing component. For example, the Adaptive Place Advisor (Thompson *et al.*, 2004) is an in-car recommender system for travel destinations that converses with the driver through spoken dialogue.

3.1 Question selection

Aha & Breslow (1997) propose an approach to question selection/ranking in NaCoDAE which takes account of the *heterogeneity* of cases that is typical in domains such as fault diagnosis. Questions are ranked according to their frequency of occurrence in the cases that are most similar to the current query. In the context of knowledge-intensive CCBR, Gupta *et al.* (2002) report that their *taxonomic* approach to question selection significantly reduces the average number of questions required for case retrieval. More recently, Gu & Aamodt (2005) describe an approach that uses meta-level relations to constrain question rankings.

Carrick *et al.* (1999) propose an approach to question selection in CaseAdvisor that is designed to address the trade-off between the quality of the information obtainable by asking a question and the cost of obtaining the information (e.g., by consulting an external information source). Question selection in CBR Strategist (McSherry, 2001a) is inspired by the *hypothetico-deductive* reasoning strategies used by doctors, such as confirming a target diagnosis or eliminating a competing diagnosis. An important benefit is that the relevance of any question can be explained in terms of the purpose for which it was selected.

Several CCBR systems have used *information gain* as a basis for question selection (Göker *et al.*, 1998; Shimazu *et al.*, 2001; Yang & Wu, 2001; Thompson *et al.*, 2004). As in *demand-driven* approaches to inductive retrieval (McSherry, 1995; Smyth & Cunningham, 1995; Friedman *et al.*, 1996), questions are selected dynamically in an interactive problem-solving dialogue, and the user's answers guide the construction of a path in a *virtual* decision tree. McSherry (2001b) shows that a

Table 1 Example CCBR (and related) systems and tools

System/Tool	Description	Features
Adaptive Place Advisor (Thompson <i>et al.</i> , 2004)	CCBR Recommender System	● Spoken dialogue interface ● Question selection based on information gain ● Dialogue operators ● Query relaxation
CaseAdvisor (Carrick <i>et al.</i> , 1999)	CCBR Tool for Fault Diagnosis	● Free text entry of problem descriptions ● Cost sensitive question selection ● Automated information gathering
CBR Strategist (McSherry, 2001a)	CCBR Tool for Fault Diagnosis	● Goal driven dialogue ● Explanation of question relevance
ExpertClerk (Shimazu <i>et al.</i> , 2001)	CCBR Tool for Knowledge Space Mentoring	● Question selection based on information gain ● Script-based indexing of cases
NaCoDAE (Aha <i>et al.</i> , 2001)	Generic CCBR Shell	● Free text entry of problem descriptions ● Ranking of questions based on frequency of occurrence in competing cases ● Dialogue inferencing
Sermo (Bridge, 2002)	Conversational Recommender System	● Dialogue grammar ● Query relaxation
ShowMe (McSherry, 2004)	Conversational Recommender System	● Explanation of query failure ● Incremental query relaxation
TCRS (Gupta <i>et al.</i> , 2002)	Taxonomic CCBR Tool	● Taxonomic representation of cases and focusing of conversations ● Dialogue inferencing
Top Case (McSherry, 2005b)	CCBR Recommender System	● Goal driven dialogue ● Explanation of question relevance ● Solution justification ● Dominance criteria for dialogue termination

simplified measure of information gain can be applied to irreducible case libraries and tends to be more effective than some other measures of question usefulness in reducing the length of CCBR dialogues. However, several authors have questioned the suitability of information gain as a basis for question selection in CCBR recommender systems (e.g., Kohlmaier *et al.*, 2001; McSherry, 2003). One alternative to information gain is the *simVar* measure proposed by Kohlmaier *et al.* (2001). Instead, McSherry (2003, 2005a,b) proposes a *goal-driven* approach to question selection in which the relevance of questions the user is asked can be explained in terms of the system's strategy of confirming a target case as the recommended case.

3.2 Dialogue inferencing

Dialogue inferencing involves the use of knowledge-intensive processes to automatically extend a partially elicited query, for example by inferring the answers to certain questions from information already provided in the course of a CCBR dialogue. As well as helping to reduce the length of

problem-solving dialogues, this is important to avoid question choices that may reduce user confidence in the system. While some commercial tools support a rule-based approach to dialogue inferencing, maintaining the manually encoded rules is a challenging problem. Aha *et al.* (1998) describe a model-based approach that significantly improves conversational efficiency while requiring only the maintenance of a relatively concise domain model. Carrick *et al.* (1999) describe the use of hierarchical task network plans to support information gathering activities in CaseAdvisor. The meta-relations used by Gu & Aamodt (2005) to guide question ranking also support a dialogue inferencing process that enables queries to be updated automatically.

3.3 Dialogue termination

An effective question selection strategy may help to avoid questions that seem irrelevant to users. However, a CCBR system should also recognize when the dialogue can be safely *terminated* while avoiding commitment to a sub-optimal solution. Often, in practice, the dialogue is discontinued when the similarity of any case reaches a predefined threshold, or maximum information gain is less than a predefined level, or the set of competing cases has been reduced to a manageable size (Doyle & Cunningham, 2000; Aha *et al.*, 2001; Kohlmaier *et al.*, 2001). However, approaches that rely on arbitrary thresholds ignore the possibility that a better solution might be found if the dialogue is allowed to continue. Instead, the recommendation dialogue in Top Case (McSherry, 2005b) is terminated only when it is certain that the outcome will be the same regardless of any user preferences that remain unknown. This enables recommendations to be *justified* on the grounds that the user's unknown preferences cannot affect the outcome. The system's dialogue termination criteria, like its question selection strategy, are defined in terms of case *dominance* with respect to the current query (McSherry, 2003).

3.4 Dialogue grammars

There is increasing interest in the study of grammars that specify the set of dialogue operators available to the system and user in a CCBR system (Gupta & Aha, 2003). The Adaptive Place Advisor (Thompson *et al.*, 2004) combines question selection based on information gain with an adaptive user model to provide personalized recommendations. Legal dialogue moves are expressed in terms of operators such as a request to remove a constraint from the user's query when there are no matching cases. Bridge (2002) instead explores the use of a dialogue grammar, inspired by *conversational analysis*, to specify the set of legal dialogue moves in a conversation, thereby ensuring local dialogue coherence. In contrast, Branting *et al.* (2004) use augmented transition networks along with a stack-based dialogue manager to identify how to respond to user speech acts. Their discourse-oriented dialogue engine also supports the integration of CCBR with other problem solving methodologies. COBBER (Gómez-Gauchía *et al.*, 2005) is a recent framework that accounts for user *variability* during a CCBR dialogue. It models conversation strategies using causal loops, and adapts to a user's cognitive state and skills.

3.5 Query relaxation

In CCBR recommender systems, the user's requirements may be treated initially as constraints that must be satisfied (Bridge, 2002). Typically the approach relies on *query relaxation* to recover from query failures that occur when there is no case that exactly matches the user's query. For example, query failures in the Adaptive Place Advisor (Thompson *et al.*, 2004) trigger a recovery process in which the selection of a constraint to be eliminated from the unsuccessful query is based on the system's current understanding of the user's personal preferences. In Ricci *et al.*'s (2002) Intelligent Travel Recommender, the user is told how many results she will get, if any, by eliminating each of the constraints in her query. McSherry (2004) instead proposes an approach to recovery from query

failure in which there is no assumption that recovery is possible by eliminating a single constraint. An explanation of what went wrong with the user's query is followed by an *incremental* relaxation process in which the user is guided in the selection of one or more constraints to be eliminated from her query.

4 Conclusions

We have focused in this commentary on contributions that highlight potential benefits of CCBR in providing more effective support for interactive problem solving. While considerable progress has been made, some important challenges remain to be addressed. For example, the need for more reliable criteria for terminating problem-solving dialogues has only recently been recognized. While recent work has emphasized the importance of CCBR systems having the ability to explain their reasoning, this remains a topic that has received relatively little attention. There has also been relatively little work on using background knowledge to guide conversations. More expressive case representations for capturing relations among objects in a decision problem are also needed to support dialogue inferencing more effectively. Finally, although a few researchers have begun to investigate CCBR approaches that proactively contribute to the problem-solving process, this remains a relatively open research topic.

Acknowledgements

Qiang Yang is supported by Hong Kong RGC grant HKUST 6187/04E.

References

- Acorn, TL and Walden, SH, 1992, SMART: support management automated reasoning technology for Compaq customer service. In *Proceedings of the Fourth Annual Conference on Innovative Applications of Artificial Intelligence*. Menlo Park, CA: AAAI Press, pp. 3–18.
- Aha, DW and Breslow, LA, 1997, Refining conversational case libraries. In *Proceedings of the Second International Conference on Case-Based Reasoning*. Berlin: Springer, pp. 267–278.
- Aha, DW, Breslow, LA and Munoz-Avila, H, 2001, Conversational case-based reasoning. *Applied Intelligence* 14(1), 9–32.
- Aha, DW, Maney, T and Breslow, LA, 1998, Supporting dialogue inferencing in conversational case-based reasoning. In *Proceedings of the Fourth European Workshop on Case-Based Reasoning*. Berlin: Springer, pp. 262–273.
- Branting, LK, Lester, J and Mott, B, 2004, Dialogue management for conversational case-based reasoning. In *Proceedings of the Seventh European Conference on Case-Based Reasoning*. Berlin: Springer, pp. 77–90.
- Bridge, D, 2002, Towards conversational recommender systems: a dialogue grammar approach. In *Workshop on Mixed-Initiative Case-Based Reasoning, ECCBR 2002*. Aberdeen: Robert Gordon University, pp. 9–22.
- Carrick, C, Yang, Q, Abi-Zeid, I and Lamontagne, L, 1999, Activating CBR systems through autonomous information gathering. In *Proceedings of the Third International Conference on Case-Based Reasoning*. Berlin: Springer, pp. 74–88.
- Doyle, M and Cunningham, P, 2000, A dynamic approach to reducing dialog in on-line decision guides. In *Proceedings of the Fifth European Workshop on Case-Based Reasoning*. Berlin: Springer, pp. 49–60.
- Friedman, JH, Kohavi, R and Yun, Y, 1996, Lazy decision trees. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press, pp. 717–724.
- Giampapa, JA and Sycara, K, 2001, Conversational case-based planning for agent team coordination. In *Proceedings of the Fourth International Conference on Case-Based Reasoning*. Berlin: Springer, pp. 189–203.
- Göker, M, Roth-Berghofer, T, Bergmann, R, Pantleon, T, Traphoner, R, Wess, S and Wilke, W, 1998, The development of HOMER: a case-based CAD/CAM help-desk support tool. In *Proceedings of the Fourth European Workshop on Case-Based Reasoning*. Berlin: Springer, pp. 346–357.
- Gómez-Gauchia, H, Díaz-Agudo, B, Gómez-Martín, PP and González-Calero, PA, 2005, Supporting conversation variability in COBBER using causal loops. In *Proceedings of the Sixth International Conference on Case-Based Reasoning*. Berlin: Springer, pp. 252–266.
- Gu, M and Aamodt, A, 2005, A knowledge-intensive method for conversational CBR. In *Proceedings of the Sixth International Conference on Case-Based Reasoning*. Berlin: Springer, pp. 296–311.

- Gupta, KM, 1998, Knowledge-based system for troubleshooting complex equipment. *International Journal of Information and Computing Science* **1**(1), 29–41.
- Gupta, KM, 2001, Taxonomic conversational case-based reasoning. In *Proceedings of the Fourth International Conference on Case-Based Reasoning*. Berlin: Springer, pp. 219–233.
- Gupta, KM and Aha, DW, 2003, A framework for incremental query formulation in mixed-initiative case-based reasoning. In *Proceedings of the ICCBR 2003 Workshops, Norwegian University of Science and Technology, Department of Computer and Information Science*, pp. 172–180.
- Gupta, KM and Aha, DW, 2004, Towards acquiring case indexing taxonomies from text. In *Proceedings of the Sixteenth International Conference of the Florida Artificial Intelligence Research Society*. Menlo Park, CA: AAAI Press.
- Gupta, KM, Aha, DW and Sandhu, N, 2002, Exploiting taxonomic and causal relations in conversational case retrieval. In *Proceedings of the Sixth European Conference on Case-Based Reasoning*. Berlin: Springer, pp. 133–147.
- Kohlmaier, A, Schmitt, S and Bergmann, R, 2001, A similarity-based approach to attribute selection in user-adaptive sales dialogues. In *Proceedings of the Fourth International Conference on Case-Based Reasoning*. Berlin: Springer, pp. 306–320.
- McSherry, D, 1995, Integrating machine learning, problem solving and explanation. In *Proceedings of Expert Systems 95*. Oxford: SGES Publications, pp. 145–157.
- McSherry, D, 2001a, Interactive case-based reasoning in sequential diagnosis. *Applied Intelligence* **14**(1), 65–76.
- McSherry, D, 2001b, Minimizing dialog length in interactive case-based reasoning. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*. San Francisco, CA: Morgan Kaufmann, pp. 993–993.
- McSherry, D, 2002a, Mixed-initiative dialogue in case-based reasoning. In *Workshop on Mixed-Initiative Case-Based Reasoning, ECCBR 2002, Robert Gordon University, Aberdeen*, pp. 1–8.
- McSherry, D, 2002b, The inseparability problem in interactive case-based reasoning. *Knowledge-Based Systems* **15**(5–6), 293–300.
- McSherry, D, 2003, Increasing dialogue efficiency in case-based reasoning without loss of solution quality. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*. San Francisco, CA: Morgan Kaufmann, pp. 121–126.
- McSherry, D, 2004, Incremental relaxation of unsuccessful queries. In *Proceedings of the Seventh European Conference on Case-Based Reasoning*. Berlin: Springer, pp. 331–345.
- McSherry, D, 2005a, Conversational CBR in multi-agent recommendation. In *Workshop on Multi-Agent Information Retrieval and Recommender Systems, IJCAI 2005 (Technical Report W20)*, pp. 331–345.
- McSherry, D, 2005b, Explanation in recommender systems. *Artificial Intelligence Review* **24**(2), 179–197.
- Muñoz-Avila, H, Aha, DW, Breslow, LA and Nau, D, 1999, HICAP: an interactive case-based planning architecture and its application to NEOs. In *Proceedings of the Eleventh Conference on Innovative Applications of Artificial Intelligence*. AAAI Press, pp. 870–875.
- Muñoz-Avila, H, Aha, DW, Nau, D, Weber, R, Breslow, LA and Yaman, F, 2001, SiN: integrating case-based reasoning with task decomposition. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence*. San Francisco, CA: Morgan Kaufmann, pp. 999–1004.
- Nguyen, T, Czerwinski, M and Lee, D, 1993, Compaq QuickSource: providing the consumer with the power of artificial intelligence. In *Proceedings of the Fifth Conference on Innovative Applications of Artificial Intelligence*. AAAI Press, pp. 142–150.
- Racine, K and Yang, Q, 1997, Maintaining unstructured case bases. In *Proceedings of the Second International Conference on Case-Based Reasoning*. Berlin: Springer, pp. 553–564.
- Ricci, F, Arslan, B, Mirzadeh, N and Venturini, A, 2002, ITR: a case-based travel advisory system. In *Proceedings of the Sixth European Conference on Case-Based Reasoning*. Berlin: Springer, pp. 613–627.
- Shimazu, H, 1999, Translation of tacit knowledge into explicit knowledge: analyses of recorded conversations between customers and human agents. In *Workshop on Exploring Synergies of Knowledge Management and Case-Based Reasoning (Technical Report WS-99-10)*. AAAI Press.
- Shimazu, H, Shibata, A and Nihei, K, 2001, ExpertGuide: a conversational case-based reasoning tool for developing mentors in knowledge spaces. *Applied Intelligence* **14**(1), 33–48.
- Smyth, B and Cunningham P, 1995, A comparison of incremental case-based reasoning and inductive learning. In *Proceedings of the Second European Workshop on Case-Based Reasoning*. Berlin: Springer, pp. 151–164.
- Thompson, CA, Göker, M and Langley, P, 2004, A personalized system for conversational recommendations. *Journal of Artificial Intelligence Research* **21**, 393–428.
- Trott, JR and Leng, B, 1997, An engineering approach for troubleshooting case bases. In *Proceedings of the Second International Conference on Case-Based Reasoning*. Berlin: Springer, pp. 178–189.
- Weber, B, Rinderle, S, Wild, W and Reichert, M, 2005, CCBR-driven business process evolution. In *Proceedings of the Sixth International Conference on Case-Based Reasoning*. Berlin: Springer, pp. 610–624.

- Weber, B, Wild, W and Breu, R, 2004, CBRFlow: Enabling adaptive workflow management through conversational case-based reasoning. In *Proceedings of the Seventh European Conference on Case-Based Reasoning*. Berlin: Springer, pp. 434–448.
- Yang, Q, Kim, E and Racine, K, 1997, CaseAdvisor: Supporting interactive problem solving and case base maintenance. In *Workshop on Practical Applications of CBR, IJCAI 1997 (Technical Report W25)*.
- Yang, Q and Wu, J, 2001, Enhancing the effectiveness of interactive case-based reasoning with clustering and decision forests. *Applied Intelligence* **14**(1), 49–64.
- Zhang, Z and Yang, Q, 1999, Dynamic refinement of feature weights using quantitative introspective learning. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*. San Francisco, CA: Morgan Kaufmann, pp. 228–233.