3C4-IOS-4a-1

Using Sensitivity Analysis for Designing Resilient Systems

Hei Chan[*1*2] Katusmi Inoue[*2]

[*1] Transdisciplinary Research Integration Center [*2] National Institute of Informatics

In this paper, we present the research results on sensitivity analysis in probabilistic graphical models, in particular Bayesian networks. We look at both the theory and application of sensitivity analysis in different domains, such as modeling (to represent the different properties and relationships between the variables), inference (to provide exact or approximate answers to user queries), and analysis (to summarize current behavior, predict future trends, and suggest actions for achieving certain targets). We also discuss using sensitivity analysis for the problem of designing systems that are resilient, such that the systems are resistant from large-scale perturbations caused by unexpected events and changes, and if their functionality is lost temporarily due to outside forces, the systems can recover gracefully and quickly to restore their functionality in the long run.

1. Introduction

The problem of systems resilience can be informally defined as the ability "to maintain its core purpose and integrity in the face of dramatically changed circumstances" [34]. In recent years, many researchers of different fields have recognized the importance of a new research discipline concerning the resilience of complex systems in the face of unexpected events in terms of time and scale (such as the 3.11 earthquake in Japan, the global economic crisis, or a new strain of virus) may cause irreversible damages to the core functionality of these agent systems. The goal of resilience research is to provide a set of general principles for building resilient systems in various domains, such that the systems are resistant from large-scale perturbations caused by unexpected events and changes, and if their functionality is lost temporarily due to outside forces, the systems can recover gracefully and quickly to restore their functionality in the long run.

The concept of resilience has appeared in various disciplines such as environmental science, materials science, sociology, ecology, disaster prevention, artificial intellignce, and so on [17, 4, 23, 2]. However, while we have seen many examples of seemingly resilient systems in various fields, researchers have not agreed on a common definition on resilience among the different domains yet [33].

We have recently developed a new research topic called "systems resilience", to provide a set of unified design principles for building resilient systems [22]. Our first step is to define a novel system model called the *SR-model* [29].

The significant aspects of our SR-model are as follows. The definition of the SR-model allows the system to change dynamically over time, such that the variables, domains, constraints, and configurations of the system can evolve based on the decisions made by agents and/or outside environmental events. The flexibility of our SR-model allows the modeling of the dynamicity of systems that is required in many domains, and is based on Constraint Satisfaction Problems (CSPs) [16]. Our SR-model enables us to measure four important properties that are central to the idea of resilience:

- *Resistance*: The ability to maintain under a certain "threshold", such that the system satisfies certain hard constraints and does not suffer from irreversible damages.
- *Recoverability*: The ability to recover to a baseline of acceptable quality as quickly and inexpensively as possible.
- *Functionality*: The ability to provide a guaranteed average degree of quality for a period of time.
- *Stabilizability*: The ability to avoid undergoing changes that are associated with high transitional costs.

In the current version of our SR-model, we assume that we have a complete knowledge on all past and current configurations of the variables in the dynamic system. However, in reality, we may only have uncertain information on some of these configurations. Therefore, we must incorporate into our models with properties which allow for probabilistic reasoning. In artificial intelligence, uncertain beliefs have often been represented using probabilistic graphical models, such as Bayesian Networks. In a Bayesian Network, the inputs are the network structure, which specify the causal relationships between variables, and the network parameters, which specify the probabilities of local events, while the outputs are the conclusions drawn from performing query inference, such as marginal probabilities.

In the next section, we will introduce Bayesian networks. Afterwards, we will introduce sensitivity analysis, which is the analysis of the relationships between changes in the inputs and the outputs of a mathematical model, and go over some of the important theoretical results of sensitivity analysis of Bayesian networks. We then show a simple example of performing sensitivity analysis on a Bayesian network. Finally, we will point out some current and future research directions in sensitivity analysis of Bayesian networks, and how it can benefit the design of resilient systems.

Contact: National Institute of Informatics, 2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo 101-8430. Hei Chan: hei@nii.ac.jp; Katsumi Inoue: inoue@nii.ac.jp.

2. Bayesian Networks

A graphical model is a graph where variables are represented by nodes, and the nodes are connected by edges in order to represent relationships between variables. The existence of an edge between two variables indicate that the two variables are directly connected, while the absence of an edge between two variables indicate they are not, and conditional independence relationships exist between them. The edges in a graphical model may be directed or undirected. Usually, a directed edge means that the parent variable is a cause of the child variable.*1 If in a graphical model, all edges are directed, and no cycles exist, the graph is called a *directed acyclic graph* (DAG).

For example, we can construct a graphical model to represent the variables used for diagnosing lung cancer. The individual's personal behavior (e.g., smoking) will be a factor in determining whether the individual has lung cancer, and will be directly connected to the lung cancer variable as its parents. Symptoms (e.g., coughing) and tests (e.g., X-ray) will depend on whether the individual has lung cancer, and will be directly connected to the lung cancer variable as its children. However, the two sets of variables above are not directly connected with each other, because their relationships are indirect through the lung cancer variable. A more elaborate model with more variables and edges can be created by incorporating the mechanism of humans developing lung cancer, such as adding a variable indicating the individual's tar deposits.

After constructing a graphical model to represent the dependencies between the variables, we need to quantify the strength of each dependency. A common way to do so is to use probabilities, giving us a *probabilistic graphical models*. For example, given two variables, X representing whether the individual is a smoker, and Y representing whether the individual has lung cancer, we can specify two probabilities, $Pr(Y = true|X = true) = 0.2$ and $Pr(Y = true|X = false) = 0.001$ to quantify the strengths of the influence of smoking on lung cancer. These probabilities can be called *local probabilities*, and are probability measures involving variables in localized structures, and can be easily estimated from data or by experts.*2

The most common probabilistic graphical models are Bayesian networks. A *Bayesian network* [24, 15] consists of a *network structure*, which is a directed acyclic graph, and a set of *network parameters*, which are probabilities grouped into tables called *conditional probability tables (CPT)*. A *directed acyclic graph* is a graph where variables are represented by nodes, which can be connected by directed edges to represent relationships between variables, and no cycles exist in the graph. The existence of a directed edge between two variables indicate that the two variables are directly connected, while the absence of an edge between two variables indicate that they are not directly connected, and conditional independence relationships exist between them. Usually, a directed edge means that the parent variable is a cause of the child variable [25].

The structure and the CPTs of a Bayesian network induce a joint probability distribution Pr. The probability of an instantiation $\mathbf{x}$ of all variables in the Bayesian network $\mathbf{X}$, is the product of all network parameters $Pr(x \mid \mathbf{u})$ where $\{x, \mathbf{u}\}$ is consistent with $\mathbf{x}$:

$$\mathrm{Pr}(\mathbf{x}) = \prod_{\{x,\mathbf{u}\}\sim\mathbf{x}} Pr(x \mid \mathbf{u}). \qquad (1)$$

Given the probability distribution Pr, we are interested in the following three types of inference questions, which we must use inference algorithms to compute the answer:

- What is the probability of evidence $\mathbf{e}$, i.e., $Pr(\mathbf{e})$?
- What it the probability of some value y given evidence $\mathbf{e}$, i.e., $Pr(y \mid \mathbf{e})$.
- What is the value of a set of variables $\mathbf{Y}$ which has the highest probability given evidence $\mathbf{e}$, i.e., $\max_{\mathbf{y}} Pr(\mathbf{y} \mid \mathbf{e})$?

Query probabilities of the form $Pr(y \mid \mathbf{e})$ can be called *global probabilities*, and they involve variables that may not be directly connected. Note that for the third inference question, we are interested in answers derived from global probabilities, such as the most likely product a customer will buy, instead of the values of the probabilities themselves.

The strengths of Bayesian networks come from three factors:

- The structure of Bayesian networks are able to represent the relationships between variables naturally, efficiently, and succinctly, and the probability values in the CPTs are meaningful values that can be easily checked with real-world scenarios;
- The structure and CPTs of Bayesian networks can be specified directly by domain experts, or obtained by learning algorithms from data; (For a survey of learning algorithms, see Chapters 17 and 18 of [15])
- Many inference algorithms, both exact and approximate (for large and complex networks), can be applied to compute the answers of the inference questions. (For a survey of exact inference algorithms, see Chapters 6, 7, 8, and 13 of [15]; for a survey of approximate inference algorithms, see Chapters 14 and 15 of [15])

Finally, the ability to update a Bayesian network according to new data, discovery of new information, or user feedback, etc, is also one of its attractive features, and this is where sensitivity analysis can be applied to modeling and inference. Moreover, sensitivity analysis can also be used to improve the process of building the Bayesian network from data. We will discuss these topics in the next section.

*1 While this is not necessarily true in some graphical models, the causality relation is imposed in some models, such as the causal Bayesian network [25].

*2 In some graphical models such as Markov random fields, the quantification does not represent probabilities, but compatibilities.

3. Sensitivity Analysis of Bayesian Networks

Sensitivity analysis is the analysis of the relationships between input changes and output changes in a mathematical model. In the case of a Bayesian network, the inputs are the local probabilities, i.e., network parameters, and the outputs are the global probabilities, such as query results, and other derived answers, such as most probable explanations. In this case, sensitivity analysis of Bayesian networks can help us answer many questions, such as [6]:

- What network parameters should we change to get a desired change in a query result? And by how much?
- Will changing certain network parameters affect the robustness of certain query results?
- How accurate should we estimate certain network parameters?

The topic of sensitivity analysis of Bayesian networks is a research topic that has been developing since the 1990s, using theoretical [20, 5, 18, 32] or empirical methods [26], and has been applied to fields such as medicine [13]. We first explain why these questions are important in the problem of data modeling and inference.

Assume that we use a Bayesian network to model data, and we can estimate each network parameter to a certain degree of accuracy, either directly by the domain experts or from data using learning algorithms. However, due to the varying amounts of available relevant data, the degrees of accuracy will be different for each parameter. Parameters for more likely events will have a higher degree of accuracy due to the large amount of available data, while parameters for less likely events will have a lower degree of accuracy. The company may want to know whether this uncertainty will affect the performance of the Bayesian network. Moreover, they may want to know whether it is beneficial to get more accurate estimates of these parameters, in terms of the costs of collecting more data versus the increased accuracy of the system. These questions can only be answered if we understand the intricate dependencies between the parameters and the query results. Therefore, sensitivity analysis can help build better Bayesian networks. When many variables are involved, a large number of parameters have to be specified, and it is often difficult and expensive to estimate all of them with great accuracy. With sensitivity analysis, we can identify parameters where a small change can greatly affect certain global query values, and those where even a large change does not affect these global query values much. This can help direct the domain experts or the learning algorithms to the parameter values that must be accurately estimated.

Suppose now we have built a Bayesian network from the current available data, and it is being used by many users. However, we may need to modify the Bayesian network due to many reasons. First, we will collect new data, which will reflect more recent demographics and trends, meaning the parameters will have to be revised to reflect them. Second, we may discover new relationships between variables. We may also receive feedback from the users or domain experts, pointing out inaccuracies in certain parts of the model. Finally, in some cases, the results given by the model may not match real-world results. Therefore, some of the parameters may need to be "tuned" accordingly, and we need to find out which parameters we need to change, and by how much, while also ensuring that other query results remain relatively robust, i.e., they do not change too much after the parameters are changed.

We now summarize some of the theoretical results developed in sensitivity analysis of Bayesian networks. We tackle our first problem, which is to find (minimum) parameter changes necessary to enforce a query constraint, in the form of $Pr(y \mid \mathbf{e}) \geq k$,[*3] for *every network parameter*. The results will give a list of suggested changes in single parameters such that the query constraint can be enforced. Notice that there may be not exist any possible changes for some parameters, and in some cases, no solution at all can be found.

Assuming that all variables are binary, for each pair of co-varying network parameters $Pr(x \mid \mathbf{u})$ and $Pr(\bar{x} \mid \mathbf{u})$ which must sum to 1, we introduce a *meta-parameter* $\tau_{x|\mathbf{u}}$, such that $Pr(x \mid \mathbf{u}) = \tau_{x|\mathbf{u}}$ and $Pr(\bar{x} \mid \mathbf{u}) = 1 - \tau_{x|\mathbf{u}}$. This way we can change each pair of co-varying parameters simultaneously.[*4] To find the necessary change in $\tau_{x|\mathbf{u}}$, we need to compute the following [7]:

- The probabilities $Pr(\mathbf{e})$ and $Pr(y, \mathbf{e})$;
- The partial derivatives $\frac{\partial Pr(\mathbf{e})}{\partial \tau_{x|\mathbf{u}}}$ and $\frac{\partial Pr(y,\mathbf{e})}{\partial \tau_{x|\mathbf{u}}}$.

We can show that the time and space complexity of computing the above values for all network parameters is $O(n2^w)$, where n is the number of variables in the network, and w is the tree-width of the network, which reflects the complexity of its structure. This time and space complexity is indeed the same as that of performing inference to compute the probability of evidence $Pr(\mathbf{e})$. Many inference algorithms can be used to find the sensitivity results. One of them is to compile an arithmetic circuit from the Bayesian network, and iterating over the circuit in two passes, an upward pass for computing the probabilities, and a downward pass for computing the partial derivatives [14].

Instead of changing only single parameters, we can also expand our search to find solutions where we change a combination of multiple parameters. The problem is, tuning multiple parameters requires computing higher-order partial derivatives in general, and is much more expensive complexity-wise. However, if we limit the parameters to those belonging to the same variable, i.e., they are in the same CPT, only first-order partial derivatives are required,

*3 We can also solve for other types of query constraints, such as $Pr(y \mid \mathbf{e}) \leq k$, $Pr(y \mid \mathbf{e}) - Pr(z \mid \mathbf{e}) \geq k$, and $Pr(y \mid \mathbf{e})/Pr(z \mid \mathbf{e}) \geq k$.

*4 If the variables are multi-valued, we can introduce a proportional scheme of co-varying parameters. Other scheme of co-varying parameters are also possible, as long as a linear relation is maintained [28].

which gives the same complexity as searching for single parameter changes, i.e., $O(n2^w)$ [8]. In this case, the solution, instead of intervals of single parameter changes, will be higher dimensional regions of multiple parameter changes. Therefore, from this solution space, we may want to suggest a "preferred" solution which minimizes the disturbance to the original parameters. One proposal is to return the solution on the "same log-odds curve", where the log-odds change in each parameter is the same.

Another important inference result we are often interested in is the *Most Probable Explanation* (MPE), defined as $\max_{\mathbf{x}} Pr(\mathbf{x} \mid \mathbf{e})$, which is the complete instantiation of all variables $\mathbf{X}$ which has the highest probability given the current evidence $\mathbf{e}$. Sensitivity analysis can also be applied to MPE, where we want to compute the allowable change in a single parameter such that the MPE remains unchanged, for every network parameter [10]. Because of the discrete nature of MPE, instead of a linear relationship between the probability of evidence and the meta-parameter $\tau_{x|\mathbf{u}}$, the MPE instantiation and its probability will change abruptly as we vary the parameter value. To find the allowable change in $\tau_{x|\mathbf{u}}$, we need to compute the following:

- The current probability of the MPE;
- Two constants, which we will call $r(\mathbf{e}, \tau_{x|\mathbf{u}})$ and $k(\mathbf{e}, \tau_{x|\mathbf{u}})$ respectively.

It turns out that the constants can be computed using a *maximizer circuit*, which can be obtained from the arithmetic circuit by changing the addition nodes into maximizer nodes. Therefore, the time and space complexity for solving this problem is also $O(n2^w)$.

The second problem we want to tackle is to bound the effects of parameter changes on any general query, *independent of the network*. For example, if a parameter changes from 0.03 to 0.05, what is the impact on a query whose current value is 0.7, without knowing the specifics of the Bayesian network? We notice that in many cases, a small absolute parameter change can induce a large absolute query change, and a small relative parameter change can also induce a large relative query change. Therefore, using absolute or relative change to quantify a parameter change is not possible to bound the amount of query change.

Assuming that we change the meta-parameter $\tau_{x|\mathbf{u}}$ by an arbitrary amount, such that the odds of x given $\mathbf{u}$, defined as $O(x \mid \mathbf{u}) = Pr(x \mid \mathbf{u})/Pr(\bar{x} \mid \mathbf{u})$, changes from $O(x \mid \mathbf{u})$ to $O'(x \mid \mathbf{u})$. Moreover, the query value of y given $\mathbf{e}$ changes as a result, from $O(y \mid \mathbf{e})$ to $O'(y \mid \mathbf{e})$. The following theorem bounds the amount of change in the query [7]:

$$|\ln O'(y \mid \mathbf{e}) - \ln O(y \mid \mathbf{e})| \leq |\ln O'(x \mid \mathbf{u}) - \ln O(x \mid \mathbf{u})|. \quad (2)$$

That is, the log-odds change (or the relative odds change) of the query is bounded by the log-odds change (or the relative odds change) of the parameter. Therefore, the log-odds change can be used to quantify a parameter change in a Bayesian network. Note that the above result is network-independent, and this bound can be computed in constant time without knowing the specifics of the network, and can serve as a preliminary recommendation when we are choosing between different parameter changes to be applied.

Finally, we may also be interested in quantifying the change in a probability distribution, instead of a single probability value. Again, we will use a measure in the same line as above, such that the measure can be used to bound the amount of change in any arbitrary query. The measure obtained, called the *Chan-Darwiche distance measure* [9], is defined as follows between two probability distributions, Pr and Pr', over the same set of variables $\mathbf{X}$:

$$CD(Pr, Pr') = \ln \max_{\mathbf{x}} \frac{Pr'(\mathbf{x})}{Pr(\mathbf{x})} - \min_{\mathbf{x}} \frac{Pr'(\mathbf{x})}{Pr(\mathbf{x})}. \quad (3)$$

We can prove that the distance measure bounds the change in any arbitrary query between the two distributions, such that:

$$|\ln O'(y \mid \mathbf{e}) - \ln O(y \mid \mathbf{e})| \leq CD(Pr, Pr'). \quad (4)$$

Therefore, this distance measure provides a worst-case bound on the change in any arbitrary query, and gives a useful quantification of the change from one probability distribution to another. On the other hand, other existing measures, such as the KL-divergence [19] and the Euclidean distance, cannot provide such a worst-case bound.

Another important property of this measure is that given two Bayesian networks that differ by only a single CPT, the distance between the probability distributions obtained from the two networks is simply the distance between the two CPTs. Therefore, the distance measure between two Bayesian networks can be easily computed in this case.

4. An Example of Sensitivity Analysis of Bayesian Networks

We now show a simple example of performing sensitivity analysis on a Bayesian network. For an example, we construct a Bayesian network to represent the decision of whether a customer buys a product, with four variables:

- *Interest* represents whether the customer is interested in this type of product.
- *Price* represents whether the price of the product is high or low.
- *History* represents whether the customer has previously bought the same type of product.
- *Buy* represents whether the customer buys the product.

We know that the customer's interest is a factor in determining whether he has previously bought the same type of product, while the price of the product and the customer's interest are both factors in determining whether he will buy this product. Therefore, we can obtain the structure as shown in Figure 1.

Now we compute query results using this Bayesian network. For example, the probability that a customer will

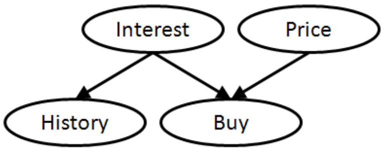

Interest	*Pr*
yes	0.25
no	0.75

Price	*Pr*
low	0.7
high	0.3

Interest	*History*	*Pr*
yes	*yes*	0.4
yes	*no*	0.6
yes	*yes*	0.05
yes	*no*	0.95

Interest	*Price*	*Buy*	*Pr*
yes	*low*	*yes*	0.5
yes	*low*	*no*	0.5
yes	*high*	*yes*	0.1
yes	*high*	*no*	0.9
no	*low*	*yes*	0.2
no	*low*	*no*	0.8
no	*high*	*yes*	0.01
no	*high*	*no*	0.99

Figure 1: An example Bayesian network.

buy the product given that the price is low and he has previously bought the same type of product is $Pr(Buy = yes \mid Price = low, History = yes) = 0.4182$.

However, after collecting customer data, we find that at least half of the customers in this case bought the product, so this probability should be higher, and the correct result should be $Pr(Buy = yes \mid Price = low, History = yes) \geq 0.5$. We now need to update the Bayesian network so that we can get the correct answer. We use SamIam, a Bayesian network software tool developed by the UCLA Automated Reasoning Group [31], to find the single parameter changes that can enforce the query constraint $Pr(Buy = yes \mid Price = low, History = yes) \geq 0.5$. Two suggestions of single parameter changes are returned, as shown in Figure 2:

1. Increase $Pr(Buy = yes \mid Interest = yes, Price = low)$ from 0.5 to ≥ 0.6125.

2. Increase $Pr(Buy = yes \mid Interest = no, Price = low)$ from 0.2 to ≥ 0.5.

We can choose to adopt either of the above parameter changes to our Bayesian network in order to satisfy our query constraint. From the two suggestions, the first one is close to the correct customer behavior, while the second one is not. So we will choose to adopt the first suggestion, and increase the probability of a customer buying a product given that he is interested in this type of product and the price is low, from 0.5 to at least 0.6125.

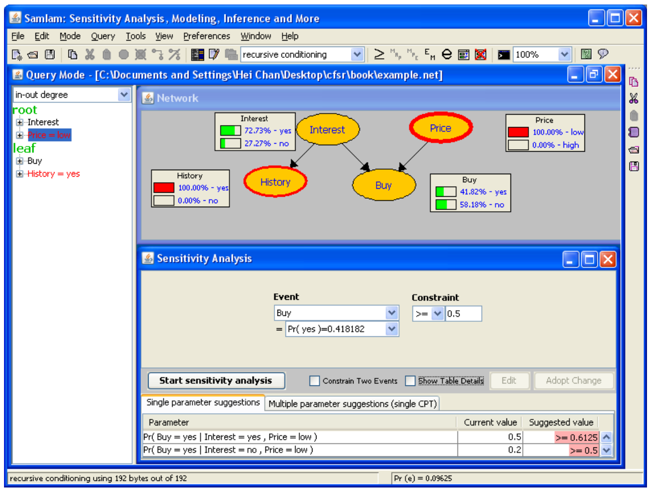

Figure 2: A screenshot of SamIam returning a list of suggestions of single parameter changes for enforcing a user-specified query constraint.

We can also ask SamIam to return multiple parameter changes to enforce the query constraint, as shown in Figure 3. The suggestion returned is to increase both parameters, $Pr(Buy = yes \mid Interest = yes, Price = low)$ from 0.5 to (at least) 0.5887, and $Pr(Buy = yes \mid Interest = no, Price = low)$ from 0.2 to (at least) 0.2635. Here, we see that the amount of increase for the parameter $Pr(Buy = yes \mid Interest = yes, Price = low)$ when we can change multiple parameters is only 0.0887 (from 0.5 to 0.5887), and thus is less than the amount of increase when we can change only single parameters, which is 0.1125 (from 0.5 to 0.6125).

5. Application on Designing Resilient Systems

Sensitivity analysis is an essential tool for designing resilient systems, by checking whether conclusions drawn from model are robust against uncertainty (e.g., estimation errors, environmental changes, unexpected events), and can be used in systems design and model debugging. For example, for model builders who design and debug models, they may ask the questions:

- What are the weak points of model that may contribute to large variations in output?

- What components we can change to improve model robustness?

While for decision makers who use and evaluate models, they may ask the questions:

- What are the causes of certain decisions being made based on the model?

- How confident are we in the decisions against uncertainty?

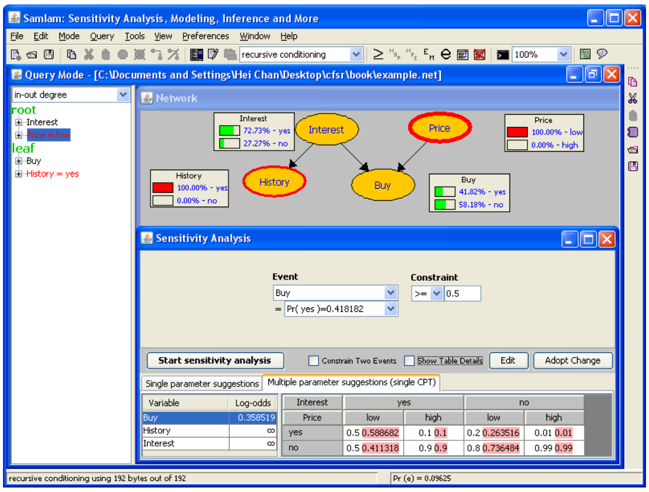

Figure 3: A screenshot of SAMIAM returning suggestions of multiple parameter changes for enforcing a user-specified query constraint.

Sensitivity analysis can help researchers and scientists build better models to represent the real world. We now present a few applications of sensitivity analysis on modeling and inference of Bayesian networks, from both a theoretical and an application point of view.

As we said earlier, in Bayesian network modeling, the structure and CPTs of Bayesian networks can be specified directly by domain experts, or obtained by learning algorithms from data. However, the former approach may be difficult if we need to construct a Bayesian network with many variables and edges. Instead, we will use learning algorithms to obtain the Bayesian network from data. To learn the structure of the Bayesian networks, we can use learning algorithms which may be information-based such as the AIC score [1], or constraint-based such as the PC algorithm [30]; To learn the CPTs of the Bayesian network, algorithms such as EM algorithm [21] can be applied.

Our sensitivity analysis results can be combined with these model learning algorithms, for measuring the deviation from the Bayesian network which is used as the data model to the data, ensuring that it faithfully represents the data, and also used to evaluate the contribution of any subset of the data to the Bayesian network. This allows the users of Bayesian networks to more thoroughly understand the constructed Bayesian network in terms of the data.

As for inference algorithms, because we may not be able to apply exact inference algorithms on large networks, approximate inference algorithms are often used to find approximate answers. In this case, our sensitivity analysis results can be used for developing new approximate inference algorithms with computable bounds, which allows the users of Bayesian networks to report the estimated errors in their data analysis results. For example, we can simplify a Bayesian network by deleting some of its edges, and obtain a new network with a simpler structure, where inference can be performed in less time. Sensitivity analysis can provide guarantees on the real query results from the approximate query results, by quantifying the amount of change from the real network to the approximate network. Recent research has shown promise in the inference method of loopy belief propagation with edge deletion [11, 12].

From an application point of view, the incorporation of sensitivity analysis into Bayesian network modeling and inference software tools can greatly benefit the users of Bayesian networks in modeling, understanding, and analyzing the data collected. First of all, an automated process of "tuning" network parameters according to user-specified constraints, feedback, and real-life results can help us build better Bayesian networks which more faithfully represent the data, and also revise them when new data or information is received. This can help reduce the costs and time that may otherwise be used to re-create new Bayesian networks, and also provide an interactive interface for the users of Bayesian networks in the modeling process.

Second, while many graphical Bayesian network software tools do a good job of providing users with visual cues indicating the current probabilities, they have not paid any attention to the change in the state of belief due to new data or information. However, it has been shown that users of these tools have difficulty visualizing probabilistic changes because of the large amount of information usually on display, the transient nature of the different probabilities, and the unfamiliarity of probability theory. Sensitivity analysis can thus be a great help for them for many purposes, such as gauging the strength of new evidence, or comparing the impact of a parameter change on different query values. By providing a measure that quantifies probability changes, we can help them visualize probabilistic changes and make better decisions.

As of now, most Bayesian network software tools do not have any implementation of sensitivity analysis. An exception is SAMIAM, which we presented in the previous section. SAMIAM includes tools that allow users to build a Bayesian network (either by hand or by learning algorithms), run different inference algorithms, and perform sensitivity analysis. We believe the implementation of new research results on sensitivity analysis in Bayesian network software tools should be an important step in improving the productivity and usability of these tools.

Finally, in Bayesian networks, the probability of evidence is linear in terms of any single variable. However, in other probabilistic graphical models with non-linear dependencies or continuous variables, such as dynamic Bayesian networks, the relations become non-linear. For example, recent sensitivity analysis results have been developed for hidden Markov models [27]. For example, in influence diagrams, which include both variable and decision nodes in order to choose a strategy of decisions to maximize pay-off, the relationship between the maximum pay-off (and the optimal strategy) is non-linear in terms of the parameters. Despite this, recent research has been applied in this domain and satisfactory results are obtained [3]. We aim to extend the results of sensitivity analysis of Bayesian networks to other probabilistic graphical models.

6. Acknowledgments

This work is supported by the "Systems Resilience" project of the Transdisciplinary Research Integration Center.

References

[1] H. Akaike, "A new look at the statistical model identification", *IEEE Transactions on Automatic Control*, Vol. 19, pp. 716–723, 1974.

[2] C. Baral, T. Eiter, M. Bjäreland, and M. Nakamura. Maintenance goals of agents in a dynamic environment: Formulation and policy construction. *Artificial Intelligence*, 172(12-13):1429–1469, 2008.

[3] D. Bhattacharjya and R. Shachter, "Sensitivity analysis in decision circuits", *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 34–42, AUAI Press, 2008.

[4] M. Bruneau. A Framework to Quantitatively Assess and Enhance the Seismic Resilience of Communities. In *Earthquake Spectra*, volume 19, 2003.

[5] E. Castillo, J. Gutiérrez, and A. Hadi, "Sensitivity analysis in discrete Bayesian networks", *IEEE Transactions on Systems, Man, and Cybernetics, Part A (Systems and Humans)*, Vol. 27, pp. 412–423, 1997.

[6] H. Chan, *Sensitivity Analysis of Probabilistic Graphical Models*, VDM Verlag Dr. Müller, 2009.

[7] H. Chan and A. Darwiche, "When do numbers really matter?", *Journal of Artificial Intelligence Research*, Vol. 17, pp. 265–287, 2002.

[8] H. Chan and A. Darwiche, "Sensitivity analysis in Bayesian networks: from single to multiple parameters", *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 67–75, AUAI Press, 2004.

[9] H. Chan and A. Darwiche, "A distance measure for bounding probabilistic belief change", *International Journal of Approximate Reasoning*, Vol. 38, pp. 149–174, 2005.

[10] H. Chan and A. Darwiche, "On the robustness of most probable explanations", *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 63–71, AUAI Press, 2006.

[11] A. Choi, H. Chan, and A. Darwiche, "On Bayesian network approximation by edge deletion", *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 128–135, AUAI Press, 2005.

[12] A. Choi and A. Darwiche, "An edge deletion semantics for belief propagation and its practical impact on approximation quality", *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI)*, pp. 1107–1114, AAAI Press, 2006.

[13] V. Coupé, N. Peek, J. Ottenkamp, and J. Habbema, "Using sensitivity analysis for efficient quantification of a belief network", *Artificial Intelligence in Medicine*, Vol. 17, pp. 223–247, 1999.

[14] A. Darwiche, "A differential approach to inference in Bayesian networks", *Journal of the ACM*, Vol. 50, pp. 280–305, 2003.

[15] A. Darwiche, *Modeling and Reasoning with Bayesian Networks*, Cambridge University Press, 2009.

[16] B. Faltings and S. Macho-Gonzalez. Open constraint programming. *Artificial Intelligence*, 161:181–208, 2005.

[17] C. Holling. Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 4:1–23, 1973.

[18] U. Kjærulff and L. Van der Gaag, "Making sensitivity analysis computationally efficient", *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 317–325, Morgan Kaufmann Publishers, 2000.

[19] S. Kullback and R. Leibler, "On information and sufficiency", *Annals of Mathematical Statistics*, Vol. 22, pp. 79–86, 1951.

[20] K. Laskey, "Sensitivity analysis for probability assessments in Bayesian networks", *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 25, pp. 901–909, 1995.

[21] S. Lauritzen, "The EM algorithm for graphical association models with missing data", *Computational Statistics and Data Analysis*, Vol. 19, pp. 191–201, 1995.

[22] K. Minami, T. Okimoto, T. Tanjo, N. Schwind, H. Chan, K. Inoue, H Maruyama. "Formalizing the Resilience of Open Dynamic System". *Joint Agent Workshop and Symposium (JAWS)*, 2012.

[23] C. M. Ozveren, A. S. Willsky, and P. J. Antsaklis. Stability and stabilizability of discrete event dynamic systems. *Journal of the ACM*, 38(3):7300–7752, 1991.

[24] J. Pearl, *Probabilistic Reasoning in Intelligence Systems*, Morgan Kaufmann, 1988.

[25] J. Pearl, *Causality: Models, Reasoning, and Inference*, Cambridge University Press, 2000.

[26] M. Pradhan, M. Henrion, G. Provan, B. Del Favero, and K. Huang, "The sensitivity of belief networks to imprecise probabilities: an experimental investigation", *Artificial Intelligence*, Vol. 85, pp. 363–397, 1996.

[27] S. Renooij, "Efficient sensitivity analysis in hidden Markov models", *International Journal of Approximate Reasoning*, Vol. 53, pp. 1397–1414, 2012.

[28] S. Renooij, "Generalised co-variation for sensitivity analysis in Bayesian networks", *Proceedings of the Sixth European Workshop on Probabilistic Graphical Models*, pp. 267–274, DECSAI Publications, 2012.

[29] N. Schwind, T. Okimoto, K. Inoue, H. Chan, T, Ribeiro, K. Minami, H. Maruyama, "Systems Resilience: A Challenge Problem for Dynamic Constraint-Based Agent Systems", *Proceedings of the Twelfth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2013.

[30] P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction, and Search*, MIT Press, 2001.

[31] UCLA Computer Science Department, Automated Reasoning Group, SamIam*: Sensitivity Analysis, Modeling, Inference and More*; URL: `http://reasoning.cs.ucla.edu/samiam/`.

[32] L. Van der Gaag and S. Renooij, "Analysing sensitivity data from probabilistic networks", *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 530–537, Morgan Kaufmann Publishers, 2001.

[33] B. H. Walker, C. S. Holling, S. C. Carpenter, and A.P. Kinzig. Resilience, adaptability and transformability. *Ecology and Society*, 9(2): 5, 2004.

[34] A. Zolli and A. M. Healy. *Resilience: Why Things Bounce Back*. Free Press, 2012.