

Modelling Multiagent Bayesian networks with Inclusion Dependencies

C.J. Butz, F. Fang

Department of Computer Science, University of Regina

Regina, Saskatchewan, Canada S4S 0A2

Email: {butz, fang11fa}@cs.uregina.ca

Abstract

Multiagent Bayesian networks (MABNs) are a powerful new framework for uncertainty management in a distributed environment. In a MABN, a collective joint probability distribution is defined by the conditional probability tables (CPTs) supplied by the individual agents. It is assumed, however, that CPTs supplied by individual agents agree on the variable domains, an assumption that does not necessarily hold in practice.

In this paper, we suggest modelling MABNs with inclusion dependencies. Our approach is more flexible, and perhaps realistic, by allowing CPTs supplied by different agents to disagree on variable domains. Our main result is that the input CPTs define a joint probability distribution if and only if certain inclusion dependencies are satisfied. Other advantages, both practical and theoretical, of modelling MABNs with inclusion dependencies are discussed.

1. Introduction

Over the past twenty years, *Bayesian networks* [3] have established themselves as an elegant practical framework for uncertainty management in a centralized and single agent setting. A Bayesian network consists of a *directed acyclic graph (DAG)* and a corresponding set of *conditional probability tables (CPTs)*. The *probabilistic conditional independencies* encoded in the DAG indicate that the product of the CPTs is a *joint probability distribution* [3].

Xiang [7, 9] extended Bayesian networks to the distributed and multiagent setting, which we will call *multiagent Bayesian networks (MABNs)*. In a MABN, the CPTs are supplied by individual agents, who are willing to cooperate to reach a common goal. There are at least two potential problems in designing a MABN: (i) the

dependency structure is *not* a DAG, and (ii) the product of the CPTs is *not* a joint probability distribution.

Xiang [8], and Wong and Butz [6], developed methods to ensure the dependency structure is a DAG. Problem (ii) can be sidestepped by assuming that variables common to multiple agents have the same domain. We believe, however, that problem (ii) has not been properly addressed, since the assumption made in the solution is not necessarily satisfied in practice.

In this paper, we consider the ramifications of allowing agents to disagree on the domain of variables. It is explicitly demonstrated that the product of the supplied CPTs is *not* a joint probability distribution, even if the input CPTs are well-defined with respect to the local domains of the individual agents. To ensure consistency, we introduce the notion of inclusion independency. Our main result is that the product of the input CPTs is a joint probability distribution if and only if certain inclusion dependencies are satisfied. Our approach possesses other salient features. Whether or not an agent needs to update the other agents with respect to some collected evidence depends on the *satisfaction* of certain inclusion dependencies. While the uncertainty community is interested in more general forms of probabilistic conditional independence [4], our proposed inclusion dependency is the *first* constraint proposed involving multiple distributions. Finally, unlike the work in [4, 5, 6] that utilized the relationship between single agent Bayesian networks and conventional relational databases, this manuscript establishes a useful relationship between *current* research in multiagent Bayesian networks with *recent* work on semantic data models in databases.

This paper is organized as follows. Inclusion dependencies, and their role in MABNs, are presented in Section 2. In Section 3, the advantages of modelling MABNs with inclusion dependencies are provided. The conclusion is presented in Section 4.

2. Inclusion Dependencies in MABNs

A *multiagent Bayesian network* (MABN) is a Bayesian network in which the CPTs are supplied by multiple agents. In this manuscript, we assume that the agents are willing to cooperate and share their knowledge to reach a common goal. Here we allow the agents in a MABN to disagree on the domains of the variables. Nevertheless, each agent can only supply CPTs that are well-defined with respect to the local domain.

Example 1. Consider a multiagent system consisting of two agents, say Agent 1 and Agent 2. Agent 1 is an expert on domain $\{a, b, c\}$, while Agent 2's expertise involves domain $\{b, c, d\}$. A MABN on $U = \{a, b, c, d\}$ can be defined by Agent 1 specifying the CPTs $p_1(a)$, $p_1(b|a)$ and $p_1(c|a)$, while Agent 2 inputs the CPT $p_2(d|b, c)$. This MABN is illustrated in Figure 1. Suppose Agent 1 and Agent 2 disagree on the domains of variables b and c . For instance, Agent 1 believes $\text{dom}(b) = \text{dom}(c) = \{0, 1, 2\}$, while Agent 2 believes $\text{dom}(b) = \text{dom}(c) = \{0, 1\}$. The supplied CPTs in Table 1 are well-defined with respect to the local domains. However, their product $\Phi(a, b, c, d)$ is a potential [3].




Figure 1. A multiagent Bayesian network (MABN)

Table 1. CPTs $p_1(a)$, $p_1(b|a)$, $p_1(c|a)$, and $p_2(d|b, c)$

a	$p_1(a)$	a b	$p_1(b a)$	a c	$p_1(c a)$	b c d	$p_2(d b, c)$
1	0.8	1 2	0.1	1 2	0.2	1 1 1	1.0
0	0.2	1 1	0.2	1 1	0.2	1 0 1	0.6
		1 0	0.7	1 0	0.6	1 0 0	0.4
		0 2	0.5	0 2	0.3	0 1 1	0.7
		0 1	0.2	0 1	0.3	0 1 0	0.3
		0 0	0.3	0 0	0.4	0 0 1	0.8
						0 0 0	0.2

The important point in Example 1 is that the product of the supplied CPTs is *not* a joint probability distribution, even though the CPTs are well-defined with respect to the local domains of the individual agents.

Our objective now is to ensure that the product of the supplied CPTs is a joint probability distribution, while at the same time allowing the agents to disagree on the

domain of the variables. We introduce the notion of *inclusion dependency* to assist us in meeting our goal.

Definition 1. Consider two distributions $\Phi_1(X)$ and $\Phi_2(XY)$. We say $\Phi_1(X)$ and $\Phi_2(XY)$ satisfy the *inclusion dependency*, if $\pi_X(\Phi_1(X)) \subseteq \pi_X(\Phi_2(XY))$, where π is the *projection operator* in relational databases [1]. That is, the configurations of variable X in the distribution $\Phi_1(X)$ must also appear as values of X in the distribution $\Phi_2(XY)$.

Example 2. Recall the MABN in Figure 1. Consider another set of supplied CPTs in Table 2. Here $p_1(b, c)$ and $p_2(d|b, c)$ satisfy the inclusion dependency, since $\pi_{b,c}(p_1(b, c)) \subseteq \pi_{b,c}(p_2(d|b, c))$, as shown in Table 3, and where $p_1(b, c) = \sum_a p_1(a) \cdot p_1(b|a) \cdot p_1(c|a)$.

Table 2. A second set of CPTs, where Agent 1 specifies $p_1(a)$, $p_1(b|a)$ and $p_1(c|a)$, while Agent 2 supplies CPT $p_2(d|b, c)$

a	$p_1(a)$	a b	$p_1(b a)$	a c	$p_1(c a)$	b c d	$p_2(d b, c)$
1	0.8	1 1	0.7	1 1	0.1	1 1 1	1.0
		0 1	0.6	0 1	0.8	1 0 1	0.6
						0 1 1	0.7
						0 0 1	0.8
						2 2 1	0.5

Table 3. The inclusion dependency holds since the configurations of variables b and c in the $p_1(b, c)$ are contained in $p_2(d|b, c)$

	b c		b c
$\pi_{b,c}(p_1(b, c)) =$	1 1	\subseteq	1 1
	1 0		1 0
	0 1		0 1
	0 0		0 0
			2 2

Here, a numbering of the variables in a DAG is called *ancestral*, if the number corresponding to any variable is lower than the numbers corresponding to all of its children. Such as, $a = 1, b = 2, c = 3, d = 4$ is one ancestral numbering of the variables in the DAG of Figure 1.

Definition 2. Consider a MABN with CPTs $p(A_1|P_1)$, $p(A_2|P_2)$, ..., $p(A_n|P_n)$, where A_1, A_2, \dots, A_n is an ancestral numbering. We say the MABN satisfies the *inclusion principle*, if

$$\pi_{P_i}(p(A_1|P_1) \cdot p(A_2|P_2) \cdot \dots \cdot p(A_{i-1}|P_{i-1})) \subseteq \pi_{P_i}(p(A_i|P_i)),$$

for $i = 2, \dots, n$.

Modelling MABNs as a Bayesian network satisfying the inclusion principle is very useful, as we explicitly demonstrate in the next section.

3. Advantages of modelling MABNs with inclusion dependencies

The practical and theoretical advantages of modelling MABNs with inclusion dependencies are following.

3.1 Practical Advantages

Inclusion dependencies possess at least two desirable characteristics in practical applications. First, they remove the questionable assumption traditionally made when *modelling* MABNs. Second, inclusion dependencies can dramatically alleviate some computation when *updating* the MABN with collected evidence in some situations.

Our main result (Theorem 1) is that the CPTs for a Bayesian network define a joint probability distribution if and only if the MABN satisfies the inclusion principle.

Theorem 1. Let $p(A_1|P_1), p(A_2|P_2), \dots, p(A_n|P_n)$ be a MABN, where A_1, A_2, \dots, A_n is an ancestral numbering. The product of these CPTs is a joint probability distribution if and only if the MABN satisfies the inclusion principle.

Proof: (\Rightarrow) Given $p(A_1, A_2, \dots, A_n) = p(A_1|P_1) \cdot p(A_2|P_2) \cdot \dots \cdot p(A_n|P_n)$. By contradiction, suppose the MABN does not satisfy inclusion principle. Then

$\pi_{P_j}(p(A_1|P_1) \cdot p(A_2|P_2) \cdot \dots \cdot p(A_{j-1}|P_{j-1})) \not\subseteq \pi_{P_j}(p(A_j|P_j))$,
for at least one $j \in [2, n]$. Since configurations with zero probability are not stored, there is a configuration $\langle A_1:a_1, A_2:a_2, \dots, A_{j-1}:a_{j-1} \rangle$ of $A_1 A_2 \dots A_{j-1}$ appearing in

$$p(A_1|P_1) \cdot p(A_2|P_2) \cdot \dots \cdot p(A_{j-1}|P_{j-1}), \quad (1)$$

but not appearing in the CPT $p(A_{j-1}|P_{j-1})$. Hence, $\langle A_1:a_1, A_2:a_2, \dots, A_{j-1}:a_{j-1} \rangle$ cannot appear in the joint probability distribution $p(U)$, nor in the marginal distribution $p(A_1, A_2, \dots, A_{j-1})$ obtained by marginalization of $p(U)$. This is a contradiction, since Equation (1) is also the definition of the marginal distribution $p(A_1, A_2, \dots, A_{j-1})$.

(\Leftarrow) Suppose the MABN $p(A_1|P_1), p(A_2|P_2), \dots, p(A_n|P_n)$ satisfies the inclusion principle. By definition, the CPT $p(A_i|P_i)$ must contain every configuration of P_i in $p(P_i)$. In other words, the joinable configurations are exactly those in $p(P_i)$. Hence, $p(A_i, P_i) = p(P_i) \cdot p(A_i|P_i)$, for $i = 2, \dots, n$. Thus, when $i = n$, the product defines a joint probability distribution.

Theorem 1 is important, since it indicates that previous work on MABNs [6, 7, 9] enforce a superfluous restriction, namely, the domains of common variables must be the same. The next example clearly shows that this restriction is unnecessary.

Example 3. Recall the MABN in Example 2. Even though the domains of variables b and c are *different* with respect to Agent 1 and Agent 2, the product, nevertheless, is still a joint probability distribution. That is,

$$p(a, b, c, d) = p_1(a) \cdot p_1(b|a) \cdot p_1(c|a) \cdot p_2(d|b, c).$$

While Example 3 explicitly demonstrates that common variables do not necessarily have to have the same domain, the next example emphasizes the importance of the MABN satisfying the inclusion principle.

Example 4. The MABN, defined by the CPTs in Table 1, does *not* satisfy the inclusion principle. That is, $\pi_{bc}(p_1(b, c)) \not\subseteq \pi_{bc}(p_2(d|b, c))$, with $\langle b:2, c:2 \rangle$ appears in $p_1(b, c)$ but not $p_2(d|b, c)$. By Theorem 1, $\Phi(a, b, c, d) = p_1(a) \cdot p_1(b|a) \cdot p_1(c|a) \cdot p_2(d|b, c)$ is potential.

The second practical advantage of modelling MABNs with inclusion dependencies involves the processing of evidence. Traditionally, when an agent collects *hard* and *soft* evidence E [2], two tasks are carried out: (i) the agent updates its knowledge base with E , and (ii) all other agents are then updated with E . On the contrary, if the MABN satisfies the inclusion principle, then step (ii) is not necessarily required.

Example 5. The MABN, defined by the CPTs in Table 5, satisfies the inclusion principle. Suppose Agent 2 collects the soft evidence $b \neq 2$ and $c \neq 2$. It performs task (i) by updating its knowledge base $p_2(d|b, c)$ using the *selection* operator σ [2], as illustrated in Table 4. However, performing task (ii) is unnecessary, since the updated MABN in Table 5 still satisfies the inclusion principle. By Theorem 1, we obtain the joint distribution $p(a, b \neq 2, c \neq 2, d) = p_1(a) \cdot p_1(b|a) \cdot p_1(c|a) \cdot p_2(d|b \neq 2, c \neq 2)$.

Table 4. Agent 2 updates $p_2(d|b, c)$ with soft evidence $b \neq 2$ and $c \neq 2$ as $\sigma_{b \neq 2 \wedge c \neq 2}(p_2(d|b, c))$

b	c	d	$p_2(d b, c)$	b	c	d	$p_2(d b \neq 2, c \neq 2)$
1	1	1	1.0	1	1	1	1.0
1	0	1	0.6	1	0	1	0.6
0	1	1	0.7	0	1	1	0.7
0	0	1	0.8	0	0	1	0.8
2	2	1	0.5				

Table 5. The MABN still satisfies the inclusion principle after Agent 2 updates itself. By Theorem 1, the product is a joint distribution

a	$p_1(a)$	a	b	$p_1(b a)$	a	c	$p_1(c a)$	b	c	d	$p_2(d b \neq 2, c \neq 2)$
1	0.8	1	1	0.7	1	1	0.1	1	1	1	1.0
			0	1	0	1	0.8	1	0	1	0.6
								0	1	1	0.7
								0	0	1	0.8

Example 5 reveals that the other agents need only be updated when the collected evidence violates the inclusion dependencies in a MABN. In these cases, it can be verified that after updating the MABN, the appropriate inclusion dependencies are again satisfied since common variables have the same domain.

3.2 Theoretical Advantages

Here we discuss two favourable features of utilizing inclusion dependencies in MABNs. One involves general forms of independencies within the Bayesian network community, while another involves the relationship between the Bayesian network and database communities.

The Bayesian network community is interested in general forms of probabilistic conditional independence [4]. The important point is that their notions all are defined with respect to a *single* distribution. On the contrary, our notion of inclusion dependency is the *first* independency proposed for single agent and multiagent Bayesian networks that involves multiple distributions. Hence, inclusion dependencies may lead to other forms of independencies holding between multiple distributions.

Finally, the work in this paper also establishes an intrinsic relationship between *current* research in Bayesian networks with that in databases. Our past research, which has been well-received by the international community [4, 5, 6], took full advantage of the inherent relationship between single agent Bayesian networks and conventional relational databases. Nowadays, however, research focuses on MABNs [7, 9] and semantic data models in databases [1]. Semantic data models usually encode *functional dependencies* and *inclusion dependencies* [1]. In this paper, we adopt the notion of inclusion dependency for modelling MABNs. Thus, this work is yet another example of how the uncertainty management and data management communities can benefit each other. Modelling MABNs with inclusion dependencies is also meritorious, since the use of inclusion dependencies is arguably closer to ways that humans organize information [1].

4. Conclusion

Based on the striking successes of applying single agent Bayesian networks for uncertainty management, MABNs are emerging as a basis for uncertainty management in a distributed setting [6, 7, 9]. These works, however, assume that variables common to multiple agents have the same domain. This assumption is not necessarily satisfied in practice.

In this paper, we have suggested modelling MABNs with inclusion dependencies. Our main result is that the

collection of CPTs supplied by the individual agents is a joint probability distribution if and only if the MABN satisfies the inclusion principle. We have also explicitly shown that inclusion dependencies may facilitate probabilistic inference during the processing of collected evidence in practice. On the theoretical side, unlike the general forms of independence [4], inclusion dependencies are the *first* constraint involving multiple distributions. Finally, in contrast to [4, 5, 6] this study establishes a relationship between the *current* work on MABNs in the uncertainty management community with the *current* research on semantic data models in the database community [1].

Reference

- [1] Abiteboul, S., Hull, R., and Vianu, V., *Foundations of Databases*, Addison-Wesley, United States of America, 1995.
- [2] Butz, C.J. and Fang, F., "Incorporating Evidence in Bayesian networks with the Select Operator," in *Proceedings of the 18th Canadian Conference on Artificial Intelligence*, Springer, Victoria, British-Columbia, 2005.
- [3] Pearl, J., *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, San Francisco, United States of America, 1988.
- [4] Wong, S.K.M., Butz, C.J. and Xiang, Y., "A Method for Implementing a Probabilistic Model as a Relational Database," in *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann, Montreal, Quebec, 1995, pp. 556-564.
- [5] Wong, S.K.M. and Butz, C.J., "Contextual Weak Independence in Bayesian networks," *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann, Montreal, Quebec, 1999, pp. 670-679.
- [6] Wong, S.K.M and Butz, C.J., "Constructing the Dependency Structure of a Multiagent Probabilistic Network," *IEEE Transactions on Knowledge and Data Engineering*, IEEE Computer Society, 2001, Vol. 13, No. 3, pp. 395-415.
- [7] Xiang, Y., "A Probabilistic Framework for Cooperative Multi-agent Distributed Interpretation and Optimization of Communication," *Artificial Intelligence*, Elsevier Science, 1996, Vol. 87, No. 1-2, pp. 295-342.
- [8] Xiang, Y., "Verification of DAG Structures in Cooperative Belief Network-Based Multi-agent Systems," *Networks*, Wiley InterScience, 1998, Vol. 31, Issue 3, pp. 183-191.
- [9] Xiang, Y., *Probabilistic Reasoning in Multiagent Systems: A Graphical Models Approach*, Cambridge University Press, Cambridge, United Kingdom, 2002.