Semantics and Knowledge Acquisition in Bayesian Knowledge-Bases

Eugene Santos Jr.
Dept of Comp Sci & Engr
University of Connecticut
Storrs, CT 06269-3155
eugene@cse.uconn.edu

Eugene S. Santos
Dept of Comp & Info Sci
Youngstown State University
Youngstown, OH 44555
santos@cis.ysu.edu

Solomon Eyal Shimony
Dept. of Comp Sci
Ben Gurion University of the Negev
84105 Beer-Sheva, Israel
shimony@cs.bgu.ac.il

Abstract

Maintaining semantics for uncertainty is critical during knowledge acquisition. 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 theoretical results concerning BKBs and how BKBs 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.

Introduction

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 (Santos 2001; Barr 1999; Knauf, Gonzalez, & Jantke 2000). 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 must provide sound semantics for the knowledge/information it does have. On top of all this, uncertainty is a primary facet of incompleteness that pervades every stage of the knowledge acquisition cycle.

Approaches to maintaining semantic consistency either (1) enforce strict local semantic assumptions or (2) require extensive modifications and recomputations over the existing knowledge-base, to accommodate new knowledge. In Bayesian networks (BNs) (Pearl 2000), the semantics of uncertainty are represented by probabilistic conditional independence. Additions made to a BN are reflected as changes in the underlying graph structure. Such changes affect the conditional independence semantics of nearly all the reachable nodes from the affected region. However, in BNs, local semantics with respect to the immediate neighbors are established directly by the knowledge engineer.

Our goal in this paper is to address the preservation of semantics during incremental knowledge acquisition under uncertainty without the local assumptions levied by BNs. In particular, we examine Bayesian Knowledge-Bases (BKBs) which are a generalization of BNs (Santos & Santos 1999). BKBs have been extensively studied both theoretically (Shimony, Domshlak, & Santos 1997; Johnson & Santos 2000; Shimony, Santos, & Rosen 2000) and for use in knowledge engineering (Santos 2001; Santos *et al.* 1999). BKBs provide a highly flexible and intuitive representation following a basic "if-then" structure in conjunction with probability theory. BKBs were designed keeping in mind typical domain incompleteness while BNs typically assume that a complete probability distribution is available. 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.

In this paper, we present new theoretical results concerning BKBs and how they can naturally and implicitly preserve semantics as new knowledge is added. In particular, the relationship between rule weights and conditional probabilities is established through stability of inferencing in BKBs.

Bayesian Knowledge Bases

The formulation of BKBs presented here is slightly different from existing definitions found in (Santos & Santos 1999) 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 $A_1, A_2, \ldots, A_k, \ldots$ be a collection of finite discrete random variables (abbrev. rvs) where $r(A_i)$ denotes the set of possible values for A_i.

Definition 1 *A* conditional probability rule (CPR)*, R, is of the form*

$$R : A_{i_1} = a_{i_1} \wedge \ldots \wedge A_{i_{n-1}} = a_{i_{n-1}} \Longrightarrow A_{i_n} = a_{i_n}$$

for some positive n where $a_{i_j} \in r(A_{i_j})$ such that $i_j \neq i_k$ for all $j \neq k$. The weight *of R is 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 $\mathsf{ant}(R)$ and $\mathsf{con}(R)$. When $n = 1$,

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 $\wedge \ldots$	$\Rightarrow$	Wastes = Present	0.77
pH $> .75 \wedge$ Ammonia = High	$\Rightarrow$	Fish Stress = Y	0.85
Ammonia = High $\wedge$ Temp = Low $\wedge \ldots$	$\Rightarrow$	Fish Stress = Y	0.36
Fish Stress = Y	$\Rightarrow$	Hungry = Not	0.92
Hungry = Not	$\Rightarrow$	Wastes = None	0.20

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

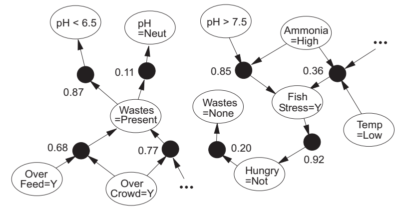

FIGURE 2. A BKB fragment from fresh-water aquarium maintenance knowledge-base as a directed graph.

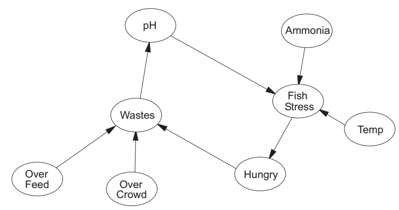

FIGURE 3. Underlying rv relationships for BKB in Figure 2.

$\mathsf{ant}(R)$ is the empty set and we write R as

$$R : \mathsf{true} \Longrightarrow A_{i_n} = a_{i_n}.$$

The weight $P(R)$ eventually corresponds to the conditional probability of R as we shall see in the next section.

Definition 2 *Given two CPRs*

$$R_1 : A_{i_1} = a_{i_1} \wedge \ldots \wedge A_{i_{n-1}} = a_{i_{n-1}} \Longrightarrow A_{i_n} = a_{i_n} \textit{and}$$

$$R_2 : A_{j_1} = a'_{j_1} \wedge \ldots \wedge A_{j_{m-1}} = a'_{j_{m-1}} \Longrightarrow A_{j_m} = a'_{j_m},$$

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 $a_{i_k} \neq a'_{j_l}$.*

Definition 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 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.

Definition 4 *A Bayesian Knowledge Base 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) $\mathsf{con}(R_1) \neq \mathsf{con}(R_2)$, and*
- *for any subset S of mutually consequent-bound CPRs of B, $\sum_{R \in S} P(R) \leq 1$.*

Figure 1 presents a sample BKB. BKBs can also be represented graphically as depicted in Figure 2 where labeled nodes represent unique specific instantiations of rvs. For example, the rv "pH" has three possible values corresponding 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. Figure 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.

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

Definition 5 *A subset S of B is said to be a* deductive set *if for each CPR R in S where*

$$R : A_{i_1} = a_{i_1} \wedge \ldots \wedge A_{i_{n-1}} = a_{i_{n-1}} \Longrightarrow A_{i_n} = a_{i_n},$$

the following two conditions hold:

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

The first condition states that the antecedents of a given CPR must be supported by the consequents of other CPRs. 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,

$$R : A_{i_1} = a_{i_1} \wedge \ldots \wedge A_{i_{n-1}} = a_{i_{n-1}} \Longrightarrow A_{i_n} = a_{i_n}$$

we say that each $(A_{i_k} = a_{i_k})$ is an *immediate ancestor* of $(A_{i_n} = a_{i_n})$ for $k = 1, \ldots, n-1$ and that $(A_{i_n} = a_{i_n})$ is an *immediate descendant* of each $(A_{i_k} = a_{i_k})$ for $k = 1, \ldots, n-1$. We 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 = false and A = true. With such a derivation, $P(S)$ becomes ill-defined as a potential joint probability.

Definition 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 $i_k = j_l$.*

Definition 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 *[(Santos & Santos 1999), 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 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 3 *[(Santos & Santos 1999), Key Theorem 4.3] For any set of mutually incompatible inferences Y in B, $\sum_{I \in Y} P(I) \leq 1$.*

Theorem 4 *[(Santos & Santos 1999), 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 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, Shimony, & Santos 2001) presents 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, Shimony et al. 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 5 *[(Rosen, Shimony, & Santos 2001), 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. While work has been done on capturing BNs as sets of rules (Poole 1997) and relaxing the conditional dependency requirements (Shimony 1993; Poole 1993; Boutilier *et al.* 1996), a total ordering on the rvs must still be maintained. BKBs do not require a total ordering 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 (Santos & Santos 1999).

Semantics

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.

In this section we present new results on how semantics is naturally preserved in BKBs during incremental knowledge acquisition without local semantic assumptions. 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 $P(\mathsf{con}(R) | \mathsf{ant}(R))$ consistent with the probability distribution(s) as defined by the current BKB. Also, this property is invariant as the BKB evolves in a stable fashion as long as R itself is not altered and continues to participate in inferences.

Deductive Set Support

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

Definition 10 *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 $\mathsf{con}(R) = \{A_{i_k} = a_{i_k}\}$.*

Definition 11 *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 (Selman & Levesque 1990).

Proposition 6 *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 12 *Given a set of CPRs S from B, the* frontier *of S is the set of all rv assignments $\{A = a\}$ such that $\{A = a\} = \mathsf{con}(R)$ 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 an inference, we can also denote by $F(S)$ the set of unique CPRs R in S whose consequents are in the frontier. Now, we consider the impact of forward chaining in our semantics for CPRs.

Definition 13 *A deductive set S is said to be* consistent *with CPR R if and only if $S \cup \{R\}$ is an inference.*

Definition 13 above implies that continuing forward chaining from S with CPR R is valid only when no inconsistencies in rv assignments can occur.

Proposition 7 *If S is consistent with R, then S is also an inference.*

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

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

Lemma 8 *$S_1 \in D_B(\mathsf{ant}(R) \cup \mathsf{con}(R), R)$ if and only if both $S_2 \in D_B(\mathsf{ant}(R), R)$ and $S_2 = S_1 - \{R\}$.*

Lemma 8 proves that there exists a one-to-one and onto mapping between deductive sets in $D_B(\mathsf{ant}(R) \cup \mathsf{con}(R), R)$ and $D_B(\mathsf{ant}(R), R)$.

Theorem 9

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

Examining Theorem 9, the fraction seems closely related to the definition of conditional probabilities where the numerator reflects $P(\mathsf{ant}(R) \cup \mathsf{con}(R))$ and the denominator, $P(\mathsf{ant}(R))$. 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.

Assignment Completeness

The inequalities found in Theorems 3 and 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 14 *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 3, 4, and 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 consistent 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 "conclude" with only consequents found in T.

Proposition 10 *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 10 states that there exists some rv assignment in $V(I_1)$ that is incompatible with $V(I_2)$.

Theorem 11 *For any set T defined above,*

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

Theorem 11 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.

Conditional Probabilities

Returning to the sets of inferences $D_B(\mathsf{ant}(R) \cup \mathsf{con}(R), R)$ and $D_B(\mathsf{ant}(R), R)$ in Theorem 9, these sets reflect inferences that support $\mathsf{ant}(R) \cup \mathsf{con}(R)$ and $\mathsf{ant}(R)$, respectively, and whose frontiers are bounded by $\mathsf{ant}(R) \cup \mathsf{con}(R)$ and $\mathsf{ant}(R)$, respectively. We now examine the relationships between the sets $D_B(\mathsf{ant}(R))$ and $D_B(\mathsf{ant}(R) \cup \mathsf{con}(R))$ to the sets $I_B(\mathsf{ant}(R))$ and $I_B(\mathsf{ant}(R) \cup \mathsf{con}(R))$.

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

Definition 15 *S is said to be* unstable *if $\{A = a\} = \mathsf{con}(R_n)$ and $\{A = a'\} \in \mathsf{ant}(R_1)$. (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 Figures 1 through 3.

Theorem 12 $D_B(\mathsf{ant}(R) \cup \mathsf{con}(R), R) = I_B(\mathsf{ant}(R) \cup \mathsf{con}(R))$.

Lemma 13 $D_B(\mathsf{ant}(R), R) \subseteq I_B(\mathsf{ant}(R))$.

Theorem 14 *If B is stable, then* $D_B(\mathsf{ant}(R), R) = I_B(\mathsf{ant}(R))$.

Combining Theorems 11, 12, and 14 above, we get the following:

Theorem 15 *If B is stable, assignment complete, and probabilistically complete, then for all $R \in B$, $P(R)$ is a conditional probability consistent with p.*

When B is not probabilistically complete, the summations $\sum_{S_1 \in D_B(\mathsf{ant}(R) \cup \mathsf{con}(R), R)} P(S_1)$ and $\sum_{S_2 \in D_B(\mathsf{ant}(R), R)} P(S_2)$ approaches $P(\mathsf{ant}(R) \cup \mathsf{con}(R))$ and $P(\mathsf{ant}(R))$, respectively, as B is completed.

Clearly, changes to B affect the various joint probabilities found in the BKB. However, from Theorem 9, 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 by using a variant on depth-first search on the graphical representation for BKBs. While problems of stability arise from cyclicity, stability does not preclude all forms of cyclicity. Figure 1 with underlying rv cyclicity is stable. Furthermore, stable BKBs properly subsume a special class of BKBs called causal BKBs (Santos & Santos 1999). Causal BKBs admit a polynomial time reasoning algorithm.

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 16 *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 16 *If R is grounded in both B and B' and was not modified, then $P(R)$ satisfies Equation 1 in both B and B'.*

Conclusions

In this paper, we presented new results regarding how Bayesian Knowledge-Bases 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. From these results, we believe BKBs to be an ideal knowledge representation for constructing knowledge-based systems.

Acknowledgments. This paper was supported in part by AFOSR Grant Nos. #940006 and F49620-99-1-0244 and the Paul Ivanier Center for Robotics and Production Management, BGU.

References

Barr, V. 1999. Applying reliability engineering to expert systems. In *Proceedings of the 12th International FLAIRS Conference*, 494–498.

Boutilier, C.; Friedman, N.; Goldszmidt, M.; and Koller, D. 1996. Context-specific independence in Bayesian networks. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*. San Francisco, CA: Morgan Kaufmann Publishers.

Johnson, G., and Santos, Jr., E. 2000. Generalizing knowledge representation rules for uncertain knowledge. In *Proceedings of the 13th International FLAIRS Conference*, 186–190.

Knauf, R.; Gonzalez, A. J.; and Jantke, K. P. 2000. Towards validation of rule-based systems – the loop is closed. In *Proceedings of the Thirteenth International FLAIRS Conference*, 331–335. AAAI Press.

Pearl, J. 2000. *CAUSALITY: Models, Reasoning, and Inference*. Cambridge University Press.

Poole, D. 1993. The use of conflicts in searching Bayesian networks. In *Uncertainty in AI, Proceedings of the 9th Conference*.

Poole, D. 1997. Probabilistic partial evaluation: Exploiting rule structure in probabilistic inference. In *IJCAI*, 1284–1291.

Rosen, T.; Shimony, S. E.; and Santos, Jr., E. 2001. Reasoning with bkbs – algorithms and complexity. Technical Report IDIS Technical Report 103, Intelligent Distributed Information Systems Laboratory, University of Connecticut.

Santos, Jr., E., and Santos, E. S. 1999. A framework for building knowledge-bases under uncertainty. *Journal of Experimental and Theoretical Artificial Intelligence* 11:265–286.

Santos, Jr., E.; Banks, S. B.; Brown, S. M.; and Bawcom, D. J. 1999. Identifying and handling structural incompleteness for validation of probabilistic knowledge-bases. In *Proceedings of the 12th International FLAIRS Conference*, 506–510.

Santos, Jr., E. 2001. Verification and validation of knowledge-bases under uncertainty. *Data and Knowledge Engineering* 37:307–329.

Selman, B., and Levesque, H. J. 1990. Abductive and default reasoning: A computational core. In *Proceedings of the AAAI Conference*, 343–348.

Shimony, S. E.; Domshlak, C.; and Santos, Jr., E. 1997. Cost-sharing heuristic for Bayesian knowledge-bases. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 421–428.

Shimony, S. E.; Santos, Jr., E.; and Rosen, T. 2000. Independence semantics for bkbs. In *Proceedings of the 13th International FLAIRS Conference*, 308–312.

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