
Decision-Theoretic Case-Based Reasoning

John S. Breese David Heckerman

Microsoft Research
One Microsoft Way
Redmond, WA, 98052-6399
(breese@microsoft.com,heckerma@microsoft.com)

Abstract

We describe a decision-theoretic methodology for case-based reasoning in diagnosis and troubleshooting applications. The system utilizes a special-structure belief network to represent diagnostic cases, with nodes representing issues, causes, and symptoms. Dirichlet distributions are assessed at knowledge acquisition time to indicate the strength of relationships between variables. During a diagnosis session, a relevant subnetwork is extracted from a belief-network database that describes a very large number of diagnostic interactions and cases. The constructed network is used to make recommendations regarding possible repairs and additional observations, based on an estimate of expected repair costs. As cases are resolved, observations of issues, causes, symptoms, and the success of repairs are recorded. New variables are added to the database, and the probabilities associated with variables already in the database are updated. In this way, the inferential behavior of system adjusts to the characteristics of the target population of users. We show how these elements work together in a cycle of troubleshooting tasks, and describe some results from a pilot system implementation and deployment.

1 Introduction

Most model-based [de Kleer and Williams, 1987] and probabilistic [Heckerman et al., 1992] methods for automated diagnosis use prespecified domain models as fundamental knowledge representations. The inference techniques used in these methods are powerful, but they rely on the availability of robust and relatively complete characterizations of potential problems, and have no explicit mechanism for feedback and learning from experience. Case-based reasoning (CBR) is

a method for problem solving that uses a database of previously encountered problem solving incidents as its core representation. The presumption in CBR is that past problem solving behavior is the best predictor of future problems and solutions. In contrast to other diagnostic approaches, CBR typically utilizes less powerful reasoning methods, but is more responsive to the problem solving environment and less modeling intensive. In this paper, we present a method that uses decision-theoretic and statistical techniques to provide normative inference and updating techniques, while retaining the case-driven modeling that is characteristic of CBR.

Typically, CBR systems provide mechanisms to (1) index into the database of incidents to retrieve potentially relevant cases, (2) apply previous cases to the solution of the current problem instance, and (3) modify and update the case database to reflect new information and results from the current problem solving incident. In the literature, many techniques for representing and indexing cases, generating solutions, and updating the database have been developed [Kolodner, 1993, Allen, 1994].

In the approach described here, the database of previous cases is represented as a belief network [Pearl, 1988]. The initial belief network is constructed manually from an expert based on a set of past problem solving experiences. The parameters in the belief network are encoded as Dirichlet distributions. At consultation time, input text describing the current problem is used to index into the relevant portion of the belief network, corresponding to Step 1 in the standard CBR process. The current implementation uses a probabilistic information-retrieval similarity metric to generate a list of potentially relevant symptoms and causes. Based on the user's selection, the relevant portion of the belief network is constructed.

Probabilistic inference methods are then used to calculate the probabilities of various problems, and cost analysis is used to suggest possible repairs and information gathering actions. We refer to the result as a troubleshooting plan. This portion of the consultation

uses the current means of the Dirichlet distributions stored in the belief network, and corresponds Step 2 in the standard CBR process.

Once a case is resolved, the belief network database is updated to reflect the new case. If the attributes observed in the case match those represented in the belief-network database, then the relevant Dirichlet distributions in the network are updated to reflect the current observations. This allows the belief-network parameters in the database to adjust to the characteristics of the target population. If attributes or observations in the current case do not match entities in the current belief network database, then the database is extended to include those variables with an initial set of human assessed parameters. These updating procedures correspond to Step 3 in CBR.

In the remainder of the paper, we describe the structure of the belief-network database (Section 2). Then, we outline a cycle of decision-theoretic troubleshooting, including steps for database access, model construction, repair planning, and case recording (Section 3). Finally, we describe experiences with a pilot implementation of these methods in Section 4.

2 Belief Network Structure

We have developed a special-structure belief network for case-based diagnostic applications. The basic organizing structure consists of nodes corresponding to causes, issues, and symptoms.

1. **Cause:** A contributing factor, activity, or configuration item—for example, “Gateway 4DX2-66” or “Printing Postscript File.”
2. **Issue:** A conflict among a set of causes. Typically, the issue is defined to occur if all the associated causes occur—for example, “Gateway 4DX2-66V” and “BusLogic BT-445S SCSI/VL-Bus Adapter” and “Windows for Workgroups 3.11.”
3. **Symptoms:** A particular behavior or malfunction caused by an issue—for example, “System will not boot,” or “Cannot connect to network server.”

The basic relationships among causes, issues, and symptoms in the belief network structure is shown in Figure 1(a) for a single issue. Issues may share common causes and symptoms as shown in Figure 1(b). For convenience, each issue is constrained to have at least one cause and at least one symptom.

Each cause is labeled as to whether it is fixable or unfixable. A fixable cause is one that can be altered in the course of troubleshooting. For example, in Figure 1(a), “Using MS/Style mouse port” is fixable with

the action (or resolution) of “Switch to the PS/2 style mouse port.” A cost is associated with each fix. In our formulation, an unfixable node corresponds to a fixable node with a very high fix cost. For example, the cause node “Windows NT 3.1” is treated as unfixable in the current model, though, in principle, one could switch operating systems at very high cost.

Symptoms and causes are labeled as being observable or unobservable. An observable node is one that the user can reasonably test or observe. For example, all the causes and symptoms in Figure 1(a) are observable. We assign a cost of observation to each observable in the database. Some causes—for example, defective hardware—may be unobservable yet fixable by replacing the component. In our formulation, symptoms are always observable.

All nodes in the belief-network database correspond to discrete variables. The uncertainties associated with the parameters in the belief-network database are encoded as Dirichlet distributions—in particular, means (probabilities for the next case) and equivalent sample sizes. The particular parameter means are as follows.

1. $\Pr(C_i)$ is the probability that cause i is present in a given case.
2. $\Pr(S_k|I_j, \bar{I}_i, \forall i \neq j)$ is the probability of symptom k given only issue j is true.
3. $\Pr(S_k|\bar{I}_j)$ is the “leak” probability of symptom k given issue j is false.

In this formulation, we assume that there is a “noisy OR” relationship [Pearl, 1988] between each symptom and the issues that cause it. This assumption makes both model assessment and inference more efficient [Heckerman, 1993, Heckerman and Breese, 1994]. Updating of the Dirichlet distributions is described in Section 3.5. Leak probabilities are not updated in the current scheme.

3 The Troubleshooting Cycle

The troubleshooting cycle is demand driven and consists of five basic steps:

1. **Database Access:** The text describing the problem is used to find potentially relevant symptoms or causes. One such variable is selected by the user to initiate the session.
2. **Belief-Network Construction:** The relevant portions of the belief network database are extracted for the following step.
3. **Belief-Network Solution:** Using the constructed belief network, the system generates recommendations for components to repair, and makes suggestions regarding additional observations.

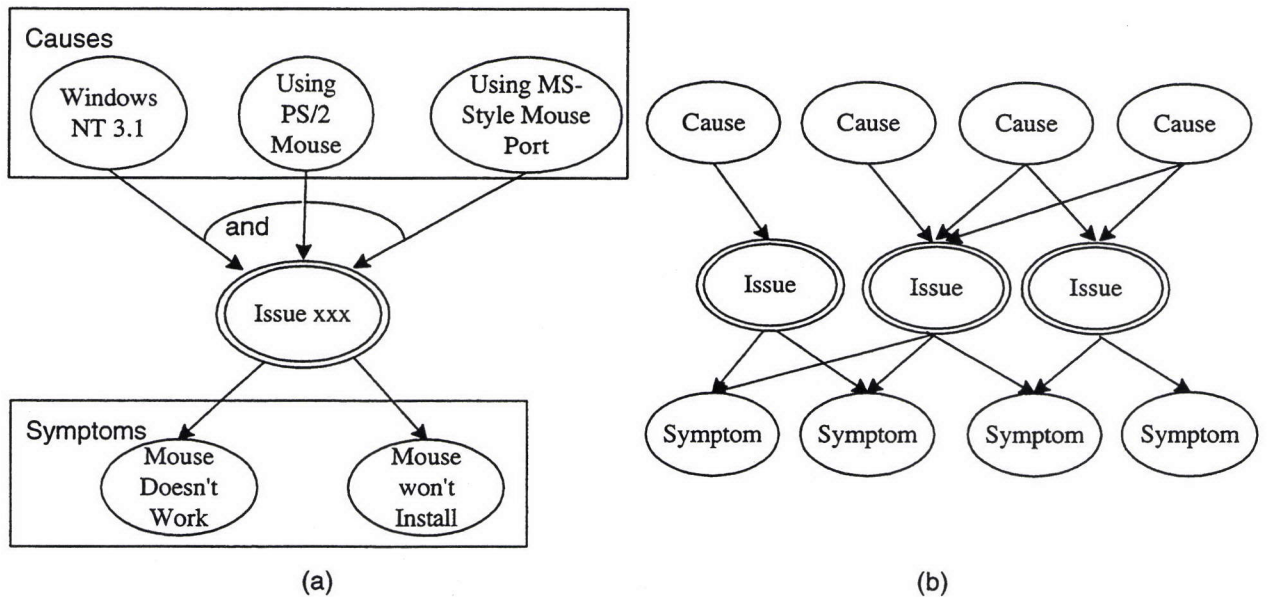


Figure 1: (a) The belief-network structure for a single issue, consisting of causes and symptoms. (b) The belief network structure for the case database, consisting of an interdependent set of issues.

4. User Execution: The user performs one or more repairs or observations, and reports the results to the system. The construction-solution-execution steps are repeated, until the problem is solved or the user suspends the session.
5. Case Recording: Upon termination, relevant Dirichlet distributions are updated. If variables not in the belief network database were observed, these can be added to the domain model, expanding the scope of issues covered.

In this section we describe each of these steps.

3.1 Database Access

A consultation is initiated by describing the problem with free text input. In the current test database, each cause, symptom and resolution is accompanied by a textual description. A full-text information-retrieval algorithm, similar to that proposed by Salton [Salton et al., 1994], is used to identify those causes or symptoms in the case database whose descriptions most closely match the textual problem description. The standard information-retrieval methods have been extended to deal with domain specific jargon and acronyms.

3.2 Model Construction

As observations are made, the system identifies the portions of the belief-network that are relevant to the observations, constructing a belief-network that is a subset of the database. In this subset, we include all

nodes that are dependent on the current set of observed nodes. In addition, because we are going to use the model to generate recommendations regarding questions to ask (i.e. perform value-of-information analysis), we also include portions of the network that would become relevant if we were to make observations.

Because we can observe any set of observable nodes in the model, the entire database is always potentially relevant to the current problem-solving task. However, as described in Section 3.3, we are using the constructed belief network in a myopic value-of-information analysis that assumes that the diagnostician will make just one observation at a time. Therefore, we need to incorporate only those portions of the network made relevant by a single observation.

This construction rule is not sound, in principle, because the best observation to make may be completely unrelated to the current set of observations. This may occur, for example, if some problem is so likely to occur *a priori* that one should always inquire about that possibility, irrespective of the currently reported symptoms. Nonetheless, this situation occurs infrequently in many domains. Furthermore, the value of information of such variables is independent of the current consultation and may be computed off line.

We identify variables that are relevant to (i.e., dependent on) current observations using the graphical test of *d-separation*. In the following definition, a path refers to any undirected path.




Figure 2: Three belief network fragments illustrating d-separation. (a) Observing a cause on a path between two subnetworks *blocks* that path. (b) Observing a symptom on a path between two subnetworks *activates* that path. (c) Observing a successor of an issue on a path between two subnetworks *activates* that path.

Definition (d-separation, Pearl, 1988) If X , Y , and Z are three disjoint subsets of nodes in a directed acyclic graph \mathcal{I} , Z is said to **d-separate** X from Y , written $\langle X, Z, Y \rangle_D$, if there is no path between a node in X and a node in Y such that the following two conditions hold: (1) every node with converging arrows is in Z or has a descendant in Z and (2) every other node is outside Z .

A path that meets conditions (1) and (2) is said to be *active*. A path that does not meet these conditions is said to be *blocked*. We illustrate the use of d-separation as a graphical indicator of conditional independence in Figure 2. In Figure 2a, we have $\langle I, \{C'\}, II \rangle_D$ but it is not the case that $\langle I, \{\}, II \rangle_D$. That is, observing C' blocks all paths between the left and right subnetworks. In Figure 2b and c, we have $\langle I, \{\}, II \rangle_D$ but not $\langle I, \{S'\}, II \rangle_D$. That is, observing S' activates all paths between the left and right subnetworks. Thus, if two portions of the network are separated by a common cause, then observing that cause means that the probabilities of nodes in one subnetwork do not depend on observations in the other subnetwork and vice-versa. Conversely, observing a symptom that separates two subnetworks (or observing a symptom that is a successor of an issue that separates two subnetworks) means that the variables in the subnetworks are dependent on one another.

The model construction algorithm proceeds as follows.

1. From currently observed causes and symptoms, add all nodes and arcs on active paths from all observed nodes. This procedure stops exploring a path when either:
 - (a) The path terminates at an unobserved symptom, or

- (b) The path has a segment of the form $I \rightarrow S \leftarrow I'$ and S is unobserved, or
- (c) The path has a segment of the form $C \rightarrow I \leftarrow C'$ and the successor symptoms of I are unobserved.

2. For each unobserved symptom in the constructed network, add all of its ancestors, if not already in the network.
3. End.

This process is illustrated in Figure 3 for an example belief-network database. Note that additional portions of the database would become relevant only if we observed more than one additional symptom. Therefore, for purposes of myopic value-of-information, the currently constructed network is sufficient.

This algorithm can be viewed as a repeated application of an efficient method for finding sets of d-separated nodes developed by Geiger et al. (1990). It also can be viewed as constructing a graph that is equivalent to the union of a set of networks that would be constructed by the knowledge-based model construction algorithm developed by Breese, where each symptom in the network constructed in Step 1 is provisionally observed [Breese, 1992].

3.3 Belief-Network Solution

The objective at this stage of problem solving is to generate a set of recommendations for repairs and observations based on the case at hand. In this paper, we adopt a scheme developed by Heckerman et al. (1994, 1995). In the following discussion, we first consider recommending just repair actions. We then address the possibility of observing additional symptoms and causes to focus troubleshooting.




Figure 3: (a) The implicit database network structure. The shaded nodes are the initial observations. (b) The constructed network after Step 1 has completed. Both the original and added nodes are shaded. (c) The final constructed network consists of the shaded nodes, where all ancestors of symptoms are added to the previous network.

Given a constructed network containing several possible fixable nodes, what order of repairs should be attempted to repair the device or system? In order to determine this sequence, we assume that there is a single *problem-defining* symptom, repair costs are independent, and each fixable cause must be observed to be abnormal immediately prior to its repair. For our constructed network, one of the observed symptoms will be designated as the problem defining node.

Let C_i^o and C_i^r denote the cost of observation and repair of a fixable cause c_i , respectively. Let p_i be the probability that c_i is in a faulty state given current evidence, calculated using the belief network constructed as described in Section 3.2. If we observe and possibly repair components in the order c_1, \dots, c_n , then the expected cost of repairing the system, denoted ECR is

$$\begin{aligned} \text{ECR} &= (C_1^o + p_1 C_1^r) \\ &\quad + (1 - p_1) \left(C_2^o + \frac{p_2}{1 - p_1} C_2^r \right) \\ &\quad + (1 - p_1 - p_2) \left(C_3^o + \frac{p_3}{1 - p_1 - p_2} C_3^r \right) \\ &\quad + \dots \\ &= \sum_{i=1}^n \left[\left(\sum_{j=i}^n p_j \right) C_i^o + p_i C_i^r \right] \end{aligned} \quad (1)$$

That is, we first observe cause c_1 incurring cost C_1^o . With probability p_1 , we find that the cause is faulty and repair it incurring cost C_1^r . With probability $1 - p_1$, we find that the cause is normal, and observe cause c_2 incurring cost C_2^o . With probability $p_2/(1 - p_1)$, we find that c_2 is faulty and repair it; and so on.

Heckerman et al. (1995) show that the optimal troubleshooting sequence under these assumptions is given by observing and possibly repairing causes in descending order of the ratio p_i/C_i^o . Ties may be broken arbitrarily.¹

This ordering assumes that there are no additional observations interleaved in the repair sequence. The additional observations, which we call *nonbase observations*, provide information about the system without changing the state of the system (i.e., without repairing causes). In order to estimate the value of these observations, we evaluate the expected cost of repair as if we can make at most one nonbase observation before executing a plan consisting of only observation–repair actions as described above.

This evaluation proceeds as follows. We assume that each observable o_i can take on exactly one of r_i possible states. We write $o_i = k$ to indicate that obser-

vation o_i takes on state k . Let \mathbf{E} denote the current information state of the troubleshooter. \mathbf{E} may include information about previous observations as well as repairs. First, we use Equation 1 to compute the expected cost of repair for the plan where no nonbase observation is made, denoted $\text{ECR}(\mathbf{E})$.

Then, for one possible nonbase observation o_i , we determine the best observation–repair sequence for every possible outcome of o_i . For each possible outcome, we again use Equation 1 to compute the expected cost of repair of the best sequence, denoted $\text{ECR}(\mathbf{E} \cup \{o_i = k\})$. Next, we compute $\text{ECO}(\mathbf{E}, o_i)$, the overall expected cost of first observing o_i and then executing the observation–repair sequence:

$$\text{ECO}(\mathbf{E}, o_i) = C_i^o + \sum_{k=1}^{r_i} \Pr(o_i = k | \mathbf{E}) \text{ECR}(\mathbf{E} \cup \{o_i = k\}) \quad (2)$$

Note that the conditional troubleshooting sequence following the observation may be different for every possible outcome of the observation. We repeat the computation of ECO for every possible nonbase observation.

If $\text{ECR}(\mathbf{E}) < \text{ECO}(\mathbf{E}, o_i)$ for any nonbase observation o_i , then we choose not to make a nonbase observation and recommend the first action in the observation–repair sequence encoded in $\text{ECR}(\mathbf{E})$. Otherwise, we recommend to observe that symptom or cause o_i with the lowest ECO. After a repair or nonbase observation has been carried out, we update the information state \mathbf{E} . If a repair action has been carried out, we must remove all observed symptoms that correspond to nodes that are descendants of the repair-action node, because the repair action may have caused these observations to change.

Although we may make additional nonbase observations in the actual troubleshooting sequence, we ignore these possible actions when selecting the next nonbase observation. Similarly, although the repair sequences are generated assuming a single fault, the iterative procedure allows us to deal with multiple fault scenarios. We have found this myopic approximation to yield good results in real-world problems [Heckerman et al., 1995].

3.4 Execution

Given the recommendations, the user performs one or more repairs or observations, and reports the results to the system. As additional variables are observed, it is necessary to reinvoke the model construction algorithm to extend the set of relevant causes, issues, and symptoms. Note that since repair actions may invalidate observations, the model could contract. Because our observation valuation procedure assumes a single additional observation at a time in the evaluation of

¹This procedure can be modified to include a recommending a service call, i.e. some more sophisticated troubleshooting agent. See Heckerman et al. (1995) for more details.

Equation 2, the model construction algorithm is guaranteed to produce the correct probabilities. If we extended the myopic procedure to include k -step lookahead, then the model construction algorithm would need to similarly be extended.

3.5 Case Recording

During a session, variables represented in the database and new variables may have been observed. At the end of a consultation, the Dirichlet distributions associated with observed nodes in the database are updated. For simplicity, we only update a distribution if both its corresponding node and the parents of that node were observed.

If a new variable was observed, then a corresponding node is added to the database, and arcs and Dirichlet distributions are assessed. If a new issue was identified, several new causes and symptoms may also be added to the database.

In order to account for dynamics in the population of problems encountered, a variable length window can be used for Dirichlet updating, where we only include the last k observations in updating the distribution. For example, a short window might be appropriate for new product introductions or revisions where we want to discount the effects of previous observations, while a longer window is appropriate for stable products and user populations.

4 Implementation

A pilot project codenamed "Aladdin," which uses the techniques described in this paper, was undertaken at Microsoft during the summer of 1994. A belief network database, consisting of approximately 200 issues, 250 causes, and 220 symptoms, was developed for warranty problems for Microsoft Windows NTtm Version 3.1 operating system. The pilot advisory system was used by Microsoft support personnel for a two-month period. The implemented version used a somewhat different set of heuristics for repair and observation recommendations than reported here. Each issue was treated independently and information gathering recommendations were based on an entropy-style information measure.

Overall satisfaction with the tool by pilot participants was very high in terms of speed of resolution and style of interaction. Unfortunately, the pilot knowledge base did not have sufficiently wide coverage for the range of issues actually encountered in on-line use, and therefore was not used extensively. Although Dirichlet parameter updating and model extension facilities were implemented and deployed, empirical results regarding effectiveness are not yet available.

In order to investigate the utility of the techniques

while controlling for coverage, we developed an experiment comparing the Aladdin system to the currently available information tools provided to support engineers. In these experiments, the correct solutions were available in both tools. The results in these test indicate that engineers using Aladdin could support products for which they had little or no training. In comparison to standard methods, the probabilistic case-based reasoning approach was faster and resulted in fewer unsolved problems, as well as being preferred by engineers.

5 Summary

We have developed a design for a case-based reasoning system that uses a belief network for runtime problem solving. The primary components of the system are an information retrieval mechanism to index into a belief-network database, a data-driven model-construction algorithm, a method for generating recommendations for observations and repairs from a belief network, and a method for updating the belief-network database as new cases are encountered. Preliminary results suggest that the method is useful as a troubleshooting aid.

Acknowledgments

The authors thank Eric Horvitz and Greg Shaw for useful comments and assistance on this work.

References

- [Allen, 1994] Allen, B. (1994). Case-based reasoning: Business applications. *Communications of the ACM*, 37(3):40-42.
- [Breese, 1992] Breese, J. S. (1992). Construction of belief and decision networks. *Computational Intelligence*, 8(4):624-648.
- [de Kleer and Williams, 1987] de Kleer, J. and Williams, B. (1987). Diagnosing multiple faults. *Artificial Intelligence*, 32(1):97-130.
- [Geiger et al., 1990] Geiger, D., Verma, T., and Pearl, J. (1990). Identifying independence in Bayesian networks. *Networks*, 20:507-534.
- [Heckerman, 1993] Heckerman, D. (1993). Causal independence for knowledge acquisition and inference. In *Proceedings of the 9th Conference on Uncertainty in Artificial Intelligence*, pages 122-127. Morgan Kaufmann.
- [Heckerman and Breese, 1994] Heckerman, D. and Breese, J. (1994). A new look at causal independence. In *Proceedings of the 10th Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann.

- [Heckerman et al., 1994] Heckerman, D., Breese, J., and Rommelse, K. (1994). Sequential troubleshooting under uncertainty. Technical Report MSR-TR-94-07, Microsoft Research, Redmond, WA.
- [Heckerman et al., 1995] Heckerman, D., Breese, J., and Rommelse, K. (1995). Troubleshooting under uncertainty. *Communications of the ACM*. To appear in the Special Issue on Uncertain Reasoning.
- [Heckerman et al., 1992] Heckerman, D., Horvitz, E., and Nathwani, B. (1992). Toward normative expert systems: Part I. The Pathfinder project. *Methods of Information in Medicine*, 31:90–105.
- [Kolodner, 1993] Kolodner, J. (1993). *Case-Based Reasoning*. Morgan Kaufman, San Mateo, Ca.
- [Pearl, 1988] Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufman, San Mateo, Ca.
- [Salton et al., 1994] Salton, G., Allan, J., and Buckley, C. (1994). Automatic structuring and retrieval of large text files. *Communications of the ACM*, 37(2):97–108.