

Hybrid Case-Based Reasoning for the Diagnosis of Complex Devices

M. P. Féret and J. I. Glasgow

Department of Computing & Information Science,
Queen's University, Kingston,
Ontario, Canada, K7L 3N6
{feret,janice}@qucis.queensu.ca

Abstract

A novel approach to integrating case-based reasoning with model-based diagnosis is presented. The main idea is to use the model of the device and the results of diagnostic tests to index and match cases representing past diagnostic situations with the current one. The initial diagnostic methodology is presented as well as the problems encountered while applying this methodology to two real-world devices. The incorporation of a case-based reasoning system is then motivated and described in detail. Experimental results show the effectiveness of both the indexing schema and the matching algorithm. The paper also discusses how and why these results can be generalized to a multiple fault situation, to other types of device models and to other applications in the field of artificial intelligence.

Introduction

This paper presents an approach to integrating case-based reasoning with a traditional diagnostic method for complex devices. This generic approach to diagnosis is based on a hierarchical decomposition of mechanical devices and uses sensor data, collected in real-time and stored in a database, to guide the search towards hypothetical diagnoses [6, 7, 8, 16]. Some of the difficulties encountered while applying a model-based reasoning (MBR) diagnostic method to two real-world devices are identified. These difficulties arose from imperfections of the device model, due to human errors or misconceptions. These imperfections lead to incorrect models which produced inadequate diagnostic performance. In this paper we further develop ideas initially introduced in [7], providing additional motivations for

*This research was supported through a contract from the Canadian Space Agency (STEAR program), a scholarship and an operating grant from the Natural Sciences and Engineering Research Council (NSERC) of Canada. We also would like to thank Spectrum Engineering Corporation Ltd., Peterborough, Ontario, Canada.

our approach, and experimental results that support the claims of the paper.

CBR has traditionally been used as a stand-alone problem-solving method [13, 14], sometimes applied to diagnostic problems (e.g. [20]). Only recently has CBR been used in association with other problem-solving paradigms [15, 10, 19]. Our approach is novel in that it uses CBR only *after* the MBR process has taken place, and in that it uses the model and the results of the MBR process to index the cases.

The hybrid CBR/MBR methodology described in this paper incorporates a critique of the results of the model-based approach in the light of past experience and provides the human operator with a means for exploring alternative hypotheses. The integration of CBR with the structural isolation process allows for a simple and effective indexing schema as well as a computationally inexpensive similarity measure for cases.

This paper initially presents the structural isolation process. It then lists some of the problems arising when trying to apply any MBR diagnostic method to complex devices. The hybrid approach combining CBR and MBR is presented along with some experimental results. The final discussion summarizes the contributions of the research.

Generic Diagnosis

Our approach to generic diagnosis has been fully implemented in the Automated Data Management System (ADMS), which has previously been described and compared to other techniques for diagnosis in [6]. This section describes the structural isolation process and outlines the problems that arose in applying it to two real-world devices, a robotic system called the Fairing Servicing Subsystem, and a Reactor Building Ventilation System [6, 7]. The first device is a robot placed at the rear of a ship to automatically replace damaged fairings on a cable which drags an underwater detec-

tion system. The fairings prevent the detection system from drifting away from the axis of the ship. The Reactor Building Ventilation System is a modified model of a ventilation system for an existing nuclear power plant.

Structural Isolation Process

Govindaraj and Su suggested that empirical constraints should direct the formation of knowledge representation in a format that reflects how experts solve problems [12]. They also observed that human diagnosis proceeds in a hierarchical manner, starting at higher levels of abstraction to generate hypotheses that guide the diagnosis at lower levels. Hierarchical, structural and conceptual data structures have already been found useful for diagnostic applications (e.g. [11, 12, 25]).

The ADMS methodology for diagnosis of mechanical devices is based on considering the device as a hierarchy of components or groups of components. This hierarchy is expressed using a frame language and constitutes the backbone of the knowledge base built for individual applications of the ADMS to a mechanical device. All recognizable components that can be diagnosed as sources of failure are situated at the bottom of the hierarchy. Each component has associated failure mechanism patterns. These patterns are conjunctions of sensor functions, testing specific conditions on the sensor data stored in a database [16]. Between these top and bottom levels, the device is decomposed into meaningful substructures, associated with test conditions indicating whether these substructures are potentially faulty.

While structural knowledge is usually available from engineering design documents and easily encoded, the knowledge about failure modes and patterns tends to be complex and of various types. In the case of sensor-based diagnosis, the queries to the database and their use in the sensor functions lead to numerous difficulties. These difficulties constitute the major differences between complex devices and electronic circuits, where failures and their consequences are straightforward to characterize.

The models used by the ADMS methodology therefore consist of a structural decomposition of the devices, along with necessary conditions for substructures and basic components being potential diagnoses. They are fault models in which all testing conditions use the real-time sensor data stored in the database through a set of predefined queries.

In the context described above, performing diagnosis involves traversing the hierarchy according to results of

necessary conditions applied to the substructures represented in the hierarchy. At a given node in the hierarchy, if there is no evidence that the substructure can be faulty, then the whole substructure can be pruned from the search space of potentially faulty components. If there is such evidence, the substructure is examined further and more local conditions are applied to subnodes of the current node. This process is known as the structural isolation process [17]. It also handles multiple faults by simultaneously investigating multiple paths in the model.

The output of the diagnostic algorithm is a ranked list of potential diagnoses from which the operator can (or not) select the final (supposedly correct) diagnosis. The ranking involves relative levels of confidence in each potential diagnosis as well as criteria related to the history of each component (such as meantime between failures, time to life expectancy, etc.).

Model-Based Diagnosis Weaknesses

Diagnosis from first-principles is believed to be NP-hard in the general case [19, 21, 22]. Many researchers have tried to focus the search for minimal diagnoses in the diagnosis from first-principle approach [2, 3, 4, 5] or to reduce the complexity of similar methods (e.g. [9, 18]).

Analyzing and compiling human diagnostic problem-solving capabilities is difficult. Misunderstandings, incorrect specifications, typos, etc. typically lead to partially incorrect models which are difficult to debug, especially when there is no simulation program available for the device. Moreover, device models are not always the most natural or efficient representation for diagnosing faulty components [23]. These knowledge acquisition and validation problems clearly weaken the reliability of model-based systems such as the ADMS and need to be addressed before the system can be used for critical, real-world applications. Fault models consisting of necessary conditions (for parts to be faulty) are simpler to express and to implement than correct behavior models. However, fault models have a well-known drawback: they only model foreseen, predictable faults. Therefore, diagnostic systems based on fault models are ineffective in the presence of unforeseen faults.

The process of human diagnostic problem-solving is often suboptimal [12, 26]. Resulting shortcomings are likely to be found in any model designed and implemented by humans. In our experience, we have found such mistakes in the experts' explanations and reasoning processes. This leads to models that are either incomplete or inconsistent because they incorporate human limitations.

Recalling past relevant decisions (diagnostic cases) is an effective way of reducing the impact of both the model's and the human's inadequacies. The next section describes the addition of a CBR system to the approach described above. This hybrid approach assists in overcoming the bottlenecks of knowledge acquisition and human reasoning imperfections that limit the capabilities of current model-based approaches to diagnosis.

CBR and Structural Isolation

The philosophy behind CBR is that "raw", unabstracted experiences can be used as a source of knowledge for problem-solving [14, 24]. A CBR system stores past experiences in the form of cases. When a new problem arises, the system retrieves the cases most similar to the current problem, then combines and adapts them to derive and criticize a solution. If the solution is not satisfactory, new cases are retrieved to further adapt it in the light of new constraints, expressed from the non-satisfactory parts of the proposed solution. The process is iterated until the proposed solution is judged acceptable. After a problem is solved, a new case can be created and stored in the casebase. The main issues to address when building CBR systems are to define an effective indexing schema, efficient retrieval and storage mechanisms, a reliable similarity measure for cases and an adaptation mechanism.

CBR is limited by the difficulty of indexing, retrieving and evaluating previous experiences. This is especially true in the case of diagnostic applications, where similar symptoms can have very different or multiple causes. Techniques that work for small problems do not necessarily handle scaled-up versions of the same problem. Model-based approaches are limited by the fact that complete and consistent models of complex devices are difficult to produce.

Researchers have previously combined CBR with other problem solving paradigms. Rajamoney and Lee use CBR to decompose a novel, large problem into smaller known ones that are then solved with a model-based reasoning (MBR) system [19]. They use CBR for a separate task than the one the MBR system is used for. Koton's CASEY system uses CBR to speed up a model-based diagnosis system by storing previous experiences and recalling them when appropriate [15]. In this system, CBR is tried first as an attempt to reason by analogy. The cases are directly derived from the MBR system and are used, once created, in isolation from the MBR system. In these two systems, the paradigms are used independently from one another. Golding and Rosenbloom use CBR to improve

the accuracy of a rule-based system [10]. Their cases denote exceptions to rules and are indexed by the rules they confirm and by the rules they contradict. Cases do not store the same information as the rules, but provide a different source of information for decision making. These systems all show that CBR does improve the performance of the whole system by either speeding up the same process without bringing new information or by storing a different type of experiential knowledge that is used to improve the accuracy of the overall system. However, all these systems have to face critical problems in retrieving and matching their cases, problems that are typical of CBR systems.

The ADMS hybrid approach to diagnosis is unique in that it uses the results of the structural isolation process to index cases. There is little overhead in retrieving relevant cases and matching is simplified since it is only applied to similar cases. We are considering CBR as a tool for assisting the operator in the final stage of a diagnostic session. Once the structural isolation has produced a list of potential diagnoses, the operator can call on the CBR component of the system to validate diagnoses or investigate other paths in the search tree.

The remainder of this section describes the content of cases, the retrieval process and the matching algorithm currently being used to evaluate case similarity.

What is a case?

In general, a case stores a fragment of a past experience. In the context of the ADMS, a case stores a past diagnostic scenario, consisting of a description of the fault that occurred (fault type, fault time, detecting sensor), the series of pruning steps used to produce a list of potential diagnoses (i.e. the tests performed during diagnosis and their values), the list of potential diagnoses produced by the structural isolation process and the correct diagnosis selected by the operator. A successful case is a case where the correct diagnosis was produced by the structural isolation process and confirmed by the human operator. A failure case is a case where the diagnosis failed to find the correct diagnosis, and for which the operator chose a component that was not in the list of proposed diagnoses.

Indexing and Storage of Cases

The structural isolation process can be seen as a rough estimate of the location of a component whose failure explains the observed symptoms. The list of potential diagnoses is used as a means of indexing the casebase, leaving the values of the associated sensor functions for

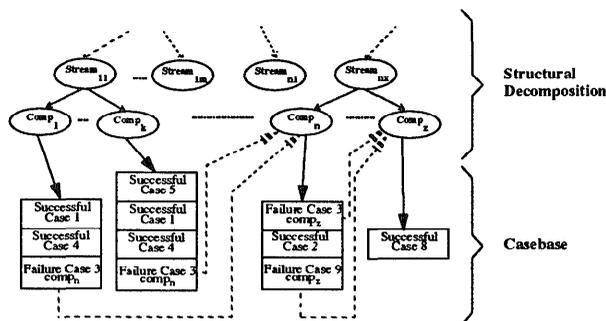


Figure 1: The casebase and the structural decomposition.

the matching step which is a finer judgement of similarity. For either a successful or a failure case, we use each potential diagnosis produced by the structural isolation process as an index for the case. Figure 1 illustrates this indexing schema. Each case is stored at the bottom of the structural decomposition under the basic components it contains as potential diagnoses. The dashed arrows originate from failure cases and point to the components representing the correct diagnoses for those cases.

This indexing schema is satisfying because of the interdependencies that exist among “neighboring” components. Such components often share the same characteristics and are likely to appear in each other’s lists of potential diagnoses. They will likely share the same cases, or cases that are very similar to each other, except for failure cases. This ensures a useful grouping of similar cases with “bridges” from one grouping to the next provided by failure cases. This indexing schema, based on the information generated by the structural isolation process, is therefore both simple and effective.

Retrieval and Matching of Cases

Cases are retrieved from the casebase to evaluate and criticize the current list of potential diagnoses. The operator can ask the ADMS to explore its casebase and to criticize or confirm a potential diagnosis or to suggest new diagnoses that were not generated by the structural isolation process. If the current potential diagnosis is supported by a previous successful case, the level of confidence in this potential diagnosis can be raised. If the correct diagnosis for the most similar case disagrees with all the suggested diagnoses, and points towards a failure case that matches sufficiently well with the current situation, the validity of the current diagnosis is lowered. The diagnosis stored in the failure case is extracted from the casebase and presented to

the user as a new potential diagnosis that can, in turn, be evaluated.

Retrieved cases are matched with the current situation, using finer criteria to evaluate similarity. Such criteria include state information from both the past case and the current situation. The measure of similarity yielded by the matching algorithm is a normalized, weighted sum of the number of sensor functions sharing the same value, the number of substructures shared in the path followed during the session represented by the past case and during the current session, the common characteristics of the symptoms, etc. All steps in the matching algorithm involve comparisons of booleans or of reals and are computationally inexpensive.

Because of the indexing method described above, components at the bottom of the hierarchy serve as pointers to cases that represent diagnostic sessions caused by similar or related failures. The matching is effective because the knowledge contained in those cases and on which the matching is based is relevant in both the current and the past cases. The matching algorithm is focused on the part of the system that is the most relevant to the current situation.

Experimental Results

The thesis of this paper and the hypothesis for our experiments is that the analysis of past experiences can aid in the diagnostic process. In particular, CBR can be used to effectively validate or critique a diagnostic decision resulting from an imperfect system model. Testing the effectiveness of the hybrid approach to diagnosis requires an initial casebase of diagnostic cases. Such a casebase could be created with the normal running of the system. For testing purposes, however, a simulator was constructed for the Fairing Servicing Subsystem to automatically generate sensor data and the corresponding sensor function boolean values, representing single faults. This simulator was used to generate initial casebases and to simulate test cases for evaluating the CBR component of the system.

Some simplifying assumptions were made in carrying out the two experiments. First, we assumed that all failure modes for components are equally probable. Although this does not reflect reality, it does present - to some degree - a worst case scenario. Secondly, we assume that the sensor information is accurate, i.e. the sensors are not included among the components that may fail.

The hybrid approach described above is implemented as an interactive process where the human operator explores the casebase indexed by the structural decomposition of the device. To test our approach,

we have implemented a non-interactive version of the same program that retrieves all the cases stored under the potential diagnoses produced by the structural isolation. We progressively degrade the model, starting from a model that produces perfect diagnosis performance, and moving towards models that contain errors. This degradation process takes place in the sensor functions, which are randomly selected and failed. A failed sensor function returns an incorrect result, falsely describing the state of the device. The model for the Fairing Servicing Subsystem contains 19 sensor functions, describing 15 failure patterns (conjunctions of sensor functions) for 100 components of 8 different component types, e.g. motors (4 failure patterns), cables (2 failure patterns). The following experiments involved running 77 simulations of faults occurring in two modules of the Fairing Servicing Subsystem containing 36 components and 7 sensors.

Experiment 1

The first experiment's goal is to measure the diagnostic performance degradation in relation to model degradation and to show the effectiveness of the indexing mechanism for CBR. The hypothesis is that the diagnostic performance becomes worse as the model degrades. Another goal of this experiment is to show how many potential diagnoses the CBR system contributes to the final result of the integrated approach.

Method. The simulations are run twice. The first pass does not use a casebase but generates one. The second pass uses the casebase containing 77 cases. A casebase containing the same cases as the ones that are currently being run is a not a good test for measuring the performance of the CBR system. However, the goal of this experiment is not to measure diagnostic performance itself but rather its degradation.

Results. Figure 2.a shows that the degradation is approximately linear in the number of failed sensor functions. Figure 2.b shows the number of potential diagnoses produced in the same experiment. In this experiment, all retrieved cases are considered similar enough to the current situation. Figure 2.b shows the influence of the CBR system in generating new potential diagnosis. The number of potential diagnoses produced by both the hybrid method and the single model-based method are linear in the number of failed sensor functions.

The linearity of the degradation in Figure 2.a is the result of a balanced use of sensor functions in the failure patterns. Failure patterns consist of 3 to 6 sensor functions, possibly negated. Figure 2.b shows the number of retrieved cases, when the casebase contains

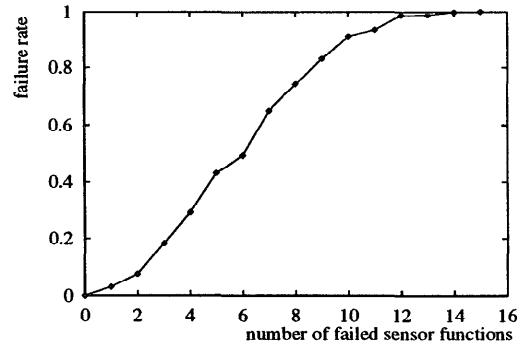


Figure 2.a) Performance Degradation from Model Degradation

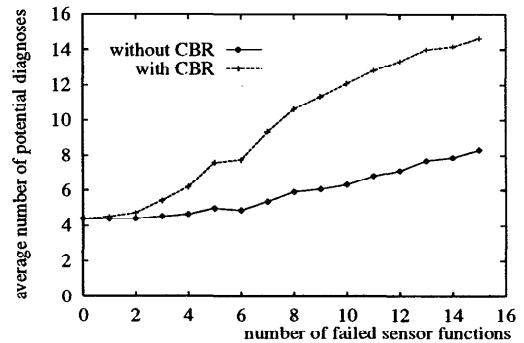


Figure 2.b) Number of Potential Diagnoses

all the cases that were generated by the simulations. This gives an indication of the computational cost of the CBR system. This cost increases linearly when the quality of the model decreases. Although significant, this cost remains within reasonable bounds. This experiment also shows the effectiveness of the retrieval process. Regardless of the quality of the model, the appropriate cases were always retrieved. Intuitively, Figures 2.a and 2.b correspond to the intuition that the worse the device model is, the less accurate it is and the more experience is required to compensate for erroneous knowledge.

Experiment 2

The hypothesis of the second experiment is that the hybrid approach, including a CBR component performs better than a the model-based approach alone, for a small additional computational cost and without over-

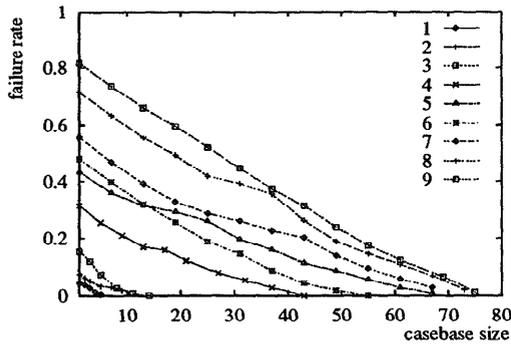


Figure 3.a) Failure Rates

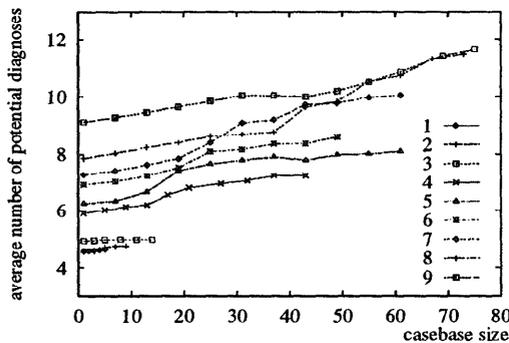


Figure 3.b) Average Number of Potential Diagnoses

whelming the operators with potential diagnoses. It also aims at showing the influence of experience (cases) on the diagnostic performance. Experiment 1 already showed that the number of retrieved matches remains within reasonable bounds.

Method. In this experiment, casebases of different sizes are used for different numbers of failed sensor functions, and for each casebase size, the 77 simulations are run on 10 different casebases. The matching algorithm described above is used to measure the similarity and the plausibility of the retrieved matches compared to the current diagnostic situation.

Results. The horizontal axes in Figure 3.a and 3.b represent the size of the casebase. Each curve represents a number of failed sensor functions. Figure 3.a shows that the improvement is linear in the size of the casebase. Figure 3.b shows that the average number of potential diagnoses increases only marginally with the size of the casebase. This experiment shows that the

performance of the hybrid diagnostic system increases linearly with the size of the casebase. The matching algorithm is itself generic and does not use any specific knowledge about the Fairing Servicing Subsystem. We therefore believe that it could be applied, with similar results, to the Reactor Building Ventilation System, or to any other complex device. Figure 3.b illustrates its effectiveness in pruning the cases that are irrelevant to the current situation. The number of potential diagnoses is almost constant as the casebase grows larger. Figure 3.a shows the failure rates decreasing linearly with the increasing size of the casebase. Better results could be achieved if the retrieval and the matching algorithms allowed for generalizations over component types. For example, most motors in the Fairing Servicing Subsystem share the same installation configuration and are monitored by similar sensors. Therefore cases related to one motor could be adapted to other motors in the device.

Closing Remarks on the Experiments

There are other possible ways to test this approach. The size of the models could be decreased by removing sensor functions. This simulates a decreasing number of sensors as opposed to a decreasing number of working sensors. This would show how important the model, even a partially incorrect model, is in the indexing. This approach could also be tested for incorrect failure patterns, or for erroneous structural knowledge. Our experience shows that these types of errors are typically easier to detect than errors in sensor functions.

Discussion

This paper presents a hybrid approach to diagnosis. It is based on a structural isolation search for potential diagnoses, enhanced by contributions from an integrated CBR component that assists the human operators in their final decision by using both successful and failure experiences. Such assistance is useful because the structural decomposition and the pruning rules associated with it are not guaranteed to be either consistent or complete and because human operators might not consider all the cues that are available to them.

Compared to an all-model-based approach to diagnosis, a hybrid approach addresses a number of problems. CBR allows the system to improve on the model built by humans. It overrules some mistakes that can be made in the design and implementation of this model. This is accomplished in a computationally inexpensive way, by using the decomposition tree

as the basis for indexing previous cases. This original indexing method allows for an accurate and effective retrieval of relevant previous cases and avoids the problems of computational complexity encountered by other hybrid CBR systems in their retrieval and matching tasks (e.g. [15, 19]).

Experimental results show that the performance gain brought by the CBR system is significant. The CBR system improves the failure rate of partially incorrect models without overwhelming the operator with potential diagnoses. These results show that considering CBR as a way to improve an existing search method is a valid approach. This solidifies the results presented in [10]. The power of Golding and Rosenbloom's method as well as ours comes mostly from the indexing schema provided by the other problem-solving algorithm. Cases are indexed by relevant pieces of knowledge, organized in a hierarchical manner. The other problem-solving algorithm (be it rules triggering or a structural isolation process) can be seen as an indexing schema that extensively uses background knowledge. Both systems can therefore be seen as constrained instances of explanation-based indexing [1].

The CBR component does not depend on the device nor on which type of model is used. The structural decomposition of the device is device dependent. Its completion or correction by the CBR component is not. In fact, this hybrid approach can also be incorporated in other model-based approaches to diagnosis, including first-principle approaches based on correct behavior device models. This is a definite advantage, both in the framework of the genericity of the ADMS as a diagnostic system and for the applicability of CBR as a complement to other problem-solving techniques. The Reactor Building Ventilation System's model is an acyclic digraph (instead of a tree for the Fairing Servicing Subsystem). We have found that both the indexing and the matching apply to the Fairing Servicing Subsystem are directly applicable to the Reactor Building Ventilation System's model, which also includes the possibility of failing sensors.

It is clear that this approach is applicable to a whole range of searching and problem-solving methods. An interesting extension to our work would be to apply this hybrid approach to other domains such as natural language parsing and understanding, planning, or game playing, where hierarchies and context information expressed by context functions, equivalent to the sensor functions, are readily available.

Cases could also effectively be used in diagnosing multiple faults. We have not tested this aspect of our approach yet. If a fault is confirmed by a case X, another case Y could also be retrieved that has case X's

final diagnosis as one of its potential, or final diagnoses. The final diagnosis of Y could be examined as a potential diagnosis for the current situation, potentially uncovering a double fault situation. Multiple faults could therefore be diagnosed using the same hybrid approach with no added complexity.

This paper describes how a hybrid model-based/case-based methodology permits the relaxation of the completeness and consistency constraints imposed by the model-based diagnosis approach, and helps overcome shortcomings in human capabilities. We show how CBR can guide the human operator in the last phase of the diagnostic process, using previous experiences indexed by the state of the device. The paper also contributes to the area of CBR, by showing that CBR is well suited to applications where it is combined with an already existing, but imperfect, method or paradigm for problem-solving.

References

- [1] R. Barletta and W. Mark. Explanation-based indexing of cases. In *Proceedings of the 6th National Conference on Artificial Intelligence*, pages 541–546, 1988.
- [2] L. Console, L. Portinale, and D. Theseider Dupre. Focusing abductive diagnosis. *AI Communications*, 4(2/3):88–97, 1991.
- [3] J. de Kleer. Diagnosis with behavioral modes. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 1324–1330, Detroit, 1989.
- [4] J. de Kleer. Using crude probability estimates to guide diagnosis. *Artificial Intelligence*, 45(3):381–391, 1990.
- [5] J. de Kleer. Optimizing focusing model-based diagnosis. In *Proceedings of the 3rd International Workshop on Principles of Diagnosis*, pages 26–29, Rosario, Washington, 1992.
- [6] M. P. Féret and J. I. Glasgow. Generic diagnosis for mechanical devices. In *Proceedings of the 6th International Conference on Applications of Artificial Intelligence in Engineering*, pages 753–768, Oxford, UK, July 1991. Computational Mechanics Publications, Elsevier Applied Science.
- [7] M. P. Féret and J. I. Glasgow. Case-based reasoning in model-based diagnosis. In *Proceedings of the 7th International Conference on Applications of*

- Artificial Intelligence in Engineering*, pages 679–692, Waterloo, Canada, July 1992. Computational Mechanics Publications, Elsevier Applied Science.
- [8] M. P. Féret, J. I. Glasgow, D. Lawson, and M. A. Jenkins. An architecture for real-time diagnosis systems. In *Proceedings of the Third International Conference on Industrial and Engineering Applications and Expert Systems*, pages 9–15, Charleston, SC, July 1990.
- [9] G. Friedrich. Theory diagnoses: A concise characterization of faulty systems. In *Proceedings of the 3rd International Workshop on Principles of Diagnosis*, pages 117–131, Rosario, Washington, 1992.
- [10] A. R. Golding and P. S. Rosenbloom. Improving rule-based systems through case-based reasoning. In *Proceedings of the 9th National Conference on Artificial Intelligence*, pages 22–27. AAAI Press, MIT Press, July 1991.
- [11] F. Gomez and B. Chandrasekaran. Knowledge organization and distribution for medical diagnosis. *IEEE Transactions on Systems, Man and Cybernetics*, 11:34–42, 1981.
- [12] T. Govindaraj and Y. L. Su. A model of fault diagnosis performance of expert marine engineers. *International Journal on Man Machine Studies*, 29:1–20, 1988.
- [13] K. J. Hammond. Chef. In Riesbeck C. and Schank R., editors, *Inside Case-Based Reasoning*. Lawrence Erlbaum Associates, 1989.
- [14] J. L. Kolodner. Improving human decision making through case-based decision aiding. *AI Magazine*, 12(2):52–68, 1991.
- [15] P. Koton. Reasoning about evidence in causal explanations. In *Proceedings of AAAI-88*, pages 256–261, 1988.
- [16] T. P. Martin, J. I. Glasgow, M. P. Féret, and T. G. Kelley. A knowledge-based system for fault diagnosis in real-time engineering applications. In *Proceedings of DEXA'91 - International Conference on Database and Expert System Applications*, pages 287–292, Berlin, Germany, August 1991.
- [17] R. Milne. Strategies for diagnosis. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-17(3):333–339, 1987.
- [18] I. Mozetic. Reduction of diagnostic complexity through model abstractions. In *Proceedings of the 1st International Workshop on Principles of Diagnosis*, pages 102–111, Stanford, CA, July 1990, 1990.
- [19] S. A. Rajamoney and H. Y. Lee. Prototype-based reasoning: An integrated approach to solving large novel problems. In *Proceedings of the 9th National Conference on Artificial Intelligence*, pages 34–39, Anaheim, CA, July 1991. AAAI Press, MIT Press.
- [20] M. Redmond. Distributed cases for case-based reasoning; facilitating use of multiple cases. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-90)*, Boston, MA, 1990. Morgan Kaufmann.
- [21] J.A. Reggia, D.S. Nau, and P.Y. Wang. A formal model of diagnosis inference. *Information Sciences*, 37:227–256, 1985.
- [22] R. Rymon. A final determination of the complexity of current formulations of model-based diagnosis (or maybe not final). Technical Report MS-CIS-91-13, LINC LAB 194, Department of Computer and Information Science, School of Engineering and Applied Science, University of Pennsylvania, Philadelphia, PA 19104-6389, 1991.
- [23] V. Sembangamoorthy and B. Chandrasekaran. Functional representation of devices and compilation of diagnostic problem-solving systems. Technical Report Tech. Rep., Ohio State University, Columbus, Ohio, 1985.
- [24] S. Slade. Case-based reasoning: A research paradigm. *AI Magazine*, 12(1):42–55, 1991.
- [25] P. Slovic, B. Fischhoff, and S. Lichtenstein. Behavioral decision theory. *Annual Review of Psychology*, 28:1–39, 1977.
- [26] W. C. Yoon and J. M. Hammer. Deep-reasoning fault diagnosis: An aid and a model. *IEEE Transactions on Systems, Man, and Cybernetics*, 18(4):659–675, 1988.