

An Information-Based model for Trust

Carles Sierra
Institut d'Investigació en Intel·ligència Artificial
Spanish Scientific Research Council, UAB
08193 Bellaterra, Catalonia, Spain
sierra@iia.csic.es

John Debenham
Faculty of Information Technology
University of Technology, Sydney
NSW, Australia
debenham@it.uts.edu.au

ABSTRACT

In this paper we present a rich decision model for intelligent agents involved in multi issue negotiations. The model, grounded on information theory, takes into account the aspects of trust and preference to devise mechanisms to manage dialogues. The model supports the design of agents that aim to take 'informed decisions' taking into account that which they have actually observed.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—Multiagent systems

General Terms

Theory

Keywords

Multiagent systems, Trust

1. INTRODUCTION

Negotiation is an information exchange process as well as a proposal exchange process. If an agent proposes a deal then it reveals information about what it is prepared to accept. If it rejects a deal then it reveals information about what is not acceptable to it. If it transmits preference information then it reveals information about its general willingness to accept deals. If it transmits information that is intended to influence its opponent's preferences for a deal then it reveals information both about what it believes its opponent knows and about what information it believes to be pertinent for that deal. Anything that an agent communicates, or fails to communicate, during a negotiation gives away (valuable) information.

The information that is exchanged both directly as part of, and indirectly because of, the exchange of proposals may not be truthful. For example, an agent may reject an acceptable

proposal in an attempt to secure a more satisfactory deal. So the information received by an agent during a negotiation may be more or less true. Such information is only useful if an agent is prepared to attach a degree of belief to it. Further the integrity of information generally will decay in time, and these degrees of belief will reflect that decay.

Our negotiating agent, α , uses ideas from information theory to evaluate its negotiation information. If its opponent, β , communicates information, the value of that communication is the decrease in uncertainty in α 's model of β . One measure of this decrease in uncertainty is *Shannon information* [10], or negative entropy. If α communicates information it evaluates that information as its expectation of the resulting decrease in β 's uncertainty about α . Any such decrease in uncertainty is seen against the continually increasing uncertainty in information because information integrity necessarily decays in time. Information theory may also be used to measure features of inter-agent relationships that extend beyond single negotiations — for example, the sharing of information in a trading pact or relationship, and the strength of trading networks that form as a result of such information sharing. α uses entropy-based inference — both maximum entropy inference and minimum relative entropy inference — to derive probability distributions for that which it does not know in line with the following principle.

Information Principle. α 's information base contains only observed facts — in the absence of observed facts, α may speculate about what those facts might be. For example in competitive negotiation, β 's utility, deadlines, and other private information will never be observable — unless β is foolish. Further, β 's motivations (such as being a utility optimizer) will also never be observable. So in competitive negotiation β 's private information is "off α 's radar" — α does not contemplate it or speculate about it.

We assume that agents may observe the values of different types of environment variables: illocutions uttered by agents, the role played by agents, the result of the execution of contracts, and, in general, agents may observe other variables of the environment that could influence their opinion about the behaviour of other agents, e.g. market data, or general information from the Internet. We assume that the interactions between agents are made within the framework of an infrastructure that fixes ontology and meaning, for instance an Electronic Institution [1]. Thus, no differences in the observation of illocutions have to be assumed. Moreover, we assume that the role of a player is public information that can be observed — this is the case for instance if we are within an electronic institution framework. How-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'05, July 25-29, 2005, Utrecht, Netherlands.
Copyright 2005 ACM 1-59593-094-9/05/0007 ...\$5.00.

ever, the perception of the execution of a contract is subjective in the sense that two deviations of behaviour can be perceived differently by two different agents. We will therefore assume that each agent is equipped with a perception function (noted in this paper as $\text{Observe}(\cdot)$) that determines which contract execution has actually occurred. Based on these assumptions, in this paper we present the architecture of an agent based on an information theory approach and then propose a trust model based on both relative and conditional entropy.

In Sections 2 and 3 we motivate the need for modelling trust and preferences in the context of multiagent negotiation. In Section 4 we introduce a representation language to support negotiation dialogues that also permits the interchange of information. Dialogues are given an information-based semantics in Section 5 that is the base for the trust model presented in Section 6. We conclude with an appropriate agent architecture to support the model and a discussion section.

2. TRUST AND NEGOTIATION

Trust is a multi-faceted concept that has received increasing attention recently [11, 13, 12, 6]. In the context of negotiation, trust represents a general assessment on how ‘serious’ an agent is about the negotiation process, i.e. that his proposals ‘make sense’ and he is not ‘flying a kite’, and that he is committed to what he signs. A lack of trust may provoke agents to breakdown negotiations, or to demand additional guarantees to cover the risk of potential defections. Therefore, in any model of trust the central issue is how to model expectations about the actual outcome at contract execution time. Contracts, when executed, may, and frequently do, yield a different result to what was initially signed. Goods may be delivered late, quality may be (perceived) different from the contract specification, extra costs may be claimed, etc. So the outcome is uncertain to some extent, and trust, precisely, is a measure of how uncertain the outcome of a contract is. Naturally, the higher the trust in a partner the more sure we are of his or her reliability. Trust is therefore a *measure of expected deviations of behaviour* along a given dimension, and in many cases for a given value (region) in that dimension (e.g. I might trust you on low-priced contracts but not on high-priced ones). In this sense, the higher the trust the lower the expectation that a (significant) deviation from what is signed occurs. Trust values can be used for three basic purposes:

- Trust permits us to select *what offer to send next*. As trust measures the expected deviation of behaviour of our opponent, we can use it as a ‘counter-balance’ by asking for even better deals.
- Trust permits us to better *select negotiation partners*. Humans (normally) prefer to live in worlds with low uncertainty about the future. The higher the trust on a partner the less probable a deception is going to happen, even if what is to be expected is not extraordinarily good. This would also explain the fact that humans tend not to explore too much once they settle down in a situation that is satisfactory enough. That is, a situation in which they have a landscape of trust that permits them to select partners in a good enough way. If I’m satisfied with my butcher I will not explore more butchers unless I am disappointed at some point.

- Trust *simplifies negotiation dialogues*. Trust is mostly built through ongoing relationships, the repetition of negotiation dialogues and contract executions means that some terms and conditions (negotiation dimensions) need not be discussed and specified again and again. A high trust in what will happen along a certain dimension enables partners to omit it from the negotiation dialogue. The higher the trust the smaller the number of dimensions of the negotiation object, and thus the shorter the negotiation dialogue.

We will propose in this paper to base our measure of trust as the negative entropy of the probability distribution of possible outcomes for a given contract given that trust measures the relationship between commitment and execution of contracts. More precisely, between signed contracts and *perceived* execution of contracts. In this way, a natural way to base our modelling of trust is on a conditional probability, P^t , between contracts given a context e as:

$$P^t(\text{Observe}(\alpha, b') | \text{Accept}(\beta, \alpha, (a, b)), e)$$

where every contract execution represents a point in that distribution.¹ A concrete relation between a signed contract and the perception of an executed one.

3. PREFERENCES AND NEGOTIATION

Rhetorics in a negotiation context represents the use of language constructs to persuade the oponent to accept our proposals. Agents use rhetorics because they want to change their opponents’ preferences or their opponents’ view of them. Rhetorical constructs are intimately linked to *social structure*² and *time*³. The *amount* and the *persistence* of preference change induced by a rhetorical construct depends strongly on the distance in some social scale between the speaker and the hearer. The larger the distance the bigger the impact. The closer, the longer the effect. A Nobel Prize winner is able to change *a lot* our views on a certain subject although getting convinced by a peer on some matter usually has *longer* impact [9]. The impact of the use of rhetoric constructs in negotiation models is the following:

1. Rhetoric permits us to *change the preferences* of the opponent. The acceptance of a proposal depends on

¹To simplify notation in the rest of the paper we will denote $P^t(\text{Observe}(\alpha, b') | \text{Accept}(\beta, \alpha, (a, b)), e)$ as $P^t(b' | b, e)$.

²The social relationship that negotiators hold determines to a great extent how probable it is that a proposal will be accepted. By exerting *power* or *authority* a negotiator can achieve the acceptance of a proposal by the opponent that would otherwise be unacceptable, for instance if coming from a peer. Thus, an important element when considering the acceptance of a proposal is to analyse the *role* of our opponent and what can we expect to occur in future interactions. Social relations are, in a sense, built and maintained to permit an accurate assessment of what to expect from others, now but *mainly* in the future.

³Time limits are strategic elements. Sometimes they even become public information, and therefore there is total certainty on the outcome of a negotiation should the limit been reached. Often, the reasons for a breakdown are kept private and therefore we can only guess whether negotiation will continue or not. Negotiations are even undertaken although one party is decided to breakdown from the very beginning, whatever the proposals are. It is important to note that Negotiation is *temporally bounded* and that Negotiation time has an associated *cost*.

the set of beliefs held by the opponent. Depending on the social relationship between the agents, certain rhetoric constructs might provoke a belief update, and therefore a change of preferences.

2. Rhetoric permits us to *refine the opponent's model of ourselves*. By issuing critiques of proposals received, or putting questions, we are giving information to our opponent that may help him/her to build a more accurate model of our preferences. We can also simply inform our opponent about our preferences to perhaps shorten the dialogue by reducing the exploration space.
3. Rhetoric permits us to *maintain the level of trust*. For instance when a negotiation is not converging as usual, a negotiator may 'remind' the opponent of the level of trust he has built in the past as encouragement to maintain it.

Another dimension that is very important in the analysis of dialogues is ontology. The contents of illocutions determine whether our assumptions about the opponent's model of the problem are correct. New values can be added to a dimension by a simple question: "Have you got yellow plastic crocodiles? — Oops, I never thought yellow crocodiles existed!". A simplifying solution is to start any negotiation process by fixing a common ontology. Alternatively, we can use ontology clarifying dialogues and then be ready to modify our models during the dialogue. This seems more complex but more challenging as well.

4. A NEGOTIATION LANGUAGE

Agent α is negotiating with an opponent β . They aim to strike a deal $\delta = (a, b)$ where a is α 's commitment and b is β 's. (Where a or b might be empty.) We denote by A the set of all possible commitments by α , and by B the set of all possible commitments by β . The agents have two languages, \mathcal{C} for communication (illocutionary based) and \mathcal{L} for internal representation (as a restricted first-order language).⁴

In this paper we assume that the illocution particle set is:

$$\iota = \{\text{Offer, Accept, Reject, Withdraw, Inform}\}$$

with the following syntax and informal meaning:

- Offer(α, β, δ) Agent α offers agent β a deal $\delta = (a, b)$ with action commitments a for α and b for β .
- Accept(α, β, δ) Agent α accepts agent β 's previously offered deal δ .
- Reject($\alpha, \beta, \delta, [info]$) Agent α rejects agent β 's previously offered deal δ . Optionally, information explaining the reason for the rejection can be given.
- Withdraw($\alpha, \beta, [info]$) Agent α breaks down negotiation with β . Extra *info* justifying the withdrawal may be given here.
- Inform($\alpha, \beta, [info]$) Agent α informs β about *info*.

⁴It is commonly accepted since the works by Austin and Searle that illocutionary acts are actions that succeed or fail. We will abuse notation in this paper and will consider that they are predicates in a first order logic meaning 'the action has been performed'. For those more pure-minded an alternative is to consider dynamic logic.

The accompanying information, *info*, can be of two basic types: (i) referring to the process (plan) used by an agent to solve a problem, or (ii) data (beliefs) of the agent including preferences. When negotiating, agents will therefore try to influence the opponent by trying to change their processes (plans) or by providing new data.

Dialogues, especially in Electronic Institutions [1], tend to be structured in order to facilitate the decision making of the participants, and because some illocutions only make sense in a particular conversation context, e.g. I cannot accept an offer that has not been made. This structuring of dialogues is what is usually called a protocol. We will not detail it here, but will assume that such a protocol exists.

Following the extensive literature on preferences, preferences are divided into two classes:

Quantitative. These preferences are usually called *soft constraints* (hard constraints are particular cases of soft constraints). A soft constraint associates each instantiation of its variables with a value from a partially ordered set. One natural interpretation of this value is the probability of choice. In general, preferences can be expressed as values within a semi-ring $\langle A, +, \times, 0, 1 \rangle$ such that A is a set and $0, 1 \in A$; $+$ is commutative, associative with 0 as its unit element; \times is associative, distributes over $+$, 1 is its unit element, and 0 is its annihilating element. Given a semi-ring $\langle A, +, \times, 0, 1 \rangle$, an ordered set of variables $V = \{v_1, \dots, v_n\}$ and their corresponding domains $D = \{D_1, \dots, D_n\}$ a soft constraint is a pair $\langle f, con \rangle$ where $con \subseteq V$ and $f: D^{|con|} \rightarrow A$ with the following intuitive meaning: $f(d_1, \dots, d_n) = \kappa$ means that the binding $x_1 = d_1, \dots, x_n = d_n$ satisfies the constraint to a level of κ .⁵

Qualitative. In many domains it is difficult to formulate precise numerical preferences, and it is more convenient to express preference relations between variable assignments: "I prefer red cars to yellow cars". The usual way to represent this relationship formally is $v = a > v = a'$, or simpler $a > a'$, meaning that we prefer the assignment of variable v to a than to a' . Also, in case of an absolute preference for a particular value in a domain, that is, when our preference is $\forall x \neq a. v = a > v = x$ we can simply write $v = a$ or just a . Also, in many cases preference relations depend on the values assigned to other variables (configuring what is called a Conditional Preference Net (CP-net) [2]). "If they serve meat, I prefer red wine to white wine". A conditional preference can be represented as $v_1 = c : v_2 = d > v_2 = d'$ (or again $c : d > d'$) meaning that we prefer d to d' in the context where c is the case. In general, any DNF over value assignments could be used as the condition. And also, other comparatives than '=' could be used.

Finally, it seems natural that the constraints have an associated certainty degree representing their degree of truth. We thus propose the following content language expressed in BNF: ($info \in \mathcal{L}$):

$$info ::= unit[\mathbf{and} info]$$

⁵As we use maximum entropy inference we have to make the simplifying assumption that domains of quantitative constraints must be finite. This means that continuous domains must be represented as a finite set of intervals, further the way in which those intervals are chosen affects the outcome. This is sometime cited as a weakness of the maximum entropy approach. In [3] it is argued to the contrary, that the choice of intervals represents our prior expectations in fine-grained detail.

<i>unit</i>	::= $K B soft qual cond$
<i>K</i>	::= $\mathbf{K}(WFF)$
<i>B</i>	::= $\mathbf{B}(WFF)$
<i>soft</i>	::= $\mathbf{soft}(f, \{V^+\})$
<i>qual</i>	::= $V=D >V=D$
<i>cond</i>	::= $\mathbf{If DNF Then qual}$
<i>WFF</i>	::= <i>any wff over subsets of variables</i> $\{V\}$
<i>DNF</i>	::= <i>conjunction</i> $[\mathbf{or DNF}]$
<i>conjunction</i>	::= $qual[\mathbf{and conjunction}]$
<i>V</i>	::= $\mathbf{v}_1 \dots \mathbf{v}_n$
<i>D</i>	::= $\mathbf{a a' b} \dots$
<i>f</i>	::= <i>any function from the domains of subsets of V to a set A. For instance a fuzzy set membership function if $A = [0,1]$</i>

Examples could be: “I prefer red wine to white wine when served meat.” as $\text{Inform}(\alpha, \beta, \text{if Food} = \text{meat then Wine} = \text{red} > \text{Wine} = \text{white})$; “I prefer more money to less money” as $\text{Inform}(\alpha, \beta, \text{soft}(\tanh, \{\text{Money}\}))$; “I reject your offer since I definitely cannot pay more than 200\$” as $\text{Reject}(\alpha, \beta, \text{Money} = 250, \text{hard}(\text{Money} < 200\$, \{\text{Money}\}))$; and “I prefer red cars to yellow cars” as $\text{Inform}(\alpha, \beta, \text{if thing} = \text{car then Colour} = \text{Red} > \text{Colour} = \text{Yellow})$

5. INFORMATION-BASED NEGOTIATION

We ground our negotiation model on information-based concepts. *Entropy*, H , is a measure of uncertainty [10] in a probability distribution for a discrete random variable X : $H(X) \triangleq -\sum_i p(x_i) \log p(x_i)$ where $p(x_i) = P(X = x_i)$. Maximum entropy inference is used to derive sentence probabilities for that which is not known by constructing the “maximally noncommittal” [8] probability distribution.

Let \mathcal{G} be the set of all positive ground literals that can be constructed using our language \mathcal{L} . A *possible world*, v , is a valuation function: $\mathcal{G} \rightarrow \{\top, \perp\}$. $\mathcal{V}|\mathcal{K} = \{v_i\}$ is the set of all possible worlds that are consistent with an agent’s knowledge base \mathcal{K} that contains statements which the agent believes are true. A *random world* for \mathcal{K} , $W|\mathcal{K} = \{p_i\}$ is a probability distribution over $\mathcal{V}|\mathcal{K}^a = \{v_i\}$, where p_i expresses an agent’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 $W|\mathcal{K}$ is:

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

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

$$(\forall \sigma \in \mathcal{L}) P_{\{W|\mathcal{K}, \mathcal{B}\}}(\sigma) \triangleq \sum_n \{p_n : \sigma \text{ is } \top \text{ in } v_n\} \quad (2)$$

From Eqn. 2, each belief imposes a linear constraint on the $\{p_i\}$. The maximum entropy distribution: $\arg \max_{\underline{p}} H(\underline{p})$,

$\underline{p} = (p_1, \dots, p_N)$, subject to $M + 1$ linear constraints:

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

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

$c_{ji} = 1$ if φ_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.

5.1 Updating from decay and experience

An important aspect that we want to model is the fact that beliefs ‘evaporate’ as time goes by. If we don’t keep an ongoing relationship, we somehow forget how *good* the opponent was. If I stop buying from my butcher, I’m not sure anymore that he will sell me the ‘best’ meat. This decay is what justifies a continuous relationship between individuals. In our model, the conditional probabilities should tend to ignorance. If we have the set of observable contracts as $B = \{b_1, b_2, \dots, b_n\}$ then complete ignorance of the opponent’s expected behaviour means that given the opponent commits to b the conditional probability for each observable contract becomes $\frac{1}{n}$ — i.e. the unconstrained maximum entropy distribution. This natural decay of belief is offset by new observations.

We define the evolution of the probability distribution that supports the previous definition of decay using an equation inspired by pheromone like models [5]:

$$P^{t+1}(b'|b) = \kappa \cdot \left(\frac{1-\rho}{n} + \rho \cdot (P^t(b'|b) + \Delta^t P(b'|b)) \right) \quad (3)$$

where κ is a normalisation constant to ensure that the resulting values for $P^{t+1}(b'|b)$ are a probability distribution. This equation models the passage of time for a conveniently large $\rho \in [0, 1]$ and where the term $\Delta^t P(b'|b)$ represents the increment in an instant of time according to the last experienced event as the following possibilities show.

Similarity based. The question is how to use the observation of a contract execution c' given a signed contract c in the update of the overall probability distribution over the set of all possible contracts. Here we use the idea that given a particular deviation in a region of the space, *similar* deviations should be expected in other regions. The intuition behind the update is that if my butcher has not given me the quality that I expected when I bought lamb chops, then I might expect similar deviations with respect to chicken. This idea is built upon a function $f(x, y)$ that takes into account the difference between acceptance probabilities and similarity between the perception of the execution x of a contract y , that is a contract for which there was an $\text{Accept}(\beta, \alpha, y)$. Thus, after the observation of c' the increment of probability distribution at time $t + 1$ is:

$$\Delta^t P(b'|b) = (1 - |f(c', c) - f(b', b)|) \quad (4)$$

where $f(x, y)$ is

$$f(x, y) = \begin{cases} 1 & \text{if } P^t(\text{Accept}(x)) > P^t(\text{Accept}(y)) \\ \text{Sim}(x, y) & \text{otherwise.} \end{cases}$$

and where Sim is an appropriate similarity function (reflexive and symmetric) that determines the indistinguishability between the perceived and the committed contract.

Entropy based. Suppose that α observes the event $(c'|c)$, the entropy based approach estimates $\Delta^t P(b'|b)$ by applying the principle of minimum relative entropy.⁶ Let:

$$(P_C^t(b_j|b))_{j=1}^n = \arg \min_{\underline{p}} \sum_{i=1}^n p_i \log \frac{p_i}{P^t(b_i|b)} \quad (5)$$

satisfying the constraint C , and $\underline{p} = (p_j)_{j=1}^n$. Then:

$$\Delta^t P(b'|b) = P_C^t(b'|b) - P^t(b'|b) \quad (6)$$

Constraint C is specified as follows in three cases: first when $c = b$, second when $c' = c \neq b$, and third when $c' \neq c \neq b$.

First, if $c = b$ then C is: $P_C^t(b'|b) = P^t(b'|b) + \nu(1 - P^t(b'|b))$, for $\nu \in [0, 1]$ — the value of ν represents the strength of α 's belief that the probability that $(b'|b)$ will occur at time $t + 1$ should increase if $(b'|b)$ occurs at time t .

Second, if $c' = c \neq b$ then constraint C is:

$$P_C^t(b|b) = P^t(b|b) + g_1(b, c)(1 - P^t(b|b))$$

for: $g_1 \in [0, 1]$, where $g_1(b, c)$ represents the strength of α 's belief that the probability that $(b|b)$ will occur at time $t + 1$ should increase if $(c|c)$ occurs at time t .

Third, if $c' \neq c \neq b$ then suppose that c' is preferred to c by α then $h(c', c) = P^t(\text{Accept}(c')) - P^t(\text{Accept}(c)) > 0$. Let $B(b)^+ = \{x \mid h(x, b) > 0\}$, ie: the set of contract executions that α prefers to b . Given a signed contract b , the prior probability that the contract execution will be preferred by α to b is: $p(b)^+ = \sum_{x \in B(b)^+} P^t(x|b)$. After observing $(c'|c)$ we wish to increase the probability that a preferred execution will occur for contract b to: $p(b \mid (c'|c))^+ = p(b)^+ + g_2(b, c, c')(1 - p(b)^+)$, where $g_2(b, c, c')$ represents the strength of α 's belief that the probability that execution of contract b at time $t + 1$ will be preferred to b should increase if $(c'|c)$ occurs at time t . Constraint C then is: $\sum_{x \in B(b)^+} P_C^t(x|b) = p(b \mid (c'|c))^+$. Similarly, if c' is not preferred to c by α then construct $B(b)^-$.

5.2 Updating from preferences

[4] describes the application of maximum entropy inference to enable α to estimate $P^t(\text{Accept}(\beta, \alpha, \delta))$ the probability that β will accept deal δ from α in response to α transmitting the illocution $\text{Offer}(\alpha, \beta, \delta)$. This distribution is derived from previously observed $\text{Offer}(\beta, \alpha, \dots)$ and $\text{Reject}(\beta, \alpha, \dots)$

⁶Given a prior probability distribution $\underline{q} = (q_i)_{i=1}^n$ and a set of constraints, the *principle of minimum relative entropy* chooses the posterior probability distribution $\underline{p} = (p_i)_{i=1}^n$ that has the least relative entropy with respect to \underline{q} , $\arg \min_{\underline{p}} \sum_{i=1}^n p_i \log \frac{p_i}{q_i}$, and that satisfies the constraints. The principle of minimum relative entropy is a generalization of the principle of maximum entropy. If the prior distribution \underline{q} is uniform, the relative entropy of \underline{p} with respect to \underline{q} differs from $-H(\underline{p})$ only by a constant. So the principle of maximum entropy is equivalent to the principle of minimum relative entropy with a uniform prior distribution.

illocutions received from β — the former indicating readiness to accept and the latter readiness to reject. α may not accept this historic readiness as being definitive now, if so then $P^t(\text{Accept}(\beta, \alpha, \delta))$ is estimated by attaching time-discounted beliefs (as sentence probabilities) to these observations, and then by calculating the maximum entropy distribution subject to those probabilities as constraints.

Suppose that α now receives preference information from β in the form of an $\text{Inform}(\beta, \alpha, [\text{info}])$ illocution, and is prepared to accept this information into its belief set \mathcal{B} as a belief with sentence probability p_{info} — this probability may decay in time. How will this new information influence α 's estimate of $P^t(\text{Accept}(\beta, \alpha, \delta))$? Preference information induces a partial ordering on the set of deals. If deal δ_1 is preferred by β to deal δ_2 then: if $\text{Accept}(\beta, \alpha, \delta_2)$ α may conclude to certainty p_{info} that $\text{Accept}(\beta, \alpha, \delta_1)$.

As described in Sec. 4, preference illocutions generally refer to particular issues within deals — e.g. “I prefer red to yellow”. In general, “I prefer deals with property Q_1 to those with property Q_2 ” becomes the following constraint on the $P^t(\text{Accept}(\beta, \alpha, \delta))$ distribution:

$$p_{\text{info}} = \frac{\sum_{\delta: Q_1(\delta)} p_{\delta}}{(\sum_{\delta: Q_1(\delta)} p_{\delta}) + (\sum_{\delta: Q_2(\delta)} p_{\delta})}$$

the posterior distribution for $P^t(\text{Accept}(\beta, \alpha, \delta))$ is calculated by applying the principle of minimum relative entropy⁶ to it subject to this constraint.

The method of representing preference information above is quite general. Although if it is used to represent a preference ordering on an issue such as “ β prefers to pay less money to more” it generates a set of constraints. If however such a constraint is assumed with $p_{\text{info}} = 1$ — ie: if it is represented in the knowledge base \mathcal{K} — then the following device is very economical. [4] describes the representation of $P^t(\text{Accept}(\beta, \alpha, \delta))$ where β is attempting to purchase something for money but with a period of warranty. There α assumes that β prefers “less money to more” and “more warranty to less”. These two preference orderings are dealt with neatly by estimating instead $P^t(\text{LimAccept}(\beta, \alpha, \delta))$ meaning “ δ is the greatest w.r.t. money and least w.r.t warranty that β will accept from α ”.

In this way, quantitative preferences over finite domains will give a finite set of linear constraints (in particular, the device above may be used to great effect when $p_{\text{info}} = 1$), and qualitative preferences including conditional preferences also yield a finite set of linear constraints.

5.3 Updating from social information

Social relationships between agents, and social roles or positions held by agents, introduce a bias, i.e. a constraint, on the admissible probability distributions. A social model can be then a set of constraints introduced in \mathcal{K} that has to be respected by the inference mechanism.

For instance, with respect to *power*, and assuming we model power as a function from agents to real values, we could model a meek agent by adding the following constraint in \mathcal{K} that establishes different degrees of acceptability for proposals according to the power of the proposer:

$$\text{Power}(\beta) > \text{Power}(\gamma) \rightarrow$$

$$P^t(\text{Accept}(\alpha, \beta, \varphi)) > P^t(\text{Accept}(\alpha, \gamma, \varphi))$$

A similar case can be made for *reputation*, which refers to

the institutional endorsement of observed trustworthiness⁷. Power and reputation are different instruments that help an agent to form an *a priori* assessment of the trustworthiness of an unknown opponent, or to modify the assessment of a known one. If α learns that her good friend γ has a high opinion of β then this may cause α to increase her trust in β and to ‘tighten up’ the distribution $P^t(b', b)$. Likewise, if α learns that β has a high reputation in a respected institution. So it is natural to represent reputation as $\text{Reputation}(\Phi, \beta)$ where Φ is an institution name.

If α receives information, Θ , such as $\text{Reputation}(\Phi, \beta)$ then Θ will either be a positive influence on α 's estimate of $P^t(b', b)$ [written Θ^+], a negative one [Θ^-], or neutral — ie a positive influence on $P^t(b, b)$ [written Θ^0]. If α receives Θ^+ then her estimate of the probability that the execution of contract b will be preferred to b becomes: $p(b | \Theta^+)^+ = p(b)^+ + g_3(b, \Theta^+)(1 - p(b)^+)$, where $p(b)^+$ is the prior probability as in Sec. 5.1, $g_3(b, \Theta^+)$ represents the strength of α 's belief that the probability that execution of contract b at time $t + 1$ will be preferred to b should increase given Θ^+ was received at time t . α revises this estimate using the principle of minimum relative entropy (Eqn. 5) subject to the constraint C : $\sum_{x \in B(b)^+} P_C^t(x|b) = p(b | \Theta^+)^+$, where $B(b)^+$ is as in Sec. 5.1. Similarly, if α receives Θ^- or Θ^0 .

6. A TRUST MODEL

6.1 Trust as conditional entropy

One way of modelling α 's trust on β is as one minus the normalised negative entropy of distribution P^t . The idea is that the more trust the less dispersion of the expected observations and therefore the closer to 1. In this way we can define the Trust that an agent α has on agent β with respect to the fulfilment of a contract (a, b) as:

$$T(\alpha, \beta, b) = 1 + \frac{1}{B^*} \cdot \sum_{b' \in B(b)^+} P^t(b'|b) \log P^t(b'|b)$$

where $B(b)^+$ is the set of contract executions that α prefers to b as defined in Sec. 5.1, $B^* = 1$ if $|B(b)^+| = 1$ and $\log |B(b)^+|$ otherwise, and β has agreed to execute b , and α systematically observes b' , for some b' that α does not prefer to b , the trust value will be 0. Trust will tend to 0 when the dispersion of observations is maximal.

And, as a general measure of α 's trust on β we naturally use the normalised negative conditional entropy of executed contracts given signed contracts:

$$T(\alpha, \beta) = 1 + \frac{\sum_{b \in B} \left[P^t(b) \cdot \sum_{b' \in B(b)^+} P^t(b'|b) \log P^t(b'|b) \right]}{B^* \cdot \sum_{b \in B} P^t(b)}$$

This formulation of trust is useful when any variation from the agreed contract is undesirable.

6.2 Trust as relative entropy

We are usually happier, and ready to trust more, if the actual execution of a contract goes in the direction of our increasing preference. We capture this idea using as a model

⁷Electronic Institutions [1] warrant, within specific limits, the *bona fides* of the players therein — it is in their interests to report anecdotal evidence of ‘good’ behaviour beyond those limits.

for trust the relative entropy⁸ between the probability distribution of acceptance and the distribution of the observation of contract execution. That is:

$$T(\alpha, \beta, b) = 1 - \sum_{b' \in B(b)^+} P^t(b') \log \frac{P^t(b')}{P^t(b'|b)}$$

and, similarly

$$T(\alpha, \beta) = 1 - \sum_{b \in B} P^t(b) \left[\sum_{b' \in B(b)^+} P^t(b') \log \frac{P^t(b')}{P^t(b'|b)} \right]$$

Finally, and as proposed in Sec. 2, the trust we place in an agent is useful to determine which agent to prefer in order to accept proposals. That is, trust is useful to assess the distribution of probability $P(\text{Accept}(\alpha, \beta, \delta) | \text{Offer}(\beta, \alpha, \delta))$. What trust does to this distribution is to impose constraints on its values. As follows:

$$P^t(\text{Accept}(\alpha, \beta, \delta) | \text{Offer}(\beta, \alpha, \delta)) >$$

$$P^t(\text{Accept}(\alpha, \gamma, \delta) | \text{Offer}(\gamma, \alpha, \delta)) \text{ if } T(\alpha, \beta) > T(\alpha, \gamma).$$

e.g. I prefer the same deal from my usual butcher than from someone I trust less.

7. AGENT ARCHITECTURE

The agent architecture developed so far is summarised in Fig. 1. If α wants to do something at time t then it may build the distributions shown in Fig. 1 “from scratch” using the full message history, however this leads to difficult belief revision problems. At each time t , this involves importing messages from the “in box” \mathcal{X} , using “import rules”, into a repository \mathcal{I}_t where they are given trust functions — or, integrity decay functions⁹, instantiating the knowledge base \mathcal{K}_t and the belief set \mathcal{B}_t , resolving inconsistencies, and applying maximum entropy inference. Alternatively, as new messages are received α may update the distributions following the approaches in Sec. 5.

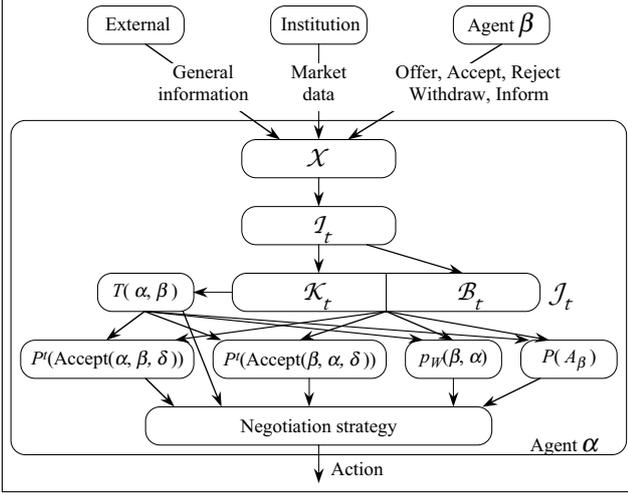
The probability distributions in Fig. 1 provide the fuel for the negotiation strategy. The $\text{Accept}(\alpha, \beta, \delta)$ illocution is perhaps dangerously simple. It is simple because it is binary — a deal is either acceptable or it isn't. The danger comes from misinterpreting the estimate $P^t(\text{Accept}(\alpha, \beta, \delta))$ as an indication of how good the deal, δ , is — that is *not* what is intended. This estimate is an estimate of α 's belief in the proposition that “ δ is an acceptable deal” — the value of that probability will be a function of *both* deal δ , *and* the integrity of the information against which δ has been evaluated. So $P^t(\text{Accept}(\alpha, \beta, \delta_1)) > P^t(\text{Accept}(\alpha, \beta, \delta_2))$ does not mean that δ_1 is an intrinsically better deal than δ_2 — it means that α is more certain that δ_1 is acceptable than δ_2 is acceptable.¹⁰ The reason for this could be that α has been unable to access reliable information against which to evaluate δ_2 . The choice of this simple, binary $\text{Accept}(\cdot)$ is justified because α will either accept an offer or it will not.

⁸Otherwise called *cross entropy* or the *Kullback-Leibler distance* — although it is not reflexive and so it is not a metric. See also Footnote 6.

⁹The values of these trust functions tend in time to the maximum entropy value of their distribution which denotes “no information”.

¹⁰We differ from game-theory-based agents in that our agent is not necessarily “utility aware”.

Figure 1: Trust and the informed agent



The $P^t(\text{Accept}(\alpha, \beta, \delta))$ distribution is derived from observations made. The significance of those observations will vary with δ — different criteria will apply to the purchase of a bar of gold to those used to purchase a used automobile. In particular, we can imagine the probability of acceptance of a deal as a composed measure:

$$P^t(\text{Accept}(\alpha, \beta, \delta)) = \kappa_1 T(\alpha, \beta, \delta) + \kappa_2 \text{Power}(\beta) + \kappa_3 \text{Reputation}(\alpha, \beta, \delta)$$

where $\kappa_1 + \kappa_2 + \kappa_3 = 1$ are constants, or the result of functions depending on the environment, represent the importance that an agent gives to each of the dimensions. A benevolent agent might have a large κ_3 , a meek agent a large κ_2 and a cautious agent a large κ_1 .

The estimation of the probability of breakdown, $p_W = P(\text{Withdraw}(\beta, \alpha)|e)$, will depend on various factors including: time — i.e. duration of the negotiation, and the “expected acceptability to β ” of α ’s offers — e.g. if α persistently sends completely unacceptable offers to β then β may well withdraw. Breakdown is discussed in Sec. 8.1.

The $P^t(\text{Accept}(\beta, \alpha, \delta))$ distribution is α ’s estimate of β readiness to accept δ . This is derived by interpreting Offer(β, α, δ) messages as indicating β ’s readiness to accept δ — to some subjective degree of belief determined by α . And by interpreting Reject(.) messages as indicating β ’s reluctance to accept. Then either maximum entropy inference or minimum relative entropy inference is used to “fill in the gaps”, and so to derive probability estimates for β ’s readiness to accept any deal.

The estimation of the probability distributions over β ’s future actions, $P(A_\beta)$, is not specifically described here, but may be achieved by applying either of the two entropy inference methods to the history of actions observed by α . To estimate all of the foregoing probability distributions α applies the *Information Principle* described in Sec. 1.

Finally, α ’s negotiation strategy should encapsulate its overall aim. If α is “greedy” then it could use an “expected acceptability to α optimising strategy”¹¹. If α strives to

¹¹If α ’s acceptance machinery happened to equate to utility then such an agent would be an expected utility optimiser.

appear “fair” then it could use some variant of an “equitable information revelation” strategy [see Sec. 8.1] that sought to equate in the expected (Shannon) value of the information in the messages sent to that which has been received. If α strives to appear “cooperative” then it could use some variant of an “issue tradeoffs” strategy [7].

8. OPERATIONAL SEMANTICS

α observes β ’s actions as well as other general information in the Institution, or marketplace — for example, α may detect that a competitor has newly arrived. All of this information is represented in α ’s knowledge base, \mathcal{K} , and belief set, \mathcal{B} . At time t , α ’s information is $\mathcal{J}_t = \mathcal{K}_t \cup \mathcal{B}_t$. α infers from \mathcal{J}_t beliefs about β ’s future actions as a probability distribution $P(A_\beta)$. We propose here that α measures the value of the information within an illocution uttered by β as the change in the entropy of $P(A_\beta)$. We denote it by $\Delta_\mu |\mathcal{J}_t^\alpha|$ where μ is the illocution received by α and $|\mathcal{J}_t^\alpha|$ is the certainty (as negative entropy) of α ’s inferred beliefs $P(A_\beta)$.

Although two agents engaged only in bilateral bargaining (i.e. they only observe each other’s actions), and using this approach, will build the model with the same observable data (exchanged illocutions) they might perceive different information gains if they represent the domains differently. One of the advantages of bargaining within normalised environments, such as *electronic institutions*, is that the ontology is determined by the environment, and therefore the chances that the representations chosen by the agents are closer are higher. Maximum entropy inference is criticised because the way in which the knowledge is represented determines the values derived. Requiring agents to interact within an Electronic Institution minimises the effect of potentially different representations.

An important question is how to determine what initial opponent’s model α should use to determine whether a proposal is going to be acceptable to β . A possibility, used here, is to assume that $\mathcal{I}_t^\beta = \text{Mirror}(\mathcal{I}_t^\alpha)$ where mirror is a function that deals with the fact that our opponent’s willingness in accepting proposals decreases as ours increases as we move along dimensions with a clear preference ordering. For instance, if $\forall x, y (x \geq y) \rightarrow (\text{Accept}(\alpha, \beta, x) \rightarrow \text{Accept}(\alpha, \beta, y)) \in \mathcal{I}_t^\alpha$ then we’ll assume that $\forall x, y (x \geq y) \rightarrow (\text{Accept}(\beta, \alpha, y) \rightarrow \text{Accept}(\beta, \alpha, x)) \in \mathcal{I}_t^\beta$.

In what follows we’ll assume that when the bargaining process starts, at time t_0 it is the case that $\mathcal{J}_t^\beta = \text{Mirror}(\mathcal{J}_t^\alpha)$ if α knows nothing about β . If α has interacted with β previously, or has access to any other information, then α can use that information to develop a better model of these beliefs. α ’s information base, \mathcal{I}_t^α is updated when an illocution arrives in the agent’s in-box, \mathcal{X} , at time t :

- Offer(β, α, δ) $\in \mathcal{X}$:
 $B(\text{Accept}(\beta, \alpha, \delta)) = f_{\text{Offer}}(T(\alpha, \beta, \delta)) \in \mathcal{I}_{t+1}^\alpha$
- Accept(β, α, δ) $\in \mathcal{X}$:
 $B(\text{Accept}(\beta, \alpha, \delta)) = 1 \in \mathcal{I}_{t+1}^\alpha$
- Reject($\beta, \alpha, \delta, [\text{info}]$) $\in \mathcal{X}$:
 $B(\text{Accept}(\beta, \alpha, \delta)) = f_{\text{Reject}}(T(\alpha, \beta, \delta)),$
 $B([\text{info}]) = f_{\text{Inform}}(T(\alpha, \beta)) \in \mathcal{I}_{t+1}^\alpha$
- Withdraw($\beta, \alpha, [\text{info}]$) $\in \mathcal{X}$:
 $\forall x B(\text{Accept}(\beta, \alpha, x)) = 0,$
 $B([\text{info}]) = f_{\text{Inform}}(T(\alpha, \beta)) \in \mathcal{I}_{t+1}^\alpha$

- $\text{Inform}(\beta, \alpha, [\text{info}]) \in \mathcal{X}$:
 $B([\text{info}]) = f_{\text{Inform}}(T(\alpha, \beta)) \in \mathcal{I}_{t+1}^\alpha$

where the functions f subscripted with illocutionary forms convert the trust values to probabilities. The first equation means that we believe, B , that agent β will sign a contract supporting his offer with the degree that our trust, T , that his dialogue behaviour indicates. The second and fourth equations contain certain beliefs, 1 and 0 respectively, in recognition of the meaning of Accept and Withdraw.

8.1 Breakdown model

In order for an agent to act sensibly in a negotiation it is necessary that it assesses both the probability that it should withdraw and the probability of breakdown of the opponent $P(\text{Withdraw}(\beta, \alpha)|e)$ — we use the abbreviation p_W when the context is unambiguous. Time spent negotiating is a cost, and the probability of breakdown will increase as the negotiation progresses (β may perceive that α is not giving away enough information and then decreases its trust on α , or β might have found a better deal in the meantime, etc.). In general it is reasonable to assume that this probability is very low initially and then it increases asymptotically towards 1. Learning the shape of this distribution from past cases seems the reasonable thing to do here, either through inductive learning or case-base reasoning.

$P(\text{Withdraw}(\beta, \alpha)|e)$ may also be a function of the attractiveness of the proposals that α has submitted. For example, if α has submitted only deals that are highly unattractive to β then β may decide that interacting with α is a waste of time. Every illocution that an agent α receives gives away information about her opponent β . That information may be valued as the increase in certainty (i.e. decrease in entropy) in $P^t(\text{Accept}(\beta, \alpha, \delta))$ as described in Sec. 5.2. Likewise, every illocution that α utters gives away information about herself. That information may be valued by α as the expected decrease in uncertainty in β 's model of α . Hence the principle of “equitable information revelation” as one strategy for managing $P(\text{Withdraw}(\beta, \alpha)|e)$ — that is, α utters illocutions that will give β comparable expected information gain as α derived from β 's previous utterance.

9. DISCUSSION

Game theory tells α that she should accept a proposal if $s_\delta > m_\delta$ where s_δ is the surplus, $s_\delta = u(\omega) - u(\pi)$ and m_δ is the margin. This is fine if everything is certain. If it is not then game theory tells α to work with a random variable, S_δ , instead. Incidentally this means that α has to be certain about her uncertainty, but that is not the issue. This means that α can consider $P(S_\delta > m_\delta)$, and the standard deviation, $\sigma(S_\delta)$, is a measure of uncertainty in the process. Then α asks “how risk averse am I”, and then is able to calculate $P(\text{accept } \delta)$.

The argument is that with uncertain information and decaying integrity, the “utility calculation” in the previous paragraph is a futile exercise. Instead we argue that it makes more sense to ask simply: “on the basis of what we actually know, what is the best thing to do?”. In support of this, we claim that α will be more concerned about the integrity of the information with which the decision is being made, than with an uncertain estimation of her utility distribution as an intermediate, and unnecessary, step in the calculation.

Further α should be concerned with “driving her acceptance machinery backwards” to identify information within the calculation whose integrity may, perhaps be increased. That is, α may view the acceptance machinery as an information management tool that may be used to proactively seek additional information to increase her certainty in her conclusions.

Finally the trust model proposed covers the requirements presented at the beginning of the paper, and is grounded on solid information based semantics.

10. REFERENCES

- [1] J. L. Arcos, M. Esteva, P. Noriega, J. A. Rodríguez, and C. Sierra. Environment engineering for multiagent systems. *Journal on Engineering Applications of Artificial Intelligence*, 18, 2005.
- [2] C. Boutilier, R. Brafman, C. Domshlak, H. Hoos, and D. Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, 21:135 – 191, 2004.
- [3] J. Debenham. Auctions and bidding with information. In P. Faratin and J. Rodríguez-Aguilar, editors, *Proceedings Agent-Mediated Electronic Commerce VI: AMEC*, pages 15 – 28, July 2004.
- [4] J. Debenham. Bargaining with information. In N. Jennings, C. Sierra, L. Sonenberg, and M. Tambe, editors, *Proceedings Third International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2004*, pages 664 – 671. ACM, July 2004.
- [5] M. Dorigo and T. Stützle. *Ant Colony Optimization*. MIT Press, Cambridge, MA, 2004.
- [6] R. Falcone and C. Castelfranchi. The socio-cognitive dynamics of trust: Does trust create trust? In *Proceedings of the workshop on Deception, Fraud, and Trust in Agent Societies*, pages 55 – 72, 2001.
- [7] P. Faratin, C. Sierra, and N. Jennings. Using similarity criteria to make issue trade-offs in automated negotiation. *Journal of Artificial Intelligence*, 142(2):205–237, 2003.
- [8] E. Jaynes. *Probability Theory — The Logic of Science*. Cambridge University Press, 2003.
- [9] M. Karlins and H. Abelson. *Persuasion*. Crosby Lockwood and Son, 1970.
- [10] D. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003.
- [11] S. Ramchurn, C. Sierra, L. Godo, and N. Jennings. Devising a trust model for multiagent interactions using confidence and reputation. *International Journal of Applied Artificial Intelligence*, 18(9–10):91–204, 2005.
- [12] J. Sabater and C. Sierra. Reputation and social network analysis in multi-agent systems. In *Proceedings of the First International Conference on Autonomous Agents and Multi-Agent systems*, pages 475 – 482, 2002.
- [13] P. Yolum and M. Singh. Achieving trust via service graphs. In *Proceedings of the Autonomous Agents and Multi-Agent Systems Workshop on Deception, Fraud and Trust in Agent Societies*. Springer-Verlag, 2003.