

The ART of IAM: The Winning Strategy for the 2006 Competition

[ART TESTBED SPECIAL SESSION]

W. T. Luke Teacy, Trung Dong Huynh,
Rajdeep K. Dash, Nicholas R. Jennings
and Jigar Patel
School of Electronics and Computer Science
University of Southampton
Southampton, SO17 1BJ, United Kingdom
{wltl,tdh,rkd,nrj}@ecs.soton.ac.uk
jpatel@mckinsey.com

Michael Luck
Department of Computer Science
King's College
London, WC2R 2LS
United Kingdom
michael.luck@kcl.ac.uk

ABSTRACT

In many dynamic open systems, agents have to interact with one another to achieve their goals. Here, agents may be self-interested, and when trusted to perform an action for others, may betray that trust by not performing the actions as required. In addition, due to the size of such systems, agents will often interact with other agents with which they have little or no past experience. This situation has led to the development of a number of trust and reputation models, which aim to facilitate an agent's decision making in the face of uncertainty regarding the behaviour of its peers. However, these multifarious models employ a variety of different representations of trust between agents, and measure performance in many different ways. This has made it hard to adequately evaluate the relative properties of different models, raising the need for a common platform on which to compare competing mechanisms. To this end, the ART Testbed Competition has been proposed, in which agents using different trust models compete against each other to provide services in an open marketplace. In this paper, we present the winning strategy for this competition in 2006, provide an analysis of the factors that led to this success, and discuss lessons learnt from the competition about issues of trust in multiagent systems in general. Our strategy, IAM, is Intelligent (using statistical models for opponent modelling), Abstemious (spending its money parsimoniously based on its trust model) and Moral (providing fair and honest feedback to those that request it).

1. INTRODUCTION

Trust constitutes an important facet of multiagent systems research since it provides a form of distributed social control within highly dynamic and open systems whereby agents

form opinions about other agents based on their past interactions, as well as from reports of other agents [8]. As a result, a number of trust models and strategies have been proposed in order to deal with distinct aspects of the interactions between agents (e.g. to deal with lying agents [1], to model and learn the behaviour of other agents [10, 9] and to fuse information from disparate sources and models [5]). This, in turn, has rendered it hard to compare trust strategies since the underlying problem addressed by these strategies has been different. Therefore, in order to provide a common platform on which "researchers can compare their technologies against objective metrics" the Agent Reputation and Trust (ART) Testbed Competition has been proposed [3].

The ART Testbed simulates an environment consisting of service providers that compete in order to provide information services. There is a fixed total number of clients who are apportioned between the service providers according to the comparative quality of their service provision. Each of these information providers needs to spend money in order to gain information. Furthermore, they can improve their quality of service by requesting (against a payment) the other agents for information. However, it is not necessary that the requested agents will provide good information. In fact, as a result of the competition between the agents, it is quite likely that the agents will provide bad information. Thus, within this competition, trust and reputation become important metrics with which to measure the reliability of other agents. This is because trust and reputation measure the certainty with which each of the other agents provides good opinions and reputation information.

In this paper, we describe the *winning* strategy, **IAM** (our research group name), of the 2006 competition. The ART testbed competition received 17 entries and our entry proved to be highly successful, beating its closest competitor in the final by approximately 28%. The aim of this paper, therefore, is to describe the various facets of the strategy which contributed to the effectiveness of the IAM agent.

The paper is organised as follows. In Section 2 we analyse the decision problems faced by a generic agent within the competition. We then go on to describe the two main com-

ponents of the IAM agent, namely its trust model and its strategy. In Section 3, we describe how the IAM agent fuses information from the disparate sources within the competition for the purposes of its trust model. Section 4 then details the spending and earning strategy of the IAM agent according to the trust model. Finally, we discuss the strategy and the competition in Section 5 and conclude in Section 6.

2. DESIGNING A COMPETITION ENTRANT

In this section, we first provide a brief description of the ART testbed competition. We then analyse the general problem faced by an agent within the competition which gives rise to the broad outlines of the design of a generic agent.

2.1 Competition Overview

The ART competition consists of a number of art appraising agents that provide appraisals about the value of the art objects of their clients, who pay them to do so. A game, within this competition, consists of a predetermined (but unknown to the agents) number of iterations. At the beginning of each game, each agent is assigned a *privately known* expertise vector that determines the variance of its error in appraising a painting when spending a certain amount of money. The vector is over the ten eras from which the paintings can come. The higher the expertise of the agent, the lower the variance of its error in assessing the painting’s true value. The agent can then spend a certain amount of money gathering *opinions* from other agents about the paintings it has been tasked with appraising. Since the other agents can provide spurious opinions, the agent also needs to build a trust model to reflect its belief of how other agents will act when providing opinions. In order to aid it to build this trust model, the agent can also buy *reputation* information from other agents, but again this is open to dishonest behaviour on the part of the reputation provider. Having decided which set of agents it will buy the opinions from, the agent then provides the game simulator with *weights* that determine how the different opinions are fused in order to provide the final *appraisals* of the paintings. This then completes one iteration of the game. The accuracy of this appraisal as compared with that of the other competing agents then determines the client share of the next iteration i.e. the number of clients each agent obtains in the next round. Thus, in each round, the agent makes money from providing appraisal to its clients and spends money in ensuring that the appraisal of the art objects are as accurate as possible. The winner of the competition is the agent with the highest bank balance at the end of the game. A much more detailed description of the competition can be obtained in [3] or at <http://www.lips.utexas.edu/art-testbed/>.

2.2 The Agent Decision Model

We now describe the decision problem faced by a generic agent at a given round in a game, which can be represented as shown in Figure 1. Note that this model is part of our analysis of the problem in which we are identifying the main functional blocks that competition agents would need to a more or less sophisticated degree. Though they may not implement the model in this way, they still require to have all this functionality in whatever architecture they do have.

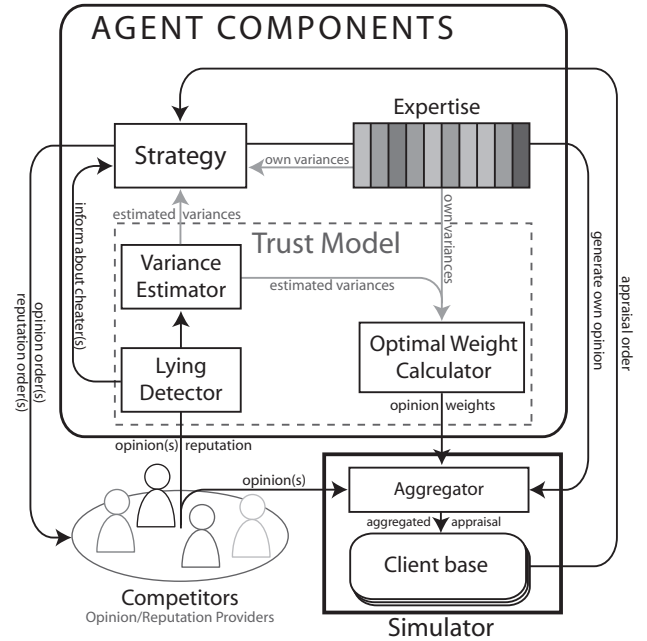


Figure 1: Agent design for the ART game

As such, it can be observed that in order to be a successful agent within the ART Testbed Competition, there are two main issues to consider. Firstly, we need to consider how to model the performance (i.e. accuracy) of each opinion provider in order to use their opinions accordingly. This is shown by the trust model component in Figure 1. Secondly, we need to formulate the strategy of the agent which will in turn regulate how it spends and earns money as well as how it provides information to its competitors. The expertise of the agent is a variable fixed by the simulator at the beginning of the game. As a result, an agent knows about its own variances in generating opinions and this will in turn influence the strategy of the agent. We now describe these two main parts of a generic agent design.

The first part is concerned with the trust model which itself consists of three main parts: (i) the lying detector which detects malicious agents, (ii) the variance estimator which estimates how much variance there is in the error of the opinions of other agents (i.e. their performance), and (iii) the optimal weight calculator which provides the optimal weights to provide to the simulator. These weights determine how the opinions an agent requested from its competitors are fused with its own opinion to generate the final appraisal and thereby its client share in the next iteration.

Now a large part of the quality of the trust model will rest on the reliability of the information the agent has gathered. The agent receives two main types of information. Firstly, it receives opinions about the paintings from its competitors (by requesting and paying for them) and it can receive opinions from the simulator by spending a certain amount of money. Secondly, it can obtain information about the behaviour of other agents, termed reputation information, by paying its competitors. Now, the quality of information obtained from the various sources is dependent on the

amount of money the agent spends on these various information providers as well as the behaviour of these sources. The provision of this information is what the strategy component of the agent determines.

An agent can spend money on requesting opinions and reputation information from its competitors, as well as generating its own opinion. Furthermore, an agent can make money by providing opinions and reputation information, as well as by providing appraisals to its client base. Thus, a successful agent needs to formulate a strategy that prudently spends money on obtaining opinions and information whilst ascertaining that its overall appraisal of the art painting of its client is comparatively good. Furthermore, the strategy determines how it interacts with its requesters when requested for opinions and reputation information.

Having briefly described the ART competition and how to construct a generic competition agent for it, we now describe the specific components of the IAM strategy. We will first describe the trust model component of IAM in the next section and then go on to detail its strategy in the subsequent section.

3. THE IAM AGENT'S TRUST MODEL

In this section, we show how, by combining its own opinion with those of others, the IAM agent can generate an overall appraisal of a painting that is more reliable than any single opinion on which it is based. That is, for a given painting with true value v , we wish to derive a combined estimate, \bar{e} , which is a function of the agent's own estimate, e_0 , and the reported third party estimates $\{e_1, \dots, e_q\}$. Moreover, our aim is to perform this combination such that the mean squared error, $E[(\bar{e} - v)^2]$, is smaller than that of each individual appraisal in $\{e_0, \dots, e_q\}$. Thus, by reporting \bar{e} during the competition, the IAM agent can increase its market share, due to the increased accuracy of its appraisals, and hence increase its revenue.

How this can best be achieved depends on the reliability of each agent's individual appraisals per era and, more specifically, on the variance, $var(e_i)$, of each estimate, e_i . Although, in the competition, each agent knows the variances of its own private appraisals, agents are not obliged to reveal these variances, nor are they obliged to reveal their private appraisals truthfully.

For these reasons, we must estimate the variances associated with each agent, and consequently adopt a three-step approach to combining appraisals. First, under the assumption that an agent reveals its appraisals truthfully, we use a Bayesian analysis to estimate an agent's variance, by modelling the factors which determine an agent's opinions. Second, we calculate a lower bound on the probability that an agent is lying about its appraisals, and use this to discard the opinions of potentially malicious agents. Finally, based on the estimated variances of agents believed to be truthful, we derive the optimal method for combining appraisals, under the constraints imposed by the competition. In the following subsections, we examine each of these steps in more detail.

3.1 Calculating Optimal Weights

The problem of combining estimates of an unknown value is one that has received much attention in the literature, for example [4] and [6]. In general, however, it can be viewed as an optimisation problem, in which we need to find a function $\bar{e} = f(e_0, \dots, e_q)$ of several estimates, $\{e_0, \dots, e_q\}$, such that \bar{e} has minimal mean squared error (MSE).

In many cases, the optimal function for achieving this is non-linear. However, one of the restrictions imposed by the competition is that estimates must be combined by the testbed, on an agent's behalf. This is achieved according to Equation 1, in which $\{w_0, \dots, w_1\}$ are weights applied to the individual opinions, which an agent can specify:

$$\bar{e} = \sum_{i=0}^q w_i \cdot e_i \quad (1)$$

Clearly, this function is linear, which constrains our search for an optimal combination. Moreover, optimality results for functions of this form are well known, and are generally referred to as Best Linear Unbiased Estimates (BLUE) [7]. Here, by unbiased we mean that the expected value of the combined estimate is equal to the value being estimated, or equivalently in this case, $E[\bar{e}] = v$. According to the design of the competition, the private appraisals of each agent are themselves unbiased estimates of v , which, although not essential, does simplify the problem of deriving a BLUE. Specifically, if we assume for the moment that private appraisals are truthfully revealed and that each $var(e_i)$ is known, a BLUE combination of appraisals can be achieved by setting each w_i as follows:

$$w_i = \frac{1/var(e_i)}{\sum_{i=0}^q 1/var(e_i)} \quad (2)$$

Unfortunately, the variance of each e_i is unknown during the competition, and so the optimal weights must be estimated, given an agent's past experience and knowledge of the environment. However, this task is simplified because each w_i only depends on the relative proportions of the variances. That is, if we multiply or divide each variance by an arbitrary constant then the weights remain unchanged, for example:

$$w_i = \frac{c/var(e_i)}{\sum_{i=0}^q c/var(e_i)} \quad (3)$$

$$w_i = \frac{c/var(e_i)}{c \sum_{i=0}^q 1/var(e_i)} \quad (4)$$

$$w_i = \frac{1/var(e_i)}{\sum_{i=0}^q 1/var(e_i)} \quad (5)$$

This has important implications for our strategy, when we consider how variances are assigned in the competition. Specifically, the variance of each agent's private opinion, for a given era, is generated by the testbed according to Equation 6, where s_i is the component of the i th agent's expertise vector specifying its knowledge of the painting's era, C_i is the amount of money spent by the i th agent to generate the appraisal, and α is a constant shared by all agents that during the competition was set to 0.5:

$$var(e_i) = v^2 \left(s_i + \frac{\alpha}{C_i} \right)^2 \quad (6)$$

Although, in general, s_i and C_i are unknown, the value of v^2 is constant for all i , and so is irrelevant. Therefore, rather than estimate $\text{var}(e_i)$ directly, it is sufficient to estimate the value of $\text{var}(e_i/v)$, which from the general properties of random variables is equal to:

$$\text{var}(e_i/v) = \left(s_i + \frac{\alpha}{C_i} \right)^2 \quad (7)$$

This is a much easier problem because we know that $\text{var}(e_i/v)$ will remain the same, even if v changes. In the following section we show how, using this fact, an agent can estimate $\text{var}(e_i/v)$ and so choose appropriate weights for each appraisal it receives.

3.2 Estimating Truthful Variances

From the previous section, we know that, for $i > 0$, $\text{var}(e_i/v)$ depends on two unknown parameters, s_i and C_i , and consequently must be estimated. One well founded way to achieve this is to use Bayesian analysis to marginalise over the unknown parameters [2]. That is, given s_i and C_i we know that $\text{var}(e_i/v)$ is equivalent to Equation 7, and from this, we can then calculate the *marginal* variance, $E[\text{var}(e_i/v)]$, as follows:

$$E[\text{var}(e_i/v)] = \sum_{s_i \in \mathcal{S}} \sum_{C_i \in \mathcal{C}} P(s_i, C_i) \left(s_i + \frac{\alpha}{C_i} \right)^2 \quad (8)$$

Here, $P(s_i, C_i)$ is the joint probability of the unknown parameters, \mathcal{S} & \mathcal{C} are their respective domains, and based on this, optimal weights can be estimated using:

$$w_i = \frac{1/E[\text{var}(e_i/v)]}{\sum_{j=0}^q 1/E[\text{var}(e_i/v)]} \quad (9)$$

In the case of an agent's own variance, the values of s_0 and C_0 are known, and so $E[\text{var}(e_0/v)]$ simplifies to $\text{var}(e_0/v)$. For third party opinions, on the other hand, we must model the probability distribution of s_i and C_i , based on available evidence. To this end, we assume that domains \mathcal{S} and \mathcal{C} are defined as follows, based on an analysis of the testbed software and competition rules:

$$\mathcal{S} = \{0.1, 0.2, \dots, 1\} \quad (10)$$

$$\mathcal{C} = \{1, 2, \dots, 10\} \quad (11)$$

This means that we can record $P(s_i, C_i)$ using a 10x10 conditional probability table (CPT) which we can update in light of an agent's experiences, and we can use this to calculate $E[\text{var}(e_i/v)]$. For example, in absence of evidence to the contrary, we may assume that all possible values of s_i and C_i are equally likely, resulting in the CPT shown in Table 1. Substituting the appropriate entries into Equation 9, this gives us a marginal variance of 0.5848. Moreover, if we hold the same beliefs for all agents then $w_i = 1/q$ for each i :

$$w_i = \frac{1/E[\text{var}(e_i/v)]}{\sum_{j=0}^q 1/E[\text{var}(e_i/v)]} = \frac{1/0.5848}{\sum_{j=0}^q 1/0.5848} \quad (12)$$

$$= \frac{1/0.5848}{q/0.5848} = \frac{1}{q} \quad (13)$$

Over time, the entries in the CPT are updated to reflect both an agent's experiences with different appraisers, and

	s_i									
C_i	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
1	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
2	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
3	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
4	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
5	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
6	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
7	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
8	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
9	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
10	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01

Table 1: Joint CPT for C_i and s_i and its initial values

the possibility that an agent may change its strategy between episodes. This is achieved by performing two operations at the end of each timestep: (1) update the CPT in light of any received opinions, and (2) add time dependent noise to account for dynamic behaviour.

To define these steps, we must first specify how we model dynamic behaviour in truth telling opinion providers. We know from the competition rules that s_i does not change for the duration of a game, but C_i can change if the opinion provider changes its policy. To account for this, we define C_i^t as the value of C_i at time $t \in \mathbb{Z}^+$. For each possible value of C_i , we then define transition probabilities that specify the conditional probability $P(C_i^{t+1} = c_1 | C_i^t = c_0)$ where c_1 and c_0 belong to \mathcal{C} . For example, we may define transition probabilities such that:

$$p(C_i^{t+1} = c_1 | C_i^t = c_0) = \begin{cases} \frac{9}{10} & \text{if } c_0 = c_1 \\ \frac{1}{90} & \text{otherwise} \end{cases} \quad (14)$$

thus assigning a probability of 0.1 to a provider changing C_i on any given time step, with equal probability of that transition being to any other possible state.

With this in mind, we can now consider the CPT update rules. First, if by the end of a time step we know that the true value of a painting is v^t , and that the opinion provided by an opinion provider is e_i^t , then we perform step 1 by calculating the posterior CPT using Bayes rule as follows:

$$P(s_i, C_i^t | p_i^t) = \frac{P(p_i^t | s_i, C_i^t) P(s_i, C_i^t)}{\sum_{s \in \mathcal{S}} \sum_{c \in \mathcal{C}} P(p_i^t | s, c) P(s, c)} \quad (15)$$

Here, p_i^t is defined as e_i^t/v^t ; $P(s, C_i^t)$ is the prior probability of the parameters, taken from the current CPT; and $P(p_i^t | s, C_i^t)$ is the data likelihood function. The latter is determined by the competition testbed, which generates estimates from a Gaussian distribution with mean v^t and variance $\text{var}(e_i^t)$. By substituting into the Gaussian p.d.f., we thus have:

$$p(p_i^t | s_i, C_i^t) = \frac{1}{\sqrt{2\pi(s_i + \alpha/C_i^t)^2}} \exp \left[-\frac{(p_i^t - 1)^2}{2(s_i + \alpha/C_i^t)^2} \right] \quad (16)$$

Finally, in step 2, we prepare the ground for the next time step by marginalising over C_i^t using the transition probabilities. This is achieved by updating the CPT for timestep

$t + 1$, as follows:

$$P(s_i, C_i^{t+1}) = \sum_{C_i^t \in \mathcal{C}} P(C_i^{t+1}, s_i, C_i^t) \quad (17)$$

$$= \sum_{C_i^t \in \mathcal{C}} P(C_i^{t+1} | s_i, C_i^t) P(s_i, C_i^t) \quad (18)$$

which, assuming we do not specify different transition probabilities depending on s_i , is equal to:

$$P(s_i, C_i^{t+1}) = \sum_{C_i^t \in \mathcal{C}} P(C_i^{t+1} | C_i^t) P(s_i, C_i^t) \quad (19)$$

3.3 Dealing with Liars

So far, we have shown how optimal weights can be estimated for agents that reveal their appraisals truthfully. However, as this is not necessarily the case, we need to be able to identify lying behaviour among agents, and act appropriately. This serves two purposes: (1) it identifies potentially fictitious appraisals, which are then eliminated from consideration before the combined appraisals are calculated; and (2) it informs the IAM agent's spending strategy, and how it behaves toward agents it believes to be malicious (see Section 4).

To this end, we test the hypothesis that an agent's last k opinions are truthful, based on what we know about how an agent's private opinions are generated. Here, k can be set to the total number of opinions generated. Here, k can be set to the total number of opinions generated from an agent, or it can be set to the number of opinions received in the last n timesteps, to allow for the possibility that an agent may change its strategy over time. More specifically, according to our assumptions, the maximum mean squared error for a truthful opinion is:

$$\max[mse] = v^2 \left(\max[s_i] + \frac{\alpha}{\min[C_i]} \right)^2 \quad (20)$$

$$= v^2 \left(1 + \frac{0.5}{0.1} \right)^2 \quad (21)$$

$$= 2.25v^2 \quad (22)$$

On this basis, the random variable e_i/v is normally distributed with mean 1 and variance 2.25. According to the standard properties of normally distributed random variables, if the last k opinions are all generated in this way, then the following statistic has a chi-squared distribution with k degrees of freedom:

$$Q_k = \frac{1}{2.25} \sum_{i=1}^k (e_i/v_i - 1)^2 \quad (23)$$

This allows us to place an upper bound on the probability that an agent's last k opinions are truthful, by comparing the actual value of Q_k to the quantiles of the chi-squared distribution.

For example, the distribution of Q_k for $k = 5$ is shown in Figure 2, in which the vertical lines mark the 0.75, 0.95 and 0.99 quantiles respectively. Now suppose that, as illustrated, Q_k is recorded as 18. As this occurs beyond the 0.99 quantile, this means that, even if all k opinions were generated with the highest possible variance for truthful opinions, the probability of Q_k occurring by chance is less than 0.01. In




Figure 2: Chi-squared distribution with 5 degrees of freedom




Figure 3: Cash flows from/to the (IAM) agent

the next section, we show how, by making such comparisons, an agent can make appropriate decisions about how to interact with its peers.

4. THE IAM AGENT'S STRATEGY

This section describes the strategy that the IAM agent adopts in order to maximise its final earnings (i.e. income less spending) in the ART game. First, we need to look at the cash flows to and from our agent in order to identify where decisions need to be made to optimise the agent's spending and earnings (see Figure 3). In brief, there are three sources and three sinks of cash for an agent participating in the competition. Figure 3 also depicts the magnitudes of these six flows of cash as decided by the competition. We will now discuss each of these cash flows in turn. We will denote with + those flows that generate income for the agent and by - those where the agent needs to spend.

1. Appraisal orders (+): At the beginning of each round in the game, the simulator generates a number of orders for each agent. The agent earns \$100 per appraisal order regardless of how it will fulfill that order (even if it decides not to provide an appraisal). Since the amount earned per order is fixed, an agent can only increase its earnings from this flow by increasing its number of orders in successive rounds. This can only be achieved if it provides a better service than its competitors (i.e. providing more accurate appraisals than

those of its competitors). This in turn requires the spending of money on 2, 3, and 4 below.

2. Generating own appraisals (–): The IAM agent can spend an arbitrary amount of money (representing the time cost associated with the appraising activity) to generate its own opinion for a particular order. The more money it spends, the more accurate the appraisal will be. However, this accuracy is limited to a pre-defined value according to the agent’s expertise in a particular art era. More mathematically, the standard deviation of the error of the opinion generated by the simulator, to which the appraiser with expertise s_i has provided C_i is given by:

$$\text{std}(e_i) = \sqrt{\text{var}(e_i)} = \left(s_i + \frac{\alpha}{C_i} \right) v \quad (24)$$

where α is a parameter set by the competition and v is the true value of the painting. Thus, it can be observed that deciding on the expenditure of money in generating one’s own opinion is not straightforward and this will be investigated in greater detail in Section 4.1.

3. Opinion requests (–): An agent can also request the opinion(s) of other agents on a particular appraisal order. Each opinion request has a fixed cost of \$10. Note that opinion providers (i.e. other agents in the game) are not obliged to provide their opinions after having received the money. In addition, they can provide opinions of low accuracy either due to having very low expertise on the concerned art era and/or deciding to spend little money on generating their opinions. Thus, the selection of opinion providers and the calculation of the weights to be provided to the simulator in order to aggregate opinions are challenging problems which an agent faces when buying opinions. Section 4.2 discusses how the IAM agent tackles these problems.
4. Reputation requests (–): In order to help agents find out about good opinion providers, the ART testbed allows agents to contact one another to ask for their evaluation about the appraising performance of a particular agent in a particular era¹.
5. Providing opinions (+): An agent can also earn by providing opinions to other agents (i.e. its competitors) when requested. The earnings for such a service are fixed at \$10 per opinion request. This service can provide extra income to boost the IAM agent’s profits and the strategy for this is specifically looked at in Section 4.3.
6. Reputation service (+): Similarly to providing its opinions, the IAM agent can also provide to other agents the reputation value of a particular agent in the game for \$1 per reputation request.

¹The performance of an agent in an art era is not necessarily equivalent to its expertise in that era. The actual performance of an agent is also determined by the amount of money it decides to spend on generating appraisals for others in addition to its expertise. For this reason, an agent can provide different levels of quality of service to different partners and, thus, the reputation values of the same agent provided by different partners can be different (assuming that the partners provide their honest evaluation as the reputation values).

From the cash flow analysis above, since providing the appraisal service to customers earns relatively much more money than providing opinions or reputation values, our strategy focuses on maximising the customer base of the IAM agent by trying to provide a highly accurate appraisal service. However, the more accurate an agent wants its appraisals to be, the more money it needs to spend (see the money flows 2, 3, and 4), and, as a result, the less profit it makes from a job. Meanwhile, spending less will result in less accurate appraisals, and in the ART game, this means fewer customers and thus less income. This is a classic tradeoff between profit and customers. Therefore, it is important to find the balance between these; i.e. spend wisely to generate accurate appraisals and to retain a high profit at the same time. In more detail, before a final appraisal is produced for a particular order, the IAM agent needs to make the following decisions:

1. how much should it spend on generating its own appraisal given its expertise on the art era of the order,
2. whether it needs to ask for external opinions for a given order and, if so, which opinion providers it should choose and
3. whether it needs to ask for the reputation values of opinion providers and how these values can be used to help it makes the decisions in (2).

Regarding using reputation values provided by other agents, since the game rules do not specify the semantics of a reputation value, except that it is a number in the range $[0, 1]$, interpreting their real meanings is very difficult (because there is no commonly shared semantics). Moreover, given the small number of participating agents in a game (5), the behaviours of the other agents can be learnt fairly quickly and thus reputation values are not of much use. For these reasons, the IAM agent does not request reputation values from other participants (and saves a small amount of money by doing so). Thus, it relies solely on the variance estimator (Section 3.2) to estimate the performance of an opinion provider.

In the remainder of this section, Sections 4.1 and 4.2 will address the first two issues above. In addition, Section 4.3 looks at how the IAM agent can increase its earnings by providing opinions and reputation values to other agents in a game.

4.1 Generating its Own Appraisals

An agent’s own appraisal for a painting is generated by the simulator from a Gaussian distribution whose mean is the painting’s true value (v) and whose deviation is determined by the agent’s expertise (s_i) and the money it spends (C_i) as in Equation 24. Therefore, given that s_i is predetermined by the simulator, an agent can only directly affect the accuracy of its appraisals by changing the amount of money it is willing to spend to generate the appraisals (i.e. C_i). However, since the painting’s true value v is factored out when the accuracy of an appraisal is calculated [3], for the sake of simplicity, it is equivalent to considering the standard




Figure 4: The cost vs. deviation trade off

deviation of an agent’s appraisals to be the following:

$$std(e_i/v) = \sqrt{var(e_i/v)} = s_i + \frac{\alpha}{C_i} \quad (25)$$

Although the more an agent spends, the less the deviation and the more accurate its appraisals, from Equation 25, it is clear that the lower bound of the deviation is s_i and it might not be worth spending more money when $std(e_i/v)$ is close enough to s_i . In this respect, Figure 4 shows the values of $std(e_i/v)$ corresponding to different amounts of spending on C_i for different values of s_i in $\{0.1, 0.2, \dots, 1.0\}$. Coincidentally, all the deviation curves in Figure 4 level off from spending of \$4 upwards, despite the value of s_i . This means, in any case, \$4 is a balance point between spending and the accuracy gained from it. Spending more than \$4 only yields a very marginal decrease of the deviation (or a marginal increase of appraisal accuracy). Given this, the IAM agent only spends \$4 for generating its own appraisals in all situations (regardless of its expertise).

4.2 Ordering External Opinions

In this section, we consider, given a particular appraisal order and having known our own expertise s_i in the order’s art era, whether external opinion(s) should be sought and, if so, from whom. From the findings in the previous sections, since it is relatively cheap (\$4 compared to the order fee of \$100) to generate its own appraisals which are reasonably close to its expertise, the IAM agent always uses its own appraisals regardless of whether it needs external opinions or not because this would provide a safety catch should other agents decide to cheat by not providing their opinions as agreed. Having estimated the performance of each opinion provider (i.e. its variance, see Section 3.2), the IAM agent can evaluate the benefit of having an opinion from a particular opinion provider to justify the decision of requesting an opinion from it. In other words, the IAM agent needs to calculate the expected variance of the final appraisal² resulting from combining its own appraisal with that of the potential opinion provider. Given that in producing the final appraisals, the simulator combines all the submitted appraisals as in Equation 1 and the weight for each appraisal w_i is defined as in Equation 5, the variance

²The expected variance of the final appraisal reflects the expected accuracy of the appraisal.

of the combined appraisal is given by:

$$var(\bar{e}) = \sum_{i=0}^q w_i^2 \cdot var(e_i) \quad (26)$$

where q is the number of external opinions (e_i) the IAM agent receives and e_0 is its own appraisal. Replacing w_i from Equation 5 gives:

$$var(\bar{e}) = \sum_{i=0}^q \left(\frac{1/var(e_i)}{\sum_{j=0}^q 1/var(e_j)} \right)^2 \cdot var(e_i) \quad (27)$$

$$= \left(\frac{1}{\sum_{j=0}^q 1/var(e_j)} \right)^2 \sum_{i=0}^q 1/var(e_i) \quad (28)$$

$$= \frac{1}{\sum_{j=0}^q 1/var(e_j)} \quad (29)$$

Factoring out the painting’s value v from both sides of Equation 29 gives us:

$$var(\bar{e}/v) = \frac{1}{\sum_{i=0}^q 1/var(e_i/v)} \quad (30)$$

Although we do not know $var(e_i/v)$ for other agents, we can estimate it (as shown in Section 3.2), and thus we can estimate $var(\bar{e}/v)$:

$$E[var(\bar{e}/v)] = \frac{1}{\sum_{i=0}^q 1/E[var(e_i/v)]} \quad (31)$$

where $E[var(e_0/v)] = var(e_0/v) = s_0 + \alpha/C_0$ (because we know our own parameters) and $E[var(e_i/v)]$ can be calculated as in Equation 8 for other agents.

Knowing how to calculate the combined variance of opinions, the IAM agent can start selecting agents it will ask for opinions by performing the following four steps. First, all other agents are sorted in ascending order, according to their estimated variances in the art era of the order (i.e. best performance first). Second, the list of agents whose opinions will be used is initialised to contain (only) the IAM agent itself (because it always uses its own appraisal). Third, the IAM agent then considers adding each of the agents into the list according to their performance order by calculating the expected combined variance of the final appraisal assuming that agent is in the list. Finally, if the new combined variance is just insignificantly improved (i.e. improvement is less than 15%³) then the selection process stops; otherwise, that agent is put into the list and the process continues with the next agent.

After the selection process is finished, the IAM agent sends opinion requests to agents selected in the list. However, since the variance estimator assumes that agents do not cheat when providing their appraisals, this process does not filter out cheating agents. Therefore, before the selection process begins, all the other agents are evaluated based on their previous interactions using the procedure as described in Section 3.3 to calculate the probability that an agent has cheated in the past. If the cheating probability is over 0.6, then the opinion provider is classified as a *potential cheater* and its opinions will not be taken into account in our agent’s

³This improvement threshold is hand-picked based on our experiments.

final appraisal. However, it will still be asked for opinions in order to confirm whether it is actually cheating. If the cheating probability is over 0.95, the opinion provider is classified as a *cheater* and no future interaction with it will be made. Moreover, an agent is also classified as a cheater if it failed to provide an appraisal after having confirmed that it will provide an opinion.

4.3 Earning from Being Honest

Beside earning money from appraising orders by the simulator, the IAM agent’s strategy is to maximise its income via providing opinions and reputation values. Its philosophy in doing so is to provide an honest and reliable service to all agents in order to maintain a good business relationship with them, which hopefully results in a regular source of income. Moreover, providing a low quality service or cheating other agents might initiate retaliatory behaviours from them. This will have a detrimental effect on the IAM agent’s own appraising business when it requires external opinions to fulfil orders in whose art era it is not an expert. Although the general strategy here is to provide a good service, the IAM agent is not naïvely benign. Rather, it looks for reciprocity from its partners (who are also its competitors). Therefore, if a cheater is identified, the IAM agent will use a retaliatory policy toward it. Following this general strategy, the IAM agent’s strategies for the earning cash flows 5 and 6 (i.e. providing opinions and reputation values for other agents, see Figure 3) are discussed next.

4.3.1 Providing opinions:

The IAM agent aims to develop its reputation as an honest agent in order to earn (good) money from providing opinions to others. Thus, it always provides honest certainty values (taking into account the money it will spend) and honest appraisals (i.e. spending the same amount of money it said it would spend). In more detail, the IAM agent always spends \$4 to generate appraisals for others (the same amount it would spend to generate appraisals for its own orders, see Section 4.1). When requested for its certainty assessment on its appraisals, the IAM agent will present the certainty value (denoted by cv) based on its expertise in the art era of that particular request and its intended spending to generate an appraisal (i.e. $C_i = \$4$):

$$cv = 1 - \frac{(1 + \alpha/C_i)s_i}{1.5} \quad (32)$$

where s_i is the expertise of the IAM agent in the concerned art era and 1.5 is the maximum deviation of the Gaussian distribution generating an agent’s appraisal value (given by Equation 25 in the worst case scenario where the minimum expertise $s_i = 1.0$ and the minimum spending $C_i = \$1$).

In case it detects that an agent cheats by providing false opinions to it or by taking the opinion fee without providing an opinion, it will retaliate against the cheating agent in all future opinion transactions by spending only a fractional amount of money (\$0.01) to generate appraisals that are extremely skewed (due to the ensuing high deviation) for the cheating agent.

4.3.2 Providing reputation values:

Agent	Affiliation	Revenue	Cost	Profit
IAM	University of Southampton	149812	18299	131583
Neil	Nanyang Technological University	116764	13741	103023
Frost	Bogazici University	120753	18176	102577
Sabatini	Universidad Carlos III de Madrid	127137	25726	101411
Joey	University of Nebraska-Lincoln	111985	19506	92479
	<i>mean</i>	<i>125290</i>	<i>19076</i>	<i>106215</i>

Table 2: Average revenue, profit & cost at end of final round games

Agent	Opinion Costs	Opinion Generation Costs	Reputation Costs
IAM	59.46	40.54	0.00
Neil	6.80	92.03	1.16
Frost	31.65	68.35	0.00
Sabatini	38.85	61.15	0.00
Joey	0.00	100.00	0.00

Table 3: Breakdown of cost percentages.

As previously mentioned, there is no defined semantics for reputation values except that their range is $[0, 1]$. We therefore define the reputation value (r_i) to provide based on the estimated variance of an agent, which we believe best reflects the performance of that agent as so far perceived:

$$r_i = \max\left\{1 - \frac{\sqrt{\text{var}(e_i/v)}}{1.5}, 0\right\} \quad (33)$$

where 1.5 is again the maximum deviation given the worst case scenario and $\sqrt{\text{var}(e_i/v)}$ is the estimated deviation of that agent’s appraisals. In this service, the IAM agent produces random reputation values in $[0, 1]$ for cheating agents.

5. COMPETITION RESULTS

In this section, we give an in depth discussion of the competition results, concentrating on its final round in particular. This consisted of 10 games with 60 time steps each, in which the top 5 agents from a series of preliminary games were pitted against each other (there were 17 agents entered into the competition). At the end of each game, the total profit earned by each agent was recorded, and then averaged over all 10 games to produce the final scores. The results of this process are shown in Table 2, which includes the average revenue and cost for each agent over the 10 games. In particular, this shows that the IAM agent won the competition by a 28% profit margin over its nearest competitor, earning the highest revenue of any agent, with below average costs.

A more detailed breakdown of each agent’s cash flow is shown in Tables 3 and 4, in terms of the types of payment that contributed to the costs and revenue in the final round games. Here, client, opinion and reputation payments refer to the total amount of each agent’s revenue that was due to payments for client appraisals, opinion generation and reputation respectively; opinion costs refer to the amount spent

Agent	Client Payments	Opinion Payments	Reputation Payments
IAM	96.09	3.89	0.03
Neil	98.63	1.37	0.00
Frost	98.45	1.52	0.03
Sabatini	88.25	11.72	0.03
Joey	96.96	3.00	0.04

Table 4: Breakdown of revenue percentages.




Figure 5: Opinion Spending

requesting third party opinions; opinion generation costs are those spent on generating opinions for third parties; and reputation costs are those spent requesting reputation about an agent’s peers.

From these results, it is clear that largest proportion of each agent’s revenue was due to payments for client appraisals, followed by opinion payments and reputation in that order. This is partially a consequence of the comparatively high price of client appraisals (\$100 compared to \$10 for opinions and \$1 for reputation). However, even taking this into account, it is clear that very little reputation was exchanged between any of the agents, reinforcing the idea that reputation is only useful in large populations with clear semantics.

On the other hand, opinion transactions did feature more prominently in the final rounds, and played an important role in determining the order of the leader table. Although this is not reflected directly in agent revenues, opinion costs accounted for the majority of the overhead in the IAM agent’s winning strategy. To understand the reason for this, suppose that two agents can co-operate by spending a total of \$10 on their opinions to appraise a painting. Assuming that, for both agents, $s_i = 1$, $\alpha = 0.5$, and the difference between their individual expenditure is y , then the best achievable variance for the appraisal (relative to a painting value) is:

$$var(\bar{e}/v) = \left(\frac{1}{[1 + 1/(10 + y)]^2} + \frac{1}{[1 + 1/(10 - y)]^2} \right)^{-1} \quad (34)$$

Based on this equation, the appraisal variance resulting from difference values of y is plotted in Figure 5. This shows that the lowest possible variance is achieved when $y = 0$ or, in other words, when the agents divide the \$10 equally between their respective opinions. This can easily be shown to be true in general, and explains why it was beneficial for the IAM agent to spend significantly on third party opinions,

Agent	MSE of Opinions Reported to IAM	No. Opinions Sent to IAM
Neil	6.20	439
Frost	142533.43	179
Sabatini	0.41	10601
Joey	1338173.81	270

Table 5: Total number of opinions received by the IAM agent compared to average opinion MSE.

Agent	Revenue	Cost	Profit
IAM	150132	18879	131253
Neil	139870	13757	126113
Joey	133417	19506	113911
Sabatini	127814	46470	81344
Frost	120833	42381	78452

Table 6: Alternative scores based on recorded number of opinion transactions.

rather than investing solely in its own direct assessments. On the other hand, agents such as Neil and Joey, who spent little or no income on third party opinions, did make it to the final round. This is perhaps an indication that, if opinions are not used wisely based on solid statistical models, then the simple strategy of using only an agent’s direct experience can produce better results.

Concentrating more specifically on the IAM agent’s performance, its success in the competition can be further explained by three other factors. First, by spending \$4 on opinion generation, it managed to spend less per client on opinion generation than any of the other four finalists, while still maintaining close to minimal opinion errors. Second, the IAM agent further reduced its costs by only buying opinions from reliable agents. This is made apparent when we compare the total number of opinions received by IAM from the other competitors, compared to the MSE of their reported opinions (Table 5). This shows a clear correlation, with IAM relying heavily on Sabatini’s accurate opinions, while mostly rejecting opinions from the other competitors, who provided misleading opinions. Third, by estimating optimal weights for received opinions based on standard statistical theory, the IAM agent was able to achieve the lowest MSE for its appraisals compared to the other finalists, and so was awarded the highest overall client market share.

In addition, the IAM agent’s share of the opinion market was second only to Sabatini. This is partially because the IAM agent itself was the highest consumer of third party opinions in terms of expenditure, most of which was spent on Sabatini’s expertise. However, there does appear to be a discrepancy in the competition data between the number of opinions received by each agent, and the number of opinions paid for. More specifically, all five finalists sent more opinions than they were paid for, perhaps due to inadequate checks for payment, before opinions were sent. Without further investigation, it is unclear whether the cause of this lies with the competition testbed, the agents themselves, or some combination of both. Nevertheless, if funds had been exchanged for all the opinion transactions that took place,

then this would have had a dramatic effect on the competition, and would have resulted in the alternative scores shown in Table 6. What this shows is that both Frost and Sabatini were sent significantly more opinions than they paid for and, had they been expected to pay for all these opinions, they would have been placed at the bottom of the leader board. However, we emphasise that this does not suggest that these alternative scores are correct, or that the official scores are inaccurate. This is particularly true if the extra opinions were unsolicited, because it would be unfair to make agents pay for unwanted opinions. Even so, this does reveal an interesting anomaly that is perhaps worth further investigation.

6. CONCLUSIONS

From the design of our successful agent for the ART competition, three lessons can be drawn that are applicable to trust assessment within multi-agent systems in general.

The first lesson concerns the conditions under which an agent can benefit by pooling information from other agents about their shared environment. In the ART competition, this aggregation of information occurs at two levels, namely at appraising art objects and at evaluating the trustworthiness of other agents. We first consider the appraisal of art objects. *In this case, an agent can benefit from third party information, if it is easier to establish such information as reliable, than it is to obtain equivalent verifiable information directly.* This lesson is especially true with regard to art appraisal, due to the way in which the MSE of an agent's opinions is calculated. That is, beyond a certain point, only marginal increases in accuracy can be achieved by spending more money on the opinion of an individual agent (Section 4.1). Thus, it is generally more economical for an agent to purchase opinions from a number of third parties than it is to invest heavily in its own opinion. The success of the IAM agent can be partially attributed to this conclusion, when we compare its performance to other agents, such as Joey, who invested little or no income on third party opinions.

The second lesson is the counter of the first, establishing when it is not beneficial to try pooling information. For example, in the competition, there is little apparent advantage to reputation sharing, as is reflected by the negligible number of reputation transactions that occurred in the final rounds. This can be attributed to the small number of agents that participated in each round (typically 5 agents), each of which had sufficient opportunity and funds to interact with all of their peers in every time-step. As a result, participants could gain reliable first-hand experience of their peers, as easily as they could gather reputation. Furthermore, the semantics of reputation were not well defined, increasing the difficulty of its interpretation and assessment. For example, suppose that reputation from one agent consists of estimated standard deviations for opinions, while for another it consists of estimated variances. Both statistics have the same range of values, $[0, \infty)$, and in both cases, lower values indicate more reliable opinions. However, their correct interpretation is still different: a variance of 10 implies a much more reliable opinion than one with a standard deviation of 10. Thus, to assess the reliability of reputation, it would first be necessary to learn the function used to create it. However, this does not imply that reputation

sharing is worthless in general. Rather, the lesson learnt is that *reputation is most valuable in cases where direct experience is relatively more difficult to acquire, and in which the semantics are clearly defined.*

The final lesson is that, *even if third party opinions carry useful information, this value is wasted if opinions are not properly applied using well founded statistical techniques.* Although trust can be viewed as a sociological concept, and inspiration for computational models of trust can be drawn from multiple disciplines, the problem of combining estimates of unknown variables (such as trustee behaviour) is fundamentally a statistical one. Thus, any mechanism for making predictions based on reputation that ignores statistical theory, does so at its peril. This claim may partially explain the relative success of agents, such as Joey and Neil, that spent little or no income on third party opinions. By their presence in the final rounds, these agents suggest that, unless appropriate statistical techniques are used (such as those used by the IAM agent), it is perhaps better to ignore reputation completely, in favour of a simpler solution.

7. REFERENCES

- [1] R. K. Dash, S. D. Ramchurn, and N. R. Jennings. Trust-based mechanism design. In *Proc. 3rd Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 748–755, 2004.
- [2] M. DeGroot and M. Schervish. *Probability & Statistics*. Addison-Wesley, 3rd edition, 2002.
- [3] 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 *Proc. 4th Int. Conf. on Autonomous agents and multiagent systems*, pages 512–518. ACM Press, 2005.
- [4] C. Genest and J. V. Zidek. Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, 1(1):114–135, 1986.
- [5] T. D. Huynh, N. R. Jennings, and N. R. Shadbolt. An integrated trust and reputation model for open multi-agent systems. *Journal of Autonomous Agents and Multi-Agent Systems*, 13(2):119–154, 2006.
- [6] J. Kittler, M. Hatef, R. P. W. Duin, and J. Matas. On combining classifiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(3):226–239, 1998.
- [7] R. H. Myers, D. C. Montgomery, and G. G. Vining. *Generalized Linear Models With Applications in Engineering and the Sciences*. Wiley, 2002.
- [8] S. D. Ramchurn, T. D. Huynh, and N. R. Jennings. Trust in multi-agent systems. *The Knowledge Engineering Review*, 19(1):1–25, March 2004.
- [9] W. T. L. Teacy. *Agent-Based Trust and Reputation in the Context of Inaccurate Information Sources*. PhD thesis, School of Electronics and Computer Science, University of Southampton, 2006.
- [10] W. T. L. Teacy, J. Patel, N. R. Jennings, and M. Luck. TRAVOS: Trust and reputation in the context of inaccurate information sources. *Autonomous Agents and Multi-Agent Systems*, 12(2):183–198, 2006.