Causal Query Elaboration in Conversational Case-Based Reasoning

David W. Aha
Navy Center for Applied Research in AI
Naval Research Laboratory (Code 5515)
Washington, DC 20375
aha@aic.nrl.navy.mil

Kalyan Moy Gupta
ITT Industries
AES Division
Alexandria, VA 22303
gupta@aic.nrl.navy.mil

Abstract

A key research focus for conversational case-based reasoning (CCBR) is *incremental query elaboration*, which is the process of maximizing the extraction of relevant problem state information throughout the querying process. Several companies and researchers have addressed this problem (e.g., by dynamically applying domain-specific plans (Carrick *et al*., 1999)). Recently, Gupta (2001) demonstrated how CCBR can be significantly enhanced through the judicious use of (1) taxonomies to represent domain information and (2) a control algorithm for focusing case retrieval. However, in that original conception, the individual taxonomies were isolated from each other, and from other information sources that could support query elaboration. This prevents information from being propagated to these taxonomies, and could inflate the length of the user's problem-solving session. In this paper, we outline and exemplify a causal query elaboration method for inter-taxonomy communication and highlight its potential benefits, which include shorter (and potentially more accurate) conversations, support for causal inferencing, and more concise case representations.

Introduction

Conversational case-based reasoning (CCBR) (Aha *et al*., 2001) is a case-based reasoning (CBR) (Aamodt & Plaza, 1994) methodology that supports incremental query refinement. It has been primarily deployed for use in help-desk (Acorn & Walden, 1992), troubleshooting (Gupta, 1998), and electronic commerce (Shimazu, 2001) applications, although it has also been suggested for use in a variety of planning tasks (Abi-Zeid *et al.,* 1999; Muñoz-Avila *et al*., 1999). CCBR defines a mixed-initiative process in which a user incrementally specifies their query by providing text annotations, answering prompted questions, or via other modalities (Shimazu *et al.,*1994; Göker *et al.,* 1998).

A key concern for CCBR is in minimizing the cognitive effort required by the user to retrieve cases of their interest (Shimazu, 2001). In addition to developing models that anticipate user needs (Göker & Thompson, 2000), cognitive load could potentially be reduced by developing general-purpose inferencing methods to reduce the amount of information that the user must provide (e.g., number of answered questions) prior to satisfactory case retrieval (Aha *et al.,* 1998). With this goal in mind, Gupta (2001) introduced a *Taxonomic CCBR* approach that re-structures cases into a set of taxonomies, one per *factor*, that, by adapting to the user's level of abstraction, can more quickly focus case retrieval.

In our initial experiments, we have found efficiency gains in using Taxonomic CCBR. However, we also recognized that further advantages could be obtained by focusing the query elaboration effort on this taxonomic framework. In particular, causal inferencing procedures could be employed that relate and propagate information among these taxonomies, and from outside information sources to these taxonomies. The potential benefits include shorter conversations, increased exploitation of domain-specific information, and more concise case representations.

This short paper summarizes our approach for extending the Taxonomic CCBR framework to support causal query elaboration. We briefly review CCBR and Taxonomic CCBR prior to introducing the proposed causal inferencing approach, and then summarize future work.

Conversational Case-Based Reasoning

Case-based reasoning (Aamodt & Plaza, 1994) is a general problem-solving methodology that involves maintaining a *case library*, where each case is a <problem, solution, outcome> triplet, and using it to solve new problems (i.e., *queries*) through processes involving case *retrieval*, *reuse*, and *revision.* This can lead to the generation of new cases, which can be incorporated through a *retention* process.

Conversational CBR (Aha *et al*., 2001) is a mixed-initiative variant of CBR in which a query is incrementally acquired from a user. During this *conversation*, the CCBR system continually updates ranked displays for both the most similar cases and their unanswered questions. Case similarity is determined by matching the query with the cases' problems, while questions can be ranked according to frequency within the top-ranked cases (Aha *et al*., 2001), by an information gain metric (Yang & Wu, 2001), or by more informed metrics (e.g., McSherry, 2001a). (McSherry (2001b) describes an evaluation of alternative methods for ranking questions.) Typically, users initialize a query at the start of a conversation by providing a textual description of their problem, although alternative modalities can be used (Shimazu *et al*., 1994; Göker *et al.,* 1998). At any point during the conversation, the user can select and answer a displayed question, or select and examine a solution for a displayed case. Inference Corporation pioneered this technology for the help desk market niche, and several researchers and companies (e.g., CaseBank Technologies) have since explored its variants. Figure 1 displays a CCBR case obtained from a dataset concerning TV troubleshooting. In this paper, we ignore the *outcome* component, assuming it here to be positive, and represent problems with a text component (i.e., a *description*) and a set of <question, answer> pairs (i.e., a set of *conditions*).

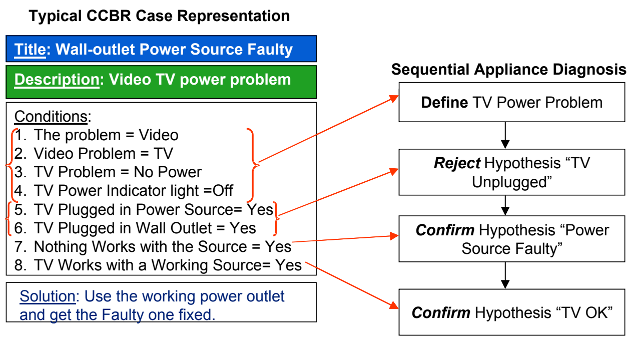

Figure 1: An example CCBR case and its interpretation for

Like most published work to date on CCBR, this paper focuses almost exclusively on case *retrieval* among the four processes mentioned above. We argue that *query elaboration* is a key component of the case retrieval process. A simplification of the typical CCBR processing loop for case retrieval is shown in Figure 2, which highlights a query elaboration step that is executed each time the query is updated (i.e., either after query initialization or after the user answers a displayed question). This step can involve any process for query updating. For example, Inference Corporation's products provided a means for users to incorporate query-elaboration rules in their applications that would fire during this step.

```
CCBR(L) =
  Query := initialize_query();
  Repeat
    Query :=
  query_elaboration(Query);
    D := update_displays(Query,L);
    Question :=
  user_select_displayed_question(D);
    A := prompt_for_answer(Question);
```

At least three research groups have explored alternative approaches for exploiting domain-specific information during query elaboration, which we refer to as *causal query elaboration*. First, Montazemi and Gupta (1996) clarified the importance of query elaboration during sequential diagnosis processes, and demonstrated how belief networks could be used to support this process so that CCBR could be adaptive to user needs. Second, Aha *et al.* (1998) argued that query elaboration rules could be more easily maintained if organized into a model, and provided evidence of this approach's utility. Third, Carrick *et al.* (1999) used pre-stored plans rather than more simplistic rules for query elaboration.

Taxonomic CCBR

Although these three forms of causal query elaboration can assist by automating information gathering activities (through distinct inferencing approaches), they only partially address the goal of minimizing cognitive load and dialog length during conversational case retrieval. In particular, they do not focus on minimizing case representation. In addition, they are not adaptive to the user's level of expertise, which would require displaying only those questions that are appropriate to a user's demonstrated level of understanding.

Gupta (2001) introduced the Taxonomic CCBR approach to address these two problems. In particular, he noted the following issues arising from ignoring abstraction relations between features (i.e., <question, answer> pairs):

- Correlation among features (at different levels of abstraction) could lead to similarity assessment errors,
- similarity cannot be assessed among features that are related by abstraction,
- redundant questions cannot be generated during query (i.e., problem description) elaboration, and
- decisional information during case representation can be lost, and representational inconsistencies could accrue during case base maintenance.

The Taxonomic CCBR methodology reduces the impact of these problems by explicitly representing abstraction relations in <question, answer> pair taxonomies. Its key observation is that, typically, CCBR cases are structured from most general to most specific <question, answer> pairs (e.g., see Figure 1), and that the more general among these are often shared with other cases. However, more than one factor might exist in a single case, where each factor focuses on a separate dimension of the problem domain. Thus, Taxonomic CCBR represents each distinct factor in the problem domain as a subsumption taxonomy of <question, answer> pairs, and cases are represented as a set of (incoming) pointers, where each points from a leaf in a distinct feature taxonomy. This distributed case representation, which is more structured than some other approaches (e.g., case retrieval nets (Lenz & Burkhard, 1996)), trades off representational flexibility so that it can more easily support query elaboration and abstraction adaptation.

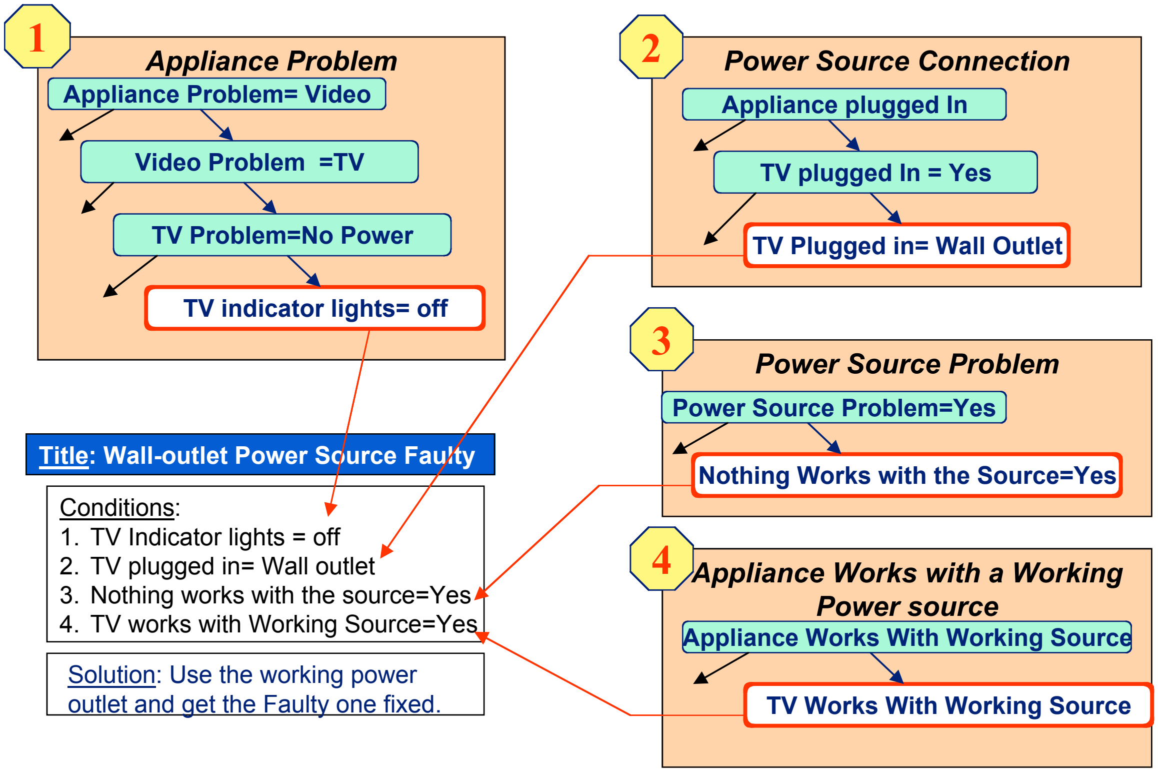

Figure 3: The example case, and its related four taxonomies, represented using a Taxonomic CCBR approach.

Gupta (2001) argued that these taxonomies could facilitate retrieval by focusing attention; if the query identifies a node (i.e., <question, answer> pair) in a taxonomy, then we can deduce that all parent nodes in that hierarchy have implicitly been answered. This can decrease the number of questions that the user needs to answer prior to case (i.e., solution) retrieval. For example, Figure 3 displays a Taxonomic CCBR representation for the case shown in Figure 1. It highlights the elimination of CCBR's *description* component (because text matching during *initialize_query*() attempts to match the user's text directly with taxonomy nodes), a reduction in the case's number of <question, answer> pairs (from 8 to 4, one per sequential diagnosis phase), and its pointers from the taxonomy's leaves. Contrasted with traditional conversational case retrieval (Figure 2), Taxonomic CCBR introduces a pre-processing step to create the taxonomies and distributed case representations. It also uses a more ambitious *query_elaboration*() step that locates the nodes in the available taxonomies that are implied by the query, and then constrains the set of displayed questions (in *update_displays*()) to those that permit downward taxonomic traversal from these nodes.

Causal Inferencing Opportunities

While the Taxonomic CCBR approach, as originally conceived, captures subsumption (i.e., abstraction) relations among features, these are not the only type of inter-feature relations that should be exploited. In particular, exploiting *causal* (e.g., event A is caused by B) and *implication* relations (i.e., conclusion A implies conclusion B) can also contribute towards query elaboration.

Existing CCBR approaches, including taxonomic CCBR, often use cases to represent decision-making or problem solving activity that may include a sequence of decision making/problem solving steps. For example, a troubleshooting case in a TV troubleshooting help-desk application includes a sequence of steps for confirming or eliminating intermediate hypotheses to identify the final root cause (see Figure 1). Similarly, a case for a product configuration application could include a sequence of configuration decisions, each step dependent on the previous one.

These types of (common) CCBR applications include feature dependencies. However, like abstraction, these inter-feature dependencies cannot be accommodated using a standard CCBR representation. Thus, some of the same problems occur as those arising from abstraction: the dependencies yield feature correlations, cause redundant questions to be displayed during query acquisition, and accrue representational inconsistencies during case library maintenance.

We propose and sketch an approach for propagating causal inferences during query elaboration. In particular, we augment the Taxonomic CCBR methodology by ***incorporating dependency relations among features (e.g., causal, implied, sequential) to improve both representational and query elaboration efficiency.*** We use two rules to guide this process:

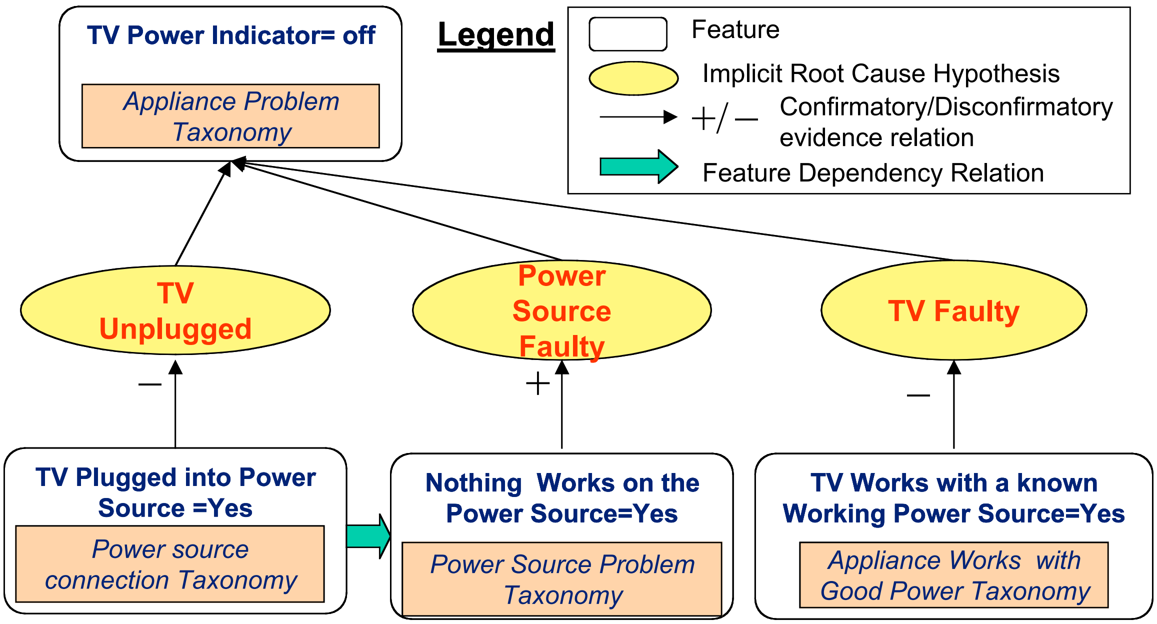

Figure 4: Dependencies among the features of the example case.

1. Relation Insertion Rule _: Causal relations are inserted only among (two) leaf nodes in distinct taxonomies.
2. Feature Inclusion Rule _: In causal relations, only the dependent features are included in the case.

The first rule prevents the inclusion of causal relations that involve interior nodes, or that relate nodes within a single taxonomy. Although we may later explore relaxations of this rule, it suffices to focus causal reasoning between nodes that represent the most specific features of a case, and it simplifies our example. The second rule, which can reduce the number of conditions associated with a case, deletes features from a case that, if true, imply other features in the case. These can be safely deleted because their dependency relation is captured by this augmentation of Taxonomic CCBR, and they are otherwise redundant.

We applied these rules to the example case in Figure 3. Figure 4 explicates the dependencies among the features of this case, showing how the three competing hypotheses (i.e., *TV unplugged*, *Power Source Faulty*, and *TV Faulty*) are rejected, confirmed, and rejected respectively. There is also a logical dependency between the feature *TV plugged into power source=Yes* and *Nothing works on the power source=Yes*. By incorporating inter-feature dependency relations using the Relation Insertion Rule, the example in Figure 3 can be represented as shown in Figure 5. This figure includes a dependency relation from the feature *TV plugged into power source=Yes* to *Nothing works on the power source=Yes*, resulting in the elimination of the former <question, answer> pair from this case due to the Feature Inclusion Rule.

This illustrates how representational gains can be realized when dependencies between features from different taxonomies are recognized and represented. This can simplify case maintenance by abstracting inter-feature dependencies, which no longer require to be repeated in multiple cases (a potential source of representational inconsistency).

To incorporate these rules, the Taxonomic CCBR approach must be modified as follows:

Search, match, and retrieval: In addition to supporting taxonomy traversal from a starting node that was identified from a user's textual description, our method for processing dependency relations will also support inter-taxonomy traversal.

Query elaboration: Our method will augment the conversation generation algorithm to accommodate the dependency relations by ordering questions within a case. Specifically, none of the dependent questions will be asked before the user answers their corresponding "dependee" questions. For example, the question "*Does anything work with the power source?*" cannot be asked until the user has answered the question "*Is the TV plugged into the power source?*" This imposes a valid ordering on questions during the conversation. Further, it suppresses questions that may not relevant to a particular stage of the information gathering process.

Summary and Future Work

We described an augmentation of the Taxonomic Conversational Case-Based Reasoning (CCBR) approach that, in addition to exploiting taxonomic relations among features in a case base, exploits other types of (causal) relations. We also described an example in which our approach would increase representational efficiency and

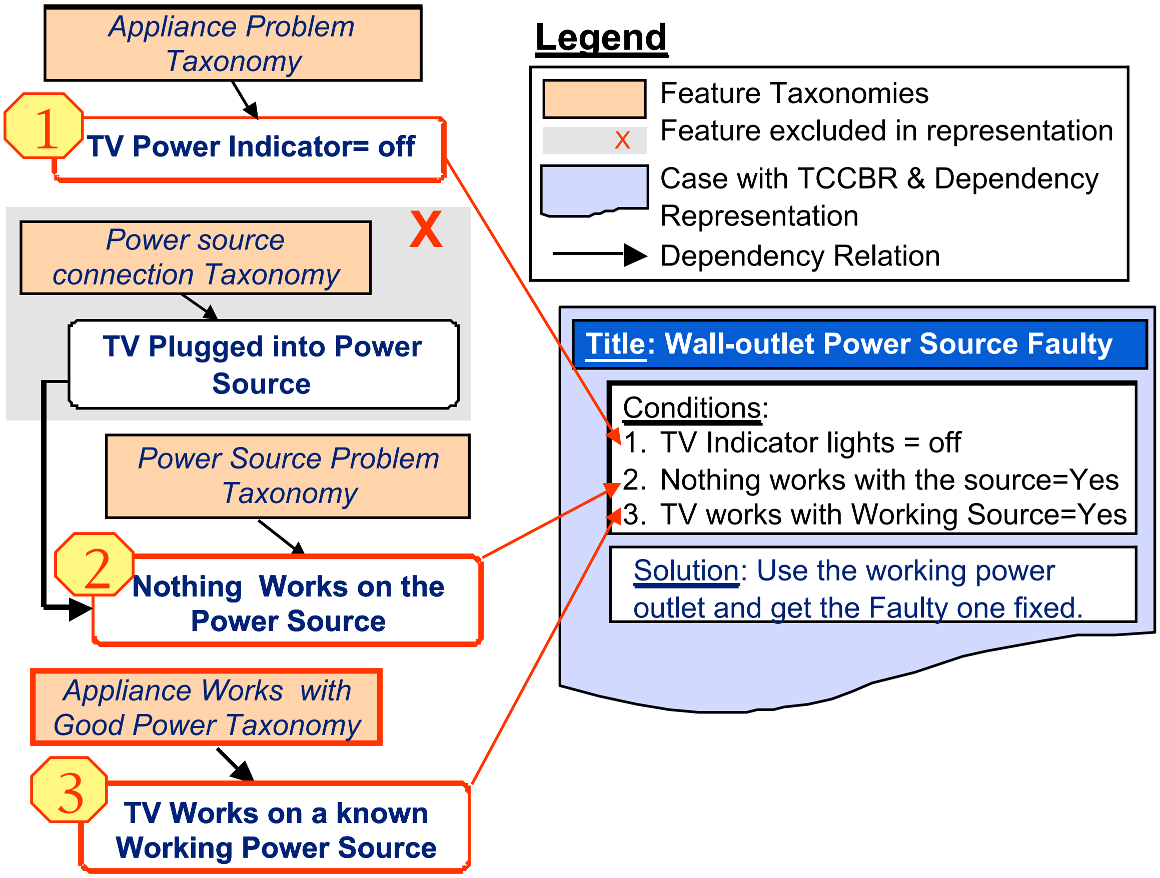

Figure 5: The example case, re-represented after accounting for inter-

improve query elaboration by exploiting feature dependency relations.

In our future work, we will formalize this augmented approach, and evaluate whether its effects (e.g., potential reductions in feature correlation and representational inconsistency) will yield significant improvements in CCBR performance (e.g., conversational efficiency). These results will be contrasted with those from using standard CCBR and Taxonomic CCBR approaches. It is possible that our proposed approach could also assist with the application of CCBR to task decomposition tasks in the context of plan authoring (Muñoz-Avila *et al.,* 1999). More generally, we intend to explore issues concerning information propagation in the context of query elaboration agents.

Acknowledgements

We thank the FLAIRS organization and Héctor Muñoz-Avila for inviting this paper and its presentation at FLAIRS'02. This research was funded by a grant from the Naval Research Laboratory.

References

Aamodt, A., & Plaza, E. (1994). Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications, 7*, 39-59.

Abi-Zeid, I., Yang, Q., & Lamontagne, L. (1999). Is CBR applicable to the coordination of search and rescue operations? *Proceedings of the Third International Conference on Case-Based Reasoning* (pp. 358-371). Seeon, Germany: Springer.

Acorn, T.L., & Walden, S.H. (1992). SMART: Support management automated reasoning technology for COMPAQ customer service. *Proceedings of the Fourth Annual Conference on Innovative Applications of Artificial Intelligence*. San Jose, CA: AAAI Press.

Aha, D.W., Breslow, L.A., & Munoz-Avila, H. (2001). Conversational case-based reasoning. *Applied Intelligence*, *14*(1), 9-32.

Aha, D. W., Maney, T., & Breslow, L. A. (1998). Supporting dialogue inferencing in conversational case-based reasoning. *Proceedings of the Fourth European Workshop on Case-Based Reasoning* (pp. 262-273). Dublin, Ireland: Springer.

Carrick, C., Yang, Q., Abi-Zeid, I., & Lamontagne, L. (1999). Activating CBR systems through autonomous information gathering. *Proceedings of the Third International Conference on Case-Based Reasoning* (pp. 74-88). Seeon, Germany: Springer.

Göker, M., Roth-Berghöfer, T., Bergmann, R., Pantleon, T., Traphöner, R., Wess, S., & Wilke, W. (1998). The development of HOMER: A case-based CAD/CAM help-desk support tool. *Proceedings of the Fourth European Workshop on Case-Based Reasoning* (pp. 346-357). Dublin, Ireland: Springer.

Göker, M., & Thompson, C.A. (2000). Personalized conversational case-based recommendation. *Proceedings of the Fifth European Workshop on Case-Based Reasoning* (pp. 99-111). Trento, Italy: Springer.

Gupta, K.M. (1998). Knowledge-based system for troubleshooting complex equipment. *International Journal of Information and Computing Science*, *1*(1), 29-41.

Gupta, K.M. (2001). Taxonomic conversational case-based reasoning. *Proceedings of the Fourth International Conference on Case-Based Reasoning* (pp. 219-233). Vancouver, Canada: Springer.

Lenz, M., & Burkhard H. D. (1996). *Case retrieval nets: Basic ideas and extensions.* In G. Görz & S. Hölldobler (Eds.) *KI-96: Advances in Artificial Intelligence*. Berlin: Springer.

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. *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence* (pp. 993-998). Seattle, WA: Morgan Kaufmann.

Montazemi A.R., & Gupta K.M., (1996). An adaptive agent for case description in diagnostic CBR systems. *Computers in Industry*, *29*(3), 209-224.

Muñoz-Avila, H., Aha, D.W., Breslow, L.A., & Nau, D. (1999). HICAP: An interactive case-based planning architecture and its application to NEOs. *Proceedings of the Eleventh Conference on Innovative Applications of Artificial Intelligence* (pp. 870-875). Orlando, FL: AAAI Press.

Shimazu, H. (2001). ExpertClerk: Navigating shoppers' buying process with the combination of asking and proposing. *Proceedings of the Seventeenth International Conference on Artificial Intelligence* (pp. 1443-1448). Seattle, WA: Morgan Kaufmann.

Shimazu, H., Shibata, A., & Nihei, K. (1994). Case-based retrieval interface adapted to customer-initiated dialogues in help desk operations. *Proceedings of the Twelfth National Conference on Artificial Intelligence* (pp. 513-518). Seattle, WA: AAAI Press.

Yang, Q., & Wu, J. (2001). Enhancing the effectiveness of interactive case-based reasoning with clustering and decision forests. *Applied Intelligence, 14*(1), 49-64.