

© [2008] IEEE. Reprinted, with permission, from [Debenham, John., Lawrence, Elaine. Automating contract negotiation. *International Conference on Information Technology*, April 2008]. This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Technology, Sydney's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to [pubs-permissions@ieee.org](mailto:pubs-permissions@ieee.org). By choosing to view this document, you agree to all provisions of the copyright laws protecting it

# Automating Contract Negotiation

John Debenham & Elaine Lawrence  
Faculty of Information Technology  
University of Technology, Sydney  
PO Box 123 Broadway  
NSW 2007, Australia  
email: {debenham, elaine}@it.uts.edu.au

## ABSTRACT

The automation of contract negotiation requires intelligent agents that can assimilate and use real-time information flows wisely. Electronic markets are information-rich with access to the Internet and the World Wide Web. A new breed of “information-based” agents are founded on concepts from information theory, and are designed to operate with information flows of varying and questionable integrity. These agents are part of a larger project that aims to make informed automated trading in applications such as eProcurement a reality.

## KEY WORDS

Intelligent agents, negotiation, argumentation, data mining.

## 1 Introduction

Electronic trading environments are awash with information, including information drawn from general resources such as the World Wide Web using smart retrieval technology [1]. Despite the substantial advances in multiagent systems and automated negotiation [2], it is perhaps surprising that negotiation in electronic business [3] remains a substantially manual procedure. We propose that rather than strive to make strategic, economically rational decisions, intelligent agents in electronic markets should capitalise on the real-time information flows and should aim to make ‘informed decisions’ that take account of the integrity of all relevant information. Traditional agent architectures, such as BDI, do not address directly the management of dynamic information flows of questionable integrity. This paper describes an agent architecture that has been designed specifically to operate in tandem with information discovery systems. This is part of our e-Market Framework that is available on the World Wide Web<sup>1</sup>. This framework aims to make informed automated trading a reality, and aims to address the realities of electronic business — “its what you know that matters”. This work does not address all of the issues in automated trading [3]. For example, the work relies on developments in: XML and semantic web, secure data exchange, value chain management and financial services. Further the design of electronic marketplaces is not described here.

<sup>1</sup><http://e-markets.org.au>

Intelligent agents that are built on an architecture designed specifically to handle real-time information flows are described in Sec. 2. The way in which these agents manage dynamic information flows is described in Sec. 3. The interaction of more than one of these agents engaging in competitive negotiation is described in Sec. 4. Sec. 5 concludes.

## 2 Information-Based 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. The central issues of trust in the execution of contracts is discussed in [4] [5]. 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.

### 2.1 The foundation of information-based agency

This section provides the rationale for the formal work that follows.

*Percepts*, the content of messages, are all that an agent has to inform it about the world and other agents. The validity of percepts will always be uncertain due to the reliability of the sender of the message, and to the period that has elapsed since the message arrived. Further, the belief that an agent has in the validity of a percept will be determined by the agent’s level of individual *caution*. The information-based agent’s *world model* is deduced from the percepts using *inference rules* that transform percepts into statements in probabilistic logic.

The integrity of percepts decreases in time. The way in which it decreases will be determined by the type of the percept, as well as by the issues listed in the previous paragraph including ‘caution’. An agent may have background knowledge concerning the expected integrity of a percept as  $t \rightarrow \infty$ . Information-based agents represent this background knowledge as a *decay limit distribution*. If the background knowledge is incomplete then one possibility for an agent is to assume that the decay limit distribution has maximum entropy whilst being consistent with the data.

All messages are valueless unless their integrity can be verified to some degree at a later time, perhaps for a cost.

To deal with this issue we employ an *institution agent* that always reports promptly and honestly on the execution of all commitments, forecasts, promises and obligations. This provides a conveniently simple solution to the integrity verification issue. The institution agent also takes care of “who owns what”. This enables the agents to negotiate and to evaluate the execution of commitments by simple message passing.

An agent’s percepts generally constitute a sparse data set whose elements have differing integrity. An agent may wish to induce tentative conclusions from this sparse and uncertain data of changing integrity. Percepts are transformed by inference rules into statements in probabilistic logic as described above. Information-based agents may employ entropy-based logic [6] to induce complete probability distributions from those statements. This logic is consistent with the laws of probability, but the results derived assume that the data is ‘all that there is to know’ — Watt’s Assumption.

Information is strategic. It has value as information. With the exception of highly cautious messages such as “make me an offer” in simple bargaining, *everything* that an agent says gives away information. This is a central issue here — the exchange of information is a strategic component of competitive interaction. Information-based agents evaluate information received as the reduction in entropy of the agent’s current world model. Information transmitted is evaluated as the agent’s expectation of the reduction of entropy in the recipient’s world model (by assuming that the recipient’s reasoning apparatus mirrors its own).

An agent will form expectations of other agents’ behaviours by observing the difference between commitments encapsulated in contracts, promises and statements of intent, and their subsequent execution. These observations may take account of whether those differences, if any, are ‘good’ or ‘bad’ for the agent.

An agent acts in response to some *need* or needs. A need may be exogenous such as the agent’s ‘owner’ needs a bottle of wine, or a message from another agent offering to trade may trigger a latent need to trade profitably. A need may also be endogenous such as the agent deciding that it owns more wine than it requires. An agent may be attempting to satisfy a number of needs at any time, and may have expectations of its future needs. A need may involve acquiring some physical object, it may also be a need to acquire information or to develop some on-going relationship. Information-based agents have a planning system that is invoked to aim to satisfy needs — the planning system is flexible and explores a number of (possibly inter-related) options at the same time, and devotes resources (if necessary) to the most promising. One component of the world model will necessarily be an estimate of the expectation that some contract will subsequently satisfy a need in an acceptable way — this will be in the form of a probability distribution over some evaluation space that measures ‘the expectation of the acceptability of the execution of a contract with respect to the satisfaction of that need’.

## 2.2 Agent Architecture

An agent observes events in its environment and represents 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*. 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 summarise and process that information. Such an agent is called *information-based*.

An agent called  $\Pi$  is the subject of this discussion.  $\Pi$  engages in multi-issue negotiation with a set of other agents:  $\{\Omega_1, \dots, \Omega_o\}$ . The foundation for  $\Pi$ ’s operation is the information that is generated both by

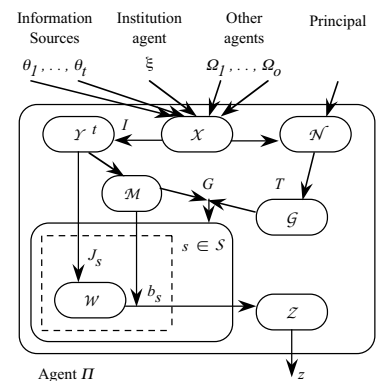


Figure 1. Architecture of agent  $\Pi$

and because of its negotiation exchanges. Any message from one agent to another reveals information about the sender.  $\Pi$  uses ideas from information theory to process and summarise its information. If  $\Pi$  *does* know its utility function *and* if it aims to optimise its utility *then*  $\Pi$  may apply the principles of game theory to achieve its aim. The information-based approach does not to reject utility optimisation — the selection of a strategy is secondary to the processing and summarising large amounts of uncertain information. In addition to the information derived from its opponents,  $\Pi$  has access to a set of information sources  $\{\Theta_1, \dots, \Theta_t\}$  that may include the marketplace in which trading takes place, and general information sources such as news-feeds accessed via the Internet. Together,  $\Pi$ ,  $\{\Omega_1, \dots, \Omega_o\}$  and  $\{\Theta_1, \dots, \Theta_t\}$  make up a multiagent system.

$\Pi$  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  $\Pi$ ’s epistemic belief in the truth of that sentence. Fig. 1 shows a high-level view of how  $\Pi$  operates. 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  $\Pi$  triggers a goal.

$\Pi$  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 €100 in exchange for a pallet of paper”, and second in response to some need,  $\nu \in \mathcal{N}$ , “we need to order some more paper”. In either case,  $\Pi$  is motivated by a need — either a need to strike a deal, or a general need to trade.  $\Pi$ ’s goals could be short-term such as obtaining some information “what is the euro / dollar exchange rate?”, 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  $\Pi$  has a trigger mechanism  $T$  where:  $T : \{\mathcal{X} \cup \mathcal{N}\} \rightarrow \mathcal{G}$ .

For each goal that  $\Pi$  commits to, it has a mechanism,  $G$ , for selecting a strategy to achieve it where  $G : \mathcal{G} \times \mathcal{M} \rightarrow \mathcal{S}$  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. 2.3 below].

**Maximum relative entropy inference.** 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. 2.3 below].

### 2.3 $\Pi$ ’s Reasoning

Once  $\Pi$  has selected a plan  $a \in \mathcal{A}$  it uses maximum entropy inference to derive the  $\{D_i^s\}_{i=1}^n$  [see Fig. 1] and minimum relative entropy inference to update those distributions as new data becomes available. *Entropy*,  $\mathbb{H}$ , is a measure of uncertainty [7] 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” [6] 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  $\Pi$ ’s language  $\mathcal{L}$ . A *possible world*,

$v$ , is a valuation function:  $\mathcal{G} \rightarrow \{\top, \perp\}$ .  $\mathcal{V}|\mathcal{K}^s = \{v_i\}$  is the set of all possible worlds that are consistent with  $\Pi$ ’s knowledge base  $\mathcal{K}^s$  that contains statements which  $\Pi$  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  $\Pi$ ’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:

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

The agent’s *belief set*  $\mathcal{B}_t^s = \{\Omega_j\}_{j=1}^M$  contains statements to which  $\Pi$  attaches a *given sentence probability*  $\mathbb{B}(\cdot)$ . 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:

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

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

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

$$g_0(\underline{p}) = \sum_{i=1}^N p_i - 1 = 0$$

where  $c_{ji} = 1$  if  $\Omega_j$  is  $\top$  in  $v_i$  and 0 otherwise, and  $p_i \geq 0$ ,  $i = 1, \dots, N$ , 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*<sup>2</sup> 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

<sup>2</sup>Otherwise called *cross entropy* or the *Kullback-Leibler* distance between the two probability distributions.

for any sentence,  $\sigma \in \mathcal{L}$ , is:

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

where  $\{p_i\} = \{W|q, C\}$ . The principle of minimum relative entropy is a generalisation of the principle of maximum entropy. If the prior distribution  $q$  is uniform, then the relative entropy of  $p$  with respect to  $q$ ,  $p||q$ , differs from  $-\mathbb{H}(p)$  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 Managing dynamic information flows

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  $\Pi$ 's ontology. We define an *ontology signature* as a tuple  $S = (C, R, \leq, \sigma)$  where  $C$  is a finite set of concept symbols (including basic data types);  $R$  is a finite set of relation symbols;  $\leq$  is a reflexive, transitive and anti-symmetric relation on  $C$  (a partial order); and,  $\sigma : R \rightarrow C^+$  is the function assigning to each relation symbol its arity. Concepts play the role of *type*, and the is-a hierarchy is the notion of subtype. Thus, type inference mechanisms can be used to type all symbols appearing in expressions. We assume that  $\Pi$  makes at least part of that ontology public so that the other agents  $\{\Omega_1, \dots, \Omega_o\}$  may communicate  $[info]$  that  $\Pi$  can understand.  $\Omega$ 's *reliability* is an estimate of the extent to which this  $[info]$  is correct.

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 in the world model. A chunk of  $[info]$  may not be directly related to one of  $\Pi$ '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  $\Pi$ .  $J_s^D([info])$  denotes the set of constraints on distribution  $D$  derived by  $J_s$  from  $[info]$ .

#### 3.1 Updating the world model with $[info]$

The procedure for updating the world model as  $[info]$  is received follows. If at time  $u$ ,  $\Pi$  receives a message containing  $[info]$  it is time-stamped and source-stamped  $[info]_{(\Omega, \Pi, u)}$ , and placed in a repository  $\mathcal{Y}^t$ . If  $\Pi$  has an active plan,  $s$ , with model building function,  $J_s$ , then  $J_s$  is applied to  $[info]_{(\Omega, \Pi, u)}$  to derive constraints on some, or none, of  $\Pi$ 's distributions. The extent to which those constraints are permitted to effect the distributions is determined by a value for the *reliability* of  $\Omega$ ,  $R^t(\Pi, \Omega, O([info]))$ , where  $O([info])$  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  $\Pi$  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:

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

where  $\rho^D \in (0, 1)$  is the decay rate. This raises the question of how to determine  $\rho^D$ . Just as an agent may know the decay limit distribution it may also know something about  $\rho^D$ . In the case of an information-overfed agent there is no harm in conservatively setting  $\rho^D$  “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  $\Omega$  and that  $\Pi$  attaches the epistemic belief probability  $R^t(\Pi, \Omega, O([info]))$  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  $\omega_i$  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  $D$ ,  $J_s^D([info])$ . If the constraints  $J_s^D([info])$  are taken by  $\Pi$  as valid then  $\Pi$  could update  $D$  to the posterior distribution  $(p_i^{[info]})_{i=1}^n$  that is the distribution with least relative entropy with respect to  $(q_i^t)_{i=1}^n$  satisfying the constraint:

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

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

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

*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  $\Pi$  that the interest rate will go up by 1%, and 1 minute ago a very unreliable agent advises  $\Pi$  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 — this is just one criterion for determining whether the  $[info]$  should be used — and  $[info]$  is only used if:

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

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:

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

for  $i = 1, \dots, n$ , and decay rate  $\rho^D$  as before.

### 3.2 Information reliability

Sec. 3.1 relies on an estimate of  $R^t(\Pi, \Omega, O([info]))$ . This estimate is constructed by measuring the ‘error’ in observed information as the error in the effect that information has on each of  $\Pi$ ’s distributions. Suppose that a chunk of  $[info]$  is received from agent  $\Omega$  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  $\Pi$  about agent  $\Omega$ ’s reliability? Consider one of  $\Pi$ ’s distributions  $D$  that is  $\{q_i^s\}$  at time  $s$ . Let  $(p_i^{[info]})_{i=1}^n$  be the minimum relative entropy distribution given that  $[info]$  has been received as calculated in Eqn. 5, and let  $(p_i^{[fact]})_{i=1}^n$  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  $D$ ,  $R^t(\Pi, \Omega, O([info]))$ .

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”  $(p_i^{[fact]})_{i=1}^n$  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  $(p_i^{[info]})_{i=1}^n$  and  $(p_i^{[fact]})_{i=1}^n$  using Eqn. 5. Then the *observed reliability* for distribution  $D$ ,  $R_D^{([info]||[fact])}$ , 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  $(p_i^{[fact]})_{i=1}^n$ :

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

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

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

where  $\text{Sem}(\cdot, \cdot) : O \times O \rightarrow [0, 1]$  measures the semantic distance [8] between two sections of the ontology, and  $\rho$  is the learning rate. Over time,  $\Pi$  notes the ontological context of the various chunks of  $[info]$  received from  $\Omega$  and over the various ontological contexts calculates the relative frequency,  $P^t(o_j)$ , of these contexts,  $o_j = O([info])$ . This leads to an overall expectation of the *reliability* that agent  $\Pi$  has for agent  $\Omega$ :  $R^t(\Pi, \Omega) = \sum_j P^t(o_j) \times R^t(\Pi, \Omega, o_j)$ .

## 4 Negotiation

For illustration  $\Pi$ ’s communication language [9] is restricted to the illocutions: Offer( $\cdot$ ), Accept( $\cdot$ ), Reject( $\cdot$ ) and Withdraw( $\cdot$ ). The simple strategies that we will describe all use the same world model function,  $J_s$ , that maintains the following two probability distributions as their world model:

- $\mathbb{P}^t(\Pi \text{Acc}(\Pi, \Omega, \nu, \delta))$  — the strength of belief that  $\Pi$  has in the proposition that she should accept the proposal  $\delta = (a, b)$  from agent  $\Omega$  in satisfaction of need  $\nu$  at time  $t$ , where  $a$  is  $\Pi$ ’s commitment and  $b$  is  $\Omega$ ’s commitment.  $\mathbb{P}^t(\Pi \text{Acc}(\Pi, \Omega, \nu, \delta))$  is estimated from:
  1.  $\mathbb{P}^t(\text{Satisfy}(\Pi, \Omega, \nu, \delta))$  a subjective evaluation (the strength of belief that  $\Pi$  has in the proposition that the expected outcome of accepting the proposal will satisfy some of her needs).
  2.  $\mathbb{P}^t(\text{Fair}(\delta))$  an objective evaluation (the strength of belief that  $\Pi$  has in the proposition that the proposal is a “fair deal” in the open market).
  3.  $\mathbb{P}^t(\Pi \text{CanDo}(a))$  an estimate of whether  $\Pi$  will be able to meet her commitment  $a$  at contract execution time.

These three arrays of probabilities are estimated by importing relevant information,  $[info]$ , as described in Sec. 3.

- $\mathbb{P}^t(\Omega \text{Acc}(\beta, \alpha, \delta))$  — the strength of belief that  $\Pi$  has in the proposition that  $\Omega$  would accept the proposal  $\delta$  from agent  $\Pi$  at time  $t$ . Every time that  $\Omega$  submits a proposal she is revealing information about what she is prepared to accept, and every time she rejects a proposal she is revealing information about what she is *not* prepared to accept. Eg: having received

the stamped illocution  $\text{Offer}(\Omega, \Pi, \delta)_{(\Omega, \Pi, u)}$ , at time  $t > u$ ,  $\Pi$  may believe that  $\mathbb{P}^t(\Omega\text{Acc}(\Omega, \Pi, \delta)) = \kappa$  this is used as a constraint on  $\mathbb{P}^{t+1}(\Omega\text{Acc}(\cdot))$  which is calculated using Eqn. 3.

#### 4.1 Negotiation Strategies

An agent's strategy  $s$  is a function of the information  $\mathcal{Y}^t$  that it has at time  $t$ . Four simple strategies make offers only on the basis of  $\mathbb{P}^t(\Pi\text{Acc}(\Pi, \Omega, \nu, \delta))$ ,  $\Pi$ 's acceptability threshold  $\gamma$ , and  $\mathbb{P}^t(\Omega\text{Acc}(\Omega, \Pi, \delta))$ . The greedy strategy  $s^+$  chooses:

$$\arg \max_{\delta} \{ \mathbb{P}^t(\Pi\text{Acc}(\Pi, \Omega, \nu, \delta)) \mid \mathbb{P}^t(\Omega\text{Acc}(\Omega, \Pi, \delta)) \gg 0 \}$$

it is appropriate when  $\Pi$  believes  $\Omega$  is desperate to trade.

The *expected-acceptability-to- $\Pi$ -optimizing strategy*  $s^*$  chooses:

$$\arg \max_{\delta} \{ \mathbb{P}^t(\Omega\text{Acc}(\Omega, \Pi, \delta)) \times \mathbb{P}^t(\Pi\text{Acc}(\Pi, \Omega, \nu, \delta)) \mid \mathbb{P}^t(\Pi\text{Acc}(\Pi, \Omega, \nu, \delta)) \geq \gamma \}$$

when  $\Pi$  is confident and not desperate to trade. The strategy  $s^-$  chooses:

$$\arg \max_{\delta} \{ \mathbb{P}^t(\Omega\text{Acc}(\Omega, \Pi, \delta)) \mid \mathbb{P}^t(\Pi\text{Acc}(\Pi, \Omega, \nu, \delta)) \geq \gamma \}$$

it optimises the likelihood of trade — when  $\Pi$  is keen to trade without compromising its own standards of acceptability.

An approach to issue-tradeoffs is described in [10]. The bargaining strategy described there attempts to make an acceptable offer by “walking round” the iso-curve of  $\Pi$ 's previous offer  $\delta'$  (that has, say, an acceptability of  $\gamma_{\delta'} \geq \gamma$ ) towards  $\Omega$ 's subsequent counter offer. In terms of the machinery described here, an analogue is to use the strategy  $s^-$ :

$$\arg \max_{\delta} \{ \mathbb{P}^t(\Omega\text{Acc}(\Omega, \Pi, \delta)) \mid \mathbb{P}^t(\Pi\text{Acc}(\Pi, \Omega, \nu, \delta)) \geq \gamma_{\delta'} \}$$

with  $\gamma = \gamma_{\delta'}$ . This is reasonable for an agent that is attempting to be accommodating without compromising its own interests. The complexity of the strategy in [10] is linear with the number of issues. The strategy described here does not have that property, but it benefits from using  $\mathbb{P}^t(\Omega\text{Acc}(\Omega, \Pi, \delta))$  that contains foot prints of the prior offer sequence — estimated by repeated use of Eqn. 3 — in that distribution more recent data gives estimates with greater certainty.

## 5 Conclusions

The key to automating trading is to build intelligent agents that are ‘informed’, that can proactively acquire information to reduce uncertainty, that can estimate the integrity of real-time information flows, and can use uncertain information as a foundation for strategic decision-making

[3]. An ‘information-based’ agent architecture has been described, that is founded on ideas from information theory, and has been developed specifically for this purpose. This work is in collaboration with “Institut d’Investigacio en Intel.ligencia Artificial”, Spanish Scientific Research Council, UAB, Barcelona, Spain.

## References

- [1] Zhang, D., Simoff, S.: Informing the Curious Negotiator: Automatic news extraction from the Internet. In Williams, G., Simoff, S., eds.: Data Mining: Theory, Methodology, Techniques, and Applications. Springer-Verlag: Heidelberg, Germany (2006) 176 – 191
- [2] Rosenschein, J.S., Zlotkin, G.: Rules of Encounter. The MIT Press, Cambridge, USA (1994)
- [3] Lawrence, E., Newton, S., Corbitt, B., Lawrence, J., Dann, S., Thanasankit, T.: Internet Commerce — Digital Models for Business. 3rd edn. John Wiley and Sons, Inc. (2003)
- [4] Sierra, C., Debenham, J.: Trust and honour in information-based agency. In Stone, P., Weiss, G., eds.: Proceedings Fifth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2006, Hakodate, Japan, ACM Press, New York (2006) 1225 – 1232
- [5] Debenham, J., Simoff, S.: An e-Market Framework for Informed Trading. In Carr, L., Roure, D.D., Iyengar, A., Goble, C., Dahlin, M., eds.: proceedings 15th International World Wide Web Conference, WWW-2006, Edinburgh, Scotland (2006)
- [6] Jaynes, E.: Probability Theory — The Logic of Science. Cambridge University Press (2003)
- [7] MacKay, D.: Information Theory, Inference and Learning Algorithms. Cambridge University Press (2003)
- [8] Li, Y., Bandar, Z.A., McLean, D.: An approach for measuring semantic similarity between words using multiple information sources. IEEE Transactions on Knowledge and Data Engineering **15** (2003) 871 – 882
- [9] Rahwan, I., Ramchurn, S., Jennings, N., McBurney, P., Parsons, S., Sonenberg, E.: Argumentation-based negotiation. Knowledge Engineering Review **18** (2003) 343–375
- [10] Debenham, J., Lawrence, E.: Intelligent agents that make informed decisions. In: proceedings 16th International Symposium on Methodologies for Intelligent Systems, ISMIS 2006, Springer Verlag (2006) 137–146