ISBN: 972-8924-16-X © 2006 IADIS

INFORMED AGENTS FOR E-MARKET TRADING

John Debenham & Simeon Simoff
Faculty of IT, University of Technology, Sydney, NSW 2007, Australia

ABSTRACT
Fully automated trading, such as e-procurement, using the Internet is virtually unheard of today. To fully automate the trading process agents need access to information from the market, the market environment and from general news sources. Here informed trading agents are built on a synthesis of the two technologies: data mining, and intelligent agents. This paper describes a demonstrable prototype e-trading system that is populated with informed agents and is available on the World Wide Web. This is part of a larger project that aims to make informed automated trading a reality.

KEYWORDS
Electronic markets, data mining, trading agents.

1 INTRODUCTION

The potential size of the electronic business market and the comparatively small amount of automated negotiation presently deployed provides a major incentive for research in automated trading. Fully automated trading, such as e-procurement, using the Internet is virtually unheard of today. Trading involves the maintenance of effective business relationships, and is the complete process of: need identification, product brokering, supplier brokering, offer-exchange, contract negotiation, and contract execution. Two core technologies are fused to create informed trading agents:

- data mining — real-time data mining technology to tap information flows from the marketplace and the World Wide Web, and to deliver timely information at the right granularity.
- trading agents — intelligent agents that are designed to operate in tandem with the real-time information flows received from the data mining systems.

This paper describes informed agents that integrate these two technologies. The work described here augments conventional plan-based agent architectures — it is *not* intended to replace them. Our agents have goals and plans, although those goals may not be utilitarian. This work describes how an agent may deal with real-time information flows that deliver information with decaying integrity. It provides a basis for proactive information acquisition to reduce uncertainty in the agent's world model. The actions of the agent are then managed in a conventional way in a BDI architecture that incorporates the world model. The e-Market Framework is available on the World Wide Web[1]. This project aims to make informed automated trading a reality, and develops further the "Curious Negotiator" framework (Simoff & Debenham 2002). This work does not address all of the issues in automated trading. For example, the work relies on developments in: XML and semantic web, secure data exchange, value chain management and financial services.

The data mining systems that have been developed for mining information both from a virtual institution and from general sources from the World Wide Web are described in Sec. 2. Intelligent agent that are built on an architecture designed specifically to handle real-time information flows are described in Sec. 3. Our work on virtual institutions has been carried out in collaboration with "Institut d'Investigacio en Intel.ligencia Artificial[2]", Spanish Scientific Research Council, UAB, Barcelona, Spain. Sec. 4 concludes.

2 DATA MINING

We have designed information discovery and delivery agents that utilise text and network data mining for supporting real-time negotiation. This work has addressed the central issues of extracting relevant information from different on-line repositories with different formats, with possible duplicative and erroneous data. That is, we have

[1] `http://e-markets.org.au`
[2] `http://www.iiia.csic.es/`

addressed the central issues in extracting information from the World Wide Web. Our mining agents understand the influence that extracted information has on the subject of negotiation and takes that in account.

Real-time embedded data mining is an essential component of the proposed framework. In this framework the trading agents make their informed decisions, based on utilising two types of information:

- information extracted from the negotiation process (i.e. from the exchange of offers), and;
- information from external sources, extracted and provided in condensed form.

The embedded data mining system provides the information extracted from the external sources. The system complements and services the information-based architecture developed in (Debenham 2004b) and (Sierra & Debenham 2005). The information request and the information delivery format is defined by the interaction ontology. As these agents operate with negotiation parameters with a discrete set of feasible values, the information request is formulated in terms of these values. As agents proceed with negotiation they have a topic of negotiation and a shared ontology that describes that topic. For example, if the topic of negotiation is buying a number of digital cameras for a University, the shared ontology will include the product model of the camera, and some characteristics, like "product reputation" (which on their own can be a list of parameters), that are usually derived from additional sources (for example, from different opinions in a professional community of photographers or digital artists). As the information-based architecture assumes that negotiation parameters are discrete, the information request can be formulated as a subset of the range of values for a negotiation parameter. For example, if the negotiator is interested in cameras with 8 megapixel resolution, and the brand is a negotiation parameter, the information request can be formulated as a set of camera models, e.g. {"Canon Power Shot Pro 1", "Sony f828", "Konica Minolta Dimage A2", "Nikon Coolpix 8400", "Olympus C-8080"} and a *preference estimate* based on the information in the different articles available. The collection of parameter sets of the negotiation topic constitutes the input to the data mining system. Continuous numerical values are replaced by finite number of ranges of interest.

The data mining system initially constructs data sets that are "focused" on requested information. From the vast amount of information available in electronic form, we need to filter the information that is relevant to the information request. In our example, this will be the news, opinions, comments, white papers related to the five models of digital cameras. Technically, the automatic retrieval of the information pieces utilises the universal news bot architecture presented in (Zhang & Simoff 2004). Developed originally for news sites only, the approach is currently being extended to discussion boards and company white papers.

The "focused" data set is dynamically constructed in an iterative process. The data mining agent constructs the news data set according to the concepts in the query. Each concept is represented as a cluster of key terms (a term can include one or more words), defined by the proximity position of the frequent key terms. On each iteration the most frequent (terms) from the retrieved data set are extracted and considered to be related to the same concept. The extracted keywords are resubmitted to the search engine. The process of query submission, data retrieval and keyword extraction is repeated until the search results start to derail from the given topic.

The set of topics in the original request is used as a set of class labels. In our example we are interested in the evidence in support of each particular model camera model. A simple solution is for each model to introduce two labels — positive opinion and negative opinion, ending with ten labels. In the constructed "focused" data set, each news article is labelled with one of the values from this set of labels. An automated approach reported in (Zhang & Simoff 2004) extends the tree-based approach proposed in (Reis, et al. 2004).

The data sets required further automatic preprocessing, related to possible redundancies in the information encoded in the set that can bias the analysis algorithms. For example, identifying a set of opinions about the camera that most likely comes from the same author, though it has been retrieved from different "opinion boards" on the Internet.

Once the set is constructed, building the "advising model" is reduced to a classification data mining problem. As the model is communicated back to the information-based agent architecture, the classifier output should include all the possible class labels with an attached probability estimates for each class. Hence, we use probabilistic classifiers (e.g. Naïve Bayes, Bayesian Network classifiers (Berthold & Hand 2003) without the min-max selection of the class output [e.g., in a classifier based on Naïve Bayes algorithm, we calculate the posterior probability $\mathbb{P}_p(i)$ of each class $c(i)$ with respect to combinations of key terms and then return the tuples $< c(i), \mathbb{P}_p(i) >$ for all classes, not just the one with maximum $\mathbb{P}_p(i)$. In the case when we deal with range variables the data mining system returns the range within which is the estimated value. For example, the response to a request for an estimate of the rate of change between two currencies over specified period of time will be done in three steps: (i) the relative focused news data set will be updated for the specified period; (ii) the model that takes these news in account is updated, and; (iii) the output of the model is compared with requested ranges and the matching one is returned. The details of this part of the data mining system are presented in (Zhang, et al. 2005). The currently used model is a modified linear model with an additional term that incorporates a news index Inews, which reflects the news effect on exchange rate. The data mining system provides parameters that define the "quality of the information", including:

- the time span of the focused data set, defined by the eldest and the latest information unit);
- estimates of the characteristics of the information sources, including reliability, trust and cost, that then are used by the information-based agent architecture.

Overall the parameters that will be estimated by the mining algorithms and provided to the negotiating agents are expected to allow information-based agents to devise more effective and better informed situated strategies. In addition to the data coming from external sources, the data mining component of the project will develop techniques for analysing agent behaviourist data with respect to the electronic institution setup.

3 TRADING AGENTS

We have designed a new agent architecture founded on information theory. These "information-based" agents operate in real-time in response to market information flows. We have addressed the central issues of trust in the execution of contracts, and the reliability of information (Sierra & Debenham 2005). Our agents understand the value of building business relationships as a foundation for reliable trade. An inherent difficulty in automated trading — including e-procurement — is that it is generally multi-issue. Even a simple trade, such as a quantity of steel, may involve: delivery date, settlement terms, as well as price and the quality of the steel. The "information-based" agent's reasoning is based on a first-order logic world model that manages multi-issue negotiation as easily as single-issue.

Most of the work on multi-issue negotiation has focussed on one-to-one bargaining — for example (Faratin, et al. 2003). There has been rather less interest in one-to-many, multi-issue auctions — (Debenham 2004a) analyzes some possibilities — despite the size of the e-procurement market which typically attempts to extend single-issue, reverse auctions to the multi-issue case by post-auction haggling. There has been even less interest in many-to-many, multi-issue exchanges.

The generic architecture of our "information-based" agents is presented in Sec. 3.1. The agent's reasoning employs entropy-based inference and is described in Sec. 3.2. The integrity of the agent's information is in a permanent state of decay, Sec. 3.3 describes the agent's machinery for managing this decay leading to a characterization of the "value" of information. Sec. 3.4 describes metrics that bring order and structure to the agent's information with the aim of supporting its management.

3.1 Information-Based Agent Architecture

The essence of "information-based agency" is now described. An agent observes events in its environment including what other agents actually do. It chooses to represent some of those observations in its world model as beliefs. As time passes, an agent may not be prepared to accept such beliefs as being "true", and qualifies those representations with epistemic probabilities. Those qualified representations of prior observations are the agent's *information*. This information is primitive — it is the agent's representation of its beliefs about prior events in the environment and about the other agents prior actions. It is independent of what the agent is trying to achieve, or what the agent believes the other agents are trying to achieve. Given this information, an agent may then choose to adopt goals and strategies. Those strategies may be based on game theory, for example. To enable the agent's strategies to make good use of its information, tools from information theory are applied to summarize and process that information. Such an agent is called *information-based*.

An agent called Π is the subject of this discussion. Π engages in multi-issue negotiation with a set of other agents: $\{\Omega_1, \cdots, \Omega_o\}$. The foundation for Π's operation is the information that is generated both by and because of its negotiation exchanges. Any message from one agent to another reveals information about the sender. Π also acquires information from the environment — including general information sources —to support its actions. Π uses ideas from information theory to process and summarize its information. Π's aim may not be "utility optimization" — it may not be aware of a utility function. If Π *does* know its utility function *and* if it aims to optimize its utility *then* Π may apply the principles of game theory to achieve its aim. The information-based approach does not to reject utility optimization — in general, the selection of a goal and strategy is secondary to the processing and summarizing of the information.

In addition to the information derived from its opponents, Π has access to a set of information sources $\{\Theta_1, \cdots, \Theta_t\}$ that may include the marketplace in which trading takes place, and general information sources such as news-feeds accessed via the Internet. Together, Π, $\{\Omega_1, \cdots, \Omega_o\}$ and $\{\Theta_1, \cdots, \Theta_t\}$ make up a multiagent system. The integrity of Π's information, including information extracted from the Internet, will decay in time. The way in which this decay occurs will depend on the type of information, and on the source from which it was drawn. Little appears to be known about how the integrity of real information, such as news-feeds, decays, although its validity can often be checked — "Is company X taking over company Y?" — by proactive action given a cooperative information source Θ_j. So Π has to consider how and when to refresh its decaying information.

Π has two languages: $\mathcal{C}$ and $\mathcal{L}$. $\mathcal{C}$ is an illocutionary-based language for communication. $\mathcal{L}$ is a first-order language for internal representation — precisely it is a first-order language with sentence probabilities optionally attached to each sentence representing Π's epistemic belief in the truth of that sentence. Messages expressed in $\mathcal{C}$ from $\{\Theta_i\}$ and $\{\Omega_i\}$ are received, time-stamped, source-stamped and placed in an *in-box* $\mathcal{X}$. The messages in $\mathcal{X}$ are then translated using an *import function* I into sentences expressed in $\mathcal{L}$ that have integrity decay functions (usually of time) attached to each sentence, they are stored in a *repository* $\mathcal{Y}^t$. And that is all that happens until Π triggers a goal.

Π triggers a goal, $g \in \mathcal{G}$, in two ways: first in response to a message received from an opponent $\{\Omega_i\}$ "I offer you €1 in exchange for an apple", and second in response to some need, $\nu \in \mathcal{N}$, "goodness, we've run out of coffee". In either case, Π is motivated by a need — either a need to strike a deal with a particular feature (such as acquiring coffee) or a general need to trade. Π's goals could be short-term such as obtaining some information "what is the time?", medium-term such as striking a deal with one of its opponents, or, rather longer-term such as building a (business) relationship with one of its opponents. So Π has a trigger mechanism T where: $T : \{\mathcal{X} \cup \mathcal{N}\} \to G$.

For each goal that Π commits to, it has a mechanism, G, for selecting a strategy to achieve it where $G : \mathcal{G} \times \mathcal{M} \to \mathcal{S}$ where $\mathcal{S}$ is the strategy library. A *strategy* s maps an information base into an action, $s(\mathcal{Y}^t) = z \in \mathcal{Z}$. Given a goal, g, and the current state of the social model m^t, a strategy: $s = G(g, m^t)$. Each strategy, s, consists of a *plan*, b_s and a *world model* (construction and revision) *function*, J_s, that constructs, and maintains the currency of, the strategy's *world model* W_s^t that consists of a set of probability distributions. A *plan* derives the agent's next action, z, on the basis of the agent's world model for that strategy and the current state of the social model: $z = b_s(W_s^t, m^t)$, and $z = s(\mathcal{Y}^t)$. J_s employs two forms of entropy-based inference:

- Maximum entropy inference, J_s^+, first constructs an *information base* $\mathcal{I}_s^t$ as a set of sentences expressed in $\mathcal{L}$ derived from $\mathcal{Y}^t$, and then from $\mathcal{I}_s^t$ constructs the world model, W_s^t, as a set of complete probability distributions [using Eqn. 2 in Sec. 3.2 below].

- Given a prior world model, W_s^u, where $u < t$, minimum relative entropy inference, J_s^-, first constructs the incremental information base $\mathcal{I}_s^{(u,t)}$ of sentences derived from those in $\mathcal{Y}^t$ that were received between time u and time t, and then from W_s^u and $\mathcal{I}_s^{(u,t)}$ constructs a new world model, W_s^t [using Eqn. 3 in Sec. 3.2 below].

3.2 Π's Reasoning

Once Π has selected a plan $a \in \mathcal{A}$ it uses maximum entropy inference to derive the $\{D_i^s\}_{i=1}^n$ and minimum relative entropy inference to update those distributions as new data becomes available. *Entropy*, $\mathbb{H}$, is a measure of uncertainty (MacKay 2003) in a probability distribution for a discrete random variable X: $\mathbb{H}(X) \triangleq -\sum_i p(x_i) \log p(x_i)$ where $p(x_i) = \mathbb{P}(X = x_i)$. Maximum entropy inference is used to derive sentence probabilities for that which is not known by constructing the "maximally noncommittal" (Jaynes 2003) probability distribution, and is chosen for its ability to generate complete distributions from sparse data.

Let $\mathcal{G}$ be the set of all positive ground literals that can be constructed using Π's language $\mathcal{L}$. A *possible world*, v, is a valuation function: $\mathcal{G} \to \{\top, \bot\}$. $\mathcal{V}|\mathcal{K}^s = \{v_i\}$ is the set of all possible worlds that are consistent with Π's knowledge base $\mathcal{K}^s$ that contains statements which Π believes are true. A *random world* for $\mathcal{K}^s$, $W|\mathcal{K}^s = \{p_i\}$ is a probability distribution over $\mathcal{V}|\mathcal{K}^s = \{v_i\}$, where p_i expresses Π's degree of belief that each of the possible worlds, v_i, is the actual world. The *derived sentence probability* of any $\sigma \in \mathcal{L}$, *with respect to* a random world $W|\mathcal{K}^s$ is:

$$(\forall \sigma \in \mathcal{L}) \mathbb{P}_{\{W|\mathcal{K}^s\}}(\sigma) \triangleq \sum_n \{\, p_n \, : \, \sigma \; is \top in \; v_n \,\} \tag{1}$$

The agent's *belief set* $\mathcal{B}_t^s = \{\Omega_j\}_{j=1}^M$ contains statements to which Π attaches a *given sentence probability* $\mathbb{B}(.)$. A random world $W|\mathcal{K}^s$ is *consistent* with $\mathcal{B}_t^s$ if: $(\forall \Omega \in \mathcal{B}_t^s)(\mathbb{B}(\Omega) = \mathbb{P}_{\{W|\mathcal{K}^s\}}(\Omega))$. Let $\{p_i\} = \{\overline{W}|\mathcal{K}^s, \mathcal{B}_t^s\}$ be the "maximum entropy probability distribution over $\mathcal{V}|\mathcal{K}^s$ that is consistent with $\mathcal{B}_t^s$". Given an agent with $\mathcal{K}^s$ and $\mathcal{B}_t^s$, *maximum entropy inference* states that the *derived sentence probability* for any sentence, $\sigma \in \mathcal{L}$, is:

$$(\forall \sigma \in \mathcal{L})\mathbb{P}_{\{\overline{W}|\mathcal{K}^s, \mathcal{B}_t^s\}}(\sigma) \triangleq \sum_n \{\, p_n \;:\; \sigma \; is \top in \; v_n \,\} \tag{2}$$

From Eqn. 2, each belief imposes a linear constraint on the $\{p_i\}$. The maximum entropy distribution: $\arg\max_{\underline{p}} \mathbb{H}(\underline{p})$, $\underline{p} = (p_1, \ldots, p_N)$, subject to $M + 1$ linear constraints:

$$g_j(\underline{p}) = \sum_{i=1}^{N} c_{ji} p_i - \mathbb{B}(\Omega_j) = 0, \quad j = 1, \ldots, M. \;\; g_0(\underline{p}) = \sum_{i=1}^{N} p_i - 1 = 0$$

where $c_{ji} = 1$ if Ω_j is $\top$ in v_i and 0 otherwise, and $p_i \geq 0, i = 1, \ldots, N$, is found by introducing Lagrange multipliers, and then obtaining a numerical solution using the multivariate Newton-Raphson method. In the subsequent subsections we'll see how an agent updates the sentence probabilities depending on the type of information used in the update.

Given a prior probability distribution $\underline{q} = (q_i)_{i=1}^n$ and a set of constraints C, the *principle of minimum relative entropy* chooses the posterior probability distribution $\underline{p} = (p_i)_{i=1}^n$ that has the least *relative entropy*[3] with respect to $\underline{q}$:

$$\{\underline{W}|\underline{q}, C\} \triangleq \arg\min_{\underline{p}} \sum_{i=1}^{n} p_i \log \frac{p_i}{q_i}$$

and that satisfies the constraints. This may be found by introducing Lagrange multipliers as above. Given a prior distribution $\underline{q}$ over $\{v_i\}$ — the set of all possible worlds, and a set of constraints C (that could have been derived as above from a set of new beliefs) *minimum relative entropy inference* states that the derived sentence probability for any sentence, $\sigma \in \mathcal{L}$, is:

$$(\forall \sigma \in \mathcal{L})\mathbb{P}_{\{\underline{W}|\underline{q}, C\}}(\sigma) \triangleq \sum_n \{\, p_n \;:\; \sigma \; is \top in \; v_n \,\} \tag{3}$$

where $\{p_i\} = \{\underline{W}|\underline{q}, C\}$. The principle of minimum relative entropy is a generalization of the principle of maximum entropy. If the prior distribution $\underline{q}$ is uniform, then the relative entropy of $\underline{p}$ with respect to $\underline{q}$, $\underline{p}\|\underline{q}$, differs from $-\mathbb{H}(\underline{p})$ only by a constant. So the principle of maximum entropy is equivalent to the principle of minimum relative entropy with a uniform prior distribution.

3.3 The agent manages information

The illocutions in the communication language $\mathcal{C}$ include information, [*info*]. The information received from general information sources will be expressed in terms defined by Π's ontology. We assume that Π makes at least part of that ontology public so that the other agents $\{\Omega_1, \ldots, \Omega_o\}$ may communicate [*info*] that Π can understand. Ω's *reliability* is an estimate of the extent to which this [*info*] is correct. For example, Ω may send Π the [*info*] that "the price of fish will go up by 10% next week", and it may actually go up by 9%.

The only restriction on incoming [*info*] is that it is expressed in terms of the ontology — this is very general. However, the way in which [*info*] is used is completely specific — it will be represented as a set of linear constraints on one or more probability distributions. A chunk of [*info*] may not be directly related to one of Π's chosen distributions or may not be expressed naturally as constraints, and so some inference machinery is required to derive these constraints — this inference is performed by model building functions, J_s, that have been activated by a plan s chosen by Π. $J_s^D([\mathit{info}])$ denotes the set of constraints on distribution D derived by J_s from [*info*].

3.3.1 Updating the world model with [*info*]

The procedure for updating the world model as [*info*] is received follows. If at time u, Π receives a message containing [*info*] it is time-stamped and source-stamped $[\mathit{info}]_{(\Omega,\Pi,u)}$, and placed in a repository $\mathcal{Y}^t$. If Π has an active plan, s, with model building function, J_s, then J_s is applied to $[\mathit{info}]_{(\Omega,\Pi,u)}$ to derive constraints on some,

[3]Otherwise called *cross entropy* or the *Kullback-Leibler* distance between the two probability distributions.

or none, of Π's distributions. The extent to which those constraints are permitted to effect the distributions is determined by a value for the *reliability* of Ω, $R^t(\Pi, \Omega, O([\textit{info}]))$, where $O([\textit{info}])$ is the ontological context of [*info*].

An agent may have models of integrity decay for some particular distributions, but general models of integrity decay for, say, a chunk of information taken at random from the World Wide Web are generally unknown. However the values to which decaying integrity should tend in time *are* often known. For example, a prior value for the truth of the proposition that a "22 year-old male will default on credit card repayment" is well known to banks. If Π attaches such prior values to a distribution D they are called the *decay limit distribution* for D, $(d_i^D)_{i=1}^n$. No matter how integrity of [*info*] decays, in the absence of any other relevant information it should decay to the decay limit distribution. If a distribution with n values has no decay limit distribution then integrity decays to the maximum entropy value $\frac{1}{n}$. In other words, the maximum entropy distribution is the default decay limit distribution.

In the absence of new [*info*] the integrity of distributions decays. If $D = (q_i)_{i=1}^n$ then we use a geometric model of decay:

$$q_i^{t+1} = (1 - \rho^D) \times d_i^D + \rho^D \times q_i^t, \text{ for } i = 1, \ldots, n \tag{4}$$

where $\rho^D \in (0, 1)$ is the decay rate. This raises the question of how to determine ρ^D. Just as an agent may know the decay limit distribution it may also know something about ρ^D. In the case of an information-overfed agent there is no harm in conservatively setting ρ^D "a bit on the low side" as the continually arriving [*info*] will sustain the estimate for D.

We now describe how new [*info*] is imported to the distributions. A single chunk of [*info*] may effect a number of distributions. Suppose that a chunk of [*info*] is received from Ω and that Π attaches the epistemic belief probability $R^t(\Pi, \Omega, O([\textit{info}]))$ to it. Each distribution models a facet of the world. Given a distribution $D^t = (q_i^t)_{i=1}^n$, q_i^t is the probability that the possible world ω_i for D is the true world for D. The effect that a chunk [*info*] has on distribution D is to enforce the set of linear constraints on D, $J_s^D([\textit{info}])$. If the constraints $J_s^D([\textit{info}])$ are taken by Π as valid then Π could update D to the posterior distribution $(p_i^{[\textit{info}]})_{i=1}^n$ that is the distribution with least relative entropy with respect to $(q_i^t)_{i=1}^n$ satisfying the constraint:

$$\sum_i \{p_i^{[\textit{info}]} \ : \ J_s^D([\textit{info}]) \text{ are all } \top \text{ in } \omega_i\} = 1. \tag{5}$$

But $R^t(\Pi, \Omega, O([\textit{info}])) = r \in [0, 1]$ and Π should only treat the $J_s^D([\textit{info}])$ as valid if $r = 1$. In general r determines the extent to which the effect of [*info*] on D is closer to $(p_i^{[\textit{info}]})_{i=1}^n$ or to the prior $(q_i^t)_{i=1}^n$ distribution by:

$$p_i^t = r \times p_i^{[\textit{info}]} + (1 - r) \times q_i^t \tag{6}$$

But, we should only permit a new chunk of [*info*] to influence D if doing so gives us new information. For example, if 5 minutes ago a trusted agent advises Π that the interest rate will go up by 1%, and 1 minute ago a very unreliable agent advises Π that the interest rate may go up by 0.5%, then the second unreliable chunk should not be permitted to 'overwrite' the first. We capture this by only permitting a new chunk of [*info*] to be imported if the resulting distribution has more information *relative to* the decay limit distribution than the existing distribution has. Precisely, this is measured using the Kullback-Leibler distance measure[4], and [*info*] is only used if:

$$\sum_{i=1}^n p_i^t \log \frac{p_i^t}{d_i^D} > \sum_{i=1}^n q_i^t \log \frac{q_i^t}{d_i^D} \tag{7}$$

In addition, we have described in Eqn. 4 how the integrity of each distribution D will decay in time. Combining these two into one result, distribution D is revised to:

$$q_i^{t+1} = \begin{cases} (1 - \rho^D) \times d_i^D + \rho^D \times p_i^t & \textit{if } \text{usable } [\textit{info}] \text{ is received at time } t \\ (1 - \rho^D) \times d_i^D + \rho^D \times q_i^t & \textit{otherwise} \end{cases}$$

for $i = 1, \cdots, n$, and decay rate ρ^D as before. We have yet to estimate $R^t(\Pi, \Omega, O([\textit{info}]))$ — that is described in Sec. 3.3.2 following.

[4]This is just one criterion for determining whether the [*info*] should be used.

ISBN: 972-8924-16-X © 2006 IADIS

3.3.2 Information reliability

We estimate $R^t(\Pi, \Omega, O([\textit{info}]))$ by measuring the error in information. Π's plans will have constructed a set of distributions. We measure the 'error' in information as the error in the effect that information has on each of Π's distributions. Suppose that a chunk of [*info*] is received from agent Ω at time s and is verified at some later time t. For example, a chunk of information could be "the interest rate will rise by 0.5% next week", and suppose that the interest rate actually rises by 0.25% — call that correct information [*fact*]. What does all this tell agent Π about agent Ω's reliability? Consider one of Π's distributions D that is $\{q_i^s\}$ at time s. Let $(p_i^{[\textit{info}]})_{i=1}^n$ be the minimum relative entropy distribution given that [*info*] has been received as calculated in Eqn. 5, and let $(p_i^{[\textit{fact}]})_{i=1}^n$ be that distribution if [*fact*] had been received instead. Suppose that the reliability estimate for distribution D was R_D^s. This section is concerned with what R_D^s should have been in the light of knowing *now*, at time t, that [*info*] should have been [*fact*], and how that knowledge effects our current reliability estimate for D, $R^t(\Pi, \Omega, O([\textit{info}]))$.

The idea of Eqn. 6, is that the current value of r should be such that, *on average*, $(p_i^s)_{i=1}^n$ will be seen to be "close to" $(p_i^{[\textit{fact}]})_{i=1}^n$ when we eventually discover [*fact*] — no matter whether or not [*info*] was used to update D, as determined by the acceptability test in Eqn. 7 at time s. That is, given [*info*], [*fact*] and the prior $(q_i^s)_{i=1}^n$, calculate $(p_i^{[\textit{info}]})_{i=1}^n$ and $(p_i^{[\textit{fact}]})_{i=1}^n$ using Eqn. 5. Then the *observed reliability* for distribution D, $R_D^{([\textit{info}]|[\textit{fact}])}$, on the basis of the verification of [*info*] with [*fact*] is the value of r that minimises the Kullback-Leibler distance between $(p_i^s)_{i=1}^n$ and $(p_i^{[\textit{fact}]})_{i=1}^n$:

$$\arg\min_r \sum_{i=1}^n (r \cdot p_i^{[\textit{info}]} + (1-r) \cdot q_i^s) \log \frac{r \cdot p_i^{[\textit{info}]} + (1-r) \cdot q_i^s}{p_i^{[\textit{fact}]}}$$

If $E^{[\textit{info}]}$ is the set of distributions that [*info*] effects, then the overall *observed reliability* on the basis of the verification of [*info*] with [*fact*] is: $R^{([\textit{info}]|[\textit{fact}])} = 1 - (\max_{D \in E^{[\textit{info}]}} |1 - R_D^{([\textit{info}]|[\textit{fact}])}|)$. Then for each ontological context o_j, at time t when, perhaps, a chunk of [*info*], with $O([\textit{info}]) = o_k$, may have been verified with [*fact*]:

$$R^{t+1}(\Pi, \Omega, o_j) = (1-\rho) \times R^t(\Pi, \Omega, o_j) + \rho \times R^{([\textit{info}]|[\textit{fact}])} \times \mathrm{Sem}(o_j, o_k)$$

where $\mathrm{Sem}(\cdot, \cdot) : O \times O \to [0, 1]$ measures the semantic distance between two sections of the ontology, and ρ is the learning rate. Over time, Π notes the ontological context of the various chunks of [*info*] received from Ω and over the various ontological contexts calculates the relative frequency, $P^t(o_j)$, of these contexts, $o_j = O([\textit{info}])$. This leads to a overall expectation of the *reliability* that agent Π has for agent Ω:

$$R^t(\Pi, \Omega) = \sum_j P^t(o_j) \times R^t(\Pi, \Omega, o_j)$$

3.4 Valuing Information

A chunk of information is valued first by the way that it enables Π to do something. So information is valued in relation to the strategies that Π is executing. A strategy, s, is chosen for a particular goal g in the context of a particular representation, or environment, e. One way in which a chunk of information assists Π is by altering s's world model W_s^t. A model W_s^t consists of a set of probability distributions: $W_s^t = \{D_{s,i}^t\}_{i=1}^n$. As a chunk of information could be "good" for one distribution and "bad" for another, we first value information by its effect on each distribution. For a model W_s^t, the *value* to W_s^t of a message received at time t is the resulting decrease in entropy in the distributions $\{D_{s,i}^t\}$. In general, suppose that a set of stamped messages $X = \{x_i\}$ is received in $\mathcal{X}$. The *information* in X at time t with respect to a particular distribution $D_{s,i}^t \in W_s^t$, strategy s, goal g and environment e is:

$$\mathbb{I}(X \mid D_{s,i}^t, s, g, e) \triangleq \mathbb{H}(D_{s,i}^t(\mathcal{Y}^t)) - \mathbb{H}(D_{s,i}^t(\mathcal{Y}^t \cup I(X)))$$

for $i = 1, \cdots, n$, where the argument of the $D_{s,i}^t(\cdot)$ is the state of Π's repository from which $D_{s,i}^t$ was derived. The environment e could be determined by a need ν (if the evaluation is made in the context of a particular negotiation) or a relationship ρ (in a broader context). It is reasonable to aggregate the information in X over the distributions used by s. That is, the information in X at time t with respect to strategy s, goal g and environment e is:

$$\mathbb{I}(X \mid s, g, e) \triangleq \sum_i \mathbb{I}(X \mid D_{s,i}^t, s, g, e)$$

and to aggregate again over all strategies to obtain the value of the information in a statement. That is, the *value of the information* in X with respect to goal g and environment e is:

$$\mathbb{I}(X \mid g, e) \triangleq \sum_{s \in \mathcal{S}(g)} \mathbb{P}(s) \cdot \mathbb{I}(X \mid s, g, e)$$

where $\mathbb{P}(s)$ is a distribution over the set of strategies for goal g, $\mathcal{S}(g)$, denoting the probability that strategy s will be chosen for goal g based on historic frequency data. and to aggregate again over all goals to obtain the (potential) information in a statement. That is, the *potential information* in X with respect to environment e is:

$$\mathbb{I}(X \mid e) \triangleq \sum_{g \in \mathcal{G}} \mathbb{P}(g) \cdot \mathbb{I}(X \mid g, e) \tag{8}$$

where $\mathbb{P}(g)$ is a distribution over $\mathcal{G}$ denoting the probability that strategy g will be triggered based on historic frequency data.

4 CONCLUSIONS

A demonstrable prototype e-Market system permits both human and software agents to trade with each other on the World Wide Web. The main contributions described are: the broadly-based and "focussed" data mining systems, and the intelligent agent architecture founded on information theory. These two technologies combine to present our vision of the trading agent of tomorrow.

REFERENCES

J. L. Arcos, et al. (2005). 'Environment Engineering for MultiAgent Systems'. *Journal on Engineering Applications of Artificial Intelligence* **18**.

A. Bogdanovych, et al. (2005). 'Narrowing the gap between humans and agents in e-Commerce: 3D electronic institutions'. In K. Bauknecht, B. Pröll, & H. Werthner (eds.), *E-Commerce and Web Technologies, Proceedings of the 6th International Conference, EC-Web 2005*, pp. 128–137, Copenhagen, Denmark. Springer-Verlag: Heidelberg, Germany.

J. Debenham (2004a). 'Auctions and Bidding with Information'. In P. Faratin & J. Rodriguez-Aguilar (eds.), *Proceedings Agent-Mediated Electronic Commerce VI: AMEC*, pp. 15 – 28.

J. Debenham (2004b). 'Bargaining with Information'. In N. Jennings, C. Sierra, L. Sonenberg, & M. Tambe (eds.), *Proceedings Third International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2004*, pp. 664 – 671. ACM.

EIDE (2005). 'Electronic Institution Development Environment: `http://e-institutor.iiia.csic.es/` '.

P. Faratin, et al. (2003). 'Using Similarity Criteria to Make Issue Trade-Offs in Automated Negotiation'. *Journal of Artificial Intelligence* **142**(2):205–237.

E. Jaynes (2003). *Probability Theory — The Logic of Science*. Cambridge University Press.

D. MacKay (2003). *Information Theory, Inference and Learning Algorithms*. Cambridge University Press.

M. Ramoni & P. Sebastiani (2003). *Intelligent Data Analysis*, chap. Bayesian methods, pp. 132–168. Springer-Verlag: Heidelberg, Germany.

D. Reis, et al. (2004). 'Automatic web news extraction using tree edit distance'. In *Proceedings of the 13'th International Conference on the World Wide Web*, pp. 502–511, New York.

C. Sierra & J. Debenham (2005). 'An Information-Based model for Trust'. In F. Dignum, V. Dignum, S. Koenig, S. Kraus, M. Singh, & M. Wooldridge (eds.), *Proceedings Fourth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2005*, pp. 497 – 504, Utrecht, The Netherlands. ACM Press, New York.

S. Simoff & J. Debenham (2002). 'Curious Negotiator'. In S. O. M. Klusch & O. Shehory (eds.), *proceedings 6th International Workshop Cooperative Information Agents VI CIA2002*, pp. 104–111, Madrid, Spain. Springer-Verlag: Heidelberg, Germany.

D. Zhang & S. Simoff (2004). 'Informing the Curious Negotiator: Automatic news extraction from the Internet'. In *Proceedings 3'rd Australasian Data Mining Conference*, pp. 55–72, Cairns, Australia.

D. Zhang, et al. (2005). 'Exchange rate modelling using news articles and economic data'. In *Proceedings of The 18th Australian Joint Conference on Artificial Intelligence*, Sydney, Australia. Springer-Verlag: Heidelberg, Germany.