System Diagnosability Analysis Using p-slop MAP

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

K. Wojtek Przytula
HRL Laboratories, LLC
Malibu, CA, 90265
wojtek@hrl.com

Abstract

Researchers have reported successful deployments of diagnosis decision support systems based on Bayesian networks. However, the methodology for evaluating the diagnosability for such systems has not been sufficiently addressed, which consequently hinders the pace of full embracement of such systems. In this paper, we propose a methodology to analyze diagnosability for diagnosing multiple faults for systems with multi-valued discrete variables. Our analysis procedure is based on computing the configurations, p-slop MAP, for the top p percent of posterior joint distribution for faults given evidence. p-slop MAP enables us to extend diagnosability measures beyond those developed in system diagnosis literature. Our analysis results can help not only design for diagnosability, but also for developing new diagnostic procedures for systems in service.

Introduction

Several researchers have reported successful deployments of diagnostic decision support systems (DDSS) based on Bayesian networks (Middleton *et al.* 1991; Heckerman, Breese, & Rommelse 1995; Darwiche 2000; Przytula & Smith 2004). However, we believe that evaluating the diagnosability for such systems has not been sufficiently addressed, especially for the cases of diagnosing multiple faults for systems with multi-valued discrete variables. For example, due to the lack of consensus, IEEE Standards on Testability and Diagnosability Characteristics and Metrics only addresses systems with no more than a single independent fault(Kaufman & Wilmering 2005). It is obvious that the inadequacy of diagnosability analysis can slow the acceptance of DDSS.

More importantly, we see several benefits for the evaluation of system diagnosability. For example, if systems are in service, the diagnosability analysis result can lead to development of new diagnostic procedures or maintenance schedules to improve diagnosability. If systems are in design, the analysis result can help in sensor placements and built-in test design to improve the availability of observables. We envision that diagnosability analysis will be incorporated into system life-cycle analysis such that the cost-effective trade-off analysis among system reliability, functionality, and diagnosability could be evaluated.

When evaluating the performance of DDSS, one normally compares the gold standard with the diagnosis recommended by DDSS for a given case. Gold standards are usually elicited from domain experts, which may be biased by recent, unusual, or memorable cases. Diagnosis generated by DDSS, on the other hand, depends on how conclusions are drawn from evidence, which are the subjects of representation and reasoning frameworks. Even though we can assume that we can elicit reliable gold standards and generate appropriate diagnosis, we may still encounter situations where we have no known cases (e.g., systems which are in design) or poorly recorded cases (e.g., repair records are collected for the purpose of billing). Therefore, there is a clear need for a systematic and automated evaluation method for system diagnosability analysis.

Analysis of system diagnosability have been studied in model-based diagnostics (MBD) (Simpson & Sheppard 1994; Sheppard & Simpson 1998; Provan 2001). Several testability and diagnosability measures have been reported. Although traditional model-based diagnosis was wide-spread in diagnosing electronic circuit systems, we found it is difficult to apply MBD to applications where system behavior is inherently uncertain and too complex to be fully represented by model-based representation languages. Diagnostic decision support systems based on Bayesian networks, on the other hand, have become popular in domains where uncertainty is prominent. However, diagnosability analysis for such systems receives much less attentions. In (Middleton *et al.* 1991), the diagnostic accuracy is evaluated for systems with binary variables (a disease present or absent) for test cases abstracted from Scientific American Medicine. In (Heckerman, Breese, & Rommelse 1995; Przytula, Dash, & Thompson 2003), Monte Carlo sampling technique is used to automatically generate test cases for model evaluation based on posterior marginal probabilities for faults given evidence. Two possible drawbacks are the use of posterior marginal probabilities for multiple faults and the difficulty in sampling cases for the tail end of the distribution.

Our approach to system diagnosability is an automated system diagnosability analysis procedure based on the com-

putation of the configurations, p-slop MAP, for the top p percent of posterior joint distribution for faults given evidence. Instead of drawing diagnosis from the posterior marginal probabilities, our procedure derives diagnosis from posterior joint probabilities of faults given evidence, which is not only suitable for single fault diagnosis, but also for diagnosing multiple faults for systems with multi-valued variables. Beyond diagnosability measures developed in system diagnosis, our procedure computes case-wise, observation-wise, and model-wise diagnosability measures such as accuracy, sensitivity, and specificity. In addition, we extend an *ambiguity* measure for single fault analysis to multi-valued multi-fault analysis and *separability* for faults given evidence based on conditional entropy. To address the possible complexity of MAP computations, we introduce the level parameter to partition the networks into relevant fragments with respect to evidence for diagnosability analysis.

In summary, we developed a novel automated system diagnosability analysis procedure, which extends current diagnosability measures for users to automatically and systematically generate the diagnosability for diagnosing fault configurations with respect to the top j percent of observations manifested by fault configurations and are diagnosed up to the top k percent of diagnoses. In the following sections, we will report our diagnostic Bayesian networks, analysis procedure, diagnosability measures, and examples of analysis.

Diagnostic Bayesian Networks

To represent diagnostic knowledge in a Bayesian network (BN), we classify each node in BN into one of following categories: target, observation, and auxiliary. A target node usually represents a diagnostic interest (e.g., the health status of a fuel injector). A target node has at least one target state, representing a failure mode (fault) of a component (e.g., a state "plugged" as a failure mode of a fuel injector), and at least one non-target state, representing a normal operational mode of a component (e.g., a state "ok" for an operational fuel injector). An observation node usually represents a symptom (e.g., observing an excessive smoking in engine exhaust), a built-in error message (e.g., the status of a power supply which is monitored by a feedback signal), or a test (e.g., measuring the voltage of a battery). An error message based observation is normally recorded in an archive when it obtains an abnormal state (e.g. power supply status is failed). When an error message, which is continuously monitored by a signal, is not recorded in an archive, one could assume that the error message is in its *default* ok state. This is to account for unreported observations (Peot & Shachter 1998). A node which is neither a target nor an observation is classified as an auxiliary node, which is usually used to represent intermediate relations between targets and observations. An observation node is further annotated with a Boolean flag, *ranked*, to specify whether a node will be ranked in the computation for recommending next observation to make. We normally annotate a test, but not an error message or a symptom, as *ranked*, since the states of symptoms and error messages are usually available before a diagnostic session is started and do not need to be recommended. We call such an annotated Bayesian network a diagnostic Bayesian network (dBN) (Lu & Przytula 2005).

Diagnosability Analysis

Diagnosability is usually defined as the ease of isolating faults in a physical system, i.e., whether a physical system can be properly diagnosed. To analyze the diagnosability of a system, we need to consider (1) whether there are sufficient observables and (2) whether observed evidence is properly reasoned upon to derive diagnosis. There have been several proposals on how the most likely diagnosis can be drawn from observed evidence (Poole & Provant 1991). In this paper, we will focus on quantifying sufficiency of observables by computing diagnosis as p-slop maximum a posteriori assignments for target nodes.

The maximum a posteriori probability assignment (MAP) is the most probable configuration for a set of target variables given the evidence. The p-slop MAP assignments are the configurations of a set of target variables which have the top p percent of posterior joint probability distribution region given the evidence. The p-slop MAP is selected as our diagnoses, because we believe that system diagnosability shall be addressed by posterior joint probability distribution for the set of target nodes rather than by posterior marginal distribution for each target node in the set. Please note that the configurations of a set of target variables given evidence are derived by $\max$ operator in p-slop MAP computation, whereas the configuration for each target node would be derived from a threshold operator (possibly $\max$ operator) over the posterior marginal probability for each target node. As a result, p-slop MAP provides a consistent view of overall status of system diagnosability.

The problem of computing MAP in BN has been shown to be $\mathbf{NP^{PP}}$ -complete (Park 2002). In addition to methods for exact inference (Park & Darwiche 2003), there are several approximate approaches for finding MAP in large Bayesian networks (Park & Darwiche 2001; Yuan, Lu, & Druzdzel 2004; Sun, Druzdzel, & Yuan 2007). We use both SMILE reasoning engine, developed by Decision Systems Laboratory at the U. of Pittsburgh and SamIam reasoning engine developed by Automated Reasoning Group at UCLA for computing p-slop MAP assignments.[1]

Analysis Procedure

Our diagnosability analysis procedure consists of three major phases: case simulation, diagnosis generation, and diagnosability scoring. In the *case simulation* phase, we assume that there is a generative BN faithfully modeling the *generative* distribution of a physical system. We use the generative BN to generate cases for faulty system, i.e., generating configurations for observation nodes by injecting faults into the system. In the *diagnosis generation* phase, we assume that there is a diagnostic BN (dBN) faithfully modeling the *diagnostic* distribution of a physical system. We instantiate simulated evidence (observation configurations) into dBN to compute diagnoses. In *diagnosability scoring* phase, the de-

[1] SMILE and SamIam are available for download at http://genie.sis.pitt.edu/ and http://reasoning.cs.ucla.edu/samiam/.

rived diagnoses are compared with their corresponding injected faults to compute diagnosability scores.

More formally, the case simulation phase will perform the following steps:

1. *fault simulation*: simulate a configuration $\mathbf{t_i} = \{t_{gh} | T_g \in \mathbf{T}\}$ from the generative BN, where each target node $T_g \in \mathbf{T}$ is put into one of its states, t_{gh}. We will denote the joint probability for configuration $\mathbf{t_i}$ as p_i.
2. *fault injection*: instantiate target nodes as $\mathbf{t_i}$ in the generative BN.
3. *evidence simulation*: simulate a set of evidence, $\mathbf{e_i} = \{\mathbf{o_{i1}}, \ldots, \mathbf{o_{ij}}, \ldots, \mathbf{o_{im}}\}$, by computing p-slop MAP from the instantiated generative BN, where $\mathbf{e_i}$ is a set of m observation configurations returned by p-slop MAP, i.e., the set of observation configurations for the top p percent of $\Pr(\mathbf{O}|\mathbf{t_i})$. We will denote the posterior joint probability for evidence $\mathbf{o_{ij}}$ as p_{ij}.

After completing the case simulation phase, we will have a set of cases $\mathbf{w}$ where each case $\mathbf{w_i} \in \mathbf{w}$ is a tuple $\langle \mathbf{t_i}, \mathbf{e_i} \rangle$.

The diagnosis generation phase will loop through each observation configurations $\mathbf{o_{ij}} \in \mathbf{e_i}$ to perform the following steps:

1. *evidence instantiation*: in the dBN, instantiate observation nodes as $\mathbf{o_{ij}}$;
2. *diagnoses generation*: generate diagnoses, $\mathbf{d_{ij}} = \{\mathbf{d_{ij0}}, \ldots, \mathbf{d_{ijk}}, \ldots, \mathbf{d_{ijn}}\}$, by computing p-slop MAP from the dBN, where $\mathbf{d_{ij}}$ is a set of n target configurations returned by p-slop MAP, i.e., the set of target configurations for the top p percent of $\Pr(\mathbf{T}|\mathbf{o_{ij}})$. We will denote the posterior joint probability for diagnosis $\mathbf{d_{ijk}}$ as p_{ijk}.

After the diagnosis generation phase, we extend $\mathbf{w_i}$ to a triple $\langle \mathbf{t_i}, \mathbf{e_i}, \mathbf{d_i} \rangle$ where $\mathbf{d_i}$ is the sets of diagnoses $\mathbf{d_{ij}}$ for each $\mathbf{o_{ij}} \in \mathbf{e_i}$.

In the diagnosability scoring phase, we compute diagnosability measures by comparing $\mathbf{t_i}$ with $\mathbf{d_{ijk}}$ and aggregate over $\mathbf{e_i}$ and $\mathbf{w_i}$ to derive target-wise diagnosability measures for $\mathbf{t_i}$ and model-wise diagnosability measures for $\mathbf{w}$.

We would like to emphasize the importance of the assumption of faithfulness for both generative BN and diagnostic dBN. It is normally the case that we will elicit the generative distribution of a system from design engineers to build a generative BN, whereas diagnostic distribution will be elicited from field experts for dBN. However, it is possible that we can only afford to build one BN. For example, when a system is in design, we may only have generative distribution; or alternatively we only have diagnostic distribution, because we purchased the system from a manufacture who would not disclose the internal design of the system. In such a case, we still can perform diagnosability analysis by replicating a dBN or the generative BN, but we need to be cautious that we do not violate the faithfulness assumption.

The *fault simulation* step in the case simulation phase could be done in several ways. One way is to systematically enumerate the configurations of the set of target nodes. Exhaustive enumerations will become infeasible for large BN, since the number of configurations for target nodes grows exponentially as the number of target nodes increases. Alternatively, one can rank the prior joint probabilities for target nodes and only perform fault simulation for the faut configurations, which have the top i percent of probability mass.

When we evaluate diagnosability for large BNs, exact MAP computations sometimes fail to generate solutions. Although approximate MAP can yield solution for large BNs, it is hard to judge the quality for the solutions of approximate p-slop MAP. In order to evaluate large BN, our approach allows users to evaluate BN by fragments. The network fragments are derived by the topological separation between the injected faults, which is a subset of all target nodes, and other nodes in our custom layered structure (Lu & Przytula 2005). Since we constrained ourselves in evaluating diagnosability to one set of injected faults at a time, all other target variables will not be instantiated. Consequently, the BN will automatically be fragmented by relevance reasoning before p-slop MAP computation in the steps of evidence simulation and diagnoses generation.

Diagnosability Measures

There have been several diagnosability measures developed in model-based system diagnosability analysis (Simpson & Sheppard 1994; Provan 2001). Although there is a IEEE standard on diagnosability measures for systems with single fault, there is no consensus about diagnosability measures for multiple faults with multi-valued variables. In this section, we define diagnosability measures to quantify diagnosability for systems modeled by dBNs with multiple faults and multi-valued variables.

Accuracy, sensitivity, specificity

Given a case $\mathbf{w_i} = \langle \mathbf{t_i}, \mathbf{e_i}, \mathbf{d_i} \rangle$ generated by our diagnosability analysis procedure, we can compute *accuracy* by comparing each diagnosis in $\mathbf{d_{ijk}}$ with the simulated fault injection $\mathbf{t_i}$. Recall that $\mathbf{t_i} = \{t_{gh} | T_g \in \mathbf{T}\}$, where each target node T_g is put into one of its states, t_{gh}. Similarly, a diagnosis $\mathbf{d_{ijk}}$ is $\{t_{gh'} | T_g \in \mathbf{T}\}$ where each target node T_g is diagnosed as being in one of its states, $t_{gh'}$. If a target node is correctly diagnosed, $t_{gh} = t_{gh'}$, we will denote it as $a_{ijkg} = 1$, and $a_{ijkg} = 0$, otherwise. Let t be the number of target nodes in $\mathbf{T}$, $t = |\mathbf{T}|$. The accuracy, a_{ijk}, for a diagnosis $\mathbf{d_{ijk}}$ is computed as $a_{ijk} = \frac{\sum_g a_{ijkg}}{t}$. The accuracy for diagnosing a fault simulation $\mathbf{t_i}$ given evidence $\mathbf{o_{ij}}$ is weighted by the likelihoods as:

$$a_{ij} = \frac{\sum_k a_{ijk} \times p_{ijk}}{\sum_k p_{ijk}}. \quad (1)$$

Next, we can compute the accuracy for diagnosing a fault simulation $\mathbf{t_i}$ given the set of evidence $\mathbf{e_i}$ as:

$$a_i = \frac{\sum_j a_{ij} \times p_{ij}}{\sum_j p_{ij}}. \quad (2)$$

Similarly, we can weight each fault simulation diagnosis accuracy with its likelihood to produce model-wise diagnosis accuracy:

$$a = \frac{\sum_i a_i \times p_i}{\sum_i p_i}. \quad (3)$$

In addition to accuracy, we can use similar formulations to compute *sensitivity* and *specificity* measures. Recall that each target node has a injected state $\mathbf{t_i}$ and a diagnosed state $\mathbf{d_{ijk}}$. If the injected state is a target (non-target) state and its diagnosed state is the same target (non-target) state, we count it as one of the correctly diagnosed defects (non-defects). The sensitivity (specificity) is defined as the ratio of correctly diagnosed defects (non-defects). All we need to do now is to substitute the definition of accuracy with sensitivity (specificity) in Equation 1-3 to compute observation-wise, case-wise, and model-wise sensitivity (specificity) measures.

Ambiguity

In system diagnosability analysis, ambiguity group refers to a set of faults that cannot be further isolated given evidence. Since ambiguity group is defined in terms of faults, we denote the set of target nodes which are instantiated into their target (fault) states in $\mathbf{t_i}$ as $\mathbf{t_i^f}$. Similarly, we denote the set of target nodes which are diagnosed in their target (fault) states in $\mathbf{d_{ijk}}$ as $\mathbf{d_{ijk}^{f'}}$. We define the *ambiguity* between $\mathbf{t_i}$ and $\mathbf{d_{ijk}}$ as follows:

$$\alpha_{ijk} = 1 - \frac{|\mathbf{t_i^f} \bigcap \mathbf{d_{ijk}^{f'}}|}{|\mathbf{t_i^f} \bigcup \mathbf{d_{ijk}^{f'}}|}, \tag{4}$$

where $\mathbf{t_i^f} \bigcap \mathbf{d_{ijk}^{f'}}$ is the set of target states for the target nodes which have the simulated target state the same as the diagnosed state, and $\mathbf{t_i^f} \bigcup \mathbf{d_{ijk}^{f'}}$ is the union of the simulate target states and the diagnosed target states. The value of α_{ijk} ranges between 0 and 1. Ambiguity $\alpha_{ijk} = 0$ ($\alpha_{ijk} = 1$) indicates that all faults are correctly (incorrectly) diagnosed.

To see how well a fault simulation $\mathbf{t_i}$ is diagnosed with respect to evidence $\mathbf{o_{ij}}$, we can average over ambiguity α_{ijk}. The ambiguity α_{ij} for $\mathbf{t_i}$ given evidence $\mathbf{o_{ij}}$ is defined as:

$$\alpha_{ij} = \frac{\sum_k \alpha_{ijk} \times p_{ijk}}{\sum_k p_{ijk}}. \tag{5}$$

Similarly, we can see how well a fault simulation $\mathbf{t_i}$ is diagnosed with respect to the set of evidence $\mathbf{e_i}$. The ambiguity α_i for $\mathbf{t_i}$ given evidence $\mathbf{e_i}$ is defined as:

$$\alpha_i = \frac{\sum_j \alpha_{ij} \times p_{ij}}{\sum_j p_{ij}}. \tag{6}$$

The model-wise ambiguity can be computed as follows:

$$\alpha = \frac{\sum_i \alpha_i \times p_i}{\sum_i p_i}. \tag{7}$$

Separability

So far, we define diagnosability measures by comparing diagnoses with simulated faults (gold-standards) and weighted by their frequency of occurrence. Separability, on the other hand, is defined directly on the posterior joint distribution of diagnoses to see how separable the distribution is. This is particularly useful when we only have field data without gold-standards. We can compute the separability by normalized conditional entropy. The separability for diagnosing $\mathbf{t_i}$ with diagnoses $\mathbf{d_{ij}}$ given evidence $\mathbf{o_{ij}}$ is computed as:

$$H_{\mathbf{d_{ij}}|\mathbf{o_{ij}}} = \frac{-\sum_k p_{ijk} \log p_{ijk}}{\log n}. \tag{8}$$

Similarly, we can check how separable it is for the posterior joint probability distribution of evidence manifested by a fault simulation $\mathbf{t_i}$. The separability for observing evidence $\mathbf{e_i}$ given fault simulation $\mathbf{t_i}$ is computed as:

$$H_{\mathbf{e_i}|\mathbf{t_i}} = \frac{-\sum_j p_{ij} \log p_{ij}}{\log m}. \tag{9}$$

Please note that both $H_{\mathbf{d_{ij}}|\mathbf{o_{ij}}}$ and $H_{\mathbf{e_i}|\mathbf{t_i}}$ are computed from the top p percent of $\Pr(\mathbf{T}|\mathbf{o_{ij}})$ and $\Pr(\mathbf{O}|\mathbf{t_i})$. $H_{\mathbf{d_{ij}}|\mathbf{o_{ij}}}$ and $H_{\mathbf{e_i}|\mathbf{t_i}}$ will become real normalized conditional entropies for $\Pr(\mathbf{T}|\mathbf{o_{ij}})$ and $\Pr(\mathbf{O}|\mathbf{t_i})$, when $p = 100$.

Example

To illustrate our diagnosability measures, we apply our diagnosability procedure to Asia network (Lauritzen & Spiegelhalter 1988). The network consists of three target nodes, four observation nodes, and one auxiliary node (see Figure 1). We duplicate Asia network into two copies one for case simulation and one for diagnosis generation.

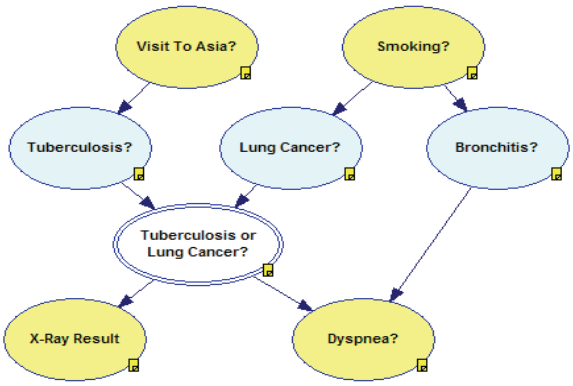

Figure 1: Asia network annotated as diagnostic BN. Target nodes (Tuberculosis?, Lung Cancer?, and Bronchitis?) are colored as light blue, observation nodes (Visit to Asia?, Smoking?, X-Ray Result, and Dyspnea?) as light yellow, and auxiliary node (Tuberculosis or Lung Cancer?) as white.

Table 1 is the list of prior joint probabilities for all configurations of targets nodes in the Asia network. We see that the multiple fault configuration $\mathbf{t_2}$ is pretty frequent, even more than the single fault configurations $\mathbf{t_3}$ and $\mathbf{t_4}$.

Assume that we would like to evaluate the diagnosability for at least the top 90% of observations manifested by $\mathbf{t_2}$ using at least the top 95% of diagnoses given the interested observations. Our procedure will first inject $\mathbf{t_2}$ onto the Asia network and simulate the evidence $\mathbf{e_2}$ for observation configurations using 90%-slop MAP. Table 2 shows the posterior joint probabilities for observation nodes given $\mathbf{t_2}$. Since we consider at least the top 90% of observations manifested by $\mathbf{t_2}$, we have $\mathbf{e_2} = \{\mathbf{o_{2,0}}, \mathbf{o_{2,1}}\}$.

Next, our procedure performs diagnosis generation phase for $\mathbf{e_2}$. Table 3 and 4 show the posterior joint probability for diagnosis given $\mathbf{o_{2,0}}$ and $\mathbf{o_{2,1}}$ separately. Since we consider at least the top 95% of generated diagnoses, we have $\mathbf{d_{2,0}} = \{\mathbf{d_{2,0,0}}, \mathbf{d_{2,0,1}}, \mathbf{d_{2,0,2}}, \mathbf{d_{2,0,3}}\}$ and $\mathbf{d_{2,1}} = \{\mathbf{d_{2,1,0}}, \mathbf{d_{2,1,1}}, \mathbf{d_{2,1,2}}, \mathbf{d_{2,1,3}}\}$.

The diagnosability measures for each diagnosis of interest are shown in Table 5. With Equation 1, 2, 4, 5, and 6, we compute observation-wise and case-wise diagnosability measures of accuracy, sensitivity, specificity, and ambiguity (see Table 6). As an example, we can read off from Table 6 and claim that the accuracy for diagnosing $\mathbf{t_2}$ with respect to at least the top 95% of diagnoses for observation $\mathbf{o_{2,0}}$ is 0.816861. Similarly, we can read off case-wise accuracy for diagnosing $\mathbf{t_2}$ with respect to at least the top 95% of diagnoses for at least the top 90% observation manifested by $\mathbf{t_2}$ is 0.793075.

To see separability, we use Equation 9 and we have separability $H_{\mathbf{d_{2,0}}|\mathbf{o_{2,0}}} = 0.830747$ and $H_{\mathbf{d_{2,1}}|\mathbf{o_{2,1}}} = 0.916433$ for how separable the diagnoses are given $\mathbf{o_{2,0}}$ and $\mathbf{o_{2,1}}$. With Equation 8, we have $H_{\mathbf{e_2}|\mathbf{t_2}} = 0.5384$ for how separable the observations are manifested by $\mathbf{t_2}$. Please recall that separability is based on the approximate computation of conditional entropy. If we do use the conditional entropy for $\Pr(\mathbf{T}|\mathbf{o_{2,0}})$ and $\Pr(\mathbf{T}|\mathbf{o_{2,1}})$ as listed in Table 3 and 4, we have 0.645823 and 0.696556. Similarly, if we compute the conditional entropy for $\Pr(\mathbf{O}|\mathbf{t_2})$, we have 0.24119.

Now, we have demonstrated how to compute diagnosability measure for $\mathbf{t_2}$, we can systematically go through each target configurations listed in Table 1 to compare their relative diagnosability or compute the model-wise diagnosability. We can also have different fault simulation strategies: for example, compute diagnosability for faults with joint prior probability above certain threshold to account for frequent fault configurations, or compute diagnosability for faults bellow certain threshold to account rare fault configurations.

Table 1: Prior joint probabilities for target nodes in the Asia network. Target nodes are abbreviated as B?:*Bronchitis?*, L?:*Lung Cancer?*, and T?:*Tuberculosis?*. Target states are abbreviated as p:present and a:absent.

$\mathbf{t_i}$	B?	L?	T?	$\Pr(\mathbf{t_i})$
$\mathbf{t_0}$	a	a	a	0.5210244
$\mathbf{t_1}$	p	a	a	0.4141476
$\mathbf{t_2}$	p	p	a	0.0311724
$\mathbf{t_3}$	a	p	a	0.0232556
$\mathbf{t_4}$	a	a	p	0.0054756
$\mathbf{t_5}$	p	a	p	0.0043524
$\mathbf{t_6}$	p	p	p	0.0032765
$\mathbf{t_7}$	a	p	p	0.0002444

Conclusion

The major contribution of our paper is a novel automated system diagnosability analysis procedure which can automatically and systematically generate the diagnosability for

Table 2: Posterior joint probabilities for observation nodes given $\mathbf{t_2}$ in the Asia network. Observation nodes are abbreviated as D?:Dyspnea?, S?:Smoking?, V?:Visit To Asia?, and X:X-Ray Result. Observation states are abbreviated as p:present, a:absent, s:smoker, ns:non_smoker, v:visit, nv:no_visit, n:normal, and ab:abnormal.

$\mathbf{o_{2,j}}$	D?	S?	V?	X	$\Pr(\mathbf{o_{2j}}\|\mathbf{t_2})$
$\mathbf{o_{2,0}}$	p	s	nv	ab	0.8319361
$\mathbf{o_{2,1}}$	a	s	nv	ab	0.0924373
$\mathbf{o_{2,2}}$	p	ns	nv	ab	0.0415968
$\mathbf{o_{2,3}}$	p	s	nv	n	0.0169782
$\mathbf{o_{2,4}}$	p	s	v	ab	0.0080638
$\mathbf{o_{2,5}}$	a	ns	nv	ab	0.0046218
$\mathbf{o_{2,6}}$	a	s	nv	n	0.0018864
$\mathbf{o_{2,7}}$	a	s	v	ab	0.0008959
$\mathbf{o_{2,8}}$	p	ns	nv	n	0.0008489
$\mathbf{o_{2,9}}$	p	ns	v	ab	0.0004031
$\mathbf{o_{2,10}}$	p	s	v	n	0.0001645
$\mathbf{o_{2,11}}$	a	ns	nv	n	0.0000943
$\mathbf{o_{2,12}}$	a	ns	v	ab	0.0000447
$\mathbf{o_{2,13}}$	a	s	v	n	0.0000182
$\mathbf{o_{2,14}}$	p	ns	v	n	0.0000082
$\mathbf{o_{2,15}}$	a	ns	v	n	0.0000009

Table 3: Posterior joint probabilities for target nodes given observation configuration $\mathbf{o_{2,0}}$ in the Asia network.

$\mathbf{d_{2,0,k}}$	B?	L?	T?	$\Pr(\mathbf{d_{2,0,k}}\|\mathbf{o_{2,0}})$
$\mathbf{d_{2,0,0}}$	p	p	a	0.4730187
$\mathbf{d_{2,0,1}}$	a	p	a	0.2452689
$\mathbf{d_{2,0,2}}$	p	a	a	0.1930688
$\mathbf{d_{2,0,3}}$	p	a	p	0.0430017
$\mathbf{d_{2,0,4}}$	a	a	p	0.0222971
$\mathbf{d_{2,0,5}}$	a	a	a	0.0160891
$\mathbf{d_{2,0,6}}$	p	p	p	0.0047777
$\mathbf{d_{2,0,7}}$	a	p	p	0.0024077

fault configurations. It provides the flexibility for users to specify the observations of interest up to the top j percent of observations manifested by fault configurations. It also allows users to specify the diagnoses for the observation of interest up to the top j percent of diagnoses. Moreover, this diagnosability analysis procedure can compute novel diagnosability measures from various perspectives: observation-wise, case-wise, and model-wise diagnosability measures.

In (Lu & Przytula 2005), we described a methodology for building complex dBN (2,147 nodes and 3,650 arcs) for locomotive diagnoses. Our practical experience of applying this analysis procedure to the locomotive dBN led us to develop the level parameter to partition the large network into smaller ones so that MAP computation is feasible. Although we are able to extend our analysis to some network fragments with the level parameter, there are still complex network fragments cannot be addressed. This calls for research in developing more efficient MAP computation algorithms.

We have also applied the analysis procedure to the the

Table 4: Posterior joint probabilities for target nodes given observation configuration $\mathbf{o_{2,1}}$ in the Asia network.

$\mathbf{d_{2,1,k}}$	B?	L?	T?	$\Pr(\mathbf{d_{2,1,k}}\|\mathbf{o_{2,1}})$
$\mathbf{d_{2,1,0}}$	a	a	a	0.3949117
$\mathbf{d_{2,1,1}}$	a	p	a	0.2866767
$\mathbf{d_{2,1,2}}$	p	p	a	0.1433383
$\mathbf{d_{2,1,3}}$	p	a	a	0.1316372
$\mathbf{d_{2,1,4}}$	a	a	p	0.0260615
$\mathbf{d_{2,1,5}}$	p	a	p	0.0130307
$\mathbf{d_{2,1,6}}$	a	p	p	0.0028957
$\mathbf{d_{2,1,7}}$	a	p	p	0.0014478

Table 5: Diagnosability measures for $\mathbf{d_{2,0}}$ and $\mathbf{d_{2,1}}$.

$\mathbf{d_{2,j,k}}$	Acc.	Sen.	Spe.	Amb.
$\mathbf{d_{2,0,0}}$	1	1	1	0
$\mathbf{d_{2,0,1}}$	2/3	1/2	1	1/2
$\mathbf{d_{2,0,2}}$	2/3	1/2	1	1/2
$\mathbf{d_{2,0,3}}$	1/3	1/2	0	2/3
$\mathbf{d_{2,1,0}}$	1/3	0	1	1
$\mathbf{d_{2,1,1}}$	2/3	1/2	1	1/2
$\mathbf{d_{2,1,2}}$	1	1	1	0
$\mathbf{d_{2,1,3}}$	2/3	1/2	1	1/2

problem of *design for diagnosability*, for which two versions of system designs are evaluated for sensor placements. Although the analysis results are informative, we believe that more research is needed to integrate the design for diagnosability into system design life cycle.

Acknowledgments

We thank the reviewers for their comments, which helped us improve our paper.

References

Darwiche, A. 2000. Model-based diagnosis under real-world constraints. *AI Magzazine* 21(2):57–73.

Heckerman, D.; Breese, J. S.; and Rommelse, K. 1995. Decision-Theoretic troubleshooting. *Communications of the ACM* 38(3):49–57.

Kaufman, M. A., and Wilmering, T. 2005. 1522 IEEE trial-use standard for testability and diagnosability characteristics and metrics. Technical report, IEEE.

Lauritzen, S. L., and Spiegelhalter, D. J. 1988. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society, Series B (Methodological)* 50(2):157–224.

Lu, T.-C., and Przytula, K. W. 2005. Methodology and tools for rapid development of large bayesian networks. In *Proceedings of the 16th International Workshop on Principles of Diagnosis (DX-05)*, 107–112.

Middleton, B.; Shwe, M.; Heckerman, D.; Henrion, M.; Horvitz, E.; Lehmann, H.; and Cooper, G. 1991. Probabilistic diagnosis using a reformulation of the INTERNIST–1/QMR knowledge base: II. evaluation of diagnostic performance. *Methods of Information in Medicine* 30(4):256–267.

Park, J., and Darwiche, A. 2001. Approximating map using local search. In *Proceedings of the 17th Annual Conference on Uncertainty in Artificial Intelligence (UAI-01)*, 403–41. San Francisco, CA: Morgan Kaufmann.

Park, J., and Darwiche, A. 2003. Solving MAP exactly using systematic search. In *Proceedings of the 19th Annual Conference on Uncertainty in Artificial Intelligence (UAI-03)*, 459–46. San Francisco, CA: Morgan Kaufmann.

Park, J. 2002. MAP complexity results and approximation methods. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence (UAI-02)*, 388–39. San Francisco, CA: Morgan Kaufmann.

Peot, M., and Shachter, R. 1998. Learning from what you don't observe. In *Proceedings of the Fourteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI–98)*, 439–446.

Poole, D., and Provant, 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.

Provan, G. 2001. System diagnosability analysis using model-based diagnosis tools. In *Proc. SPIE*, volume 4389, 93–101.

Przytula, K. W., and Smith, S. 2004. Diagnostic assistant based on graphical probabilistic models. In *SAE*.

Przytula, K. W.; Dash, D.; and Thompson, D. 2003. Evaluation of bayesian networks used for diagnostics. In *2003 IEEE Aerospace Conference*.

Sheppard, J. W., and Simpson, W. R. 1998. *Research Perspectives and Case Studies in System Test and Diagnostics*. Kluwer Academic Publishers.

Simpson, W. R., and Sheppard, J. W. 1994. *System Test and Diagnosis*. Klwer Academic Publishers.

Sun, X.; Druzdzel, M. J.; and Yuan, C. 2007. Dynamic weighting A* search-based MAP algorithm for bayesian networks. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI 2007)*, 2385–2390.

Yuan, C.; Lu, T.-C.; and Druzdzel, M. 2004. Annealed MAP. In *Proceedings of the 20th Annual Conference on Uncertainty in Artificial Intelligence (UAI-04)*, 628–63. Arlington, Virginia: AUAI Press.

Table 6: Observation-wise and case-wise diagnosability measures for $\mathbf{t_2}$.

	Acc.	Sen.	Spe.	Amb.
$\mathbf{d_{2,0}}$	0.816861	0.74782	0.954942	0.259689
$\mathbf{d_{2,1}}$	0.579001	0.368502	1.0	0.631498
$\mathbf{d_2}$	0.793075	0.709888	0.859626	0.29687