Finding Explanations in Bayesian Networks

Changhe Yuan
Department of Computer Science and Engineering
Mississippi State University
Mississippi State, MS 39762
cyuan@cse.msstate.edu

Tsai-Ching Lu
HRL Laboratories, LLC
Malibu, CA, 90265
tlu@hrl.com

Abstract

Maximum a Posteriori (MAP) and Most Probable Explanation (MPE) are popular approaches of finding explanations for given observations in Bayesian networks. One of their shortcomings is that they have to find a complete assignment to a set of target variables. However, it is often the case that only few of the target variables are most relevant in explaining the observations. In this paper, we formulate the problem of finding the most relevant variables as *explanatory* MAP (eMAP), which considers all subsets of the target variables and finds a partial assignment that maximizes a chosen quality measure. We consider two quality measures for eMAP: Bayes Factor (or Weight of Evidence) and likelihood function. We then illustrate the proposed methodology on a circuit diagnosis problem in literature.

Introduction

Bayesian networks (BNs) (Pearl 1988) offer a compact and intuitive graphical representation of uncertain relationships among random variables in a domain. They have proven their value in many disciplines during the last two decades, including a variety of decision problems in medical diagnosis, prognosis, therapy planning, machine diagnosis, user modeling, natural language interpretation, planning, vision, robotics, data mining, fraud detection, and many others. Bayesian networks provide principled approaches for finding the most likely state of a system given new observations, typically formulated as Maximum a Posteriori (MAP) or Most Probable Explanation (MPE) problems. MAP is defined as follows. Let $\mathbf{X}$ be the set of MAP nodes whose most probable configuration is of our interest. Let $\mathbf{E}$ be the set of evidence nodes and $\mathbf{e}$ be their states. The remainder of the nodes is denoted by $\mathbf{Y}$. Then, the standard MAP query $MAP(\mathbf{X}, \mathbf{e})$ is defined as

$$MAP(\mathbf{X}, \mathbf{e}) \triangleq \hat{\mathbf{x}} = \arg\max_{\mathbf{x}} \sum_{\mathbf{y}} P(\mathbf{X}, \mathbf{Y}|\mathbf{e}). \quad (1)$$

MPE is defined analogously except that $\mathbf{Y}$ is empty. In other words, MPE is interested in all the unobserved variables. Both MAP and MPE return a complete assignment to $\mathbf{X}$ as the best explanation for the evidence. In diagnostic setting, $\mathbf{X}$ typically contains fault variables in a system. The number of fault variables can be as large as several hundreds in a large model. In such case, even the best solution by MAP or MPE may have a very low probability. It is doubtful whether such an unlikely event constitutes a good explanation.

In a real world problem, it is often the case that only few of the target variables are most relevant in explaining the evidence. Take a medical domain as an example. When a patient goes to the hospital with certain symptoms, a physician can typically identify the only few diseases that the patient may have. All other diseases are regarded as unrelevant and ruled out from further consideration. Expert knowledge in this case helps in making the initial diagnosis. The question is whether we can formulate the problem formally so that we can develop computational solutions for solving it.

In this paper, we formulate the problem of identifying the most relevant target variables as *explanatory* MAP (eMAP), which considers all subsets of the target variables and finds a partial assignment that maximizes a chosen quality measure. Instead of restricting to any single measure, we compare two choices: Bayes factor (or Weight of Evidence) and likelihood function. We illustrate the method on circuit diagnosis problem in literature by comparing against several existing approaches.

Related Work

Besides MAP and MPE, there are several other existing approaches for finding explanations in Bayesian networks. We briefly review some of the approaches in this section.

Shimony (1993) defines explanations in Bayesian networks as the most probable independence-based assignment that is complete and consistent with respect to the evidence nodes. Roughly speaking, an explanation is a truth assignment to the variables "relevant" to evidence nodes. Only ancestors of evidence nodes can be relevant. An ancestor of a given node is irrelevant if it is independent of that node given the values of other ancestors.

Henrion and Druzdzel (1991) assume that the system has a set of pre-specified scenarios as potential explanations. Although they do not require the truth values of all propositions be specified, they order the explanations based on the posterior probability of the scenarios.

de Campos et al. (de Campos, Gamez, & Moral 2001) proposed yet another approach for finding explanations in

Bayesian networks. Their approach is a two-step procedure. They first solve a K-MPE problem and find the K most probable explanations in a Bayesian network given the evidence. Then, they simplify the explanations by greedily removing unimportant variables one by one. If the removal of a literal does not reduce the likelihood of an explanation or only reduce the likelihood within a certain factor, the literal is regarded as unimportant and can be removed from the explanation.

There are also some approaches that are not specific to Bayesian networks. Gärdenfors (1988) defines that X is an explanation of E relative to a state of belief $K =< W, P >$, where W is a set of possible worlds and P is a distribution on W and $E \in K$, if

1. $P_E^-(E|X) > P_E^-(E)$, and
2. $P(X) < 1 (i.e., X \notin K)$.

Gärdenfors orders all legitimate explanations using *explanatory power* (EP). The EP of X with respect to E is defined as

$$EP(X, E) = \frac{P_E^-(E|X)}{P_E^(E)}. \quad (2)$$

Since $P_E^-(E)$ is a constant for given E, EP essentially orders the explanations according to the likelihood of evidence given the explanations.

Chajewska and Halpern (1997) assume that a world is a pair (w, C) consisting of a truth assignment w and a causal structure C. They also assume that the explanandum E is one of the observations O and define P_E^- as follows:

$$P_E^-(w, C) = P'(C) P_C(w|O - \{E\}). \quad (3)$$

They place only partial orderings on explanations based on a pair of numbers $(EP(X, E), P_E^-(X))$: the explanatory power and the prior of X.

All these approaches realize the need to find variables that are most relevant to the evidence. However, as we will show later, they may still find explanations that contain too many variables, some of which may be unimportant. We propose a new approach for the problem in the next section.

Explanatory MAP

Based on the assumption that typically only few of the target variables are most relevant in explaining the evidence in a system, we formulate identifying the most relevant fault variables as *explanatory* MAP query $eMAP(\mathbf{X}, \mathbf{e})$ which takes the same input as MAP problem: the set of MAP nodes $\mathbf{X}$, the set of evidence nodes $\mathbf{E}$, and the remaining nodes $\mathbf{Y}$, and solves the following optimization problem.

$$eMAP(\mathbf{X}, \mathbf{e}) \triangleq \hat{\mathbf{x_k}} = \arg\max_{\mathbf{x_k}, \mathbf{X_k} \subseteq \mathbf{X}} f(\mathbf{x_k}, \mathbf{e}), \quad (4)$$

where $f(\mathbf{x_k}, \mathbf{e})$ is an evaluation function that measures the degree of a partial assignment to K variables $\mathbf{X_k}$ in $\mathbf{X}$ in explaining $\mathbf{e}$. In other words, we are searching for a solution simultaneously over the subspaces of potential MAP variables and their corresponding configurations. Suppose there are n MAP variables and each has d states, the size of the search space of eMAP would be $\sum_{i=1}^{n} \binom{n}{i} d^i$, which is much larger than d^n of MAP.

One element missing from the formulation is the evaluation function f. A naive selection for f is the posterior probability of the partial assignment, which is obviously not a viable choice due to the fact that the probabilities in comparisons are not at the same footing. Assignments with fewer variables always dominate those with more variables. In this paper, we compare two quality measures: likelihood function and Bayes factor (Good 1985). Both likelihood and Bayes factor provide "common" grounds for comparing all the partial assignments of $\mathbf{X}$.

Likelihood Function

The first choice for f is the likelihood function $P(\mathbf{e}|\mathbf{X_k})$. By using this measure, we basically formulate eMAP as a model selection problem: Each partial assignment of the MAP variables serves as a model candidate, and we want to find the model that mostly likely generates the data, in our case, the evidence. Since adding variables that are conditionally independent of evidence to an explanation does not reduce the likelihood of the explanation, we should penalize more complex explanations as in model selection. We introduce a penalty term as in AIC (Akaike 1974) for eMAP. In the general case, AIC is defined as follows:

$$AIC = -2 \ln L(M) + 2K(M),$$

where $L(M)$ is the likelihood of the data given model M and $K(M)$ is the number of parameters. It is equivalent to using $P(\mathbf{e}|\mathbf{X_k}) exp(-K)$ as the quality measure in eMAP, where K is the number of variables. $P(\mathbf{e}|\mathbf{X_k})$ is the likelihood, and $exp(-K)$ is an exponential prior penalizing the complexity of the model. However, when K is small, change in $exp(-K)$ when K changes may dominate the whole measure. In order to minimize the effect of the prior, we use $exp(-cK)$ instead, where c is small constant far less than 1. This prior is only intended to help discriminating explanations with different number of variables when their likelihood scores are similar.

Bayes Factor

Another choice for f in Eqn. 4 is Bayes factor (BF) (Good 1985). Suppose we have data D assumed to have arisen under one of the two hypotheses H_1 and H_2 according to a probability density $P(D|H_1)$ and $P(D|H_2)$. The Bayesian factor is given by

$$BF = \frac{P(D|H_1)}{P(D|H_2)}. \quad (5)$$

Bayes factor is similar to likelihood ratio, although in Bayesian setting we have to average likelihood over the parameters. The logarithm of BF is called the *Weight of Evidence* (WoE) given by D for H_1 over H_2 (Good 1985). A value of $BF > 1$ means that the data indicate that H_1 is more likely than H_2. Jeffreys (1961) gave the following scale for interpretation of BF:

In our problem, there are many explanations that compete with each other, so we use a generalized version of Bayes

BF	Strength of evidence
$< 1{:}1$	Negative (supports H_2)
1:1 to 3:1	Barely worth mentioning
3:1 to 12:1	Positive
12:1 to 150:1	Strong
$> 150{:}1$	Very strong

Table 1: Interpretation of Bayes factor (BF).

factor (Fitelson 2001)

$$BF = \frac{P(D|H_1)}{P(D|\neg H_1)}. \quad (6)$$

Using Bayes rule, we can transform the above definition into the following form:

$$BF = \frac{P(H_1|D)(1 - P(H_1))}{P(H_1)(1 - P(H_1|D))}. \quad (7)$$

In our problem, D is simply the evidence $\mathbf{e}$, and Hs are the explanations $\mathbf{x_k}$. We can use inference algorithms to get both $P(\mathbf{x_k})$ and $P(\mathbf{x_k}|\mathbf{e})$ in a Bayesian network.

Example

We illustrate our method using a concrete example, which originally appeared in (Poole & Provan 1991). We also compare our explanations against those of several existing approaches.

The Model

Consider a simple circuit as in Figure 1 (Poole & Provan 1991). A, B, C and D are gates that are faulty if they are closed, and the prior probabilities that the gates break independently are as follows:

$$P(A) = 0.016;$$
$$P(B) = 0.1;$$
$$P(C) = 0.15;$$
$$P(D) = 0.1.$$

We observe that there is current flowing through the circuit.

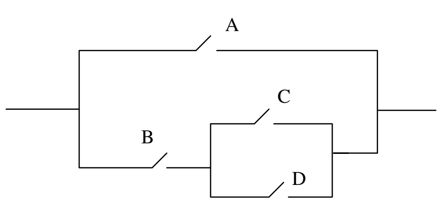

Figure 1: A circuit.

The circuit can be translated into an equivalent Bayesian network model as in Figure 2. Nodes A, B, C and D correspond to the gates in the circuit and each has two states "defective" and "ok". Node $Current$ represents the input to the circuit, with two states "yes" and "no". Other nodes are intermediate or final outputs, which can take state "current" or "ok". An intermediate output node takes state "current" if and only if one of its parent is in state "current" and the gate preceding it is "defective". The $Total\ Output$ node takes "current" if any of its parents is in state "current". The observation that current flows through the circuit is equivalent to the evidence that node $Current$ is in state "yes" and node $Total\ Ouput$ is in state "current". Given the evidence, the task is to diagnose the system and find the best explanations for the observations.

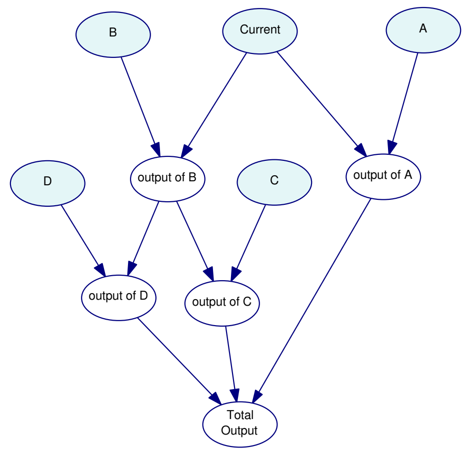

Figure 2: An equivalent BN model.

Based on our knowledge of the system, there are three scenarios which may lead to the observation:

1. A is defective;
2. B and C are defective;
3. B and D are defective;

Of course, any combination of the above three scenarios can also produce the observation. But the above three explanations are *minimal* in the sense that any explanation subsumed by them cannot fully explain the evidence. In the following sections, we compare the explanations found by multiple approaches.

To simplify notation in later discussions, we use OX to stand for variable $Ouput\ of\ X$, e.g., OA for $Output\ of\ A$. TO stands for $Total\ Output$. We use a variable and its negation to stand for the variable's two states.

Explanations by MPE and MAP

Most Probable Explanation (MPE) returns the most likely configuration of all the unobserved nodes as the explanation. The best explanation by MPE for the example is

$$P(\neg A \wedge B \wedge C \wedge \neg D \wedge \neg OA \wedge OB \wedge OC \wedge \neg OD) = 0.3395.$$

The MPE solution subsumes one of the minimal explanations asserting that B and C are defective. However,

OA, OB, OC and OD are not really necessary for the explanation, because removing them still results in a valid explanation.

The single MPE solution may not provide enough insight to the problem of a system. This problem can be addressed by solving the K-MPE problem that finds the K most probable explanations. If we set K equal to 3, K-MPE will output the following two other likely explanations.

$$P(A \wedge \neg B \wedge \neg C \wedge \neg D \wedge OA \wedge \neg OB \wedge \neg OC \wedge \neg OD) = 0.2816;$$

$$P(\neg A \wedge B \wedge \neg C \wedge D \wedge \neg OA \wedge OB \wedge \neg OC \wedge OD) = 0.2138.$$

In comparison, Maximum a Posteriori (MAP) focuses on the nodes that may be faulty. In the example model, A, B, C and D are the fault variables. The best explanation returned by MAP is

$$P(\neg A \wedge B \wedge C \wedge \neg D) = 0.3395.$$

Similarly, by allowing finding K-MAP solutions, we can also find two other likely explanations:

$$P(A \wedge \neg B \wedge \neg C \wedge \neg D) = 0.2816;$$
$$P(\neg A \wedge B \wedge \neg C \wedge D) = 0.2138.$$

MAP explanations are much more concise than those of MPE. However, for a large system with tens or hundreds of fault variables, even MAP solutions would contain too many variables. Some of the variables may not be necessary for the validity of the explanations.

Explanations By Shimony's Approach

Based on the definition of "relevance" in (Shimony 1993), $(\neg A, \neg OA, \neg TO)$ is an independence-based assignment with regard to evidence $(\neg Current, \neg TO)$, because once OA is observed at state "current", OC or OD become context specific independent (Boutilier *et al.* 1996) of TO. Such independence does not result from d-separation (Pearl 1988), which is the inherent property of the graphical structure, but rather the parameterizations of the model. Similarly, $(\neg B, \neg OB, \neg C, \neg OC, \neg TO)$ and $(\neg B, \neg OB, \neg D, \neg OD, \neg TO)$ are two other independence-based assignments. The probabilities of these three assignments are as follows:

$$P(\neg A, \neg OA, \neg TO) = 0.016;$$
$$P(\neg B, \neg OB, \neg C, \neg OC, \neg TO) = 0.015;$$
$$P(\neg B, \neg OB, \neg D, \neg OD, \neg TO) = 0.010.$$

Therefore, $(\neg A, \neg OA, \neg TO)$ is the independence-based MAP assignment in Shimony's approach.

This approach also finds an explanation that subsumes one of the three minimal explanations. Since it uses posterior probability to rank the explanations, it biases towards explanations with fewer variables and finds a more concise explanation than MPE and MAP. The drawback of this approach is that its solution may still contain unnecessary variables. As pointed out in (Chajewska & Halpern 1997), for each evidence node, the explanation must include an assignment to all the nodes in at least one path from the evidence to the root, because for each relevant node, at least one of its parents must be relevant. Furthermore, the irrelevance condition is quite strong and only achieves significant pruning in limited contexts.

Explanations by de Campos et al.'s Approach

de Campos et al.'s approach first finds the top K MPE solutions and then simplify the explanations by removing variables that do not or only slightly reduce the likelihood of evidence given the explanations. The top three MPE explanations are already listed in a previous section. It turns out that the likelihood of evidence is 1 given any of the three explanations.

After simplifying the explanations, de Campos et al.'s approach will obtain the following three explanations, which turns out to be the three minimal scenarios listed before.

$$P(e|\neg A) = 1;$$
$$P(e|\neg B, \neg C) = 1;$$
$$P(e|\neg B, \neg D) = 1.$$

One drawback of this approach is that likelihood does not distinguish among the above three explanations and treat them as equally good. Another drawback is that $P(e|X)$ does not always change monotonically as we remove variables from an explanation. Therefore, the simplifying procedure may be stuck in a local maximum. de Campos et al. recommend relaxing the criteria and allow $P(e|X)$ reduce by a small factor at each step. Another drawback, as the authors pointed out themselves, simplifying the explanations may result in fewer than K explanations.

Explanations by Gärdenfors' Approach

The probability of the evidence in this example equals 0.0391. There are many explanations that have likelihood greater than this value. Indeed, the likelihood of evidence is 1 given any of the three minimal explanations and any of the explanations subsuming one of the three minimal explanations. Since Gärdenfors orders the explanations essentially in the order of likelihood, all these explanations are regarded as equally good based on Gärdenfors' approach.

Explanations by eMAP

Finally, let us look at explanations by eMAP. We consider two quality measures separately: likelihood function and Bayes factor.

If we use likelihood function as the quality measure, eMAP will also find many explanations with the same quality, including the three minimal explanations and all explanations subsuming them. As we discussed before, we address the problem by introducing prior $P(K) \propto exp(-cK)$ to bias towards simpler models. eMAP with this prior would identify $P(e|\neg A) = 1$ as the top choice and the following explanations slightly worse.

$$P(e|\neg B, \neg C) = 1;$$
$$P(e|\neg B, \neg D) = 1;$$
$$P(e|\neg A, B) = 1;$$
$$P(e|\neg A, \neg B) = 1;$$
$$P(e|\neg A, C) = 1;$$
$$P(e|\neg A, \neg C) = 1;$$
$$P(e|\neg A, D) = 1;$$
$$P(e|\neg A, \neg D) = 1.$$

Likelihood provides no further ability to discriminate between these explanations because they all consist of the same number of variables.

Now, let us consider Bayes factor. The major difference between Bayes factor and likelihood function is that Bayes factor not only factors in the ratio between the posterior and prior probabilities but also the magnitudes of the probabilities. When using Bayes factor as the quality measure, eMAP returns the following top explanation.

$$BF(\neg A) = 42.56;$$

The explanations immediately following the above explanation are:

$$BF(\neg B, \neg C) = 40.82;$$
$$BF(A, \neg B, \neg C) = 40.43;$$
$$BF(\neg A, D) = 39.87;$$
$$BF(\neg A, B) = 39.87;$$
$$BF(\neg A, C) = 38.65;$$
$$BF(\neg B, \neg C, D) = 38.5;$$
$$...$$
$$BF(\neg B, \neg D) = 33.99;$$
$$...$$

According to Jeffreys' interpretation of BF, these explanations all enjoy strong support from the evidence. Explanation $(\neg A)$ ranks on the top because its prior and posterior probabilities are both high. Explanation $(\neg B, \neg D)$ does not rank very high in the list due to its low prior probability. In other words, Bayes factor factors in both prior and posteriori probabilities such that explanations which cannot be distinguished by likelihood measure can be ranked using Bayes factor. This is especially important in sequential diagnosis as illustrated in the next section.

Sequential Diagnosis

In sequential diagnosis, we are normally presented with some initial evidence for a defective system and start ranking diagnostic hypotheses according to how well they can explain the evidence. Next, we rank tests based on their ability to differentiate a set of focus hypotheses, so called focus faults in (Lu & Przytula 2006). We then gather more evidence by performing the selected tests to refine our hypotheses. We typically go through several iterations of information gathering and hypotheses refining until we are confident enough to conclude our final diagnosis.

For our example, given the initial evidence, we have discussed the initial ranking of hypotheses for each method in the earlier sections. In this section, we assume that our focus hypothesis will be the top hypothesis only. We also assume that we have a silver-bullet test for verifying the health of A and the test is ranked as the top test to perform. If the test result confirms that A is okay, i.e. $\Pr(A) = 1$, we can include $A =$ "ok" as our additional evidence and continue the diagnosis.

We first look at eMAP with Bayes factor measure. In the first iteration, $(\neg A)$ would be our diagnosis hypothesis. If our test results show that indeed A is faulty, we found the problem of the system. Otherwise if A is okay, we have to include the fact that A is okay as our additional evidence and continue the diagnosis. In the next step, we found the following top explanations

$$BF(\neg B, \neg C) = 115.88;$$
$$BF(\neg B, \neg C, D) = 98.65;$$
$$BF(\neg B, \neg D) = 73.33;$$
$$BF(\neg B, C, \neg D) = 66.1;$$
$$BF(\neg B, \neg C, \neg D) = 45.39.$$

Therefore, $(\neg B, \neg C)$ is our next hypothesis to test.

If we use likelihood function as the quality measure for eMAP, the first step is the same. We will test hypothesis $(\neg A)$ first because of its simplicity. If A is okay, the next step is a little problematic because likelihood cannot distinguish between explanations $(\neg B, \neg C)$ and $(\neg B, \neg D)$. Other approaches using likelihood discussed above have the same problem.

It is possible to use Most Probable Explanations (MPE) or Maximum a Posteriori (MAP) Assignments to generate multiple-fault hypotheses. Both MPE and MAP ranks the explanations containing $(\neg B, \neg C)$ as their top choice. Therefore, we need to test this hypothesis first. However, it is typically harder to test hypotheses with more faults. For comparison purpose, we assume users choose to override the recommendation and test hypothesis $(\neg A)$ first. If A is okay, the following top explanations will be generated.

$$P(\neg B \wedge \neg C \wedge D) = 0.5745;$$
$$P(\neg B \wedge C \wedge \neg D) = 0.3617;$$
$$P(\neg B \wedge \neg C \wedge \neg D) = 0.0638.$$

Shimony's approach works pretty well for this example. No matter what results testing hypothesis containing $\neg A$ yields, the ordering of the other two top explanation will not change because they are independent from the first hypothesis.

Other Applications of eMAP

eMAP has the capability to automatically identify the variables that are most relevant in explaining new observations. It has potentially a wide range of applications in diagnostic settings. Besides generating multiple-fault hypotheses, we discuss other possible applications.

First, given a diagnostic hypothesis, we typically have the opportunity to gather additional information about the state of the world before taking further actions. Value of information is a concept that captures this task (Howard 1966). Existing diagnostic systems often calculate myopic value of information and recommend a single best test to perform. However, it is possible that each single test has value of information less than its cost, but there exists some set of tests whose value of information is greater than their cost (Heckerman, Horvitz, & Middleton 1993). In this case, myopic value of information will incorrectly stop any information gathering because it believes no test is cost effective. Optimal non-myopic value of information has to consider all possible test combinations. Such an analysis is intractable unless only a small number of tests is under consideration. We can apply eMAP in this setting to select one or multiple subsets of the tests and consider different test combinations only within these test sets.

Second, eMAP can be used to perform model evaluation. After we build a model, it is important to evaluate the performance of the model before applying it to real problems. One approach is to use real cases to evaluate the model. However, real cases are typically difficult and expensive to obtain. The few cases that we manage to get are also not enough to evaluate the model thoroughly. eMAP can be applied in this setting to simulate faults and generate test cases as in (Lu & Przytula 2007). For this purpose, we can use eMAP to simulate the most probable observations given faults as our test cases. We can also use eMAP to find the most probable explanations for simulated test cases.

Conclusion and Discussion

In this paper, we propose a novel approach for finding the best explanations for new observations in Bayesian networks. Based on our analysis on an example model, eMAP demonstrates the capability to automatically identify the most relevant variables in explaining the evidence. eMAP using Bayes factor can factor in both prior and posterior probabilities properly in ranking the potential explanations and was able to distinguish between explanations that likelihood cannot distinguish. eMAP has huge potential to many diagnosis problems, such as sequential diagnosis, fault simulation, model evaluation, and value of information. In our future work, we plan to design and implement efficient exact and approximate algorithms for eMAP, and, based on that, develop methodologies for multiple-fault diagnosis.

Acknowledgements

All experimental data have been obtained using SMILE, a Bayesian inference engine developed at the Decision Systems Laboratory and available at `http://genie.sis.pitt.edu`.

References

Akaike, H. 1974. A new look at the statistical model identification. *IEEE Transactions on Automatic Control* 19(6):716723.

Boutilier, C.; Friedman, N.; Goldszmidt, M.; and Koller, D. 1996. Context-specific independence in bayesian networks. In *UAI*, 115–123.

Chajewska, U., and Halpern, J. Y. 1997. Defining explanation in probabilistic systems. In *Proceedings of the Thirteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI–97)*, 62–71. San Francisco, CA: Morgan Kaufmann Publishers.

de Campos, L. M.; Gamez, J. A.; and Moral, S. 2001. Simplifying explanations in Bayesian belief networks. *International Journal of Uncertainty, Fuzziness Knowledge-Based Systems* 9(4):461–489.

Fitelson, B. 2001. *Studies in Bayesian confirmation theory*. Ph.D. Dissertation, University of Wisconsin, Madison, Philosophy Department.

Gärdenfors, P. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press.

Good, I. 1985. Weight of evidence: A brief survey. *Bayesian Statistics* 2:249–270.

Heckerman, D.; Horvitz, E.; and Middleton, B. 1993. An approximate nonmyopic computation for value of information. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15(3):292–298.

Henrion, M., and Druzdzel, M. J. 1991. Qualitative propagation and scenario-based schemes for explaining probabilistic reasoning. In Bonissone, P.; Henrion, M.; Kanal, L.; and Lemmer, J., eds., *Uncertainty in Artificial Intelligence 6*. New York, N. Y.: Elsevier Science Publishing Company, Inc. 17–32.

Howard, R. A. 1966. Information value theory. *IEEE Transactions on Systems Science and Cybernetics* SSC2(1):22–26.

Jeffreys, H. 1961. *Theory of Probability*. Oxford University Press.

Lu, T.-C., and Przytula, K. W. 2006. Focusing strategies for multiple fault diagnosis. In *Proceedings of the 19th International FLAIRS Conference (FLAIRS-06)*, 842–847.

Lu, T.-C., and Przytula, K. W. 2007. System diagnosability analysis using p-slop MAP. In *Proceeding of the 20th International FLAIRS Conference (FLAIRS-07)*, To appear.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA: Morgan Kaufmann Publishers, Inc.

Poole, D., and Provan, G. M. 1991. What is the most likely diagnosis? In Bonissone, P.; Henrion, M.; Kanal, L.; and Lemmer, J., eds., *Uncertainty in Artificial Intelligence 6*. New York, N. Y.: Elsevier Science Publishing Company, Inc. 89–105.

Shimony, S. E. 1993. The role of relevance in explanation I: Irrelevance as statistical independence. *International Journal of Approximate Reasoning* 8(4):281–324.