Learning Whom to Trust: Using Graphical Models for Learning about Information Providers

Tracking Number 732

ABSTRACT

In many multi-agent systems, information is distributed among potential providers that vary in their capability to report useful information and in the extent to which their reports may be biased. This paper describes a series of graphical models for learning about information providers in such settings. It shows how graphical models can be used to (1) simultaneously learn the reporting strategies that agents use and learn their capabilities; (2) weigh the benefits of different combinations of information providers; and (3) calculate the expected error of selecting different combinations of information providers. These models can cope with agents that vary in their capabilities and strategies, and whose capabilities may change over time. An agent using one such graphical model was evaluated using the ART-testbed against the contestants submitted to the latest competition held at AAMAS 2008. This agent was able to outperform all of the contestants. Further experiments show that graphical models can accurately model synthetic agents that use complex strategies to decide how to report information, and determine how to combine these reports to minimize error.

1. INTRODUCTION

A series of advancements in communication technology have enabled the emergence of large, distributed systems in which information is dispersed among many possible providers. The production and procurement of reliable information is a vital aspect of multi-agent interaction in these systems. The environments in which these systems exist contain uncertainty, and agents that interact in these systems differ in their capability to generate useful information about the world. Some agents may deliberately misreport information to gain an advantage. Learning which information providers to trust is an important problem in such environments.

For illustrative purposes, consider online marketplaces, such as eBay, where users often purchase from a particular seller only once. The reliability of sellers and the quality of the goods they are selling are not known until after an interaction has taken place. In such cases it is beneficial to obtain information about the qualities of these goods from reviewers before purchasing them. However, reviewers vary in their capability to evaluate different products, and may provide biased reports when, for example, a reviewer happens to be a competitor of the seller. It is crucial in such settings to learn to distinguish between the individual capabilities of information providers and the degree to which they report truthfully. Another example is the task of predicting the winner of an election by combining the results of different poles.

This paper addresses the problem of learning about information providers in settings in which agents need to make decisions, and the outcome of these decisions is directly affected by the quality of the information they can obtain. It proposes several graphical models for describing the strategies used by different information providers whose capabilities are unknown, and that vary in the extent to which they provide truthful information. Graphical models offer a unified approach both for reasoning about the strategies employed by multiple information providers and for deciding how to optimally combine their reports. Second, they can simultaneously learn about the strategies different providers may use to report information as well as their capabilities for generating useful information. Third, they can represent prior knowledge of the modeler about information providers and the world. Fourth, they explicitly represent conditional independence in the world, which facilitates the representation of information providers and inference in the model.

We begin by presenting a graphical model that can be used to learn the capabilities of truthful agents, and then show how the model can be extended to capture the strategies of agents who may not be reporting truthful information. Our model also allows for agents who use multiple strategies. For both truthful and strategic agents, we show empirically that our model can adapt to agents whose strategies vary over time. We then present a method for optimally choosing between multiple information providers of varying strategies and capabilities.

The task of choosing between information providers is formalized as a constrained optimization problem. We show how to convert this problem into a maximum a posteriori computation on our graphical model. Because our model can distinguish between the capabilities of provider agents and the extent to which they report truthfully, it can learn to make complex but "sensible" combinations, such as choosing two information providers, neither of which are truthful, but whose biases cancel each other out.

We evaluated these models by using them to make decisions in the Agent Reputation Testbed (ART) testbed [2], an open competition forum for agent designs that model trust and reputation. The model's performance is compared to the contestants in the ART competition held at AAMAS 08, as well as to synthetic agents that used various mixture models to decide whether to report truthful information. The experiments we ran simulated the exact conditions of the competition. The agents using the graphical models described in this paper scored highest among all of the contestants. In separate experiments we show that agents using our models outperform the finalists of this competition (in the long run) when com-

Cite as: Title, Author(s), *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. XXX-XXX.

peting against synthetic agents that insert bias into their reports. We also report results from an experiment that measured the ability of our models to accurately estimate agents that use complex strategies to decide how to report information.

1.1 The ART Testbed

In the ART testbed, agents need to appraise the value of paintings from different eras. Agents vary in their expertise for appraising paintings from different eras, and painting values are uniformly distributed between 1,000 to 100,000 units. Agents may purchase opinions from other agents about their assigned paintings. Agents cannot produce estimates for their own paintings through direct observation, and so may only produce estimates of paintings for opinions requested by others. The system provides a noisy estimation for the true value of paintings. The extent of this noise depends on agents' expertise about the era of the painting and the amount they paid to the system. The noise for an agent j that pays c_j units to the system to appraise a painting of value V of era e will be drawn from a normal distribution with zero mean and standard deviation $\left(S_{j,e} + \frac{\alpha}{c_j}\right) \cdot V$. Here, α is a known constant, and $S_{j,e}$ is the expertise of provider j for a painting of era e.

For each painting that is assigned to them, agents need to report a weighting of the opinions that were acquired from other agents. The final appraisal, computed by the system, is the weighted average of these opinions. At each round of the competition, agents receive a set of paintings for appraisals, referred to as their "market share". The value of these paintings and the reported opinions are revealed at the end of the round, after all agents have submitted their appraisals. The size of the market share for each agent is based on the accuracy of their appraisals, as well as the performance of other agents. Thus, agents receive more paintings, and thus the potential for more profit, for producing more accurate appraisals. The number of rounds in the competition averages at about 1,000 and is not known in advance. Agents incur a reward for each round that depends on the size of their market share.

2. MODELING AGENTS' STRATEGIES

We will focus on two classes of information providers. Agents that report truthful information do not distort the estimates they obtain from the system about painting values, while agents that report strategically may use any type or types of strategies to bias these estimates.

2.1 Truthful Information Providers

Figure 1 describes a Bayesian network [6] for the ART testbed which models information providers that report truthful information. This model is shown from the point of view of a single agent, referred to as an "appraiser", that is reasoning about several information providers. Each node in a Bayesian network represents a variable in the ART domain and contains a local probability distribution over its domain given its parents in the network. The topology of the network describes the conditional independencies that hold between variables in the domain. A Bayesian network defines a complete joint probability distribution over its nodes that can be decomposed as the product of the local probability distribution of each node given its parents.

As described in Section 1.1, the error of the estimate, for agents that request an estimate, from the ART system depends on their expertise for an era as well as the amount paid to generate the estimate. In the model we collapse the expertise of a provider agent with the amount it pays to generate an estimate into a single variable denoted σ that is termed the "capability" of the provider agent. The nodes σ_1 and σ_2 in the network represent the capabilities of

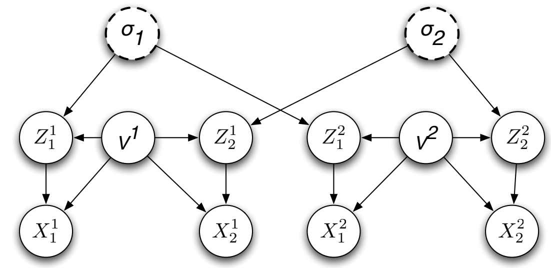

Figure 1: A Bayesian Network describing two information providers with capabilities σ_1 and σ_2. Dashed nodes represent variables whose distribution is unknown.

two information providers.[1] The values of σ_1 and σ_2 are not observed

An "interaction" is the process by which an appraiser agents needs to estimate the value of a single painting in its market share. The network of Figure 1 describes two such interactions. The nodes V^1 and V^2 represent two painting assigned for appraisal in these interactions. The nodes X_1^1, X_2^1 represent the estimates of the two provider agents, and nodes Z_1^1 and Z_2^1 are error terms representing the deviation between the estimate that both provider agents received from the system and the true value V^1. The nodes X_1^2, X_2^2, Z_1^2 and Z_2^2 represent equivalent nodes relating to the second interaction.

An edge between two nodes in a Bayesian network implies that the source node is a direct cause of the target node. For example, the parents of X_2^1 are the nodes Z_2^1 and V^1. This represents the fact that in the first interaction, the estimate X_2^1 of the second provider agent depends on the true value of the painting V^1 as well as its error Z_2^1. The parents of Z_2^1 are σ_2 and V^1, representing the fact that the error of the estimate of the second provider for the value of V^1 depends on its capability as well as on the true value of the painting. The parents of the nodes relating to the model of the first provider and the second interaction follow a similar pattern.

Each node in a Bayesian network is assigned a probability distribution given every value of its parents. The distribution over the node X_2^1 is deterministic and assigns probability 1 to each value that equals $V^1 + Z_2^1$. The probability of Z_2^1 is drawn from a normal distribution with mean 0 and standard deviation $\sigma_2 \cdot V^1$. The node σ_2 is thus a parameter specifying the distribution over Z_2^1. This node is marked by a dashed outline to represent the fact that the distribution over σ_2 is unknown to the appraiser. The value of V^1 is drawn from a uniform distribution, as specified by the rules of ART. The distributions over the nodes relating to the first provider and the second interaction are similar.

For expository convenience, this network shows two interactions with two information providers. In general, there will be nodes $\left((X_1^i, \ldots, X_K^i), V^i\right)$ for each interaction i. These represent the opinions of K provider agents about the value of painting V^i in interaction i. This network can be expanded to includes nodes for different types of providers. The size of the network will grow linearly in the number of provider types.

To model the strategy employed by the information providers, the appraiser agent must estimate the joint distribution over the capability of the provider agents given the history that is observed by

[1]For simplicity the network is presented for a single era. The era of each painting is known to all agents in ART, and in practice we learn a separate capability for each provider agent for every era.

the appraiser. This history consists of the values of the paintings for all interactions, as well as the estimates from the provider agents. This is denoted as

$$\left\{\left((X_1^1,\ldots,X_K^1),V^1\right),\ldots,\left((X_1^N,\ldots,X_K^N),V^N\right)\right\}$$

For the remainder of the paper we will use notation $V^{1,N}$ to refer to the set of nodes $(V^1,\ldots,V^N)$, and notation $X_{1,K}^{1,N}$ to refer to the set of nodes $\left((X_1^1,\ldots,X_1^N),\ldots,(X_K^N,\ldots,X_K^N)\right)$. We will use a similar notion for $Z_{1,K}^{1,N}$ and $\sigma_{1,K}$.

2.2 Learning Agents' Capabilities

The purpose of the network in Figure 1 is to infer the capabilities of the information providers given a history of interactions with these agents. To do so we need to compute $P(\sigma_{1,K} \mid X_{1,K}^{1,N}, V^{1,N})$, which is the posterior probability over the capability of the provider agents given the observations incurred by the appraiser. This factor depends on the values of the error nodes $Z_{1,K}^{1,N}$, which are unobserved. We sum over all possible values of the error nodes and get

$$P(\sigma_{1,K} \mid X_{1,K}^{1,N}, V^{1,N}) = \int_{Z_{1,K}^{1,N}} P(\sigma_{1,K} \mid Z_{1,K}^{1,N}, X_{1,K}^{1,N}, V^{1,N}) \cdot P(Z_{1,K}^{1,N} \mid X_{1,K}^{1,N}, V^{1,N}) \quad (1)$$

From the network we can infer first that the node Z_j^i, representing the error for provider j, in interaction i depends only on X_j^i and V^i (the value of the painting at interaction i, and the estimate of provider j). Therefore we can decompose the joint probability distribution $(Z_{1,K}^{1,N} \mid X_{1,K}^{1,N}, V^{1,N})$ as follows:

$$P(Z_{1,K}^{1,N} \mid X_{1,K}^{1,N}, V^{1,N}) = \prod_{i}^{N}\prod_{j}^{K} P(Z_j^i \mid X_j^i, V^i) \quad (2)$$

Because we have that $Z_j^i = X_j^i - V^i$, each of the resulting factors is easy to compute. We can also infer from the network that each node σ_j, representing the capability of provider j, depends only on $Z_j^{1,N}$, $X_j^{1,N}$, $V^{1,N}$ (the error term of j at interaction i, and the sequence of estimates for j for all interactions), so we can write that

$$P(\sigma_{1,K} \mid Z_{1,K}^{1,N}, X_{1,K}^{1,N}, V^{1,N}) = \prod_{j} P(\sigma_j \mid Z_j^{1,N}, X_j^{1,N}, V^{1,N}) \quad (3)$$

The term $P(\sigma_j \mid Z_j^{1,N}, X_j^{1,N}, V^{1,N})$ represents the posterior distribution over the capability of provider j given the prior interactions with that agent and the painting values. Using Bayes rule, we expand this term to get

$$P(\sigma_j \mid Z_j^{1,N}, X_j^{1,N}, V^{1,N}) \propto P(X_j^{1,N} \mid Z^{1,N}, V^{1,N}, \sigma_j) \cdot P(\sigma_j \mid Z_j^{1,N}, V^{1,N}) \quad (4)$$

From the network, we can infer that the nodes $X_j^{1,N}$, representing the combined estimates of provider j, are conditionally independent from each other given the nodes Z_j^i, and V^i. Therefore we can write that

$$P(X_j^{1,N} \mid Z^{1,N}, V^{1,N}, \sigma_j) = \prod_{i} P(X_j^i \mid Z_j^i, V^i) \quad (5)$$

Because of the deterministic relationship between X_j^i, Z_j^i and V_j^i these factors are easy to compute.

To compute the probability $P(\sigma_j \mid Z_j^{1,N}, V^{1,N})$ we first infer that the node σ_j is conditionally independent from the nodes $V^{1,N}$ given $Z_j^{1,N}$, so we can write

$$P(\sigma_j \mid Z_j^{1,N}, V^{1,N}) = P(\sigma_j \mid Z_j^{1,N}) \quad (6)$$

Computing the factor $P(\sigma_j \mid Z_j^{1,N})$ requires learning, because the distribution over the nodes $\sigma_{1,K}$ is unknown. To learn this distribution we use the fact that the Inverse Gamma (IG) distribution is a conjugate prior to the normal distribution with known mean and unknown variance. We set the distribution of each σ_j to follow an IG with parameters α_j and β_j. The value of these parameters represents the prior knowledge of the appraiser agent of the capability of provider j. Given that the appraiser agent engages in N rounds of play, the posterior distribution over $(\sigma_j)^2$ is also an inverse gamma with the following parameters

$$(\sigma_j)^2 \sim IG\left(\alpha_j + \frac{N}{2}, \beta_j + \frac{1}{2}\sum_{i=0}^{N}(Z_j^i - 0)^2\right) \quad (7)$$

2.3 Non-truthful Providers

In the real world, the values reported by provider agents may not be truthful. As an example from ART, provider agents may attempt to deceive an appraiser and report the wrong estimate for paintings to attain a larger market share. In addition, sophisticated agents may report their true estimate some of the time in order to avoid detection. We capture such agents using a mixture of several possible reporting strategies: overestimating (or underestimating) the estimate obtained from the system by an unknown constant, reporting the true estimate, or reporting an estimate from a normal distribution whose mean and variance is unknown. The probability that the mixture assigns to each individual strategy is also unknown. A Bayesian network describing this model is shown in Figure 2, from the point of view of an appraiser agent. For expository convenience we only include the nodes relating to a single interaction and provider agent.[2] Dashed nodes represent variables whose distribution is unknown.

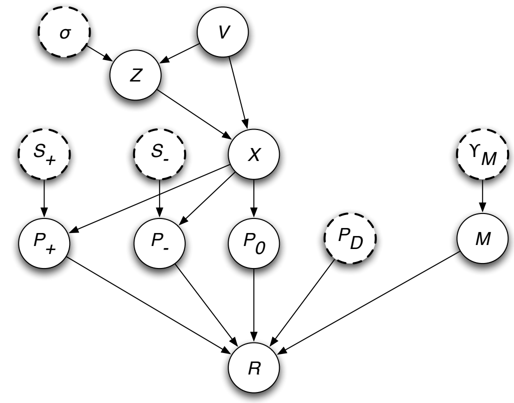

Figure 2: A Bayesian Network of a Provider Agent using a Mixture Strategy. Dashed nodes represent variables whose distribution is unknown.

The nodes σ, Z and X represent the capability of the provider agent, its error rate the estimate it receives for a painting V. The node P_0 represents a truthful report by the provider agent of the estimate X. This strategy is represented by a deterministic probabil-

[2]As in the case of truthful providers, the network can be expanded to account for multiple providers, and it will grow linearly in the number of provider types.

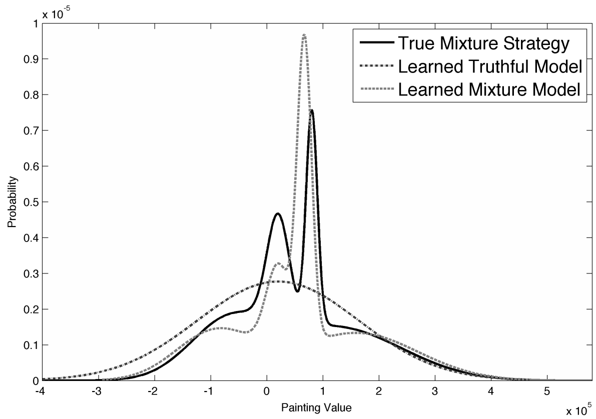

Figure 3: Learned distributions for one painting

ity distribution encoded in P_0 that assigns value 1 to X. The node P_+ represents a report by the provider agent that multiplies its estimate of X by a positive constant value represented by the node S_+. This strategy is specified by assigning a deterministic distribution over P_+ that assigns probability 1 to the value of $X_i \cdot S_+$. Similarly, the node P_- represents a report by provider agent that multiplies its estimate of X by a negative constant value represented by the node S_-. The distribution of P_- is assigned in a similar fashion to P_+. The values of both S_+ and S_- are drawn from normal distributions with a known variance and unknown mean. The node P_D represents a strategy that reports a random appraisal. It is drawn from a normal distribution whose variance and mean are both unknown.

The node R represents the report that the provider agent submits to the appraiser. The parents of R are the report nodes P_0, P_+, P_-, P_D and the node M. The domain of M ranges over the possible report nodes. The probability distribution of R is deterministic, and will assign probability 1 to the value of the report node P_0 when the value of M equals P_0 (and similarly for the other report nodes). In this way the node R serves as a multiplexer whose role is to choose among the values of the individual report nodes P_0, P_+, P_- and P_D depending on the value of M. The distribution of R is a multinomial, and the node γ_M specifies the parameters for this distribution. These parameters represent the beliefs of the appraiser agent over how the provider agent chooses between its individual reporting strategies. The values γ_M are unknown, and are drawn from a Dirichlet prior.

In this model, the distributions over the providers' capabilities, the extent to which they misreport their estimate and the mixture over the individual strategies are unknown. We define prior distributions over these parameters that represent background knowledge about the model. For example, the prior of S_- is a normal distribution that is bounded from above by 0, because it represents the mean of a negative constant.

To compute the posterior distribution over agents' capabilities in this model, we need to be able to learn the distributions over the nodes P_0, P_+, P_D and γ_M. The complexity of the model inhibits the use of an analytical solution, like in the network of Figure 1. Therefore, we use Gibbs sampling to estimate these parameters.

We demonstrate the use of the mixture model for learning the strategy of a provider agent that is employing the following mixture strategy:

- With probability 0.45 value of X_j is multiplied by 6.
- With probability 0.25 the value of X_j is multiplied by -3.5.

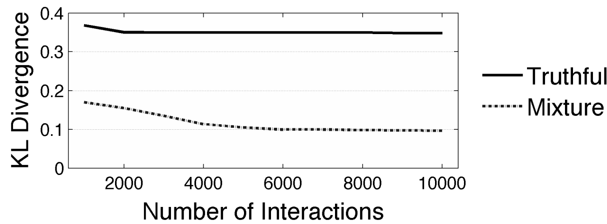

Figure 4: Model accuracy for varying number of interactions.

- With probability 0.15 the value of X_j (reported truthfully).
- With probability 0.15 a normal distribution with mean 80,000 and standard deviation 10,000.

The estimation error for this agent is distributed normally with a mean of zero and a standard deviation of 1.0.

Figure 3 compares the true mixture strategy of a provider agent to the strategy learned by two possible models. One in which the provider agent is assumed to be truthful, and one which uses the network of Figure 2 to model a mixture strategy. The figure is presented for a single painting of value of 20,000. Both of the models learned from 100,000 interactions with the provider agent.

The complexity of the network of Figure 2 means that computing the posterior distribution over the unknown parameters cannot be done in closed form, as in the case of truthful information providers. Instead, we employ Gibbs sampling to compute these quantities, through the WINBUGS system, a publicly available inference toolkit [4].[3]

The true mixture strategy for the provider agent is shown as the solid curve in the figure. The modes that can be seen in the distribution represent the different individual strategies that are used by the agent. The dotted curve shows the mixture strategy that is learned by the appraiser that models strategic agents. The dashed curve shows the truthful model that is learned by the network of Figure 1.

Figure 4 plots the average Kullback-Leibler (KL) divergence between the true and estimated values of paintings for varying number of interactions. Different sets of paintings were sampled for each point in the graph. As shown in this graph the divergence of the mixture model was significantly lower than that of the truthful model for all number of interactions. The divergence of both models decreases as the number of interactions increase, but the decrease is significantly more pronounced for the mixture model.

3. COMBINING INFORMATION FROM PROVIDERS

In this Section we show how to use the model of agents' strategies to make decisions in the ART test-bed. We first show how an appraiser agent can directly use the Bayesian network to weigh the potential contribution of the reports of different information provider agents. Then we show how the network can inform a process for optimally choosing between different providers.

3.1 Truthful Information Providers

[3]Note that considerable more interactions were used here than the average 1,000 in the ART competitions. In this section we do not mean to evaluate performance of the model in the test-bed, but to show that in expectation, the model is able to learn a good approximation to a provider agent employing a mixture strategy.

	True Weights	Learned Weights $N = 10,000$	Learned Weights $N = 1,000$
w_1	0.7024	0.6999 (sd=0.0097)	0.6953 (sd=0.022)
w_2	0.1756	0.1768 (sd=0.0053)	0.1765 (sd=0.022)
w_3	0.07805	0.07831 (sd=0.0057)	0.07712 (sd=0.015)
w_4	0.04390	0.04503 (sd=0.0020)	0.05089 (sd=0.0066)

Table 1: True versus learned weights for four truthful agents with capability values 0.25, 0.5, 0.75, 1.0, having the respective weights w_1, w_2, w_3, w_4.

We first consider information providers that report their true estimates. Let D denote a set of provider agents chosen by the appraiser for interaction i, and let W_d^i denote the weights for an individual provider $d \in D$ for interaction i. The appraiser agent's combined estimate of a painting at i, denoted A_D^i, is the weighted average of the opinions of the provider agents in D.

$$A_D^i = \sum_{d \in D} W_d^i \cdot X_d^i \tag{8}$$

The error of appraisal when combining the estimates of the providers in D, denoted E_D^i, is the difference between the true painting value and the combined estimate:

$$E_D^i = V^i - A_D^i \tag{9}$$

In the case where provider agents are truthful, we follow an approach used by Teacy *et al.* [10], and use the best linear unbiased estimate (BLUE) for the value of the weights, computed as

$$W_d^i = \sum_{d' \in D} \frac{1/(\sigma_d)^2}{1/(\sigma_{d'})^2} \tag{10}$$

Recall that σ_d is the capability of provider agent d, which is unobserved by the appraiser agent. Therefore, we use the expectation $E[(\sigma_d)^2]$ of the capability for each provider d as an estimate for $(\sigma_d)^2$. As stated in Section 2.2, the distribution over $(\sigma_d)^2$ is an inverse gamma distribution, and its expectation can be computed as:

$$E[(\sigma_d)^2] = \frac{2\beta_d + \sum_{i=0}^{N} (Z_d^i)^2}{2\alpha_d + N - 2} \tag{11}$$

To show this technique in practice, consider four provider agents with capabilities 0.25, 0.5, 0.75 and 1.0, respectively. We can compute the BLUE for the weights relating to these agents using Equation 10, and compare it to the BLUE using the learned capabilities of these agents through simulation. Results of the evaluation are shown in Table 1 for $N = 1,000$ and $N = 10,000$ interactions. Eight trials were run for each data point. On average, the learned weights had a percent error of 3.2% when $N = 10,000$, and an 11% error when $N = 1000$. As a baseline, we got a 199.1% percent error when assigning uniform weights. The learned weights all fell within one standard deviation of the optimal value except for w_4 when N equaled 1000.

3.2 Non-truthful Information Providers

For the case in which provider agents may report strategically there is no closed form solution for computing the weights. Solving Equation 10 for the case in which provider agents use a mixture strategy is a highly non-linear optimization problem. We cannot solve this problem analytically. Instead, We choose to approximate the weight values by envisioning what their ideal values would look like. To this end, we extend the Bayesian network of Figure 2 to explicitly represent the weights attributed to provider agents, and perform inference on the network to approximate their values from a sampled set of future interactions.

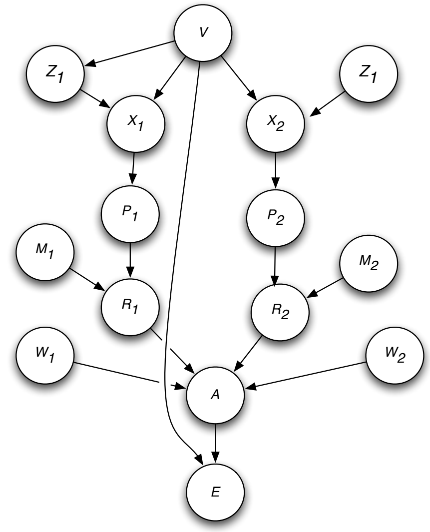

Figure 5: A Bayesian network for weighting the reports of information providers.

The node A represents the combined appraisal which is the weighted opinions of the information providers

$$A_D^i = \sum_{d \in D} W_d^i \cdot R_d^i \tag{12}$$

Figure 5 shows such a network. For ease of presentation, we have collapsed the set of nodes indicating the possible strategies used by the reporting agents to two nodes P and S. The nodes W_1 and W_2 represent the weights attributed by the appraiser to the opinions of these agents. The prior distribution over W_1 is uniform [0,1]. The distribution over W_2 is deterministic and assigns probability 1 to the value $1 - W_1$. In this way the value of the weights associated with two providers is made to sum to 1. Nodes representing the entire set of D information providers can be added, and the size of the network will grow linearly with the number of provider agents.

Weighing the potential contributions of the provider agents in D corresponds to solving a constrained optimization problem for finding the set of weights W_D^i that minimize the expected error:

$$\min_{W_D^i} E[E_D^i \mid W_D^i] = 0, \text{ such that } \sum_{d \in D} W_d^i = 1 \tag{13}$$

Ideally, we wish the value of E_D in the network to be 0. (In the network of Figure 5, this corresponds to setting the node E to equal 0). We therefore ask, what combination of weights is most likely to give an error of 0? this amounts to computing the weights that maximize $P(E_D = 0 \mid W_D)$. Because the prior distribution of the weights is uniform, this is the same as computing the weights that maximize $P(W_D \mid E_D = 0)$. The term $P(W_D \mid E_D = 0)$ is the likelihood over the weights W_D given the value of the paintings in the sample, and that the error of estimation when combining the opinions of the providers in D is zero. Computing the weights that maximize this likelihood correspond to performing a minimal a posteriori (MAP) computation in the graphical model. To compute $P(W_D \mid E_D = 0)$ in the network, we again use Gibbs sampling, but first set the values of the nodes S_+, S_-, P_D, and σ to equal the mean of their posterior distributions.

3.3 Using the Model to Choose Among Information Providers

We now show how it is possible to use the Bayesian networks of Figure 1 and Figure 2 to choose among different provider agents. To compute the expected benefit of the appraiser, we need to reason about the future market share difference for different combinations of provider agents. Reasoning about future market share is difficult, because it is affected by the combined appraisals of the other agents in the system, which are unknown. We approximate the benefit for an arbitrary interaction by setting its value to be the difference between the known expectation over V computed without information, the expected error from using a possible subset of providers, and the cost for obtaining that information, denoted opn_{cost}. For a set of D of provider agents, we define the expected benefit to the appraiser agent as

$$EU_D \quad = \quad E[V] - E_D - opn_{cost} \cdot |D| \qquad (14)$$

We use a soft-max function to allow the appraiser agent to explore different combinations of information providers. The likelihood of selecting a subset of agents is correlated with their expected benefit. The probability that a set D of provider agents will be chosen by the appraiser is defined as:

$$\frac{e^{EU_D/EU_{D^*}\cdot\tau}}{\sum_{D'\in 2^K} e^{EU'_D/EU_{D^*}\cdot\tau}} \qquad (15)$$

Here, 2^K denotes the set of all possible provider agent combinations, and EU_{D^*} denotes the largest expected benefit for all combinations of provider agents. The term τ varies the amount of exploration that is employed by the appraiser. As the number of interactions grows, τ increases and the appraiser will be less likely to explore. One drawback is that the computational complexity of computing Equation 15 is exponential in the number of possible provider agents. This did not present a problem in ART, because an appraiser agent can only ask for information from up to three information providers.

3.4 Demonstration of Model

We now demonstrate how to use this technique to learn how to combine the reports of three provider agents $\{a_1, a_2, a_3\}$, with the following parameters. A provider a_1 with low error of estimation ($Z_1 \sim N(0, 0.01)$), that consistently multiplies its estimate by a factor of 4. A provider a_2 with low error of estimation ($Z_1 \sim N(0, 0.01)$), that consistently multiplies its estimate by a factor of (-4). A provider a_1 with high error of estimation ($Z_1 \sim N(0, 4)$), that reports its estimate truthfully. We learned from 10,000 interactions with these providers. The following table lists the learned weights for each possible set of providers and combined error rate for estimating a held out test-set of paintings. The last line in the table is a baseline for the case in which no information was obtained from any provider.

Providers	w_1	w_2	w_3	Error
$\{a_1, a_2, a_3\}$	0.5133	0.4867	0.0	13350
$\{a_1, a_2\}$	0.5133	0.4867		**13350**
$\{a_1, a_3\}$	0.2245		0.7755	79470
$\{a_2, a_3\}$		0.2178	0.7822	79420
$\{a_1\}$	1.0			205500
$\{a_2\}$		1.0		209800
$\{a_3\}$			1.0	87440
$\emptyset$				51490

As these results show, each provider agent is useless when chosen in isolation. (The error obtained from any single provider was always greater than selecting none of the providers). However, the model was able to learn that combining information from both providers a_1 and a_2 results in the lowest amount of error. This is because the aggregate reports from both of these agents cancel the bias they introduce individually. The model was also able to learn that there is no benefit to adding the report from a_3 to a combination that includes the joint reports of a_2 and a_1.

4. EMPIRICAL WORK

In the following section we describe several sets of experiments were run using the ART-testbed described in Section 1.1. In all experiments, we compared the performance of an agent using this model to agents that were submitted to the last ART competition held in AAMAS 2008. We made no attempt to change the strategies employed by these agents.

4.1 Using our Model in the Art test bed

The following describes the process by which an appraiser makes a decision in ART. At each round of the competition our agent proceeded as follows.

1. Compute weights for different combinations of provider agents as described in Section 3.1 and Section 3.2.

2. Compute the expected utility for different combinations of provider agents as described in Section 3.3.

3. Probabilistically choose a subset of provider agents by using Equation 15.

4. Report the chosen agent subset and weightings to the system.

5. Observe the true values of the paintings and update model parameters.

4.2 Proof of Competitiveness

We compared a version of our graphical model agent to the eleven competitors entered in the AAMAS 2008 competition, employing precisely the same game parameters used in the competition. A facet of the 2008 competition was that agents' expertise for different eras changed over time. The extent of this change was varied by running three scenarios. The first scenario had low dynamics, in which each agent's expertise was changed by 0.05 for one of the eras. The second competition scenario had medium dynamics, in which each agent's expertise was changed by 0.1 for three of the eras. The third competition scenario had high dynamics, in which each agent's expertise was changed by 0.3 for three of the eras. The three competition scenarios were each run three times, with 90 rounds and 20 possible eras. The average market share for each agent at each round was 20 units, and the agents received a reward of 100 per assigned painting. The cost of obtaining an appraisal from a provider agent was 10 units, and each agent was allowed to ask for a maximum of three opinions per round. [4] We refer to the agent using the network of Figure 1 as the Graphical Model agent with Truthful Discounting (GMTD). Because provider agents' expertise about eras may change over time, we refined the update procedure over the inverse gamma distribution described in equation 7 to include a discount factor λ, which decreased the effects of

[4]We do not list the more minor game parameters, which were set to the same value as in the competition.

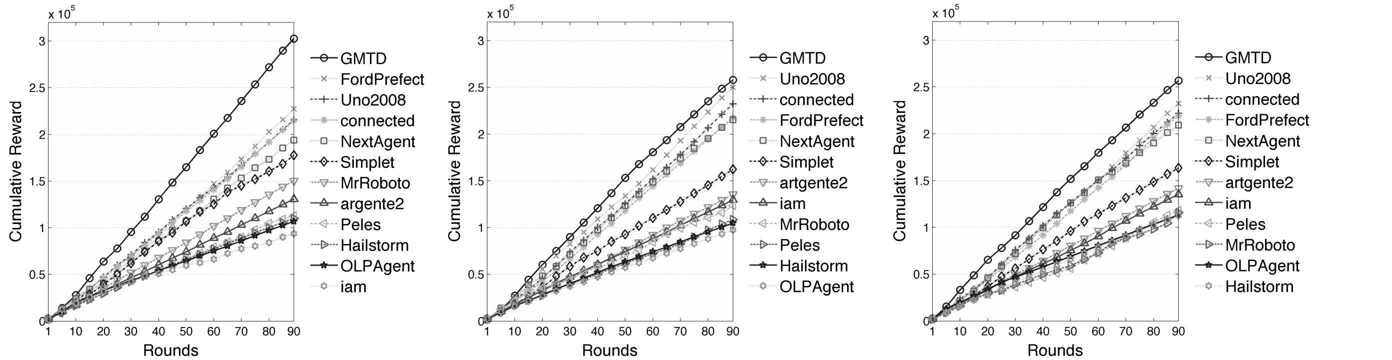

(a) (b) (c)

Figure 6: Comparison of 2008 AAMAS ART-Testbed competitors with our GMTD agent in the first competition scenario with (a) low dynamics, (b) medium dynamics and (c) high dynamics.

past rounds.

$$\sigma_j \sim IG\left(\alpha_j + \frac{\sum_{i=1}^{N}\lambda^i}{2}, \beta_j + \frac{1}{2}\sum_{i=1}^{N}(Z_j^i - 0)^2\lambda^{N-i}\right) \quad (16)$$

For the competition the GMTD agent employed a discount factor $\lambda = 0.96$ for all three scenarios.[5] The exploration factor τ, defined in equation 15 was set to 100 for all three scenarios. This had the effect of increasing the rate of exploration, to further adjust to agents' changing expertise over time. Figure 6 compares between the GMTD agent and the ART 2008 contestants for the three scenarios. As shown by the figure, for all of these scenarios, the GMTD agent received a larger cumulative reward than any of the contestants. [6]

4.3 Evaluation against Synthetic Agents

We will now compare the ability of an agent using the Bayesian network of Figure 2 to that of the top finalists in ART competition held at AAMAS 2008 for capturing synthetic agents of varying reporting strategies. Our agent, termed the Graphical Mixture Model Agent (GMMA), was designed to learn a mixture of four individual strategies, in a similar fashion to what is described in Section 2.3. We used the following synthetic provider agents in this experiment: agent a_P that reports estimates that are multiplied by a positive factor of 4 and had low variance of estimation error ($\sigma^2 = (0.06)^2$), agent a_N that reports estimates that are multiplied by a negative factor of -2 and had low variance of estimation error ($\sigma^2 = (0.06)^2$), agent a_0 that reported its true estimate but had high variance of estimation error ($\sigma^2 = (1.55)^2$).

We ran a series of experiments in which ART simulations positioned each of the competitions top three finalists (named "Uno2008", "FordPrefect", and "connected") as well as the GMMA agent against the three information providers a_P, a_N, a_0. Each of the three finalists, as well as the GMMA agent, was individually evaluated against a_P, a_N, a_0. We do not report the performance for these synthetic agents because their purpose was to compare between the abilities of the GMMA agent and the top finalists.

[5]This factor was determined empirically by using simulations using parameters of past ART competitions.

[6]We note that it is possible for agents in ART to ask for reputation information regarding the reliability of different agents in the system. The transaction logs of the competition show that the vast majority of contestants, our agents included, did not use this mechanism.

In this experiment, the GMMA agent used an exploration factor τ, from equation 15, that depended on the number of times it obtained information from the provider agent from which it had the most observation, denoted as t. Specifically, this factor was set to $\tau = 0.1^8 \cdot t^9$, where t was the maximum number of observations obtained for any opinion provider, divided by a constant of 4,000. In this way, the GMMA agent decreased its exploration rate based on the number of interactions with the most popular information provider.

We ran one scenario for this experiment in which the number of rounds was 5000, the number of eras set to 1, and where the expertise of agents did not change. All other parameters were set to those used in the AAMAS 2008 competition. Figure 7(a), Figure 7(b) and Figure 7(c) show performance, error and market share when comparing between the top contestants and the GMMA agent. Results show that after an acclamation period, the GMMA agent was able to outperform the previous finalists of the ART testbed. During the initial rounds the GMMA agent did not perform as well as the other agents. However, the GMMA agent was able to learn that combining the reports of the synthetic agents whose bias canceled each other would increase its reward, where each individual agent would have yielded a high variance and low benefit. The initial poor performance of the GMMA agent highlights the trade-off between the ability to model complex models, and the need for more interactions to be able to learn such complex models.

5. RELATED WORK

A variety of approaches towards investigating trust and reputation in multi-agent systems exist in the literature. An older version of the Uno2008, the 2008 AAMAS competition winner, by Munoz and Murillo [5] used a non-Bayesian approach based on combining several threshold factors, such as a lying and error rate, to determine which agents to query for opinion information.

The winner of the ART 2006 competition was an agent designed by Teacy *et al.* [9] that used a discrete distribution to model the capabilities of provider agents and determined whether an agent was untruthful using a Chi-squared test. In further work, Teacy *et al.* proposed a Bayesian approach that uses a similar technique to that proposed in Section 2.2 for modeling agents assumed to report truthfully. This work used a reward-based model for learning to combine between different provider agents. An agent using this model was able to beat the contestants of the ART 2007 competi-

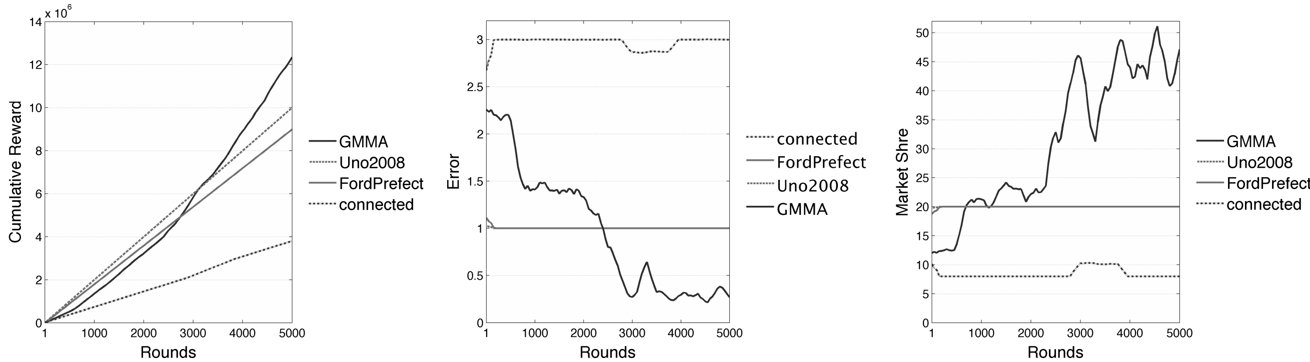

(a) (b) (c)

Figure 7: Performance of the GMMA agent as compared with the top three finalist from the AAMAS 2008 ART testbed competition for playing against opponents with positive, negative, and no bias. Results measure (a) cumulative reward, (b) error, and (c) market share.

tion. This agent was not officially entered in ART 2008.[7]. Other works [3, 8] have also used Bayesian modeling to infer the reliability of information from different provider agents.

None of these models distinguish between agents that are incapable of generating useful information and agents that intentionally do not report useful information. They also do not consider the possibility that the expertise of agents may change over time, as was the case in the 2008 competition. As we have shown in Section 3, it is crucial to be able to separate the two to optimally combine the reports of agents that use sophisticated strategies. Our work is also distinct in that we use the same graphical model for both learning about and combining information from different provider agents.

Lastly, model-free approaches such as Q-learning have been used in several works for learning to select between information providers [7, 1]. These approaches do not learn an explicit model of the strategies other agents use to report information, and do not capture reporting strategies that change over time.

6. CONCLUSION AND FUTURE WORK

This paper showed that graphical models offer a unified approach for simultaneously learning about information reporting strategies, agent capabilities, and how to combine information from multiple sources to minimize error. Our model was able to capture a mixture of agent strategies, as well as agents that changed their reported strategies. The graphical model outperformed all agents in the previous year's AAMAS 2008 ART testbed competition. A more sophisticated mixture model was shown to to learn complex provider strategies with varying bias, and was able to combine multiple biased opinions into a final appraisal with low error. In future work, wish to explore conditions in which multiple, possibly competing agents, share their learned information. Lastly, we intend to use this model to learn the information exchange strategies of humans in strategic environments.

7. REFERENCES

[1] K. K. Fullam and K. S. Barber. Dynamically learning sources of trust information: experience vs. reputation. In *AAMAS '07*, pages 1–8, New York, NY, USA, 2007. ACM.

[2] K. K. Fullam, T. B. Klos, G. Muller, J. Sabater, A. Schlosser, Z. Topol, K. S. Barber, J. S. Rosenschein, L. Vercouter, and M. Voss. A specification of the agent reputation and trust (art) testbed: experimentation and competition for trust in agent societies. In *AAMAS '05*, pages 512–518, New York, NY, USA, 2005. ACM.

[3] A. Jøsang and R. Ismail. The beta reputation system. In *In Proceedings of the 15th Bled Electronic Commerce Conference*, Bled, Slovenia, 2002.

[4] D. Lunn, A. Thomas, N. Best, and D. Spiegelhalter. Winbugs – a bayesian modelling framework: concepts, structure, and extensibility. *Statistics and Computing*, 10::325–337, (2000).

[5] V. Munoz and J. Murillo. Agent uno: Winner in the 2nd spanish art competition. *Inteligencia Artificia*, 12(39):19–27, 2008.

[6] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, 1988.

[7] K. Regan, R. Cohen, and P. Poupart. The advisor pomdp: A principled approach to trust through reputation in electronic markets. *Proceedings of Privacy, Security and Trust Conference*, November 2005.

[8] W. T. Teacy, J. Patel, N. R. Jennings, and M. Luck. Travos: Trust and reputation in the context of inaccurate information sources. *AAMAS 06*, 12(2):183–198, 2006.

[9] W. T. L. Teacy, G. Chalkiadakis, A. Rogers, and N. R. Jennings. Sequential decision making with untrustworthy service providers. In *AAMAS 08*, pages 755–762, Richland, SC, 2008. International Foundation for Autonomous Agents and Multiagent Systems.

[10] W. T. L. Teacy, T. D. Huynh, R. K. Dash, N. R. Jennings, M. Luck, and J. Patel. The art of iam: The winning strategy for the 2006 competition. *The 10th International Workshop on Trust in Agent Societies*, 2007 2007.

[7]The API for ART 2008 is not backwards compatible so we could not present a direct comparison between our model and the winners of the 2007 or 2006 competition.