

# Distributed Multi-agent Probabilistic Reasoning With Bayesian Networks

Yang Xiang

Department of Computer Science, University of Regina  
Regina, Saskatchewan, Canada S4S 0A2, yxiang@cs.uregina.ca

**Abstract.** Main stream approaches in distributed artificial intelligence (DAI) are essentially logic-based. Little has been reported to explore probabilistic approach in DAI. On the other hand, Bayesian networks have been applied to many AI tasks that require reasoning under uncertainty. However, as commonly applied, a single-agent paradigm is assumed in Bayesian networks.

This paper extends multiply sectioned Bayesian networks (MSBNs) for single-agent systems into a framework for multi-agent distributed interpretation systems. Each cooperative agent is represented as a Bayesian subnet that consumes its own computational resource, gathers its own evidence, and can answer queries. Unlike in single-agent systems where evidence is entered one subnet at a time, multiple agents may acquire evidence in parallel. We add to the set of single-agent MSBN operations new belief propagation operations which, when performed, regain global consistency. Due to the inter-agent 'distance' and the associated communication cost, global consistency can not be maintained constantly. We show that, when the proposed operations are followed, between two successive communications, the answers to queries from an agent are consistent with all local evidence gathered so far, and are consistent with all global evidence gathered up to the last communication.

**Keywords:** knowledge representation and integration, approximate reasoning, probabilistic reasoning, Bayesian networks, distributed artificial intelligence, distributed reasoning.

## 1 Introduction

Probabilistic reasoning in Bayesian networks (BNs), as commonly applied, assumes a single-agent paradigm: A single processor accesses a single global network representation, updates the joint probability distribution over the network variables as evidence becomes available, and answers queries. Concurrency, as applied to BNs, primarily aims at performance and decentralization of control [10, 4], but not at modeling inference among multiple agents with multiple perspectives. The resultant individual concurrent element is thus 'fine-grained', e.g., a node in a BN [10] or a clique in the junction tree representation of a BN [4].

The single-agent paradigm is inadequate when uncertain reasoning is performed by elements of a system between which there is some 'distance', which may be spatial, temporal, or semantic (elements are specialized differently) [1].

Such systems pose special issues that must be addressed. A multi-agent view is thus required where each agent is an autonomous intelligent subsystem. Each agent holds its own partial domain knowledge, accesses some information source, and consumes some computational resource. Each agent communicates with other agents to achieve the system's goal cooperatively.

Distributed artificial intelligence (DAI), a subfield of artificial intelligence, addresses the problems of designing and analyzing such 'large-grained' coordinating multi-agent systems [2, 3]. Main stream approaches, e.g., blackboard systems, contract nets, and open systems are essentially logic-based. To our best knowledge, little has been reported to explore probabilistic approach in DAI.

This paper reports our pilot study to apply probabilistic approach to distributed multi-agent reasoning. Our representation is based on multiply sectioned Bayesian networks (MSBNs)[12], which were developed for single-agent-oriented, modular knowledge representation and more efficient inference[11]. We show that the modularity of MSBNs allows a natural extension into a multi-agent reasoning formalism.

We address the use of MSBNs in distributed interpretation, a subclass of problems in DAI. As defined originally by Lesser and Erman [7], an *interpretation* system accepts evidence from some environment and produces higher level descriptions of objects and events in the environment. A *distributed* interpretation system is needed when sensors for collecting evidence are distributed, and communication of all evidence to a centralized site is undesirable. Examples of such systems include sensor networks, medical diagnosis by multiple specialists, trouble-shooting of complex artifacts and distributed image interpretation.

Section 2 briefly introduce BNs and single-agent MSBNs. Section 3 presents the semantic extension of single-agent MSBNs to multi-agent MSBNs. Each cooperative agent is represented as a Bayesian subnet that consumes its own computational resource, gathers its own evidence, and can answer queries. Section 4 discusses new technical issues raised due to the extension, which are addressed in subsequent sections.

Section 5 extends the evidence entering operation proposed originally [12] to include both specific evidence and virtual evidence.

Unlike in single-agent systems where evidence is entered one subnet at a time, multiple agents may acquire evidence asynchronously in parallel. Section 6 addresses this and adds new belief propagation operations to the set of single-agent MSBN operations.

Section 7 shows that global consistency is guaranteed when the proposed operations are followed. Due to the inter-agent 'distance' and the associated communication cost, global consistency can not be maintained constantly. It is also shown that, when the proposed operations are followed, between two successive communications, the answers to queries from an agent are consistent with all local evidence gathered so far, and are consistent with all global evidence gathered up to the last communication.


Section 8 discusses the role of causal independence in our framework.

## 2 Single Agent Oriented MSBNs

A BN [10, 8, 6, 4] is a triplet  $(N, E, P)$ .  $N$  is a set of nodes (variables).  $E$  is a set of arcs such that  $D = (N, E)$  is a directed acyclic graph (DAG). For each node  $A_i \in N$ , the strengths of the causal influences from its parent nodes  $\pi_i$  are quantified by a conditional probability distribution  $p(A_i|\pi_i)$ . The basic dependency assumption embedded in BNs is that a variable is independent of its non-descendants given its parents.  $P$  is a joint probability distribution (JPD). For a BN with  $\alpha$  nodes,  $P$  can be specified by the following product due to the assumption:  $P = p(A_1 \dots A_\alpha) = \prod_{i=1}^{\alpha} p(A_i|\pi_i)$ .

In the following, we briefly introduce single-agent oriented MSBNs. For a formal presentation of MSBNs, see [12]. A MSBN consists of a set of interrelated Bayesian subnets. Each subnet represents one subdomain in a large domain. Each subnet shares a non-empty set of variables with at least one other subnet. The interfacing set between each pair of subnets must satisfy a *d-sepset* condition such that, when the pair is isolated from the MSBN, the interfacing set renders the two subnets conditionally independent.

The overall structure with which the subnets of a MSBN are organized is subject to a constraint called *soundness of sectioning*. Without the condition, a MSBN is subject to loss of information when later transformed into its secondary representation: a linked junction forest. A sufficient condition for sound sectioning is to construct a MSBN with a *hypertree* structure. The hypernodes in the hypertree are subnets of the MSBN. The hyperlinks are interfacing sets between subnets. Conceptually, the hypertree structured MSBN is built by adding one subnet (hypernode) to existing ones at a time. Only one interfacing set (hyperlink) to an existing subnet is explicitly stored. A hypertree structured MSBN guarantees that each subnet renders the rest of the MSBN conditionally independent. Figure 1 depicts a hypertree structured MSBN. Each box represents a subnet. Boundaries between boxes represent interfacing sets. The superscripts of subDAGs indicate a possible order of construction.



**Fig. 1.** A MSBN with a hypertree structure.

Once a hypertree structured MSBN is constructed, it is then converted into a *linked junction forest* of the identical hypertree structure. Each hypernode in the hypertree is a *junction tree* (clique tree, join tree) converted from a subnet in the hypertree structured MSBN. The conversion of a subnet to a junction tree

is subject to a *separability* constraint to guarantee non-distorted belief propagation. Each hyperlink of the hypertree is a set of *linkages* converted from the corresponding interfacing set in the hypertree structured MSBN.

Parallel to the transformation of the graphical structure, there is a corresponding transformation of the probability tables in the MSBN to *belief tables* in the junction forest. The conversion is subject to a *supportiveness* constraint such that belief can be propagated without blocking. The overall conversion guarantees a *joint system belief* of the linked junction forest can be assembled from belief tables distributed in the forest, which is equivalent to the joint probability distribution of the MSBN.

To answer queries by efficient local computation in a linked junction forest (LJF), the LJF must be made consistent. A LJF is *Locally consistent* if all junction trees (JTs) are internally consistent, i.e., when marginalized onto the same set of variables different belief tables in a junction tree yield the same marginal distribution. A LJF is *boundary consistent* if each pair of linked JTs are consistent with respect to their interfacing set. A LJF is *globally consistent* iff it is both locally consistent and boundary consistent.

A set of operations are developed to achieve consistency in a LJF during evidential reasoning: **BeliefInitialization** establishes initial global consistency. **DistributeEvidence** causes an outward belief propagation within a JT, and brings the JT internally consistent after evidence on variables in a single clique has been entered. **CollectEvidence** causes an inward belief propagation within a JT. **UnifyBelief** brings a JT internally consistent after evidence on variables in multiple cliques has been entered. **EnterEvidence** updates belief in a JT in light of new evidence, and brings the JT internally consistent again by calling either **DistributeEvidence** or **UnifyBelief**. **UpdateBelief** updates the belief of a JT  $T$  relative to an adjacent JT, and brings  $T$  internally consistent. **DistributeBelief** initiated at a JT  $T$  causes an outward belief propagation in the LJF radiating from  $T$ . **ShiftAttention** allows the user to enter multiple pieces of evidence into a JT of current attention, and, when the user shifts attention to a target JT, maintains consistency along the hyperpath in the hypertree structured forest from the current JT to the target.

### 3 Representing Multiple Agents in MSBNs

A MSBN represents a large domain by representing a set of subdomains. From the viewpoint of reasoning agents, a MSBN represents multiple perspectives of a single agent. For example, PAINULIM [11] consists of three subnets which represents a neurologist's three different perspectives of the neuromuscular diagnostic domain: clinical, EMG and nerve conduction perspectives.


In a multi-agent system, each agent can be considered as holding its own perspective of the domain. This partial perspective may be over a subdomain, a period of time, or a geographical area. The modular representation of MSBN thus allows a natural extension to multi-agent system, with a modification of the semantics: Instead of representing *one* agent's *multiple* perspectives of a domain, a multi-agent MSBN represents *multiple* agents in a domain each of which holds

one perspective of the domain. Each subnet corresponds to one such perspective. We extend the example by Lauritzen and Spiegelhalter [6] to illustrate this:


**Example 1** Dyspnoea ( $\delta$ ) may be due to tuberculosis ( $\tau$ ), lung cancer ( $\iota$ ) or bronchitis ( $\beta$ ). A recent visit to Asia ( $\nu$ ) increases the chances of  $\tau$ , while smoking ( $\zeta$ ) is a known risk factor for both  $\tau$  and  $\iota$ . After an initial diagnosis based on the above information, to further discriminate between  $\tau$  and  $\iota$ , a clinician may request lab tests from a radiology lab and a biology lab. Radiology lab has two relevant tests for  $\tau$  and  $\iota$ : X-ray ( $\chi$ ) and laminagraphy ( $\alpha$ ). Biology lab has two relevant tests as well: sputum test ( $\rho$ ) and biopsy ( $o$ ). Lab reports describe their impression upon  $\tau$  and  $\iota$  based on the results of the test(s).

The above fictitious example involves three agents: a clinical doctor, a radiologist, and a biologist. Diagnosis of a patient with dyspnoea is their common goal. But each agent has its own perspective of the same patient. The distance between them is semantic.

A multi-agent MSBN can be constructed as a decision aid in such a domain. The MSBN consists of three agents: clinical, radiological, and biological subnets. The three agents may process evidence in parallel. Though it makes sense that the radiologist and biologist may not be involved until the clinician's initial diagnosis has been reached, and the patient has to visit them in sequence, considering the time delay required to develop lab results, the radiological and biological subnets may well receive evidence and be posed with queries at the same time. The resultant multi-agent MSBN is illustrated in Figure 2. The transformed LJF is shown in Figure 3.



**Fig. 2.** A multi-agent MSBN representing three medical specialists diagnosing a patient with dyspnoea.  $D^1$ : clinical subnet,  $D^2$ : radiological subnet,  $D^3$ : biological subnet.



**Fig. 3.** The linked junction forest for the multi-agent MSBN in Figure 1. Sepsets between cliques are shown in solid lines. Linkages are shown in dotted lines.

## 4 Consistency Issues in Multi-agent MSBNs

The natural extension of MSBNs to multi-agent systems implies that all the technical constraints applicable to the construction of single-agent MSBNs must be followed in the construction of multi-agent MSBNs. In addition, evidential reasoning in multi-agent systems raises new issues regarding consistency, which must be addressed. To appreciate the issues, we first review how consistency is maintained in single-agent MSBNs:

Initial global consistency is obtained by **BeliefInitialization** (Section 2). The agent focuses attention on one subdomain at a time. Therefore, evidence is entered *one subnet at a time*. As the agent shifts attention to each subnet and enters evidence, **ShiftAttention** (Section 2) maintains (local) consistency along the hyperpath in the hypertree structured forest. It is proven [12] that such local consistency is actually at the global level, i.e., answers to queries at the current JT is consistent with the evidence acquired in the entire LJF.

### 4.1 How to regain global consistency?

Notice that **ShiftAttention** being able to maintain local consistency at the global level depends directly on the fact that evidence is always entered at the current subnet and nowhere else. In a multi-agent system, multiple agents may acquire evidence asynchronously in parallel. **ShiftAttention** can no longer be used to maintain global consistency. It must be replaced by new operations which we shall refer to as *communication*. We propose the communication operations in Section 6 and prove their property in Section 7.

### 4.2 What is the consistency level between communications?

**ShiftAttention** maintains directly only local consistency. But the local consistency happens to be at the global level in a single-agent MSBN! In a multi-agent MSBN, local consistency at a JT means one thing, and the global consistency means another. We can no longer shoot two birds with one stone.

Even with new communication operations, due to the inter-agent 'distance' and the associated communication cost, global consistency can not be maintained constantly. A question that must be answered is, between two successive communications which regain global consistency, what is the consistency level of each agent after acquiring additional evidence? We prove a theorem to answer this question in Section 7.

## 5 Entering Virtual Evidence

Entering virtual evidence (described below) to a directed tree has been discussed in [8]. Undirected tree representations (junction tree or LJF) provide more flexible inference mechanism. But to our best knowledge, virtual evidence has not been treated formally in such representations. Evidence entering operations for such representations only considered entering specific evidence (described below) into a junction tree [5] or a LJF [12]. To be complete, in this section, we extend **EnterEvidence** to include entering virtual evidence into a LJF. This result is

applicable to both single-agent MSBNs and multi-agent MSBNs. Our presentation assumes some familiarity with the technical details of MSBN construction. Readers who are interested in results pertinent to only multi-agent reasoning can skip this section, and assume the previous **EnterEvidence** operation for the rest of the paper.

A *piece* of evidence is obtained by one agent’s observation, made at one time, of a set of variables contained in its subnet. Different agents may acquire pieces of evidence from their local sources asynchronously in parallel.

Pearl [10] classifies evidence into *specific* and *virtual*. Specific evidence on a variable is obtained by a direct observation of the value of the variable. Virtual evidence “corresponds to judgments based on undisclosed observations that are outside the network but have a bearing on variables in the network”. For example, an X-ray image of chest may be regarded as normal or abnormal. Since there is no universally agreeable way to determine mechanically the status of an X-ray image, a radiologist must interpret the pattern subjectively.

Each piece of evidence is represented by an evidence function. See [12] on the evidence function for specific evidence. We first consider a piece of virtual evidence that involves only a single variable. Suppose the virtual evidence has a direct bearing on a variable  $A$ . Pearl [10] suggests to code such evidence into a BN by adding a child node  $e$  to  $A$  and assessing the distribution  $p(e|A)$ . Following this representation, we analyze its impact on the LJF. Since  $e$  is contained in only one subnet in the MSBN, it is sufficient to analyze its impact on the subnet/JT that contains  $e$ . We shall denote them by  $S$  and  $T$ , respectively.

During moralization and triangulation, since the node  $e$  is a leaf with a single parent, no link is added between  $e$  and other nodes in  $S$ . Hence the only clique that  $e$  will participate is  $C_0 = \{e, A\}$ . During construction of  $T$ , since the only possible variable that  $C_0$  can share with other cliques is  $A$ ,  $C_0$  can always be configured as a leaf in  $T$  with an arbitrary neighbour clique  $C$  such that  $A \in C$ .

Suppose  $T$  is internally consistent before  $e$  is observed. After  $e$  is observed, *DistributeEvidence* can be called on  $C_0$ . The updated belief table (BT) on  $C$  is  $B'(C) = B(C) * B'(A)/B(A)$  where  $B(A)$  and  $B'(A)$  stand for the previous and the updated sepset BTs, respectively. Since  $e$  is observed as *True*,  $B'(A) = \sum_e B'(C_0) = \sum_e B'(eA) = B'(eA) = p(e|A) * B(A)$ . Substituting  $B'(A)$  in the above equation, we have  $B'(C) = B(C) * p(e|A)$ . This means that virtual evidence can be represented and entered in a LJF in exactly the same way as specific evidence, and we do not need to create the clique  $C_0$  at all in the first place. We summarize the above analysis in the following definition:

**Definition 2 (Evidence function)** *Let  $S = (N, E, P)$  be a subnet in a MSBN. Let  $X \subset N$  be a set of variables involved in a piece of evidence, and let the space of  $A \in X$  be  $\mathcal{A}$ . An evidence function  $f : \bigcup_{A \in X} \mathcal{A} \rightarrow [0, 1]$  is defined by the following rule:*

1. *If  $A$  is directly observable, for each  $a \in \mathcal{A}$  assign to  $f(a)$  either 1 or 0 depending on whether  $a$  is still a possible outcome of  $A$ .*

2. If  $A$ 's value is not directly observable, but is interpreted by an autonomous interpreter based on an observation  $e$ , then assign to  $f(a)$  the subjective conditional probability  $p(e|a)$  of the interpreter.

Let  $f : X \rightarrow Z$  be a function. Let  $Y \subset X$  be a subset of the domain  $X$ . We define the *restriction* of  $f$  restricted to  $Y$  as a function  $f_Y : Y \rightarrow Z$  according to the rule: for each  $y \in Y$ ,  $f_Y(y) = f(y)$ .

The extended **EnterEvidence** operation that handles both specific and virtual evidence is defined as follows:

**Operation 3 (EnterEvidence)** *Let  $T$  be a JT in a LJF. Let  $X$  be a set of variables in  $T$  that is involved in a piece of evidence  $E$  represented by an evidence function  $f : \bigcup_{A \in X} \mathcal{A} \rightarrow [0, 1]$ . When **EnterEvidence** is initiated at  $T$  to enter  $E$ , the following are performed: (1) For each  $A \in X$ , a belief universe  $U = (C, B(C))$  such that  $A \in C$  is arbitrarily selected, and  $B(C)$  is multiplied by  $f$  restricted to  $\mathcal{A}$ . (2) If  $X \subseteq C$ , i.e., only a single universe is involved in the above step, **DistributeEvidence** is called in  $U$ , otherwise **UnifyBelief** is called in any universe. **EnterEvidence** is associated with JTs.*

## 6 Added Operations for Regaining Global Consistency

As discussed in Section 4, parallel evidence entering at multiple agents renders the single-agent MSBN operation for maintaining global consistency invalid. In order to regain global consistency in a multi-agent LJF, we extend the inward-outward belief propagation method in a single junction tree [4] to a LJF. Jensen's method propagates belief through a single information channel (a unique path exists between any two cliques in a junction tree). Belief propagation in a LJF must be performed over multiple linkages. Fortunately, the latter problem has been solved in single-agent MSBNs with the operation **UpdateBelief**.

Following the above idea, we add two new operations **CollectNewBelief** and **CommunicateBelief** to regain global consistency in a multi-agent LJF. **CollectNewBelief** causes an inward belief propagation in the LJF. **CommunicateBelief** calls **CollectNewBelief** and **DistributeBelief** to propagate evidence obtained from multiple agents (JTs) inward first and then outward to the entire LJF.

**Operation 4 (CollectNewBelief)** *Let  $T$  be a JT in a LJF. Let **caller** be either the LJF or a neighbour JT. When **CollectNewBelief** is called in  $T$ , the following are performed: (1)  $T$  calls **CollectNewBelief** in all neighbours except **caller** if **caller** is a neighbour. (2) After each neighbour being called has finished **CollectNewBelief**,  $T$  updates its belief with respect to the neighbour by **UpdateBelief**. **CollectNewBelief** is associated with JTs.*

**Operation 5 (CommunicateBelief)** *When **CommunicateBelief** is initiated at a LJF  $F$ , the following are performed: (1) A JT  $T$  in  $F$  is arbitrarily selected. (2) **CollectNewBelief** is called in  $T$ . (3) When  $T$  has finished **CollectNewBelief**, **DistributeBelief** is called in  $T$ . **CommunicateBelief** is associated with the LJF.*



## 7 Consistency After and Between Communications

We answer the two questions raised in Section 4. Proofs of theorems are omitted due to the limited space. First, we show that the operations proposed in Section 6, when performed, guarantees global consistency among multiple agents:

**Theorem 6 (Multi-agent consistency)** *Let  $F$  be a supportive, globally consistent and separable LJF converted from a MSBN of hypertree structure. Let  $Z$  be a subset of JTs of  $F$ . After the following operations,  $F$  is globally consistent: (1) For each JT in  $Z$ , use **EnterEvidence** to enter finite pieces of evidence into the JT. (2) Use **CommunicateBelief** to communicate belief among JTs.*

**CommunicateBelief** involves global computation of multiple agents and information exchange over 'distance'. Due to the cost involved, **CommunicateBelief** can not be performed frequently. Therefore, each agent may acquire multiple pieces of evidence between two successive **CommunicateBelief** operations, and may have to answer queries before the next **CommunicateBelief** can be performed. The second question we address is: what is the consistency level of these answers to queries. Theorem 7 shows that, between two successive communications, a JT is consistent with all local evidence acquired so far, and is consistent with all global evidence acquired up to the last communication. This is the best that one can expect.

**Theorem 7 (Semi-up-to-date)** *Let  $F$  be a supportive and separable LJF converted from a MSBN of hypertree structure. Let  $Z$  be a subset of JTs of  $F$ .*

*After a **CommunicateBelief** in  $F$  followed by a finite number of **EnterEvidence** to each JT in  $Z$ , the marginal distributions obtained in a JT  $T \in Z$  are identical as would be obtained if only the **EnterEvidence** operations in  $T$  were performed after the **CommunicateBelief**.*

## 8 Discussion

In a MSBN, the interface between any pair of subnets satisfies a d-sepset condition such that, to make neighbours up-to-date, it is sufficient to pass the up-to-date distribution on the interfacing set between them and nothing else.

Each subnet is transformed into a single JT. The tree structure maintains a single information path between any pair of nodes in the JT such that local belief updating can be performed efficiently. The price paid is the multiple paths between neighbour JTs: Passing evidence from one JT to a neighbour requires computation proportional to the number of linkages between them.

Such built-in preference of the efficiency of local computation over the efficiency of communication is consistent with the general assumption in multi-agent intelligent systems. In DAI, we assume that the distance between agents prevents constant communication. Therefore, each agent must rely on its local resource in its problem solving, and can communicate with other agent only occasionally.

One of the major concerns in DAI is how to balance the need to maintain as much as possible global consistency and the need to reduce the traffic of

communication. As argued by Pearl [10], a tree structure makes use of causal independence and allows the most efficient information passage among a group of elements. Our study on MSBNs highlights the MSBNs that are organized into a *hypertree* structure. This structural preference is a more general case of the exploration of the causal independence in singly connected BNs [9] and in the JT representation of multiply connected BNs [4].

If we call an element which can render a pair of elements conditionally independent a 'causal mediator', we see an increase of complexity of the internal structures of causal mediators in the three cases. In singly connected BNs, a causal mediator is an internal node in the network. In the JT representation of multiply connected BNs, a causal mediator is an internal clique, a group of variables. In MSBNs/LJFs of hypertree structure, a causal mediator is a subnet/JT. As argued by Pearl [10], conditional independence should not be viewed as a restrictive assumption for mathematical convenience, nor as an occasional grace of nature for which we must passively wait, but rather as a mental construct that we should actively create. The progression of probabilistic reasoning techniques from singly connected BNs, to the JT representation of multiply connected BNs, and to MSBNs/LJFs is just one example of such endeavor.

## References

1. A.H. Bond and L. Gasser. An analysis of problems and research in dai. In A.H. Bond and L. Gasser, editors, *Readings in Distributed Artificial Intelligence*, Morgan Kaufmann, 3–35, 1988.
2. A.H. Bond and L. Gasser, editors. *Readings in Distributed Artificial Intelligence*. Morgan Kaufmann, 1988.
3. L. Gasser and M.N. Huhns, editors. *Distributed Artificial Intelligence, Volume II*. Morgan Kaufmann, 1989.
4. F.V. Jensen, S.L. Lauritzen, and K.G. Olesen. Bayesian updating in causal probabilistic networks by local computations. *Comput. Stat. Quarterly*, (4):269–282, 1990.
5. F.V. Jensen, K.G. Olesen, and S.K. Andersen. An algebra of bayesian belief universes for knowledge-based systems. *Networks*, 20:637–659, 1990.
6. S.L. Lauritzen and D.J. Spiegelhalter. Local computation with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society, Series B*, (50):157–244, 1988.
7. V.R. Lesser and L.D. Erman. Distributed interpretation: a model and experiment. *IEEE Transactions on Computers*, C-29(12):1144–1163, 1980.
8. R.E. Neapolitan. *Probabilistic Reasoning in Expert Systems*. John Wiley, 1990.
9. J. Pearl. Fusion, propagation, and structuring in belief networks. *Artificial Intelligence*, (29):241–288, 1986.
10. J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.
11. Y. Xiang, B. Pant, A. Eisen, M. P. Beddoes, and D. Poole. Multiply sectioned bayesian networks for neuromuscular diagnosis. *Artificial Intelligence in Medicine*, 5:293–314, 1993.
12. Y. Xiang, D. Poole, and M. P. Beddoes. Multiply sectioned bayesian networks and junction forests for large knowledge based systems. *Computational Intelligence*, 9(2):171–220, 1993.

This article was processed using the L<sup>A</sup>T<sub>E</sub>X macro package with LLNCS style