Using Bayesian Network to aid Negotiations among Agents

Bikramjit Banerjee, Sandip Debnath, & Sandip Sen
Mathematical & Computer Sciences Department
University of Tulsa
600 South College Avenue
Tulsa, OK 74104
{bikram,debnath,sandip}@euler.mcs.utulsa.edu

Abstract

Relationship between agents in a society can be represented as a network structure where the nodes and connections represent the causal links among different actions. Such a network of causality and influences can be represented by a Bayesian network where the topology of the network together with the conditional and prior probabilities represent an agent's view of the influence of different factors on outcomes of agent interactions. We propose the use of Bayesian Networks to aid an agent in its negotiation with other agents. We define and use the concept of a *negotiation context* in which negotiation takes place between two agents. We provide a decision mechanism by which an agent can take actions to create a favorable negotiation context in addition to choosing a negotiation offer that is likely to be accepted by the other agent.

Introduction

Self-interested agents in Multiagent systems (MASs) may at times find themselves in a situation where local goals conflict with that of other agents. At other times the agents may have to exchange resources in a mutually profitable manner. The term negotiation refers to a broad class of techniques which can be used by self-interested agents to arrive at coordinated action sequences. In this particular paper, we will consider the scenario of bargaining between agents where one agent is trying to sell another agent a particular good. We posit that in addition to the selling price offered, the seller has a number of direct and indirect influences on the decision process by which the buyer decides whether or not to buy the good at the offered price. We define the *negotiation context* as the variable assignments to all factors that influence the decision of a party to negotiation.

We propose to to use a Bayesian network model to represent the influences of different factors on agent decisions. An agent's knowledge of such causal factors and their relative importance is captured in the topology of the network as well as the prior and conditional probability assignments. Initial, approximate knowledge of an agent can be further refined based on actual negotiation experiences. If values of all the factors are known, then the actual decision taken by another agent given these factors can be used to update the conditional probabilities at the outcome nodes. If some of the factors values are not known, the decision taken and the values of the known factors can be used to update either the conditional probabilities at the outcome nodes or the prior probabilities of the unknown factors. In this paper we focus on the decision mechanism that allows a modeling agent to use its knowledge represented as a Bayesian network to choose actions that set the negotiation context that will maximize the chance of its offer to be accepted by the other agent involved in the negotiation.

Model-based approaches have recently received increased attention in MAS research. Several probabilistic mechanisms have been developed to model agents (Gmytrasiewicz & Durfee 1995; Zeng & Sycara 1997). Some of these models have also been effectively used to explore opponents' strategies (Carmel & Markovitch 1998). Our work complements these efforts by using a rich representation scheme that allows the modeler to succinctly capture direct and indirect influences of various factors on the behavior of other agents. This is particularly appropriate in the context of negotiation and bargaining between agents which is rarely based on a single dimension.

Let us consider an example where an agent A needs to sell some goods to another agent B (e.g a second hand car). Now A and B will negotiate the price of the car. The financial capacity of agent B influences its **reservation price**, i.e., the upper limit of payment that B is unwilling to exceed. The agent A may find it useful to ask agent C to recommend the car to agent B if B is known to value C's recommendation. C's recommendation, in turn, can depend on other beliefs on part of C. Whether C will recommend A's car depends on C's faith in A, the fairness of A's offer and C's own perception of the condition of the car. In (Zeng & Sycara 1997) bargainers had beliefs only regarding each other's reservation prices, which they tried to model, in order to settle to a point within the zone of agreement. This assumes the existence of such an overlapping zone. However an intermediate agent (such as C here) may persuade one of the bargainers (B here) to upgrade the perceived utility of the car and as a result increase its reservation price even when a zone of agreement did not exist initially[1]. We consider the viewpoint of only the seller, A, who models both B and C, and updates its beliefs regarding both, in order to decide whether mediation by C is conducive to its purpose, as also whether the price-offer to B should be modified or not.

Bayesian Network

A Bayesian network is a graphical model that encodes relationships among variables of interest. A Bayesian network (Jensen 1996) consists of a set of *variables* and a set of *directed edges* between variables. Each variable has a finite set of mutually exclusive states. The variables together with the directed edges form a *directed acyclic graph*. Each variable A with parents $B_1,....,B_n$ a has a conditional probability table $P(A|B_1,....B_n)$ associated with it. There are four major reasons why we chose Bayesian networks to represent the belief structure:

- Bayesian networks can readily handle incomplete data sets. This is because Bayesian networks offer a way to encode the correlations among the input variables.
- Bayesian networks allow one to learn about causal relationships. This is useful to gain understanding about a problem domain. In addition it allows to make predictions in the presence of interventions.
- Bayesian networks in conjunction with Bayesian statistical techniques facilitate the combination of domain knowledge and data.
- Bayesian networks in conjunction with Bayesian methods offers an efficient and principled approach to avoiding over fitting of data.
- Bayesian networks offers a method of updating the belief or in other words the probability of occurrence of the particular event for the given causes.

Negotiation and belief update

Negotiations among self-interested agents in a multiagent environment may reveal their desired strategies to achieve individual goals. In most domains, strategies of agents are strongly coupled in the sense that an agent can influence actions of other agents, i.e. an agent A can create certain situations that can encourage another agent B to act in a way that reveals its beliefs and perceptions. Such situations can be reached by negotiations. Now, in an open environment, an effective approach for a self-interested agent is to be aware of the other agents' beliefs and the information about the importance that other agents attach to their beliefs, so that it can plan its actions. Actions taken from certain subsets of its action set in given situations, or a particular ordering of the same action set may reveal the true nature of the agents more effectively. And this true nature can be assessed by updating their beliefs about other agents. The problem in hand is to plan the actions of the modeling agent A in such a manner that either A achieves its goal or, in case of failure in a step of negotiation, learns more about others' beliefs to plan an effective course for future negotiations.

Our Bayesian Network based negotiation framework provides both an action selection mechanism that increases likelihood of success in negotiation as well as allowing for updating beliefs about relationships between several factors and agent decisions based on the decisions taken by other agent in a given negotiation context. . In order to develop such a model we assume that the other agent B is not changing its policy during the negotiations, and that the modeling agent has a reasonably good estimate of the factors or events in the environment that affect the decisions taken by the other agent. Though we restrict our discussion to negotiation between two agents in this paper, our framework can be used in situations where more than two agents are involved in a negotiation.

Example

We use an example of selling a second hand car in order to explain the procedure for action selection, where agent A quotes a price for its car and negotiation continues until B accepts A's offer. We assume B has a genuine desire to buy a car very similar to A's, and that price is the primary issue being negotiated. So A offers a price depending on its prior belief of B's capacity of payment (reservation price), and its evaluation of A's car. A may call in a third agent C, and based on the belief of its rapport with C and C's evaluation of its car, A may ask C to recommend its car to B. In a real-life negotiation situation, B may not accept A's initial offer and may counter-offer a different price. Given B's rejection of A's offer, based on its current beliefs, A will update its beliefs about B and C. The next offer may be selected based on such updated beliefs and using the same offer-choice mechanism as used before. In the process, A may also add new nodes (and connections) and/or delete some existing nodes (and connections), e.g he may choose to offer C other services to obtain a more favorable recommendation; or if A feels that C is going to give a poor impression of his car anyway, he may choose to eliminate C from the scenario and look for a better reference. So the network of influences is a dynamic one and is influenced by the past history of negotiation. In each offer cycle, A works with the recreated and/or updated network.

In this paper, we concentrate solely on the decision mechanism that allows A to set the negotiation context and the selling price based on its knowledge of the factors that influence B and C's decisions. So in our particular example the (seller) agent A is left with taking any of these two actions:

[1] We assume that the reservation price for an agent is not the upper limit of its payment-capacity, but the maximum that it is willing to pay in the deal.

- A sets the price of its car. The price range may be divided into various parts, the lower limit being A's reservation price. Initially A chooses one of these subranges as a result of its initial computations on its a priori beliefs.
- A asks the agent C to recommend its car to the agent B.

This is a simplified situation where the negotiation context is that of asking C to recommend (this may not be necessary if A believes for example that C will recommend the car anyway). In reality the negotiation context will be much richer including perhaps obtaining blue book value of the car, putting together a folder with the maintenance history of the car, may be changing the car tires or some other cosmetic improvement to the car to enhance its saleability, etc. Interestingly enough, doing all of this is not perhaps worth it and will be dictated by whether any such action is likely to significantly improve B's reservation price for the car.

The agent C's recommendation will depend on its relationship with and faith in A and its own evaluation of A's car. So according to A, C is influenced by the following factors :

- The agent C's faith in A.
- The agent C's own evaluation about the car.
- The agent A's request to recommend his car.

Agent B's ultimate decision depends on its payment-capacity and recommendation by C as well as the importance of C's recommendation to B.

- The agent B's payment-capacity.
- The price offered by A.
- The car's evaluation done by the agent B itself.
- The recommendation of the agent C.
- The importance of C's recommendation to B.

The above beliefs and actions may not be comprehensive, but serve the purpose of illustration. The Bayesian Network, based on the above setup, is shown in figure 1. The agent A starts at the sink node (labeled B) and chooses the desired value of the node (i.e B accepts the offer). The combination of values of the nodes A_2 (directly controllable by A, as it is one of its actions) and C_3 (indirectly controllable by A, as it has one of the actions of A, viz. A_1 as one of its ancestors) that maximizes the probability of occurrence of $B = accept$ for all possible value combinations of B_1, B_2 and B_3 (all uncontrolled nodes) are calculated. But since A has only indirect influence on $C3$, i.e., C's recommendation, it cannot guarantee any particular recommendation of C. It can, however, calculate the maximum probability of $C3$ taking any value given its own choice of asking or not asking for C's recommendation. This means that to calculate the probability distribution of B's choices, we have to recursively calculate the probability distribution of C's recommendation options. Such recursive calculations are used to ultimately choose values for

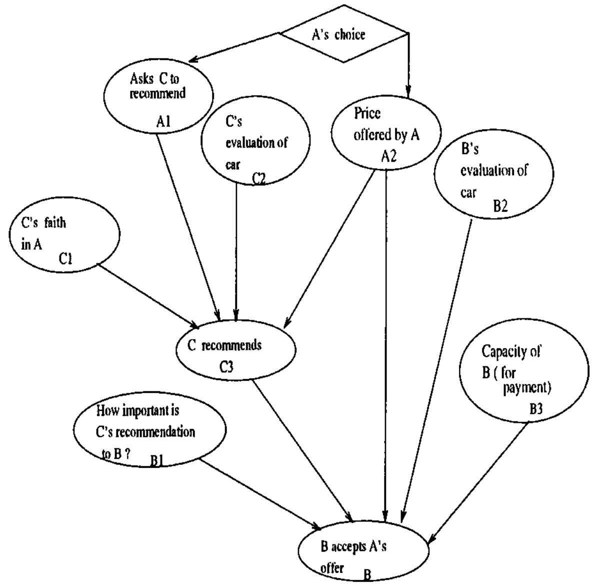

Figure 1: A Bayesian network for negotiation

nodes $A1$ (whether to ask C for recommendation) and $A2$ (what price to offer to B).

The Algorithm

Now, we present the algorithm that generates the negotiation context assignment given the target node decision choice. It is to be called by a top-level procedure that supplies the node B and its desired value (i.e. B accepts A's offer) as arguments.

This algorithm is a recursive one. Nodes in the network can be categorized, relative to the current target-node (i.e. the sink of the portion of the network currently under consideration) into three classes - D (Directly controllable), I (Indirectly controllable) and U (Uncontrollable) node. These classes were exemplified in the previous section. With these notations, the algorithm is presented in Figure 2.

The algorithm sets the negotiation context to maximize the probability of the target node taking on the target-value, and returns the corresponding probability. This algorithm does not work well for the case when a node both directly and indirectly controls the target node. For example, the node A2 directly controls B and also indirectly controls it through C3. We are currently working on a fix to this problem.

Remarks

There are two sets of recursive calls - the first set evaluates the maximum probability all possible assignments to the I nodes (I nodes are assumed to be independent, and we ignore the case where they may have common ancestors), and the next call proceeds upward along the network with the particular assignment to the I nodes that maximizes the desired value of the target node. In

```
Procedure Action-choice(target-node, target-value)
/* Returns maximized Pr(target-node = target-value)
Side effect is to assign values to action-nodes
for this maximization */
{
  D ← ∅; I ← ∅; U ← ∅
  for each node n such that (n, target-node)
  is a directed edge in the network
  {
    If n is directly controlled node then
      D ← D ∪ n
    else if ∃ a directed path from
      any action node to n then
      I ← I ∪ n
    else
      U ← U ∪ n
  }
  If (D ≠ ∅) or (I ≠ ∅) then
  /* Let the nodes in D be {D_1, D_2, ..., D_m}
  and all possible combinations of their values
  be {d_1, d_2, ..., d_{m_1}}
  Let the nodes in I be {I_1, I_2, ..., I_n}
  and all possible combinations of their values
  be {s_1, s_2, ..., s_{n_1}}
  Let the nodes in U be {U_1, U_2, ..., U_p}
  and all possible combinations of their values be
  {u_1, u_2, ..., u_{p_1}} */
  {
    max ← 0;
    for i ← 1 to m_1
       for j ← 1 to n_1
          {
            sum ← 0;
            for k ← 1 to p_1
               sum ← sum +
               Pr(target-node =
               target-value | d_i s_j u_k)
               * Pr(u_k) *
               ∏_{l=1}^{n} Action-choice
               (I_l ,value of I_l as in s_j);
            if (sum > max) then
            {
               max ← sum;
               a ← i;
               b ← j;
            }
          }
       for each v ∈ D, set its value as in d_a
       for each u ∈ I
       {
          Action-choice(u, value of u in s_b);
       }
    return max;
  }
}
```

Figure 2: The Bayesian Network based algorithm for setting negotiation context.

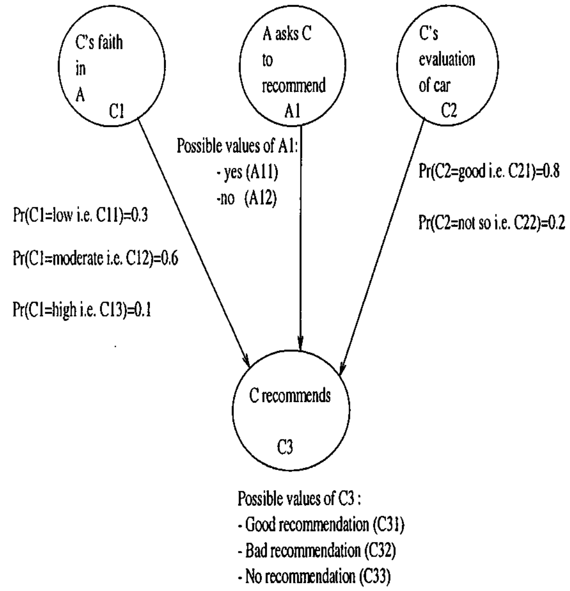

Figure 3: A portion of the Bayes network

any set of recursive calls, each of the I nodes is made the new target node.

We show the calculations for a small portion of the network for simplicity. Consider the portion shown in figure 3. Clearly C_1 and C_2 are uncontrollable nodes and hence belong to the group U while A_1 is directly controllable and belongs to group D. There is no I node in this instance. Now we compute $Pr(C_{31} \mid A_{11})$ assuming C_{31} is the value for C_3 returned by the algorithm with B as the target node (this means that A would desire a value of C_{31} for the node C_3 as that maximizes the probability of the probability of B accepting A's offer). The corresponding probability is given by

$$Pr(C_{31} \mid A_{11}) = \sum_{i=1}^{3} \sum_{j=1}^{2} Pr(C_{1i} A_{11} C_{2j}) * Pr(C_{31} \mid C_{1i} A_{11} C_{2j})$$

where

$$Pr(C_{1i} A_{11} C_{2j}) = Pr(C_{1i}) * Pr(A_{11}) * Pr(C_{2j})$$

assuming that these are independent belief/action nodes. Similarly we can compute $Pr(C_{31} \mid A_{12})$. With sample probability assignments as shown in figure 4 these values turn out to be

$$Pr(C_{31} \mid A_{11}) = 0.50111, \;\; Pr(C_{31} \mid A_{12}) = 0.04264.$$

So A should choose action A_{11}, i.e., it should ask C to recommend its car to B. In all these computations we have ignored the directed edge (A_2, C_3) since for calculation with B as the target node, A_2 would affect B both directly and indirectly (through C_3) and our algorithm does not adequately handle this scenario as we have mentioned before.

With the actions decided, A performs them and notes the result. If B still bargains, A updates the belief-node

C1 A1 C2	P(C31 | C1 A1 C2)	P(C32 | C1 A1 C2)	P(C33 | C1 A1 C2)
C11 A11 C21	0.55	0.4	0.05
C11 A11 C22	0.001	0.95	0.049
C11 A12 C21	0.01	0.009	0.981
C11 A12 C22	0.0001	0.0009	0.999
C12 A11 C21	0.75	0.21	0.04
C12 A11 C22	0.009	0.9	0.091
C12 A12 C21	0.08	0.001	0.919
C12 A12 C22	0.00001	0.00009	0.9999
C13 A11 C21	0.98	0.019	0.001
C13 A11 C22	0.01	0.45	0.54
C13 A12 C21	0.099	0.0001	0.9009
C13 A12 C22	0.000001	0.000009	0.99999

Figure 4: A sample conditional probability assignment for node C_3

values using the standard posterior probability calculation method in a Bayesian Network, and proceeds with the algorithm all over again. This continues until B accepts A's offer or gives up.

Future work

We intend to perform experiments with dynamic networks. In particular, we are interested in updating the topology of the network based on experience. Based on the result of each iteration of the negotiation, the modeling agent can delete or insert nodes and connections as well as alter the probabilities in the network. Additionally, new actions may be considered or old actions may be eliminated before embarking on to the next round of negotiation. In this paper we have considered modeling by only one party to the negotiation. We plan to study the convergence of the system when both the buyer and the seller model each other using the same approach. We also plan to modifying our algorithm so that nodes with both direct direct and indirect influences on the decision node can be adequately dealt with.

Very little work currently exists on explicitly choosing actions that aid in the model building process. We plan to investigate mechanisms that will allow the modeling agent to choose actions to persuade other agents to make decisions that will reveal maximal information about the relationships between their decisions and the influencing factors.

References

Carmel, D., and Markovitch, S. 1998. How to explore your opponent's stratey (almost) optimally. In *Proceedings of the Third International Conference on Multiagent Systems*, 64–71. Los Alamitos, CA: IEEE Computer Society.

Gmytrasiewicz, P. J., and Durfee, E. H. 1995. A rigorous, operational formalization of recursive modeling. In *First International Conference on Multiagent Systems*, 125–132. Menlo Park, CA: AAAI Press/MIT Press.

Jensen, F. V. 1996. *An Introduction to Bayesian Networks*. New York, NY: Springer Verlag.

Zeng, D., and Sycara, K. 1997. Benefits of learning in negotiation. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, 36–41. Menlo Park, CA: AAAI Press/MIT Press.