

An Ontology-based Bayesian Network Approach for Representing Uncertainty in Clinical Practice Guidelines

Hai-tao Zheng, Bo-Yeong Kang, and Hong-Gee Kim*

Biomedical Knowledge Engineering Laboratory
Dentistry College, Seoul National University
28 Yeongeon-dong, Jongro-gu, Seoul, Korea

Abstract. Clinical Practice Guidelines (CPGs) play an important role in improving the quality of care and patient outcomes. Although several machine-readable representations of practice guidelines implemented with semantic web technologies have been presented, there is no implementation to represent uncertainty with respect to activity graphs in clinical practice guidelines. In this paper, we are exploring a Bayesian Network(BN) approach for representing the uncertainty in CPGs based on ontologies. Based on the representation of uncertainty in CPGs, when an activity occurs, we can evaluate its effect on the whole clinical process, which, in turn, can help doctors judge the risk of uncertainty for other activities, and make a decision. A variable elimination algorithm is applied to implement the BN inference and a validation of an aspirin therapy scenario for diabetic patients is proposed.

1 Introduction

Clinical Practice Guidelines (CPGs) play an important role in improving the quality of care and patient outcomes; therefore, the task of clinical guideline-sharing across different medical institutions is a prerequisite to many EMR (Electronic Medical Record) applications including medical data retrieval [18], medical knowledge management [7], and clinical decision support systems (CDSSs) [13]. To facilitate clinical guideline-sharing, GLIF (GuideLine Interchange Format) and SAGE (Standards-based Sharable Active Guideline Environment) have been the focus of extensive research [12]. GLIF is a semantic web based standard for representing clinical guidelines [15] and SAGE is an interoperable guideline execution engine, which encodes the content of the clinical guideline to an ontology representation, and executes the ontology through the functions of a CIS (clinical information system) [17].

Most previous approaches using GLIF and SAGE are designed to proceed from one step to the next only if there is no uncertain data in the former step [13]. However, this expectation is unrealistic in practice. For example, a guideline, which requires risk factors for heart disease to be assessed, needs to proceed

* Corresponding author: hgkim@snu.ac.kr

even when the information about this item is uncertain. In the clinical process, uncertain data can be (1) data stemming from unreliable sources (e.g., a patient can not remember the results of his/her last glucose test); (2) data not obtainable (e.g., no historical data on familial diabetes); and (3) data not yet collected (e.g., levels of serum glucose today) [14]. If data represented in CPGs is uncertain, the activities that handle these uncertain data become uncertain as well. For instance, in CDSS systems, when using the diabetes clinical guideline, it is necessary to get the family history for evaluating the risk of insulin therapy. However, in the real hospital environment, clinicians cannot easily obtain all the needed data for his/her health care activity. Based on these issues, the goal of this paper is to construct an approach to represent the uncertainty in CPGs and help doctors judge the risk of these uncertainties in the clinical process. Uncertainty in CPGs means that activity graphs that CPGs are composed of contain uncertain activities.

As a model for uncertainty, Bayesian Networks (BNs) occupy a prominent position in many medical decision making processes and statistical inference [11, 3, 2]. However, there have been few reports on applying BNs to the representation of uncertainty in CPGs. Therefore, to address this issue, we propose an ontology-based representation of uncertainty in CPGs by using BNs.

In this paper, we first introduce BNs, then we describe the use of BNs for the medical domain, and review previous work on applying semantic web technology to model CPGs in section 2; Section 3 elaborates the mechanism of encoding uncertainty into a CPG ontology; Section 4 describes a scenario validation based on BN inference; Section 5 discusses the conclusions and future work.

2 Background and Related Work

2.1 Bayesian Network

There are several models that are used to represent uncertainty, such as fuzzy-logic, BNs, etc. Generally, a BN of n variables consists of a DAG (Direct Acyclic Graph) of n nodes and a number of arcs. Nodes X_i in a DAG correspond to random variables, and directed arcs between two nodes represent direct causal or influential relations from one variable to the other. The uncertainty of the causal relationship is represented locally by the CPT (Conditional Probability Table). $P(X_i|pa(X_i))$ associated with each node X_i , where $pa(X_i)$ is the parent set of X_i . Under the conditional independence assumption, the joint probability distribution of $\mathbf{X} = (X_1, X_2, \dots, X_n)$ can be factored out as a product of the CPTs in the network, namely, the chain rule of BN: $P(\mathbf{X}) = \prod_i P(X_i|pa(X_i))$. With the joint probability distribution, BNs support, at least in theory, any probabilistic inference in the joint space. Besides the power of probabilistic reasoning provided by BNs themselves, we are attracted to BNs in this work for the structural similarity between the DAG of a BN and activity graphs of CPGs: both of them are directed graphs, and direct correspondence exists between many nodes and arcs in the two graphs. Moreover, BNs can be utilized to represent the uncertainty visually, provide inference effectively and facilitate human understanding

of CPGs. Considering the advantages of BNs, we apply BNs to represent the uncertainty in CPGs.

2.2 Bayesian Networks for the Medical Domain

Because BNs occupy a prominent position as a model for uncertainty in decision making and statistical inference, it has been applied to many medical decision support systems [11, 3, 2]. Atoui [3] adopted a decision making solution based on a BN that he trained to predict the risk of a cardiovascular event (infarction, stroke, or cardiovascular death) based on a set of demographic and clinical data. Aronsky [2] presented the development and the evaluation of a BN for the diagnosis of community-acquired pneumonia and he showed that BNs are an appropriate method to detect pneumonia patients with high accuracy. With respect to clinical guidelines, Mani [11] proposed BNs for the induction of decision tables and generated the guideline by these tables. However, although these methods focus on predicting some feature or risk of disease by using BN inference, there has been no implementation to represent the uncertainty with respect to activity graphs in CPGs and to reason on the uncertainty to provide the probabilities of target activities, which is the focus of our approach.

2.3 Semantic Web for Clinical Practice Guideline

A representational form of clinical guideline knowledge, which promotes completeness and minimizes inconsistency and redundancy, is essential if we want to implement and share guidelines for computer-based applications. Semantic Web technology offers such sharable and manageable methodology for modeling CPGs. GLIF [15] and SAGE [17] are two good examples. For creation and maintenance of implementable clinical guideline specifications, an architecture is presented in [8]. This architecture includes components such as a rules engine, an OWL-based classification engine and a data repository storing patient data. Moreover, approaches for modeling clinical guidelines are discussed and they show that guideline maintenance is tractable when a CPG is represented in an ontology. Here, we apply an ontology to represent the uncertainty in CPGs because it is more extensible and maintainable than other methods such as relational databases.

3 Encoding Uncertainty into a CPG Ontology

Figure 1 depicts the overall procedure of the proposed method. Firstly, the original CPG is encoded into an ontology model that contains uncertainty features using BNs. For this, we propose a formal model of CPG Ontology to represent uncertainty and an algorithm to construct the CPTs (Conditional Probability Tables) of the BN. The CPG ontology can be shared and utilized in different clinical information systems. Then, when a user provides his/her observed evidence in the clinical process, the BN inference engine will load the CPG ontology

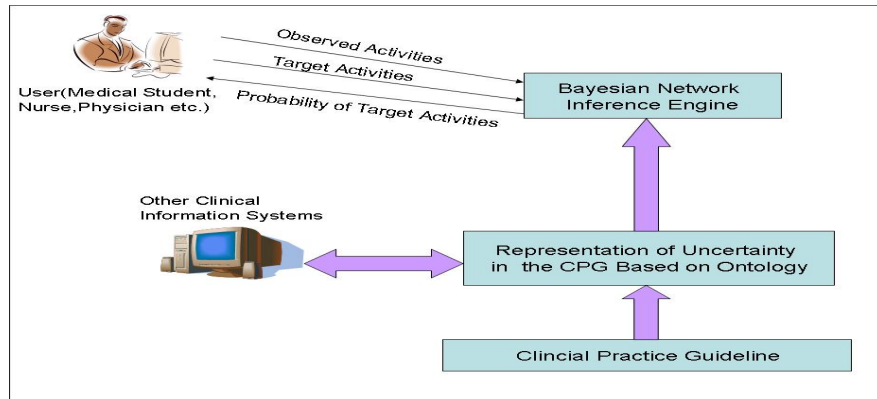


Fig. 1. The framework

as a BN and mark the nodes that are observed by the user in the BN. Based on the observed evidence, the BN inference engine can reason out the probabilities of target activities asked by the user. Given the reasoning results, the user can judge the risk of unobserved activities and make a further decision.




Fig. 2. Clinical practice guideline of aspirin therapy for diabetic patients (ASA means aspirin therapy)

3.1 Clinical Practice Guideline Ontology

CPGs typically include multiple recommendation sets represented as an activity graph that show the recommended activities during a clinical process and [4]. An activity graph describes the relationship between activities in the recommendation set as a process model. In this article, we use a single recommendation set in the SAGE diabetes CPG [1], which is an activity graph of aspirin therapy for diabetic patients, to illustrate how we represent the uncertainty in CPGs

based on ontology (Fig. 2). Typically, an activity graph contains three kinds of activities, i.e., context activity, decision activity, action activity. Each activity graph segment within a guideline begins with a context activity node that serves as a control point in guideline execution by specifying the clinical context for that segment. A decision activity node in the SAGE guideline model represents clinical decision logic by listing alternatives (typically subsequent action activity nodes), and specifying the criteria that need to be met to reach those nodes. An action activity node encapsulates a set of work items that must be performed by either a computer system or persons.

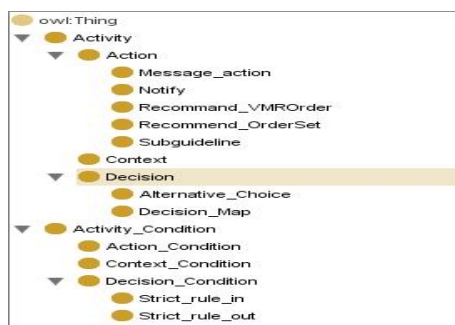


Fig. 3. Classes representation for clinical practice guideline

To represent activities in CPGs, we create the activity class that represents all the nodes in activity graph as shown in Figure 3. Because there are three kinds of activities, we construct an action class, a context class, and a decision class as sub classes of the activity class in the ontology respectively. In CPGs, activities may include internal conditions that restrict their execution. For example, for the decision activity “Yes;check for ASA(aspirin therapy) contraindications” (Fig. 2), there are many internal conditions, such as checking presence of family history, checking presence of hypertensive disorder etc., to make sure the ASA contraindications will be checked correctly. We encode these internal conditions of activity as an activity condition class in the ontology (Fig. 3).

A CPG Ontology with uncertainty features is defined as follows:

Definition 1. (CPG Ontology) CPG Ontology $O := \{C, I, Ps, cinst\}$, with an activity class set C , an activity instance set I , a property set Ps , and an activity class instantiation function $cinst : C \rightarrow 2^I$.

In CPG ontology, the activity instance set I represents the set of real activities that belong to activity classes accordingly. The property set Ps is proposed to represent the different attributes of activities in order to encode the features of the BN into ontology. The property set Ps is defined as follows:

Definition 2. (Properties for uncertainty representation) Property Set $Ps := \{cause, hasCondition, hasState, isObserved, hasPriorPro Value, hasCondiPro Value\}$,

has a property function $cause : I \rightarrow I$, a property function $hasCondition : I \rightarrow I$, a property function $hasState : I \rightarrow Boolean$, a property function $isObserved : I \rightarrow Boolean$, a property function $hasPriorProValue : I \rightarrow Float$, and a property function $hasCondiProValue : I \rightarrow Float$.

In CPGs, if the criteria associated with an activity node are satisfied, it will be successfully executed, which will cause the execution of subsequent nodes in the activity graph. Therefore, the relationship between activities is called the *cause* relationship. For example, in Figure 2, the context activity “Patient 21 years or older” causes the decision activity “Check for Aspirin therapy”. To represent this relationship in the ontology, we construct the object property *cause* whose domain and range are activity class and activity condition class. The *hasCondition* property is proposed as inverse properties of the *cause* property, which describes the “parent” activities of an activity that *cause* its execution. For instance, the decision activity “Check for aspirin therapy” has the property *hasCondition* with value “Patient 21 years or older” activity that causes its execution. With the *hasCondition* property, users can easily figure out all the conditions that cause the execution of any activity. The *cause* property plays the role of “directed arc” and all the activity instances play the role of “node” in the DAG of BN. Another property, the *hasState* property, which has a boolean value range, is denoted as the state of the activity instance; the *isObserved* property shows if the activity instances have been observed or not.

Prior probability and conditional probability are two features that represent the uncertainty level of nodes in BNs. To encode prior probability and conditional probability of activity instances into the ontology respectively, *hasPriorProValue* property and *hasCondiProValue* property are employed. Let A , B be the instances of the activity class representing two concrete activities. We interpret $P(A = a)$ as the prior probability that a value a is a state of instance A and $P(B = b|A = a)$ as the conditional probability that when A has state a , B has state b . For example, when A is activity “Patient 21 years or older”, B is activity “Check for Aspirin therapy”, $P(A = true) = 0.5$ can be expressed as follows:

```
<Context rdf:ID="Patient_21_yo_or_older">
  <hasPriorProValue
    rdf:datatype="http://www.w3.org/2001/XMLSchema#float"
    >0.5</hasPriorProValue>
  <hasState
    rdf:datatype="http://www.w3.org/2001/XMLSchema#boolean"
    >true</hasState>
  <cause rdf:resource="#Check_for_Aspirin_therapy"/>
</Context>
```

The conditional probability $P(B = true|A = true) = 1.0$ can be expressed as follows:

```
<Decision rdf:ID="Check_for_Aspirin_therapy">
  <hasCondition>
    <Context rdf:ID="Patient_21_yo_or_older">
      <hasState rdf:datatype="http://www.w3.org/2001/XMLSchema#boolean"
```

```

    >true</hasState>
  </Context>
</hasCondition>
<hasState rdf:datatype="http://www.w3.org/2001/XMLSchema#boolean"
>true</hasState>
<hasCondiProValue rdf:datatype="http://www.w3.org/2001/XMLSchema#float"
>1.0</hasCondiProValue>
<cause rdf:resource="#Check_fo_age_older_than_40"/>
</Decision>

```

3.2 Construction of Conditional Probability Tables




Fig. 4. Algorithm for constructing the conditional probability tables

In this section, we introduce an algorithm utilized to construct the CPTs of activity instances. After creating the properties to represent the uncertainty in the ontology, another important work is the construction of the CPTs for the BN because BN inference is based on the CPTs of each node in BN. For building the CPTs of activity instance in CPGs, we propose the encoding algorithm according to the features of CPGs. Each activity X_i in a CPG η has a corresponding activity instance X_{oi} in CPG Ontology O , i.e., we mark the corresponding activity instance by adding the letter “o” to the activity variable.

This algorithm provides principles to initialize the CPTs of BNs. The instances in the definition mean the activity instances in the CPGs. Firstly, we assign prior probabilities to activity instances, i.e., only when the activity instances have no “parents” in BN, they have prior probabilities. When an activity A causes the set of activities $\{B_1, B_2, \dots, B_n\}$ simultaneously, the conditional probability $P(B_i|A) = 1.0, (i = 1, \dots, n)$ (Fig. 4(a)); when an activity A causes one of the activities $\{B_1, B_2, \dots, B_n\}$, the conditional probability $P(B_i|A) = 1.0/n, (i = 1, \dots, n)$ (Fig. 4(b)); when a set of activities $\{A_1, A_2, \dots, A_n\}$ cause activity B together, then $P(B|A_1, A_2, \dots, A_n) = 1.0$ (Fig. 4(c)); when one of the activities $\{A_1, A_2, \dots, A_n\}$ can cause activity B , then $P(B|\overline{A_1}, \overline{A_2}, \dots, \overline{A_n}) = 0.0$ (Fig. 4(d)).

With the initialization of CPTs, we have finished constructing the BN from an ontology that represents the uncertainty in CPGs, namely, the activity graphs

containing uncertain activities are represented using a BN. When a BN inference engine loads this ontology, the ontology will be converted to a BN for BN inference. All the instances of the activity class and the activity condition class are translated into the node of the BN whose properties are also converted from properties of these instances in ontology accordingly. In the BN, an arc is drawn between nodes if the corresponding two activity instances are related by a *cause* property, with the direction from the activity instance that has *cause* property to the value of this property. CPTs in the BN are also easily obtained from property *hasCondiProValue* and property *hasPriorProValue* of corresponding activity instances.

Algorithm 1 Construct CPTs(CPG η , CPG Ontology O)

```

for each activity  $X_i$  in  $\eta$  do
  if  $X_i$  can be successfully executed then
    Set property hasState of activity instance  $Xo_i$  in O with value true
  end if
  if there is no activity that causes the execution of  $X_i$  then
    Set property hasCondition of activity instance  $Xo_i$  in O with value null
    Set property hasPriorProValue of  $Xo_i$  with value 0.5
  end if
end for
if activity  $A$  in  $\eta$  cause the execution of activities  $\bigcap_{i=1}^n B_i$  then
  Set property hasCondiProValue of activity instance  $Bo_i$  in O with value 1.0
  Set activity instance  $Ao$  as the value of property hasCondition of activity instance  $Bo_i$  in O
  Set property hasState of activity instance  $Ao$  and  $Bo_i$  in O with value true
end if
if activity  $A$  in  $\eta$  cause the execution of activities  $\bigcup_{i=1}^n B_i$  then
  Set property hasCondiProValue of activity instance  $Bo_i$  in O with value  $1.0/n$ 
  Set activity instance  $Ao$  as the value of property hasCondition of activity instance  $Bo_i$  in O
  Set property hasState of activity instance  $Ao$  and  $Bo_i$  in O with value true
end if
if activities  $\bigcap_{i=1}^n A_i$  in  $\eta$  cause the execution of activities  $B$  then
  Set property hasCondiProValue of activity instance  $Bo$  in O with value 1.0
  Set activity instances  $Ao_1, \dots, Ao_n$  as the value of property hasCondition of activity instance  $Bo$  in O
  Set property hasState of activity instance  $Ao_i$  and  $Bo$  in O with value true
end if
if activities  $\bigcup_{i=1}^n A_i$  in  $\eta$  cause the execution of activities  $B$  then
  Set property hasCondiProValue of activity instance  $Bo$  in O with value 0.0
  Set activity instances  $Ao_1, \dots, Ao_n$  as the value of property hasCondition of activity instance  $Bo$  in O
  Set property hasState of activity instance  $Ao_i$  with value false
  Set property hasState of activity instance  $Bo$  with value true
end if

```

4 A Scenario Validation Based on Bayesian Network Inference

We apply the variable elimination algorithm [9, 5] to perform BN inference. To verify the feasibility of our approach, a scenario of aspirin therapy for a diabetic patient is proposed. Based on this scenario, we apply our ontology-based BN approach to represent the uncertainty in CPGs and carried out the BN inference based on this BN.

4.1 Bayesian Network Inference

There are a lot of algorithms that manipulate BNs to produce posterior values [16, 10]. The variable elimination algorithm [9, 5] and the bucket elimination algorithm [6] are focused on algebraic operations. Since algebraic schemes like variable and bucket elimination compute marginal probability values for a given set of variables that is suitable for inference on observed evidence, we apply the variable elimination algorithm to implement the BN inference on the uncertainty of CPGs.

We assume all random variables have a finite number of possible values. Set of variables are denoted in bold; for instance, \mathbf{X} . The set of all variables that belong to \mathbf{X} but do not belong to \mathbf{Y} is indicated by $\mathbf{X} \setminus \mathbf{Y}$. The expression $\sum_{\mathbf{X}} f(\mathbf{X}, \mathbf{Y})$ indicates that all variables in \mathbf{X} are summed out from the function $f(\mathbf{X}, \mathbf{Y})$. Denoted by $P(X)$ is the *probability density of X*: $P(x)$ is the probability measure of the event $\{X = x\}$. Denoted by $P(X|Y)$ is the probability density of X conditional on values of Y .

Given a BN, the event E denotes the *observed evidence* in the network. Denoted by \mathbf{X}_E is the set of observed variables. Inferences with BNs usually involve the calculation of the posterior marginal for a set of query variables \mathbf{X}_q . The posterior of \mathbf{X}_q given E is:

$$P(\mathbf{X}_q|E) = \frac{P(\mathbf{X}_q, E)}{P(E)} = \frac{\sum_{\mathbf{X} \setminus \{\mathbf{X}_q, \mathbf{X}_E\}} P(\mathbf{X})}{\sum_{\mathbf{X} \setminus \mathbf{X}_E} P(\mathbf{X})} \quad (1)$$

The detail of variable elimination algorithm can be found in [9, 5].

4.2 A Validation of an Aspirin Therapy Scenario for Diabetic Patients

We demonstrate the validity of our approach by applying an experiment to the CPG of aspirin therapy for diabetic patients (Fig. 2). Let us consider a scenario:

Scenario 1 *A user (medical student, nurse or physician etc.) is trying to apply aspirin therapy for a diabetic patient using the diabetes CPG. When he/she tries to check the aspirin risk factors, he/she can get a few observed evidence, such as observations of hypertensive disorder, tobacco user finding, hyperlipidemia, and myocardial infarction. In this case, the user wants to evaluate target activities*

that he is concerned about in this CPG. In this way, he hopes the results can help him understand the effect of the observed evidence on the target activities during the whole clinical process.

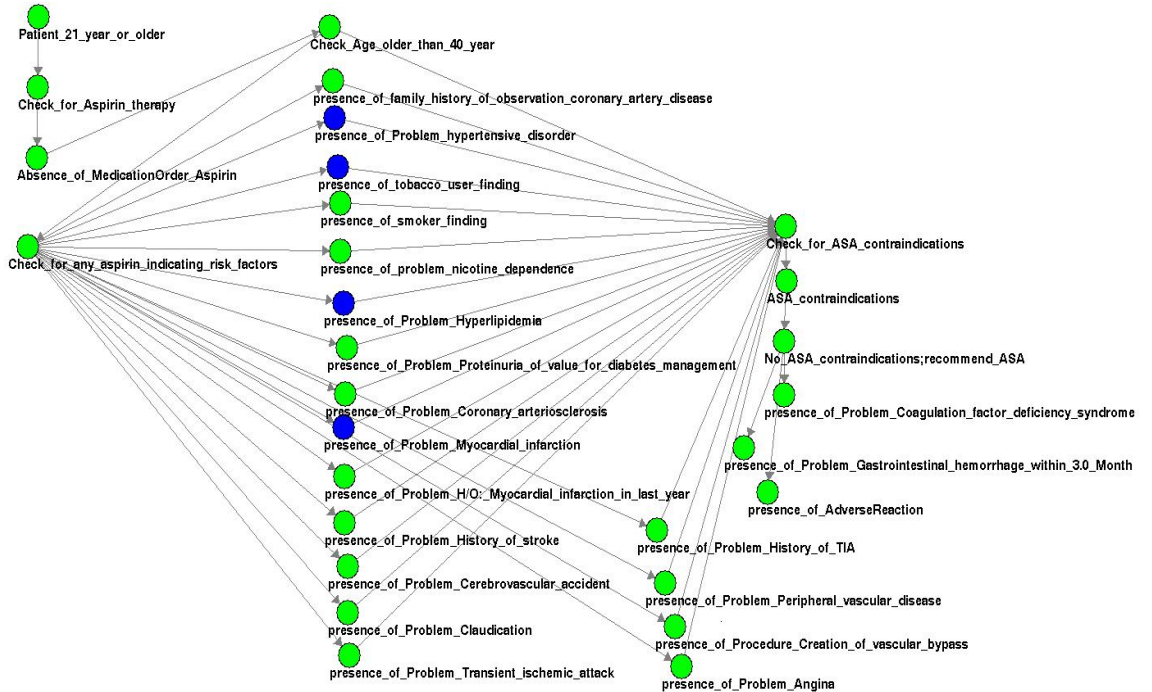


Fig. 5. An ontology based Bayesian network of aspirin therapy for diabetic patients derived from figure 2 (blue nodes are the observed ones)

In the scenario, the CPG of aspirin therapy for diabetic patients is used. Since there are some uncertain activities with respect to the activity graph in this CPG, we can apply our ontology-based BN approach to represent this uncertainty. Details are described in Section 3. Figure 5 shows the ontology-based BN representing the uncertainty in the CPG of aspirin therapy for diabetic patients.

After loading the ontology-based BN, the BN inference engine can process BN inference when the user provides his/her observed evidence, such as observations of hypertensive disorder, tobacco user finding, hyperlipidemia, and myocardial infarction in this scenario (Fig. 5). If the user selects the target activities, the BN inference engine can calculate the probability of them by using the variable elimination.

For example, after the user got the observed evidence of some aspirin risk factors, he wants to know the probability of activity “No ASA (aspirin therapy) contraindications; recommend ASA” to help him to judge whether or not his ob-

servations of aspirin risk factors are adequate. In the BN inference engine, since the activity instance “presence of problem hypertensive disorder” is observed, its property *isObserved* is set *true* and the property *hasState* is set *false*. Similarly, the activities instances “presence of problem myocardial infarction”, “presence of tobacco user finding”, and “presence of problem hyperlipidemia” are also set in the same manner. With CPTs in this BN, equation 1 (Section 4.1) is applied to calculate the probability of activity instance “No ASA contraindications; recommend ASA” :

$$P(\mathbf{X}_q|E) = \frac{P(\mathbf{X}_q, E)}{P(E)} = \frac{\sum_{\mathbf{X} \setminus \{\mathbf{X}_q, \mathbf{X}_E\}} P(\mathbf{X})}{\sum_{\mathbf{X} \setminus \mathbf{X}_E} P(\mathbf{X})} = 0.775$$

where $\mathbf{X}_q = \{ \text{“No ASA contraindications; recommend ASA”} \}$, and $E = \{ \text{“presence of problem hypertensive disorder”} = \textit{false}, \text{“presence of problem myocardial infarction”} = \textit{false}, \text{“presence of tobacco user finding”} = \textit{false}, \text{“presence of problem hyperlipidemia”} = \textit{false} \}$.

In another case, when the user wants to get the uncertain degree of activity instance “presence of problem coagulation factor deficiency syndrome”, he can choose this target activity instance based on the observed evidence E . Through BN inference, we can obtain:

$$P(\mathbf{X}_q|E) = \frac{P(\mathbf{X}_q, E)}{P(E)} = 0.6425$$

where $\mathbf{X}_q = \{ \text{“presence of problem coagulation factor deficiency syndrome”} \}$ and E is the same as the above case.

The results in the two cases show high probabilities for the target activities, which suggest the user can make a decision to go ahead based on the observed evidence. When we consult several medical experts with this scenario, their opinions are coincident with these results, which shows the feasibility of our approach.

5 Conclusion and future work

In this paper, we contribute an ontology-based BN approach to represent the uncertainty in CPGs. With this uncertain representation in ontology, computers can: (1) calculate the uncertainty of target activities in CPGs; (2) remind users of the missing important data or event items, which should be observed in the clinical process; (3) simulate the clinical process under uncertain situations, which can be applied to e-learning systems in medical schools.

In the future, we are planning to combine our approach with a real CIS environment and apply uncertain clinical data to our application. A more comprehensive evaluation based on real clinical data should also be carried out.

Acknowledgements

This study was supported by a grant of National Project for Information Technology Advancement, Ministry of Information and Communication, and the

Interoperable EHR Research and Development Center(A050909), Ministry of Health & Welfare, Republic of Korea.

References

1. Sage diabetes guideline. <http://sage.wherever.org/cpgs/diabetes/diabetes.html/phtml.html>.
2. D. Aronsky and P. J. Haug. Diagnosing community-acquired pneumonia with a bayesian network. In *Proc AMIA Symp*, pages 632–636, 1998.
3. H. Atoui, J. Fayn, F. Gueyffier, and P. Rubel. Cardiovascular risk stratification in decision support systems:a probabilistic approach. application to phealth. *Computers in Cardiology*, 33:281–284, 2006.
4. J. Campbell, S. Tu, J. Mansfield, J. Boyer, J. McClay, C. Parker, P. Ram, S. Scheitel, and K. McDonald. The sage guideline model:a knowledge representation framework for encoding interoperable cpgs. *Stanford Medical Informatics Report SMI-2003-0962*, 2003.
5. F. G. Cozman. Generalizing variable elimination in bayesian networks. In *Workshop on Probabilistic Reasoning in Artificial Intelligence*, pages 27–32, 2000.
6. R. DECHTER. Bucket elimination: A unifying framework for probabilistic inference. In *M. I. Jordan, editor, Learning in Graphical Models*, MITPress, pages 75–104, 1999.
7. G. Hripcsak. Writing arden syntax medical logic modules. *Comput Biol Med*, 24:331–363, 1994.
8. V. Kashyap, A. Morales, and T. Hongsermeier. Creation and maintenance of implementable clinical guideline specifications. In *ISWC 2005*, 2005.
9. Z. N. L and P. D. Exploiting causal independence in bayesian network inference. *Artificial Intelligence Research*, pages 301–328, 1996.
10. S. L. Lauritzen and D. J. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical SocietyB*, 50(2):157–224, 1988.
11. S. Mani and M. J. Pazzani. Guideline generation from data by induction of decision tables using a bayesian network framework. *JAMIA supplement*, pages 518–522, 1998.
12. A. H. Morris. Developing and implementing computerized protocols for standardization of clinical decisions. *Annal of Internal Medicine*, 132(5):373–383, 2000.
13. M. Musen, S. Tu, A. Das, and Y. Shahar. Eon:a component-based approach to automation of protocol-directedtherapy. *J Am Med Inform Assoc*, 2:367–388, 1996.
14. L. Ohno-Machado. Representing uncertainty in clinical practice guidelines. In *An Invitational Workshop: Towards Representations for Sharable Guidelines*, March 2000.
15. L. Ohno-Machado, S. N. Murphy, D. E. Oliver, R. A. Greenes, and G. O. Barnett. The guideline interchange format: A model for representing guidelines. *Journal of the Americal Medical Informatics Association*, 5(4):357–372, Jul 1998.
16. J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, SanMateo, California, 1988.
17. P. Ram, D. Berg, S. Tu, G. Mansfield, Q. Ye, R. Abarbanel, and N. Beard. Executing clinical practice guidelines using the sage execution engine. *Medinfo*, pages 251–255, 2004.
18. P. Stoufflet, L. Ohno-Machado, S. Deibel, D. Lee, and R. Greenes. Geode-cm:a state-transition framework for clinical management. In *20th Annu Symp Comput Appl Med Care*, page 924, 1996.