Agent-Encapsulated Bayesian Networks and the Rumor Problem

(Extended Abstract)

Scott Langevin
Department of Computer Science and Engineering
University of South Carolina
langevin@cec.sc.edu

Marco Valtorta
Department of Computer Science and Engineering
University of South Carolina
mgv@cec.sc.edu

Mark Bloemeke
LogicBlox, Inc.
Atlanta, GA
mark.bloemeke@logicblox.com

ABSTRACT

We present a multiagent organization for data interpretation and fusion in which each agent uses an encapsulated Bayesian network for knowledge representation, and agents communicate by exchanging beliefs (marginal posterior probabilities) on shared variables. We call this organization an Agent-Encapsulated Bayesian Network (AEBN) system. Communication of probabilities among agents leads to rumors, i.e. potential double counting of information. We propose a new and correct method to compensate for rumors in AEBN systems by passing extended messages that contain joint probabilities.

Categories and Subject Descriptors

I.2.3 [**Artificial Intelligence**]: Deduction and Theorem Proving—*Uncertainty, "fuzzy," and probabilistic reasoning*; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Coherence and coordination, Intelligent agents, Multiagent systems*

General Terms

Algorithms, Design, Theory

Keywords

Communication protocols, Distributed problem solving, Knowledge representation, Reasoning (single and multi-agent)

1. AGENT ENCAPSULATED BN SYSTEMS

An Agent Encapsulated Bayesian Network (AEBN) [1] is an agent that utilizes a Bayesian network for its internal representation of the world. How the agent utilizes this representation in decision support or goal based planning is unimportant so long as the world view is updated based only on local observations and observations received in the form of probabilistic messages from communicating agents.

In an AEBN system the agents communicate through the transmission of probability distributions on shared variables. The topology of the communication in the multiagent system forms a DAG structure. The Bayesian network of each agent can be divided into three distinct sets of variables: I, those about which other agents have better knowledge; L, those that are used only within the agent; and O, those that this agent has the best knowledge of and that other agents may want. This effectively produces two classes of variables in the agent: its local variables, L, and its shared variables, I and O.

The mechanism for integrating the view of the other agents on a shared variable is to simply replace the agent's current belief in the variable with that of the communicating agent. For this reason, all communication in the AEBN system occurs through the passing of messages that essentially contain the "correct" views on some shared variables. When an agent receives one of these messages, it modifies its internal model so that its local distribution either becomes consistent with the other agents' view or becomes inconsistent by entering a zero probability configuration.

To update an agent's distribution $P(V)$ with new evidence $Q(E_1, E_2, ..., E_n)$ for some set of variables $\{E_1, E_2, ..., E_n\} = I$ one calculates the joint probability $P(V)$, dividing by the marginal probability $P(I)$, and multiplying it by the new distribution of $\{E_1, E_2, ..., E_n\}$. This corresponds to the application of Jeffrey's rule,

$$Q(I) = Q(E_1) \cdot Q(E_2) \cdot ... \cdot Q(E_n), \tag{1}$$

thus obtaining:

$$Q(V) = P(V \backslash I|I) \cdot Q(I) = \frac{P(V)}{P(I)} \cdot Q(I). \tag{2}$$

In the case in which the input variables are not independent in the receiving agent, Equation 1 does not hold. (See [6, Section 5] for a detailed discussion on this point.) Lemma 1 in [6] allows the replacement of Equation 2 by:

$$Q^*(V) = P(V \backslash I|I) \cdot Q(I) = \frac{P(V)}{P(I)} \cdot Q_I^*(I), \tag{3}$$

where Q_I^* is the I_1-projection of probability distribution P on the set of all distributions defined on I and having $Q(E_i), i = 1, ..., n$, as their marginals. In practice, $P(V)$ could be updated to $Q^*(V)$ using the *big clique algorithm* of [6, 2], *lazy big clique algorithm* of [3], or the *wrapper methods* of [5].

Cite as: Agent-Encapsulated Bayesian Networks and the Rumor Problem (Extended Abstract), Scott Langevin, Marco Valtorta and Mark Bloemeke, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. XXX-XXX.

Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Thus a mechanism similar to that already used for updating probabilities in a Bayesian network adjusts the world view of the agent, $P(V)$, into a conditional probability table $P(O|I)$. Note that this table is calculated using the local observations of the agent: $P(O|I) = \sum_L P(O, I, L)/P(I)$. It then combines that table with the external view of the inputs, $Q(I)$, to allow the calculation of the new values for the output variables $Q(O)$.

Given this view of the purpose of each agent in the overall system, an agent system may be considered an expansion of the Bayesian network formalism to a DAG where the distribution of the variables of one agent is obtained by conditioning on its input variables. This is not strictly the case for two reasons. First, when input variables are not independent in the receiving agent, then the calibration equation 2 must be replaced by the formally identical, but substantially and computationally more complex equation 3.

Second, the oracular assumption imposes the additional constraint that, in the agent system, unlike a Bayesian network, all parents are not affected by their descendants. More precisely, the only variables that may affect the variables in an agent are (1) those in the agent itself and (2) those in a preceding agent. In order to provide a formal definition of "preceding agent," we define a *communication graph*, where the nodes represent agents and directed edges represent the flow of messages labeled with the shared variables that are being communicated. Let A_i and A_j be two distinct agents, let V_i, V_j be the sets of variables in agent A_i and A_j, respectively, and let $W_i \subseteq V_i$, $W_j \subseteq V_j$. Then if there is no directed path in the communication graph from A_j to A_i, any changes (whether by observation or by intervention) in the state of the variables in W_j do not affect the state of the variables in W_i. This is a very strong condition on the distribution of the variables in different agents of the agent system. This is *not* a symmetric relation, and therefore cannot be represented by any independence relation, since every independence relation is symmetric. In an AEBN, there is no possibility for a variable in an agent to be affected by a descendent agent.

2. RUMOR PROBLEM

Redundant influences arise in a communication graph whenever the combination of messages received by an agent causes the belief in some shared variable to be over included. Consider as an example a four-agent system (figure 1(a)), where a supervisor agent fuses reports from two observer agents, each of which reports information from a single sensor agent. We can see the problem centers on the fact that the supervisor agents view of the world, held in its Bayesian network, is doubly influenced by the initial sensor reading (S), through the reports (L1 and L2).

It is sufficient to identify all redundant influences by examining all pairwise node disjoint paths in the communication graph. Once they have been identified, a new graph, known as the redundancy graph, is constructed. The redundancy graph has the same nodes and edges as the communication graph, but its edge labels are expanded by adding the names of the variables that have redundant influence flowing through the edge. Figure 1(b) shows the redundancy graph for the ROSE example. We propose a method of compensating for redundant influences where agent communication has been expanded to pass joint probabilities along the appropriately labeled links in the redundancy graph, without any change in the local Bayesian networks of each agent. A data structure called the Redundancy Filter Tree is used to remove redundant influences from incoming messages before an agent performs belief propagation in its local Bayesian network. Space limitations prevent us from providing the detailed algorithm and its proof of correctness, but see [4].

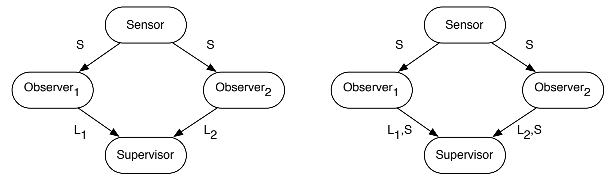

(a) Communication graph (b) Redundancy graph

Figure 1: Redundantly Observed Sensor Example.

3. CONCLUSIONS

The elimination of rumors in probabilistic agent communication is a difficult, longstanding problem that limits the applicability of graphical probabilistic models for knowledge representation in multiagent systems. Xiang [7] showed that, under a number of postulated assumptions, a correct probabilistic solution requires the topology of the communicating agents to be a tree. We replaced some assumptions with an oracular assumption, which states that the probability distribution of a variable published by an agent cannot be changed by the agents that subscribe to it. Under this new assumption, it is possible to compensate for rumors.

4. REFERENCES

[1] M. Bloemeke. *Agent Encapsulated Bayesian Networks.* PhD thesis, Department of Computer Science, University of South Carolina, 1998.

[2] Y.-G. Kim, M. Valtorta, and J. Vomlel. A prototypical system for soft evidential update. *Applied Intelligence*, 21(1), July–August 2004.

[3] S. Langevin and M. Valtorta. Performance evaluation of algorithms for soft evidential update in Bayesian networks: First results. In *Proceedings of the Second International Conference on Scalable Uncertainty Management (SUM-08)*, pages 284–297, Naples, Italy, October 2008.

[4] S. Langevin, M. Valtorta, and M. Bloemeke. Agent-encapsulated bayesian networks and the rumor problem, Technical Report CSCE TR-2010-001. Technical report, Department of Computer Science and Engineering, University of South Carolina, February 2010.

[5] R. Pan, Y. Peng, and Z. Ding. Belief update in Bayesian networks using uncertain evidence. In *ICTAI*, pages 441–444. IEEE Computer Society, 2006.

[6] M. Valtorta, Y.-G. Kim, and J. Vomlel. Soft evidential update for probabilistic multiagent systems. *International Journal of Approximate Reasoning*, 29(1):71–106, January 2002.

[7] Y. Xiang. *Probabilistic Reasoning in Multiagent Systems: A Graphical Models Approach.* Cambridge University Press, Cambridge, 2002.