



ELSEVIER

Available at
www.ComputerScienceWeb.com
POWERED BY SCIENCE @ DIRECT®

INTERNATIONAL JOURNAL OF
**APPROXIMATE
REASONING**

International Journal of Approximate Reasoning 33 (2003) 71–94

www.elsevier.com/locate/ijar

Implicitly preserving semantics during incremental knowledge base acquisition under uncertainty

Eugene Santos Jr.^{a,*}, Eugene S. Santos^b,
Solomon Eyal Shimony^c

^a *Department of Computer Science and Engineering, University of Connecticut, UTEB,
191 Auditorium Rd, U-155, Storrs, CT 06269-3155, USA*

^b *Department of Computer and Information Sciences, Youngstown State University,
Youngstown, OH 44555, USA*

^c *Department of Computer Science, Ben Gurion University of the Negev, 84105 Beer-Sheva, Israel*

Received 1 December 2001; accepted 1 November 2002

Abstract

New knowledge is incrementally introduced to an existing knowledge base in a typical knowledge-engineering cycle. Unfortunately, at most given stages, the knowledge-base is incomplete but must still satisfy sufficient consistency conditions in order to provide sound semantics. Maintaining semantics for uncertainty is of primary concern. We examine Bayesian knowledge bases (BKBs), which are a generalization of Bayesian networks. BKBs provide a highly flexible and intuitive representation following a basic “if-then” structure in conjunction with probability theory. We present new theoretical and algorithmic results concerning BKBs and how they can naturally and implicitly preserve semantics as new knowledge is added. In particular, equivalence of rule weights and conditional probabilities is achieved through stability of inferencing in BKBs. Furthermore, efficient algorithms are developed to guarantee stability of BKBs during construction. Finally, we examine and prove formal conditions that hold during the incremental construction of BKBs.

© 2002 Elsevier Science Inc. All rights reserved.

* Corresponding author. Tel.: +1-860-486-5955; fax: +1-860-486-4817.

E-mail addresses: eugene@cse.uconn.edu (E. Santos Jr.), santos@cis.ysu.edu (E.S. Santos), shimony@cs.bgu.ac.il (S.E. Shimony).

Keywords: Bayesian knowledge-bases; Probabilistic semantics; Knowledge acquisition; Uncertainty; Knowledge engineering

1. Introduction

Knowledge acquisition is an inherently sequential process. The elicitation, encoding, and testing of knowledge by human knowledge engineers follows a necessary cycle in order to obtain the required knowledge critical to constructing a usable knowledge-based system. Thus, new knowledge is incrementally introduced to the existing knowledge base as the cycle progresses [2,3,8,10,11,16,19,22–24]. Unfortunately, it is rarely the case that complete knowledge is ever available except in very specific and often simplistic domains. New knowledge is often discovered and uncovered during construction as well as even after the knowledge-based system has been fielded. Hence, at any stage, the knowledge-base is actually incomplete but must still satisfy sufficient consistency conditions in order to provide sound semantics among the knowledge/information it does have.

Uncertainty is a primary facet of incompleteness that pervades every stage of the knowledge acquisition cycle. It is well known that the problem inherent in managing uncertainty lies with how multiple sources of uncertainty interact. Without a sound and consistent semantics of uncertainty, the resulting interactions are ad hoc, unpredictable, and often counter-intuitive. Thus, the *key* difficulty during incremental knowledge acquisition lies in preserving the semantics of the knowledge-base as new knowledge is introduced. This is especially important to human knowledge engineers who themselves are attempting to maintain their own consistent internal picture of the target domain.

Approaches to maintaining semantic consistency during acquisition under uncertainty can either (1) enforce strict local semantic assumptions or (2) require extensive modifications and recomputations over the existing knowledge-base, to accommodate new knowledge. The former can be accomplished by restricting the acquisition of new knowledge in such a way that existing semantic conditions/assumptions are never violated. For the latter, the impacts of the new knowledge is essentially propagated throughout the existing knowledge-base in an effort to maintain consistency. Systems such as assumption-based truth maintenance systems rely on extensive updates throughout their knowledge base, which can be an extremely expensive process [7]. Thus, a critical *goal* for knowledge engineering is to have an approach that guarantees precise and intuitive local semantics while minimizing the maintenance expense of global semantic consistency.

In Bayesian networks (BNs) [13,14], the semantics of uncertainty are represented by probabilistic conditional independence which can be directly re-

lated to notions of causality. BNs require that conditional (in)dependence be based on a directed acyclic graph of random variables.¹ Additions made to a BN are reflected as changes in the underlying graph structure. Such changes, should they occur in the interior of the graph, affect the conditional independence semantics of not just the new knowledge introduced, but also nearly all the reachable nodes from the affected region. Local semantics with respect to the immediate neighbors are established directly by the knowledge engineer. Thus, we find that BNs trivially preserve the original local semantics for the knowledge engineer as the knowledge-base evolves and little additional maintenance computations are needed.

Our goal in this paper is to further address the preservation of semantics during incremental knowledge acquisition under uncertainty. In particular, we examine Bayesian knowledge bases (BKBs) which are a generalization of BNs [25]. BKBs have been extensively studied both theoretically [9,17,29,30] and for use in knowledge engineering [21–24]. BKBs provide a highly flexible and intuitive representation following a basic “if-then” structure in conjunction with probability theory. Furthermore, BKBs were designed keeping in mind typical domain incompleteness to retain semantic consistency as well as soundness of inference in the absence of complete knowledge. BNs, on the other hand, typically assume a complete probability distribution is available from the start. Also, BKBs have been shown to capture knowledge at a finer level of detail as well as knowledge that would be cyclical (hence disallowed) in BNs.

Probabilistic models exhibiting significant local structure are common. In such models, explicit representation of that structure as done in BKBs, is advantageous, as the resulting representation is much more compact than the full table representation of the conditional probability tables (CPTs) in a BN. For example, consider the following setting: X , a binary variable, is known to be true if any of the variables Y_i is true, for $1 \leq i \leq n$, and X is false with probability p otherwise. The global structure here is that X depends on all the Y_i and in a BN one might represent this with a set of arcs $\{(Y_i, X) | 1 \leq i \leq n\}$. The representation of the distribution in the “standard” form of a CPT would require $O(2^n)$ entries. However, the (partially) given distribution also exhibits “local” structure, as when Y_i is known to be true for some i , X no longer depends on the value of Y_j for $j \neq i$. The size of the representation of the conditional probabilities in terms of rules is only $O(n)$. Although work has been done on representing local structure using other methods, such as local decision trees and default tables [4], rules have significant advantages in size of the

¹ Nodes in a BN represent random variables and the arcs represent direct conditional dependence between the random variables. The graphical concept of d -separation is the basis of determining conditional independence between any sets of random variables in a BN [5,6,12–14,20,28,31].

representation, as well as their better explainability. For example, contrasting rules with decision trees as a representation of local structure, every decision tree is compactly representable as a set of rules, while the reverse is not necessarily true—the decision tree may be exponentially larger than the set of rules [1]. Although rule-based systems for representing an exact distribution exist (e.g. [15]), these systems are a (compact) notational variant of Bayes' networks, and are thus less flexible than BKBs, as they do not allow for incompleteness or cyclicity.

Given the ability of BNs to easily maintain local semantics for the knowledge engineer and the relationship between BNs and BKB, can BKBs also provide such capabilities for incremental knowledge acquisition in light of BKBs' added representational power? In particular, we consider the following modifications to the knowledge base, and how they may affect the semantics: (1) adding and deleting rules, and (2) changing the rule weights ("conditional probabilities"). One related issue, called "completeness" of a BKB, that we address is: given a partial BKB, that represents an incompletely specified distribution, what are the conditions for guaranteeing that there *exists* a set of rules that, after they are added, makes the BKB a completely specified and consistent distribution?

In this paper, we present new theoretical and algorithmic results on BKBs, and how the model can naturally and implicitly preserve semantics as new knowledge is added, by addressing the above problems. Additionally, equivalence of rule weights and conditional probabilities is achieved through stability of inferencing in BKBs. Furthermore, efficient algorithms are developed to guarantee stability of BKBs during construction. Finally, we present formal conditions concerning the incremental construction and completeness of BKBs.

We begin in Section 2 by formally describing the BKB representation and inferencing mechanisms. Section 3 then presents our results—the semantics of rule probabilities in BKBs, followed by stability as a sufficient condition for preserving the semantics during incremental knowledge acquisition and construction of BKBs.

2. Bayesian knowledge bases

In this section, we provide the formal definition for *Bayesian knowledge bases* to represent and reason over uncertain information based on a sound probabilistic framework. The formulation presented here is slightly different from existing definitions found in [22,25,30] but is equivalent. This formulation helps better emphasize the incremental nature of knowledge acquisition in order to provide better intuitions concerning our results in the next section.

Let $A_1, A_2, \dots, A_k, \dots$ be a collection of finite discrete random variables (rvs) where $r(A_i)$ denotes the set of possible values for A_i .

Definition 2.1. A conditional probability rule (CPR), R , is of the form

$$R : A_{i_1} = a_{i_1} \wedge A_{i_2} = a_{i_2} \wedge \cdots \wedge A_{i_{n-1}} = a_{i_{n-1}} \Rightarrow A_{i_n} = a_{i_n}$$

for some positive n , where $a_{i_j} \in r(A_{i_j})$ such that $i_j \neq i_k$ for all $j \neq k$. Rules have an associated *weight*, denoted by $P(R)$.

The left-hand side of R is said to be the *antecedent* of R and the right-hand side the *consequent* of R . We denote these, respectively, by $\text{ant}(R)$ and $\text{con}(R)$. When $n = 1$, $\text{ant}(R)$ is the empty set and we write R as follows:

$$R : \text{true} \Rightarrow A_{i_n} = a_{i_n}.$$

The weight $P(R)$ will be shown to correspond to the conditional probability of R as we shall see in the next section.

Definition 2.2. Given two CPRs

$$R_1 : A_{i_1} = a_{i_1} \wedge A_{i_2} = a_{i_2} \wedge \cdots \wedge A_{i_{n-1}} = a_{i_{n-1}} \Rightarrow A_{i_n} = a_{i_n}$$

and

$$R_2 : A_{j_1} = a'_{j_1} \wedge A_{j_2} = a'_{j_2} \wedge \cdots \wedge A_{j_{m-1}} = a'_{j_{m-1}} \Rightarrow A_{j_m} = a'_{j_m}.$$

We say that R_1 and R_2 are mutually exclusive if there exists some $1 \leq k < n$ and $1 \leq l < m$ such that $i_k = j_l$ and $a_{i_k} \neq a'_{j_l}$.

In essence, Definition 2.2 states that R_1 and R_2 must differ by at least one rv value assignment in their antecedents.

Definition 2.3. R_1 and R_2 are said to be consequent-bound if (1) for all $k < n$ and $l < m$, $a_{i_k} = a'_{j_l}$ whenever $i_k = j_l$, and (2) $i_n = j_m$ but $a_{i_n} \neq a'_{j_m}$.

Proposition 2.1. If R_1 is consequent-bound with R_2 , then R_1 and R_2 are not mutually exclusive.

Consequent-boundedness simply indicates that the difference between R_1 and R_2 only occurs in the consequents of both CPRs. Intuitively, R_1 and R_2 are opposing rules to apply when both antecedents are satisfiable. Sets of mutually consequent-bound CPRs represent the possible values the single rv in the consequents can attain given satisfiable preconditions.

Now we can define a BKB as follows.

Definition 2.4. A BKB B is a finite set of CPRs such that

- for any distinct R_1 and R_2 in B , either (1) R_1 is mutually exclusive with R_2 or (2) $\text{con}(R_1) \neq \text{con}(R_2)$, and
- for any subset S of mutually consequent-bound CPRs of B

$$\sum_{R \in S} P(R) \leq 1.$$

Fig. 1 presents a sample BKB. BKBs can also be represented graphically [25] as depicted in Fig. 2 where labeled nodes represent unique specific instantiations of rvs. For example, the rv “pH” has three possible values: “<6.5”, “neutral”, and “>7.5”. These correspond to the three labeled nodes in the graph. Each CPR is represented by a darkened node where the parents of the node are the antecedents of the CPR and the child of the node denotes the consequent. Fig. 3 shows the underlying rv relationships in our BKB example. While such a cycle is problematic in BNs, it is allowable in the BKB framework. We will occasionally use the graphical description to further provide intuitions on key ideas through out the paper.

Inferencing over BKBs is conducted similarly to “if-then” rule inferencing. Thus, sets of CPRs collectively form inferences.

Wastes = Present	\Rightarrow	pH < 6.5	0.87
Wastes = Present	\Rightarrow	pH = Neut	0.11
Over Feed = Y \wedge Over Crowd = Y	\Rightarrow	Wastes = Present	0.68
Over Crowd = Y \wedge ...	\Rightarrow	Wastes = Present	0.77
pH > .75 \wedge Ammonia = High	\Rightarrow	Fish Stress = Y	0.85
Ammonia = High \wedge Temp = Low \wedge ...	\Rightarrow	Fish Stress = Y	0.36
Fish Stress = Y	\Rightarrow	Hungry = Not	0.92
Hungry = Not	\Rightarrow	Wastes = None	0.20

Fig. 1. A sample BKB fragment for fresh water aquarium management.

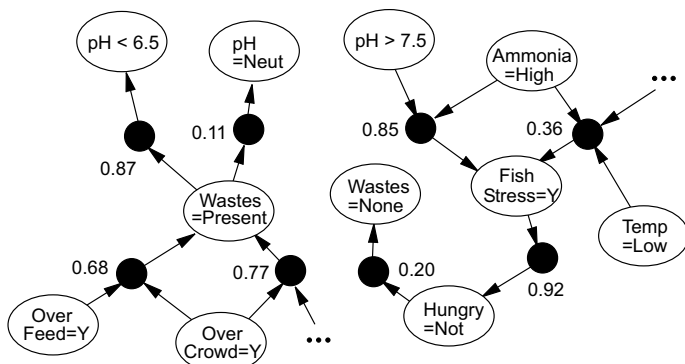


Fig. 2. A BKB fragment from fresh-water aquarium maintenance knowledge-base as a directed graph.




Fig. 3. Underlying rv relationships for BKB in Fig. 2.

Definition 2.5. A subset S of B is said to be a deductive set if for each CPR R in S where

$$R : A_{i_1} = a_{i_1} \wedge A_{i_2} = a_{i_2} \wedge \cdots \wedge A_{i_{n-1}} = a_{i_{n-1}} \Rightarrow A_{i_n} = a_{i_n},$$

the following two conditions hold:

- For each $k = 1, \dots, n-1$ there exists a CPR R_k in S such that $\text{con}(R_k) = \{A_{i_k} = a_{i_k}\}$.
- There does not exist some $R' \in S$ where $R' \neq R$ and $\text{con}(R') = \text{con}(R)$.

The first condition states that the antecedents of a given CPR must be supported by the consequents of other CPRs which corresponds to standard forward chaining in rule bases. The second condition imposes that there is a unique chain for supporting a particular rv assignment.

Notation. Given any $S \subseteq B$, $V(S)$ represents the set of rv assignments found in S and $H(S)$ represents the random variables that occur in S . $H(B)$ denote the finite set of random variables that occur in B .

Let $\Delta(B)$ represent the set of all possible sets of rv assignments to $H(B)$ such that if $T \in \Delta(B)$, then for each rv $A \in H(B)$, there exists at most one rv assignment to A in T . Furthermore, T is said to be a *complete assignment* if for each rv $A \in H(B)$, there exists exactly one rv assignment to A in T .

Given a set $S \subseteq B$, we define $P(S)$ as

$$P(S) = \prod_{R \in S} P(R).$$

For any CPR R in S

$$R : A_{i_1} = a_{i_1} \wedge A_{i_2} = a_{i_2} \wedge \cdots \wedge A_{i_{n-1}} = a_{i_{n-1}} \Rightarrow A_{i_n} = a_{i_n},$$

we say that each $(A_{i_k} = a_{i_k})$ is an *immediate ancestor* of $(A_{i_n} = a_{i_n})$ for $k = 1, \dots, n-1$ and that $(A_{i_n} = a_{i_n})$ is an *immediate descendant* of each $(A_{i_k} = a_{i_k})$ for $k = 1, \dots, n-1$. Thus, we can define this recursively with respect to the CPRs in a given set S for *ancestor* and *descendant*.

One typical problem with forward chaining in rule bases is the possibility of deriving inconsistent rv assignments. For example, we might derive both $A = \text{false}$ and $A = \text{true}$. With such a derivation, $P(S)$ becomes ill-defined as a potential joint probability.

Definition 2.6. We say that R_1 is compatible with R_2 if for all $k \leq n$ and $l \leq m$, $a_{i_k} = a'_{j_l}$ whenever $i_k = j_l$.

Compatibility guarantees that both R_1 and R_2 can be simultaneously satisfied which is the basis for forming valid inferences.

Definition 2.7. A deductive set I is said to be an inference over B if the following two conditions hold:

- I consists of mutually compatible CPRs;
- no $A_{i_k} = a_{i_k}$ is an ancestor of itself in I .

$P(I)$ is said to be the probability of inference I . Furthermore, an inference I over B is said to be complete if $H(I) = H(B)$.

Clearly, an inference I induces the set of rv assignments $V(I)$. The following theorem establishes that for each set of rv assignments V , there exists at most one inference I over B such that $V = V(I)$.

Theorem 2.2 (Santos and Santos 1999 [25], Corollary 4.4). *If I_1 and I_2 are two inferences over B where $V(I_1) = V(I_2)$, then $I_1 = I_2$.*

The collection of inferences from B can now define a probability distribution. This is established as follows.

Definition 2.8. Two inferences I_1 and I_2 are said to be compatible if for any $R_1 \in I_1$ and $R_2 \in I_2$, R_1 is compatible with R_2 . Otherwise, I_1 and I_2 are incompatible.

Furthermore, we extend the definition of compatibility between a CPR and a set of CPRs and vice versa.

Theorem 2.3 (Santos and Santos 1999 [25], Key Theorem 4.3). *For any set of mutually incompatible inferences Y in B*

$$\sum_{I \in Y} P(I) \leq 1.$$

Theorem 2.4 (Santos and Santos 1999 [25], Key Theorem 4.4). *Let I_0 be some inference. For any set of mutually incompatible inferences $Y(I_0)$ such that for all $I \in Y(I_0)$, $I_0 \subseteq I$*

$$\sum_{I \in Y(I_0)} P(I) \leq P(I_0).$$

The above two theorems establish the relationship among the inferences and with the joint probabilities that are induced by the inferences.

Definition 2.9. Let f be a function from $\Delta(B)$ to $[0, 1]$. f is said to be consistent with B (denoted $B \models f$) if for each complete inference $I \subseteq B$, $P(I) = f(V(I))$.

Hence, the structure of inferences in BKBs allows us to construct a partial joint probability distribution based on the available inferences which can then be extended to a complete distribution. Since BKBs are by nature designed to handle incomplete information, there is potentially a “missing mass” of probabilistic information not explicitly accounted for in the BKB, thus resulting in the possibility of multiple probability distributions that are fully consistent with the BKB.

Rosen et al. [17] present a constructive algorithm to automatically derive a single probability distribution. They basically examine a single interpretation of the “missing mass.” Assuming that no information is available concerning said mass, they distribute the mass uniformly across the unspecified distribution regions. This specific distribution is called the *default distribution* of B . Hence, there exists a discrete probability distribution, p over $H(B)$ that is consistent with B , i.e., $B \models p$.

From this, the following relationship between probability distributions and inferences in B is also derived.

Theorem 2.5 (Rosen et al. [18], Corollary 1). *For any inference I from B , $p(V(I)) = P(I)$.*

As we can see, unlike BNs, BKBs are organized at the individual rv assignment level instead of simply with the rvs alone. Furthermore, BKBs do not require a total ordering (causal) of the rvs or apriori complete distribution as are needed in BNs. This makes BKBs more flexible and capable of handling cyclical information while fully subsuming BNs [17,25,30].

3. Semantics

As we mentioned earlier, changes to a knowledge base occur throughout its life-cycle. The process of incremental knowledge acquisition identifies new knowledge that must be correctly introduced into the knowledge-base. For BKBs, such changes take the form of adding new CPRs, adding or removing antecedents in existing CPRs, changing the probability value of a CPR, and deleting CPRs if they are found to be incorrect.

Changes of this nature in typical probabilistic knowledge-bases, if not carefully done, can lead to potentially drastic alterations of the semantics for existing knowledge. Alterations could radically transform probability distributions. In this section, we present new results on how semantics is naturally preserved in BKBs during incremental knowledge acquisition. Our focus here is to examine the value $P(R)$ associated with a CPR R with respect to the changing probability distribution of the BKB. We will formally prove that $P(R)$ corresponds to the conditional probability $P(\text{con}(R)|\text{ant}(R))$ consistent to the probability distribution(s) as defined by the current BKB. Also, this property is invariant as the BKB evolves as long as R itself is not altered and continues to participate in inferences. Furthermore, this will lead to new formal theoretical and algorithmic results concerning the construction of BKBs.

3.1. Properties of minimal supports

Let $T = \{(A_{i_1} = a_{i_1}), (A_{i_2} = a_{i_2}), \dots, (A_{i_n} = a_{i_n})\}$ be a consistent set of rv assignments, i.e., $i_j \neq i_k$ whenever $j \neq k$.

Definition 3.1. A deductive set S is said to support T if for each $\{A_{i_k} = a_{i_k}\} \in T$, there exists some CPR R in S such that $\text{con}(R) = \{A_{i_k} = a_{i_k}\}$.

Definition 3.2. A deductive set S is said to be minimal with respect to T if S supports T and there does not exist a deductive set $S' \subset S$ that also supports T .

Clearly, T may have many minimal supports each representing different forward chaining possibilities found in B . Minimal supports are also considered to be explanations for T [27].

Proposition 3.1. If S is minimal with respect to T and S is an inference, then there does not exist an inference $S' \subset S$ that also supports T .

In this case, we also say that S is a *minimal inference* with respect to T .

Definition 3.3. Given a set of CPRs S from B , the frontier of S is the set of all rv assignments $\{A = a\}$ such that $\{A = a\} = \text{con}(R)$ for some $R \in S$ and $\{A = a\}$ has no descendants in S . We denote this set by $F(S)$.

Basically, the frontier of S represents rv assignments that have not participated in forward chaining. In the case, that S is a deductive set, we can also denote by $F(S)$ the set of unique CPRs R in S whose consequents are in the frontier.

Lemma 3.2. *If $S \neq \phi$ is an inference, then $F(S)$ is not empty.*

Proof. Since $S \neq \phi$ is an inference, by Definition 2.7, there must exist some rv assignment in S that has no descendants. \square

Lemma 3.3. *If S is a minimal deductive set supporting T , then $F(S) \subseteq T$.*

Proof. Assume that $\{A = a\}$ is in $F(S)$ but not in T . By Definition 2.5, let R be the unique CPR in S such that $\text{con}(R) = \{A = a\}$. Since $\{A = a\}$ has no descendants, we can safely remove R from S resulting in a subset S' of S . Clearly, S' is also a deductive set and S' supports T . However, S is not minimal. Contradiction. \square

Now, we consider the impact of forward chaining in our semantics for CPRs.

Definition 3.4. A deductive set S is said to be consistent with CPR R if and only if $S \cup \{R\}$ is an inference.

Definition 3.4 above implies that continuing forward chaining from S with CPR R is valid only when no inconsistencies in rv assignments can occur.

Proposition 3.4. *If S is consistent with R , then S is also an inference.*

We can now derive the following theorem relating the CPR probability to deductive sets.

Notation. $D_B(T, R)$ is the set of all minimal deductive sets (inferences) supporting T and consistent with R .

Lemma 3.5. *If $S \in D_B(\text{ant}(R) \cup \text{con}(R), R)$, then $R \in S$.*

Proof. Since S is an inference supporting $\text{ant}(R) \cup \text{con}(R)$, there exists some CPR R' in S such that $\text{con}(R') = \text{con}(R)$. Assume $R' \neq R$. This implies that R' and R are mutually exclusive. Thus, there exists rv assignments $\{A = a\}$ in

$\text{ant}(R)$ and $\{A = a'\}$ in $\text{ant}(R')$ such that $a \neq a'$. By Definition 2.5, there exists some CPR $R'' \in S$ such that $\text{con}(R'') = \{A = a\}$. However, R'' is not compatible with R' . This implies that S is not in $D_B(\text{ant}(R) \cup \text{con}(R))$. Contradiction. Therefore, $R' = R$. \square

Lemma 3.6. $S_1 \in D_B(\text{ant}(R) \cup \text{con}(R), R)$ if and only if both $S_2 \in D_B(\text{ant}(R), R)$ and $S_2 = S_1 - \{R\}$.

Proof. (\Rightarrow) Let $S_1 \in D_B(\text{ant}(R) \cup \text{con}(R), R)$. From Lemma 3.5, $R \in S_1$. From Lemma 3.3, $F(S) \subseteq \text{ant}(R) \cup \text{con}(R)$. Since $R \in S_1$, $\text{ant}(R) \cap F(S)$ is empty. From Lemma 3.2, $F(S)$ is not empty. Thus, $F(S) = \text{con}(R)$. Let $S' = S_1 - \{R\}$. Clearly, S' supports $\text{ant}(R)$ and is consistent with R . From Proposition 3.4, S' is an inference. Furthermore, $F(S') \subset \text{ant}(R)$.

Assume S' is not in $D_B(\text{ant}(R), R)$. Thus, there exists some CPR R' in S' such that $S' - \{R'\} \in D_B(\text{ant}(R), R)$. Since $F(S') \subseteq \text{ant}(R)$, this implies that either $\text{con}(R') \subset \text{ant}(R)$ or $\text{con}(R')$ is an ancestor of some $\{A = a\} \in \text{ant}(R)$. Hence, $\text{con}(R') \in V(S' - \{R'\})$. However, removing R' implies that no CPR R'' exists in $S' - \{R'\}$ such that $\text{con}(R'') = \text{con}(R')$. Thus, $S' - \{R'\}$ is not a deductive set. Contradiction.

(\Leftarrow) Let $S_2 \in D_B(\text{ant}(R), R)$. It follows that $S_2 \cup \{R\} \in D_B(\text{ant}(R) \cup \text{con}(R), R)$. \square

Lemma 3.6 proves that there exists a one-to-one and onto mapping between deductive sets in $D_B(\text{ant}(R) \cup \text{con}(R), R)$ and $D_B(\text{ant}(R), R)$.

Theorem 3.7

$$P(R) = \frac{\sum_{S_1 \in D_B(\text{ant}(R) \cup \text{con}(R), R)} P(S_1)}{\sum_{S_2 \in D_B(\text{ant}(R), R)} P(S_2)}. \quad (1)$$

Proof. Let $S_1 \in D_B(\text{ant}(R) \cup \text{con}(R), R)$. From Lemma 3.5, $R \in S_1$. We can re-write $P(S_1)$ as

$$P(S_1) = P(R) \sum_{R' \in S_1 - \{R\}} P(R') = P(R) P(S_1 - \{R\}).$$

Combined with Lemma 3.6,

$$\frac{\sum_{S_1 \in D_B(\text{ant}(R) \cup \text{con}(R), R)} P(S_1)}{\sum_{S_2 \in D_B(\text{ant}(R), R)} P(S_2)} = \frac{P(R) \sum_{S_1 \in D_B(\text{ant}(R) \cup \text{con}(R), R)} P(S_1 - \{R\})}{\sum_{S_1 \in D_B(\text{ant}(R) \cup \text{con}(R), R)} P(S_1 - \{R\})}.$$

Dividing out the common terms leaves us with $P(R)$. \square

Examining Theorem 3.7, the fraction seems closely related to the definition of conditional probabilities where the numerator reflects $P(\text{ant}(R) \cup \text{con}(R))$ and the denominator, $P(\text{ant}(R))$. In the next sections, we will be formally studying the relationship between the fraction in the above theorem and conditional probabilities. In particular, we will be formally identifying when such situations/conditions occur. Given this relationship, we will then examine the impact on a BKB's ability to manage incompleteness with respect to the semantics of conditional probabilities. As we will see in the next subsections, the semantics of BKBs are naturally defined and preserved during knowledge engineering.

3.2. Assignment completeness

The inequalities found in Theorems 2.3 and 2.4 reflect the incompleteness of information that may occur in a BKB. While a consistent distribution exists, there may be more than one such distribution. In this subsection, we examine a special class of BKBs.

Definition 3.5. B is said to be assignment complete if for every complete assignment $T \in \Delta(B)$, there exists a complete inference $I \subseteq B$, such that $V(I) = T$.

For this subsection, we only consider assignment complete BKBs and further assume that the sum of the probabilities of all complete inferences in B is 1 (also called probabilistically complete). Clearly, B defines a unique joint probability distribution p where $B \models p$. It follows from Theorems 2.3–2.5 that $p(T)$ is the sum of all complete inferences I over B such that $T \subseteq V(I)$. We now prove that $p(T)$ can be computed by summing carefully selected inferences (not necessarily complete) that are compatible with T .

Notation. $I_B(T)$ denotes the set of all inferences over B such that for each inference $I \in I_B(T)$, I is minimal with respect to T .

Intuitively, $I_B(T)$ represents all inferences that “concludes” with only consequents found in T .

Proposition 3.8. *Given any two distinct inferences I_1 and I_2 from $I_B(T)$, I_1 is incompatible with I_2 .*

In other words, Proposition 3.8 states that there exists some rv assignment in $V(I_1)$ that is incompatible with $V(I_2)$.

Theorem 3.9. *For any set T defined above*

$$p(T) = \sum_{I \in I_B(T)} P(I).$$

Proof. Let I be some inference in $I_B(T)$. Let $E(I)$ denote the set of complete inferences that are supersets of I . From Proposition 3.8, given I_1 and I_2 from $I_B(T)$ such that $I_1 \neq I_2$, $E(I_1) \cap E(I_2)$ is empty.

Since B is assignment complete

$$\bigcup_{I \in I_B(T)} E(I) = Q(T),$$

where $Q(T)$ is the set of all complete inferences in B that are compatible with T . Thus, from Theorem 2.5

$$\sum_{I \in I_B(T)} P(I) = \sum_{I \in I_B(T)} \sum_{J \in E(I)} P(J) = \sum_{I \in Q(T)} P(I)$$

and the last summation is equal to $p(T)$. \square

Theorem 3.9 demonstrates that for our special class of assignment complete BKBs, the joint probability, $p(T)$, can be calculated directly from the set of inferences in $I_B(T)$. In the following subsection, we take this observation and examine the relationship to conditional probabilities discussed earlier.

3.3. Conditional probabilities

Returning to the sets of inferences $D_B(\text{ant}(R) \cup \text{con}(R), R)$ and $D_B(\text{ant}(R), R)$ in Theorem 3.7, these sets reflect inferences that support $\text{ant}(R) \cup \text{con}(R)$ and $\text{ant}(R)$, respectively, and whose frontiers are bounded by $\text{ant}(R) \cup \text{con}(R)$ and $\text{ant}(R)$, respectively. We now examine the relationships between the sets $D_B(\text{ant}(R))$ and $D_B(\text{ant}(R) \cup \text{con}(R))$ to the sets $I_B(\text{ant}(R))$ and $I_B(\text{ant}(R) \cup \text{con}(R))$.

Let $S = \{R_1, R_2, \dots, R_n\}$ be a set of CPRs in B such that $\text{con}(R_i) \in \text{ant}(R_{i+1})$ for $i = 1, \dots, n-1$.

Definition 3.6. S is said to be unstable if $\{A = a\} = \text{con}(R_n)$ and $\{A = a'\} \in \text{ant}(R_1)$ (Note that a and a' need not be distinct.). B is said to be stable if it does not have any unstable subsets.

In graph-based terms, for unstable sets there exists a directed path between $\{A = a\}$ and $\{A = a'\}$ in the BKB. This does not preclude cycles in the underlying rv graph such as the BKB in Figs. 1–3.

Theorem 3.10. $D_B(\text{ant}(R) \cup \text{con}(R), R) = I_B(\text{ant}(R) \cup \text{con}(R))$.

Proof. (\subseteq) Let $I \in D_B(\text{ant}(R) \cup \text{con}(R), R)$. From Lemma 3.5, $R \in I$. From Definition 3.4, I is an inference. From Proposition 3.1, I is minimal with respect to $\text{ant}(R) \cup \text{con}(R)$. Thus, $I \in I_B(\text{ant}(R) \cup \text{con}(R))$.

(\supseteq) Let $I \in I_B(\text{ant}(R) \cup \text{con}(R))$. Assume I is not in $D_B(\text{ant}(R) \cup \text{con}(R), R)$. Clearly, since I is an inference, R is in I . Thus, $I \cup \{R\}$ is an inference. However, I is not a minimal deductive set with respect to $\text{ant}(R) \cup \text{con}(R)$. There exists some deductive set I' that is minimal with respect to $\text{ant}(R) \cup \text{con}(R)$ such that $I' \subset I$. From construction of $I_B(\text{ant}(R) \cup \text{con}(R))$, I is a minimal inference with respect to $\text{ant}(R) \cup \text{con}(R)$. Thus, I' cannot be an inference. Contradiction. \square

Lemma 3.11. $D_B(\text{ant}(R), R) \subseteq I_B(\text{ant}(R))$.

Proof. Let $I \in D_B(\text{ant}(R), R)$. From Lemma 3.3, $F(I) \subseteq \text{ant}(R)$. From Proposition 3.4, I is an inference. From Proposition 3.1, I is minimal with respect to $\text{ant}(R)$. Thus, $I \in I_B(\text{ant}(R))$. \square

Theorem 3.12. If B is stable, then $D_B(\text{ant}(R), R) = I_B(\text{ant}(R))$.

Proof. Let B be stable. From Lemma 3.11, we only need to prove subset equality in the other direction.

Let $I \in I_B(\text{ant}(R))$. Assume I is not in $D_B(\text{ant}(R), R)$. This implies that either (1) I is not a minimal deductive set with respect to $\text{ant}(R)$ or (2) $I \cup \{R\}$ is not an inference.

Case 1. I is not a minimal deductive set with respect to $\text{ant}(R)$. There exists some deductive set I' that is minimal with respect to $\text{ant}(R)$ such that $I' \subset I$. From construction of $I_B(\text{ant}(R))$, I is a minimal inference with respect to $\text{ant}(R)$. Thus, I' cannot be an inference. Contradiction.

Case 2. $I \cup \{R\}$ is not an inference. This implies that there exists some CPR R' in I such that R is not compatible with R' . Without loss of generality, let $\text{con}(R) = \{A = a\}$ and $\text{con}(R') = \{A = a'\}$ where $a \neq a'$. By construction of $I_B(\text{ant}(R))$, $\{A = a'\}$ is the ancestor of some rv assignment $\{B = b\}$ in $\text{ant}(R)$ in I . In $I \cup \{R\}$, $\{B = b\}$ is an immediate ancestor of $\{A = a\}$. Thus, $\{A = a'\}$ is the ancestor of $\{A = a\}$. However, B is stable. Contradiction.

Therefore, $D_B(\text{ant}(R), R) = I_B(\text{ant}(R))$. \square

Combining Theorems 3.9, 3.10 and 3.12, we get the following.

Theorem 3.13. If B is stable, assignment complete, and probabilistically complete, then for each $R \in B$, $P(R)$ is a conditional probability consistent with p .

When B is not probabilistically complete, the summations

$$\sum_{S_1 \in D_B(\text{ant}(R) \cup \text{con}(R), R)} P(S_1)$$

and

$$\sum_{S_2 \in D_B(\text{ant}(R), R)} P(S_2)$$

approach $P(\text{ant}(R) \cup \text{con}(R))$ and $P(\text{ant}(R))$, respectively, as B is completed.

Clearly, changes to B affect the various joint probabilities found in the BKB. However, from Theorem 3.7, such changes do not affect the original semantics imposed by the knowledge engineer on the individual CPRs unless they themselves are altered. *As long as the BKB is stable, the semantics correspond to conditional probabilities throughout a BKB's life-cycle.*

The check for stability in a BKB can be accomplished in polynomial time. Using the graphical representation for BKBs (e.g., Fig. 2), the following algorithm (a variant of DFS), determines stability:

Algorithm 3.1

Input: BKB B

Output: Decision on B 's stability

Initialize stack s to be empty

For every rv assignment node v in B do begin

For every node u in B , unmark u

Push all immediate descendants of v onto s

While stack s is not empty do begin

$u = \text{pop}(s)$

*If u is an rv assignment node and $H(u) = H(v)$ then
return (“unstable”)*

If u is unmarked then do begin

Push all immediate descendants of u onto s

Mark u

end

end

end

return (“stable”)

In Fig. 2, the BKB is stable although the underlying rv relationships (see Fig. 3) has an rv cycle.

Finally, assume that B is modified to B' during an incremental knowledge acquisition state and both B and B' are stable. Due to incompleteness, some CPRs may not participate in any inference over B . We say that such CPRs are *ungrounded*.

Definition 3.7. R is said to be grounded if there exists an inference I over B such that $R \in I$. Otherwise, we say that R is ungrounded.

Theorem 3.14. *If R is grounded in both B and B' and was not modified, then $P(R)$ satisfies Eq. (1) in both B and B' .*

Proof. Follows from above results. \square

From this theorem, the initial probability assigned to R is semantically preserved via Eq. (1) during changes to the knowledge-base as long as R itself is not modified and remains grounded. However, even if R becomes ungrounded at some point, the semantics for R is restored once R is again grounded.

3.4. Properties for assignment completeness

In the previous section, we observed that assignment-complete BKBs have some desirable properties, the foremost being the correct intuitive probabilistic semantics of rule weights and probability of partial assignments. Although some of the semantics properties hold for incomplete BKBs, it is a crucial issue whether there exists a way of completing a current partial BKB. Formally, the question is: given a partial BKB, is there a set of CPRs that can be added to it, resulting in an assignment-complete BKB, i.e. is the partial BKB *completable*?

Definition 3.8. We say that B is assignment completable if there exists an assignment complete BKB B' such that $B \subseteq B'$.

There exist cases where B is an inextensible BKB that is not assignment complete. That is, there exists a complete assignment to the variables of the BKB, $H(B)$, that has no corresponding complete inference in B , but no new CPRs can be added to B (based on the existing rv assignments in $V(B)$). This situation occurs when there is some rv assignment, say $\{A = a\}$, such that for all CPRs R in B that have this assignment as a consequent, i.e. $\text{con}(R) = \{A = a\}$, the rule antecedent $\text{ant}(R)$ is not consistent with the rv assignments in $H(B)$. This implies the “incompleteness” part. If, in addition, every CPR R' with the same consequent as R , i.e. $\text{con}(R') = \{A = a\}$ and $\text{ant}(R')$ consistent with $V(B)$ cannot be added to B because of mutual exclusion, the result will be an incompletable BKB.

For example, consider the following two CPRs:

$$A = a \wedge B = b \Rightarrow C = c$$

and

$$C = c \Rightarrow A = a'.$$

The rv assignment $\{A = a', B = b, C = c\}$ will never have an associated inference. Observe that to get the above rv assignment, one must use the rule $C = c \Rightarrow A = a'$ —because any other candidate rule addition that achieves

$A = a'$ must contradict $C = c$, and thus cannot be used. But in order to achieve $C = c$, a rule must be added that is mutually exclusive with $A = a \wedge B = b \Rightarrow C = c$, and does *not* have an assignment to A in the antecedent (due to requirement of acyclicity of an inference). The only way to get mutual exclusion under these conditions is to have some assignment to B in the antecedent (say $B = b'$ where $b' \neq b$), but that is not consistent with the desired rv assignment.

Definition 3.9. B is said to be inextensible if there does not exist a BKB B' such that $B \subset B'$ and $V(B) = V(B')$.

Proposition 3.15. *If B is assignment complete, then B is inextensible.*

Let $S = \{R_1, R_2, \dots, R_n\}$ be a set of CPRs in B such that $\text{con}(R_i) \in \text{ant}(R_{i+1})$ for $i = 1, \dots, n-1$ (That is, S is a *chain* of rules—a path in a proof graph.).

Definition 3.10. S is said to be strongly unstable if all following conditions hold:

- S is unstable,
- R_2, \dots, R_n are mutually consistent, and
- $\text{con}(R_1) \cup \text{ant}(R_1)$ is consistent with $\bigcup_{i=2}^n \text{con}(R_i)$.

B is said to be weakly stable if it does not have any strongly unstable subsets.

Intuitively, strongly unstable sets are more likely to participate in inferences since $\{R_2, \dots, R_n\}$ are mutually consistent.

Lemma 3.16. *If B is assignment complete, then B is weakly stable.*

Proof. Let B be assignment complete. Assume B is not weakly stable. Let $S = \{R_1, R_2, \dots, R_k\}$ be a strongly unstable set of CPRs such that for any other strongly unstable set CPRs S' , $|S| \leq |S'|$. Let $\{A_0 = a_0, A_1 = a_1, \dots, A_{k-1} = a_{k-1}, A_0 = a'_0\}$ be the sequence where $\{A_{i-1} = a_{i-1}\} \in \text{ant}(R_i)$ for $i = 1, \dots, k$, $\{A_i = a_i\} = \text{con}(R_i)$ for $i = 1, \dots, k-1$, and $\text{con}(R_k) = \{A_0 = a'_0\}$. Let M be the set of all rv assignments in $\text{ant}(R_1)$ that are inconsistent with $\bigcup_{i=2}^k \text{con}(R_i)$.

We now construct a complete assignment as follows:

$$V = \text{con}(R_k) \cup \bigcup_{i=2}^k \text{ant}(R_i) \cup \{\text{ant}(R_1) - M\}$$

For any rv not assigned in V , arbitrarily choose an assignment found in B .

Since B is assignment complete, let I be the complete inference for V . Because of mutual exclusion, $\{R_2, \dots, R_k\} \subseteq I$. If $R_1 \in I$, this implies that $a'_0 = a_0$

and M is empty. However, this implies that $\{A_0 = a_0\}$ is an ancestor of itself in I . Thus, R_1 is not in I .

Let R'_1 be the CPR in I where $\text{con}(R'_1) = \{A_1 = a_1\}$. R'_1 must be mutually consistent with V and $\{R_2, \dots, R_k\}$. Now, since $\text{ant}(R_1) - M$ is consistent with V and R_1 must be mutually exclusive with R'_1 , this implies that there exists some $\{B = b\} \in \text{ant}(R_1)$ and $\{B = b'\} \in \text{ant}(R'_1)$ where $b \neq b'$. Thus, $\{B = b\} \in M$. Therefore, $B = A_i$ and $b' = a_i$ where $\{A_i = a_i\}$ is a descendant of $\{A_1 = a_1\}$ in S . However, $\{A_i = a_i\}$ is now an ancestor of itself in I . Contradiction.

Therefore, B is weakly stable. \square

Lemma 3.17. *If B is inextensible, then for any complete assignment T to $H(B)$, for each $\{A = a\} \in T$, there exists a CPR $R \in B$ such that $\text{con}(R) = \{A = a\}$ and $\text{ant}(R) \subset T$.*

Proof. Assume that there exists some complete assignment T to $H(B)$ and some $\{A = a\} \in T$ such that no CPR R exists in B where $\text{con}(R) = \{A = a\}$ and $\text{ant}(R) \subset T$.

Let R be any CPR such that $\text{con}(R) = \{A = a\}$ and $\text{ant}(R) \subset T$. This implies that there exists CPR $R' \in B$ such that $\text{con}(R') = \text{con}(R)$ and R' is mutually consistent with R . Clearly, any such R' is not consistent with T . Thus, $\text{ant}(R')$ must contain some $\{B = b'\}$ where $\{B = b\} \in T$ and $b \neq b'$.

However, if $\text{ant}(R) = T - \{A = a\}$, then R is mutually exclusive with any R' found above. Contradiction. \square

Lemma 3.18. *If B is inextensible and weakly stable, then B is assignment complete.*

Proof. Let $T = \{A_1 = a_1, \dots, A_n = a_n\}$ be a complete assignment to $H(B)$. From Lemma 3.17, for each $\{A_i = a_i\}$ there exists a CPR $R \in B$ such that $\text{con}(R) = \{A_i = a_i\}$ and $\text{ant}(R) \subset T$. Let S be the maximal set of CPRs in B consistent with T . Thus, $V(S) = T$.

Clearly, all CPRs in S are mutually consistent. Assume that S is not an inference. This implies that for some $\{A_k = a_k\}$, it is an ancestor of itself in S . However, this implies that S is strongly unstable. Therefore S is a complete inference for T . \square

The following necessary and sufficient condition for assignment completeness follows from the lemmas above.

Theorem 3.19. *B is assignment complete if and only if B is inextensible and weakly stable.*

Theorem 3.20. *If B is assignment completable, then B is weakly stable.*

Proof. Let $B' \supseteq B$ be some assignment complete BKB. From Lemma 3.16, B' is weakly stable. Clearly, any subset of B' must also be weakly stable. Therefore, B is weakly stable. \square

Theorem 3.20 provides us a necessary condition for assignment completeness. Thus, we must guarantee that the BKB is weakly stable at each step during the construction process. As we mentioned earlier, testing for stability in a BKB can be done in polynomial time, however, determining weak stability is much more difficult given the additional conditions.

Theorem 3.21. *Deciding weak stability of a BKB is CO-NP complete.*

Proof. Membership in CO-NP is obvious, since a certificate is a strongly unstable set of CPRs, and deciding set weak stability can be done in polynomial time. We show that the problem is hard by reduction from CSP.

Let C be a constraint satisfaction problem over n variables v_1, \dots, v_n , with a set of constraints C_1, \dots, C_m . The following polynomial-time construction builds BKB K , as follows.

1. For all constraints C_i , construct one node X_{C_i} (denoting that the constraint is not violated). Also, add one node X_{C_0} , denoting a dummy constraint.
2. For each variable v_j , add a variable X_{v_j} , and one BKB node for each possible domain value d_l of v_j (the nodes are denoted by X_{v_j, d_l}). These nodes denote possible assignments to the BKB variable X_{v_j} , mirroring the CSP.
3. For all constraints (i.e. all C_i , i from 1 to m), construct one CPR for each allowed tuple t_j in C_i , (a k -ary constraint) as follows: Let $t_j = (v_{l_1} = d_{j_1}, v_{l_2} = d_{j_2}, \dots, v_{l_k} = d_{j_k})$ be any permissible variable assignment to variables participating in constraint C_i . Construct the rule

$$X_{C_{i-1}} \wedge X_{v_{l_1}, d_{j_1}} \wedge \dots \wedge X_{v_{l_k}, d_{j_k}} \Rightarrow X_{C_i}.$$

4. Finally, add the CPR R : $X_{C_m} \Rightarrow X_{C_0}$.

Claim. K is weakly stable iff C has no solutions.

Proof of Claim. (\Rightarrow) Assume that K is weakly stable and that C has a solution. Without loss of generality, let that solution be $v_j = d_j$ for all j from 1 to n . Now, since the solution violates no constraints, there is a tuple in every constraint C_i consistent with the variable assignment, and thus exactly one rule in K corresponding to that assignment for the constraint C_i . This set of CPRs, together with the CPR R , form an unstable set—contradiction.

(\Leftarrow) Assume that C has no solutions and that K is not weakly stable. Now, K must have an unstable set of CPRs, however, by construction, the only

possible weakly unstable set contains rule R as well as a sequence of rules of the form $X_{C_{i-1}} \wedge \dots \wedge \dots \Rightarrow X_{C_i}$ for all i from 1 to n . Weak instability requires all these latter CPRs to have consistent antecedents, which induces a solution to problem C —a contradiction, and proving the claim.

Thus, deciding weak stability is CO-NP complete. \square

3.5. Discussion

In prior work on BKBs, it was simply assumed that the probability values attached with each CPR semantically represented the conditional probability of the consequent conditioned on the antecedents [25]. Furthermore, strict independence assumptions were applied in order to achieve semantics similar to BNs based solely on conditional dependency modeling [30]. From our new results, we have proven that the values associated with CPRs are *inherently* conditional probabilities when the BKBs are stable. This further strengthens our semantics assumptions for BKBs by demonstrating that all probabilistic information in such BKBs are soundly and consistently derived. Furthermore, weak stability is a necessary condition for completability of BKBs.

While problems of stability arise from cyclicity, stability does not preclude all forms of cyclicity. Fig. 1 with underlying rv cyclicity is stable. Furthermore, stable BKBs properly subsume a special class of BKBs called causal BKBs [25]. Causal BKBs admit a polynomial time reasoning algorithm.

Finally, testing for stability in a BKB can be done in polynomial time (see Algorithm 3.1). Furthermore, the complexity of such checks can be further reduced if they are performed incrementally as new CPRs are introduced. In particular, when a new CPR is introduced, we only need to do the graph traversal for nodes that are ancestors of the CPR and compare against the set of nodes which are descendants from the CPR.

4. Conclusions

Maintaining correct semantics of a knowledge-base, through changes required during the knowledge engineering cycle, is a challenging problem. The task is particularly difficult in knowledge-bases that must capture uncertainty, as the consistency requirements are compounded by the necessity of adhering to the requirements of the uncertainty calculus—in our case the axioms of probability theory.

In this paper, we presented new results regarding how BKBs naturally capture and preserve uncertainty semantics, especially during incremental knowledge acquisition. In particular, we demonstrated that by using the BKB model, the numerical values of uncertainty assigned to each conditional probability rule (BKB's “if-then” rule equivalents) implicitly correspond to

conditional probabilities in the target probability distribution being constructed. This is achieved without levying explicit semantics assumptions on the values but by properly guaranteeing stability in inferencing for the BKB. Furthermore, we also demonstrated that the semantics are preserved in a BKB while changes are made during incremental knowledge acquisition. Hence, the initial value and semantics assumed by the knowledge engineer remains constant as the BKB changes and grows. We were also able to derive new results concerning the completeness of BKBs based on weak stability during inferencing. This ensures that all relevant inferences that are needed in the BKBs target domain can be methodically captured.

From these results, we believe BKBs to be an ideal knowledge representation for constructing knowledge-based systems. While our discussions have focused on the typical knowledge engineering cycle of human updating and correction to the knowledge-base, BKBs can also serve as excellent frameworks for systems that must automatically update their knowledge in dynamic environments through data-mining and machine learning techniques. For future work, we are examining approaches for distributed problem solving via BKBs where multiple processes/agents each possess BKBs and must cooperate by propagating conditional probability rules. This propagation can effectively address problems in establishing common context between agents, negotiate requests, and evaluate various probability measures of success and goal satisfaction. In particular, such a distributed system of BKBs can be applied to domains such as mission planning, manufacturing scheduling, and cooperative workspaces.

Acknowledgements

This research was supported in part by AFOSR Grant no. F49610-99-1-0059. A preliminary version of this work appears in [26].

This research was supported in part by the Paul Ivanier, Center for Robotics and Production Management, Ben-Gurion University.

References

- [1] H. Aizenstein, A. Blum, R. Khardon, E. Kushilevitz, L. Pitt, D. Roth, On learning read- k -satisfy- j DNF, *SIAM Journal on Computing* 27 (6) (1998) 1515–1530.
- [2] V. Barr, Applying reliability engineering to expert systems, in: *Proceedings of the Twelfth International FLAIRS Conference*, 1999, pp. 494–498.
- [3] E.J. Bass, S.T. Ernst-Fortin, R.L. Small, Knowledge base development tool requirements for an intelligent monitoring aid, in: *Proceedings of the Tenth International FLAIRS Conference*, 1997, pp. 412–416.

- [4] C. Boutilier, N. Friedman, M. Goldszmidt, D. Koller, Context-specific independence in BNs, in: *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann Publishers, San Francisco, CA, 1996.
- [5] G.F. Cooper, The computational complexity of probabilistic inference using Bayesian belief networks, *Artificial Intelligence* 50 (1) (1990) 393–405.
- [6] P. Dagum, M. Luby, Approximating probabilistic inference in Bayesian belief networks is NP-hard, *Artificial Intelligence* 60 (1) (1993) 141–153.
- [7] J. DeKleer, An assumption-based TMS, *Artificial Intelligence* 28 (1986) 163–196.
- [8] A.J. Gonzalez, D.D. Dankel, *The Engineering of Knowledge-Based Systems: Theory and Practice*, Prentice-Hall, Inc, Englewood Cliffs, NJ, 1993.
- [9] G. Johnson, E. Santos Jr., Generalizing knowledge representation rules for uncertain knowledge, in: *Proceedings of the Thirteenth International FLAIRS Conference*, 2000, pp. 186–190.
- [10] R. Knauf, A.J. Gonzalez, K. P. Jantke, Towards validation of rule-based systems—the loop is closed, in: *Proceedings of the Thirteenth International FLAIRS Conference*, AAAI Press, New York, 2000, pp. 331–335.
- [11] R. Knauf, I. Philippow, A.J. Gonzalez, Towards validation and refinement of rule-based systems, *Journal of Experimental and Theoretical Artificial Intelligence* 12 (4) (2000) 421–431.
- [12] D. Madigan, J. York, Bayesian graphical models for discrete data, *International Statistical Review* 63 (1995) 215–232.
- [13] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Morgan Kaufmann*, Morgan Kaufmann, San Mateo, CA, 1988.
- [14] J. Pearl, *CAUSALITY: Models, Reasoning, and Inference*, Cambridge University Press, 2000.
- [15] D. Poole, Probabilistic horn abduction and Bayesian networks, *Artificial Intelligence* 64 (1) (1993) 81–129.
- [16] S. Ram, S. Ram, Design and validation of a knowledge-based system for screening product innovations, *IEEE Transactions on Systems, Man, and Cybernetics (Part A)* 26 (2) (1996) 213–221.
- [17] T. Rosen, S.E. Shimony, E. Santos Jr., Reasoning with BKBs—algorithms and complexity, in: *Working Notes of the Sixth International Symposium on Artificial Intelligence and Mathematics*, 2000.
- [18] T. Rosen, S.E. Shimony, E. Santos Jr., Reasoning with BKBs—algorithms and complexity, Technical Report IDIS Technical Report 103, Intelligent Distributed Information Systems Laboratory, University of Connecticut, 2001.
- [19] K. Sandahl, Transferring knowledge from active expert to end-user environment, *Knowledge Acquisition* 6 (1) (1994) 1–22.
- [20] E. Santos Jr., On the generation of alternative explanations with implications for belief revision, in: *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 1991, pp. 337–347.
- [21] E. Santos Jr., Verification and validation of knowledge-bases under uncertainty, *Data and Knowledge Engineering* 37 (2001) 307–329.
- [22] E. Santos Jr., D.O. Banks, S.B. Banks, Mack: a tool for acquiring consistent knowledge under uncertainty, in: *Proceedings of the AAAI Workshop on Verification and Validation of Knowledge-Based Systems*, 1997, pp. 23–32.
- [23] E. Santos Jr., S.B. Banks, S.M. Brown, D.J. Bawcom, Identifying and handling structural incompleteness for validation of probabilistic knowledge-bases, in: *Proceedings of the Twelfth International FLAIRS Conference*, 1999, pp. 506–510.
- [24] E. Santos Jr., H.T. Gleason, S.B. Banks, Bval: Probabilistic knowledge-base validation, in: *Proceedings of the AAAI Workshop on Verification and Validation of Knowledge-Based Systems*, 1997, pp. 13–22.

- [25] E. Santos Jr., E.S. Santos, A framework for building knowledge-bases under uncertainty, *Journal of Experimental and Theoretical Artificial Intelligence* 11 (1999) 265–286.
- [26] E. Santos Jr., E.S. Santos, S.E. Shimony, Semantics and knowledge acquisition in Bayesian knowledge-bases, in: *Proceedings of the Fifteenth International FLAIRS Conference*, 2002, pp. 524–528.
- [27] B. Selman, H.J. Levesque, Abductive and default reasoning: a computational core, in: *Proceedings of the AAAI Conference*, 1990, pp. 343–348.
- [28] S.E. Shimony, Finding MAPs for belief networks is NP-hard, *Artificial Intelligence* 68 (1994) 399–410.
- [29] S.E. Shimony, C. Domshlak, E. Santos Jr., Cost-sharing heuristic for Bayesian knowledge-bases, in: *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 1997, pp. 421–428.
- [30] S.E. Shimony, E. Santos Jr., T. Rosen, Independence semantics for BKBs, in: *Proceedings of the Thirteenth International FLAIRS Conference*, 2000, pp. 308–312.
- [31] N.L. Zhang, D. Poole, Exploiting causal independence in Bayesian network inference, *Journal of Artificial Intelligence Research* 5 (1996) 301–328.