

Information-Based Agency

Carles Sierra

Institut d'Investigacio en Intel.ligencia Artificial
Spanish Scientific Research Council, UAB
08193 Bellaterra, Catalonia, Spain
sierra@iii.csic.es

John Debenham

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

Abstract

Successful negotiators look beyond a purely utilitarian view. We propose a new agent architecture that integrates the utilitarian, information, and semantic views allowing the definition of strategies that take these three dimensions into account. Information-based agency values the information in dialogues in the context of a communication language based on a structured ontology and on the notion of commitment. This abstraction unifies measures such as trust, reputation, and reliability in a single framework.

1 Introduction

In this paper we introduce an agency framework grounded on information-based concepts [MacKay, 2003]. It presents an agent architecture that admits a game-theoretical reading and an information-theoretical reading. And, what is more interesting, permits a connection between these two worlds by considering negotiation as both *cooperative*, in the sense that all interaction involves the exchange of information, and *competitive*, in the sense that agents aim at getting the best they can. This approach contrasts with previous work on the design of negotiation strategies that did not take information exchange into account, but focused on the similarity of offers [Jennings *et al.*, 2001; Faratin *et al.*, 2003], game theory [Rosenschein and Zlotkin, 1994], or first-order logic [Kraus, 1997].

We assume that a multiagent system $\{\alpha, \beta_1, \dots, \beta_o, \xi, \theta_1, \dots, \theta_t\}$, contains an agent α that interacts with negotiating agents, β_i , information providing agents, θ_j , and an *institutional agent*, ξ , that represents the institution where we assume the interactions happen [Arcos *et al.*, 2005]. Institutions give a normative context to interactions that simplify matters (e.g an agent can't make an offer, have it accepted, and then renege on it). We will describe a *communication language* \mathcal{C} based on an *ontology* that will permit us both to structure the dialogues and to structure the processing of the information gathered by agents. Agents have an *internal language* \mathcal{L} used to build a probabilistic *world model*.

We understand agents as being built on top of two basic functionalities. First, a *proactive machinery*, that transforms

needs into *goals* and these into *plans* composed of *actions*. Second, a reactive machinery, that uses the received messages to obtain a new world model by updating the probability distributions in it. Agents summarise their world models using a number of measures (e.g. trust, reputation, and reliability [Sierra and Debenham, 2006]) that can then be used to define strategies for “exchanging information” — in the sense developed here, this is the only thing that an agent can do.

We introduce the communication language in Section 2, the agent architecture in Section 3, some summary measures based on the architecture in Section 4, and some associated interaction strategies in Section 5.

2 The Multiagent System

Our agent α has two languages: \mathcal{C} is an illocutionary-based language for communication, and \mathcal{L} is a language for internal representation. \mathcal{C} is described following, and \mathcal{L} in Section 3.

2.1 Communication Language \mathcal{C}

The shape of the language that α uses to represent the information received and the content of its dialogues depends on two fundamental notions. First, when agents interact within an overarching institution they explicitly or implicitly accept the *norms* that will constrain their behaviour, and accept the established sanctions and penalties whenever norms are violated. Second, the dialogues in which α engages are built around two fundamental actions: (i) passing information, and (ii) exchanging proposals and contracts. A contract $\delta = (a, b)$ between agents α and β is a pair where a and b represent the activities that agents α and β are respectively responsible for. *Contracts* signed by agents and *information* passed by agents, are similar to norms in the sense that they oblige agents to behave in a particular way, so as to satisfy the conditions of the contract, or to make the world consistent with the information passed. Contracts and Information can then be thought of as normative statements that restrict an agent's behaviour.

Norms, contracts, and information have an obvious temporal dimension. Thus, an agent has to abide by a norm while it is inside an institution, a contract has a validity period, and a piece of information is true only during an interval in time. The set of norms affecting the behaviour of an agent define the *context* that the agent has to take into account.

The communication language that α needs requires two fundamental primitives: $\text{Commit}(\alpha, \beta, \varphi)$ to represent, in φ ,

what is the world α aims at bringing about and that β has the right to verify, complain about or claim compensation for any deviations from, and $\text{Done}(a)$ to represent the event that a certain action a^1 has taken place. In this way, norms, contracts, and information chunks will be represented as instances of $\text{Commit}(\cdot)$ where α and β can be individual agents or institutions, \mathcal{C} is:

$a ::= \text{illoc}(\alpha, \beta, \varphi, t) \mid a; a \mid$

Let context In a End

$\varphi ::= \text{term} \mid \text{Done}(a) \mid \text{Commit}(\alpha, \varphi) \mid \varphi \wedge \varphi \mid$

$\varphi \vee \varphi \mid \neg\varphi \mid \forall v.\varphi_v \mid \exists v.\varphi_v$

$\text{context} ::= \varphi \mid \text{id} = \varphi \mid \text{prolog_clause} \mid \text{context}; \text{context}$

where φ_v is a formula with free variable v , *illoc* is any appropriate set of illocutionary particles, ‘;’ means sequencing, and *context* represents either previous agreements, previous illocutions, or code that aligns the ontological differences between the speakers needed to interpret an action a .

For example, we can represent the following offer: “If you spend a total of more than €100 in my shop during October then I will give you a 10% discount on all goods in November”, as:

$\text{Offer}(\alpha, \beta, \text{spent}(\beta, \alpha, \text{October}, X) \wedge X \geq \text{€}100 \rightarrow$
 $\forall y. \text{Done}(\text{Inform}(\xi, \alpha, \text{pay}(\beta, \alpha, y), \text{November})) \rightarrow$
 $\text{Commit}(\alpha, \beta, \text{discount}(y, 10\%)))$

Note the use of the institution agent ξ to report the payment.

In order to define the language introduced above that structures agent dialogues we need an ontology that includes a (minimum) repertoire of elements: a set of *concepts* (e.g. quantity, quality, material) organised in a is-a hierarchy (e.g. platypus is a mammal, australian-dollar is a currency), and a set of relations over these concepts (e.g. price(beer, AUD)).² We model ontologies following an algebraic approach [Kalfoglou and Schorlemmer, 2003] as:

An ontology is a tuple $\mathcal{O} = (C, R, \leq, \sigma)$ 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. \leq is a traditional *is-a* hierarchy, and R contains relations between the concepts in the hierarchy.

The concepts within an ontology are closer, semantically speaking, depending on how far away are they in the structure defined by the \leq relation. Semantic distance plays a fundamental role in strategies for information-based agency. How signed contracts, $\text{Commit}(\cdot)$ about objects in a particular semantic region, and their execution $\text{Done}(\cdot)$, affect our decision making process about signing future contracts on nearby semantic regions is crucial to model the common sense that human beings apply in managing trading relationships. A measure [Li *et al.*, 2003] bases the *semantic similarity* between two concepts on the path length induced by \leq (more

¹Without loss of generality we will assume that all actions are dialogical.

²Usually, a set of axioms defined over the concepts and relations is also required. We will omit this here.

distance in the \leq graph means less semantic similarity), and the *depth* of the subsumer concept (common ancestor) in the shortest path between the two concepts (the deeper in the hierarchy, the closer the meaning of the concepts). Semantic similarity could then be defined as:

$$\text{Sim}(c, c') = e^{-\kappa_1 l} \cdot \frac{e^{\kappa_2 h} - e^{-\kappa_2 h}}{e^{\kappa_2 h} + e^{-\kappa_2 h}}$$

where l is the length (i.e. number of hops) of the shortest path between the concepts, h is the depth of the deepest concept subsuming both concepts, and κ_1 and κ_2 are parameters scaling the contribution of shortest path length and depth respectively.

Given a formula $\varphi \in \mathcal{C}$ in the communication language we define the vocabulary or *ontological context* of the formula, $O(\varphi)$, as the set of concepts in the ontology used in it. Thus, we extend the previous definition of similarity to sets of concepts in the following way:

$$\text{Sim}(\varphi, \psi) = \max_{c_i \in O(\varphi)} \min_{c_j \in O(\psi)} \{\text{Sim}(c_i, c_j)\}$$

The following does not depend on this particular definition.

3 Agent Architecture

Agent α receives all messages expressed in \mathcal{C} in an in-box \mathcal{X} where they are time-stamped and sourced-stamped. A message μ from agent β (or θ or ξ) is then moved from \mathcal{X} to a *percept repository* \mathcal{Y}^t where it is appended with a subjective belief function $\mathbb{R}^t(\alpha, \beta, \mu)$ that normally decays with time. α acts in response to a message that expresses a *need*. A need may be exogenous such as a need to trade profitably and may be triggered by another agent offering to trade, or endogenous such as α deciding that it owns more wine than it requires. Needs trigger α ’s goal/plan proactive reasoning described in Section 3.1, other messages are dealt with by α ’s reactive reasoning described in Section 3.2.

3.1 Proactive Reasoning

α ’s goal/plan machinery operates in a climate of changing uncertainty and so typically will pursue multiple sub-goals concurrently. This applies both to negotiations with a potential trading partner, and to interaction with information sources that either may be unreliable or may take an unpredictable time to respond. Each plan contains constructors for a *world model* \mathcal{M}^t that consists of probability distributions, (X_i) , in first-order probabilistic logic \mathcal{L} . \mathcal{M}^t is then maintained from percepts received using *update functions* that transform percepts into constraints on \mathcal{M}^t — described in Section 3.2.

The distributions in \mathcal{M}^t are determined by α ’s plans that are determined by its needs. If α is negotiating some contract δ in satisfaction of need χ then it may require the distribution $\mathbb{P}^t(\text{eval}(\alpha, \beta, \chi, \delta) = e_i)$ where for a particular δ , $\text{eval}(\alpha, \beta, \chi, \delta)$ is an evaluation over some complete and disjoint *evaluation space* $E = (e_1, \dots, e_n)$ that may contain hard (possibly utilitarian) values, or fuzzy values such as “reject” and “accept”. This distribution assists α ’s strategies to decide whether to accept a proposed contract leading to a probability of acceptance $\mathbb{P}^t(\text{acc}(\alpha, \beta, \chi, \delta))$. For example,

$\mathbb{P}^t(\text{eval}(\alpha, \beta, \chi, \delta) = e_i)$ could be derived from the subjective estimate $\mathbb{P}^t(\text{satisfy}(\alpha, \beta, \chi, \delta) = f_j)$ of the expected extent to which the execution of δ by β will satisfy χ , and an objective estimate $\mathbb{P}^t(\text{val}(\alpha, \beta, \delta) = g_k)$ of the expected valuation of the execution of δ possibly in utilitarian terms. This second estimate could be derived by proactive reference to the $\{\theta_i\}$ for market data. In a negotiation α 's plans may also construct the distribution $\mathbb{P}^t(\text{acc}(\beta, \alpha, \delta))$ that estimates the probability that β would accept δ — we show in Section 3.2 how α may derive this estimate from the information in β 's proposals.

α 's plans may construct various other distributions such as: $\mathbb{P}^t(\text{trade}(\alpha, \beta, o) = e_i)$ that β is a good person to sign contracts with in context o , and $\mathbb{P}^t(\text{confide}(\alpha, \beta, o) = f_j)$ that α can trust β with confidential information in context o .

The integrity of percepts decreases in time. α may have background knowledge concerning the expected integrity of a percept as $t \rightarrow \infty$. Such background knowledge is represented as a *decay limit distribution*. If the background knowledge is incomplete then one possibility is for α to assume that the decay limit distribution has maximum entropy whilst being consistent with the data. Given a distribution, $\mathbb{P}(X_i)$, and a decay limit distribution $\mathbb{D}(X_i)$, $\mathbb{P}(X_i)$ decays by:

$$\mathbb{P}^{t+1}(X_i) = \Delta_i(\mathbb{D}(X_i), \mathbb{P}^t(X_i)) \quad (1)$$

where Δ_i is the *decay function* for the X_i satisfying the property that $\lim_{t \rightarrow \infty} \mathbb{P}^t(X_i) = \mathbb{D}(X_i)$. For example, Δ_i could be linear: $\mathbb{P}^{t+1}(X_i) = (1 - \nu_i) \times \mathbb{D}(X_i) + \nu_i \times \mathbb{P}^t(X_i)$, where $\nu_i < 1$ is the decay rate for the i 'th distribution. Either the decay function or the decay limit distribution could also be a function of time: Δ_i^t and $\mathbb{D}^t(X_i)$.

3.2 Reactive Reasoning

In the absence of in-coming messages the integrity of \mathcal{M}^t decays by Eqn. 1. The following procedure updates \mathcal{M}^t for all percepts expressed in \mathcal{C} . Suppose that α receives a message μ from agent β at time t . Suppose that this message states that something is so with probability z , and suppose that α attaches an epistemic belief $\mathbb{R}^t(\alpha, \beta, \mu)$ to μ — this probability reflects α 's level of personal *caution*. Each of α 's active plans, s , contains constructors for a set of distributions $\{X_i\} \in \mathcal{M}^t$ together with associated *update functions*, $J_s(\cdot)$, such that $J_s^{X_i}(\mu)$ is a set of linear constraints on the posterior distribution for X_i . Examples of these update functions are given in Section 3.4. Denote the prior distribution $\mathbb{P}^t(X_i)$ by \vec{p} , and let $\vec{p}(\mu)$ be the distribution with minimum relative entropy³ with respect to \vec{p} : $\vec{p}(\mu) = \arg \min_{\vec{r}} \sum_j r_j \log \frac{r_j}{p_j}$ that

³Given a probability distribution \vec{q} , the *minimum relative entropy distribution* $\vec{p} = (p_1, \dots, p_I)$ subject to a set of J linear constraints $\vec{g} = \{g_j(\vec{p}) = \vec{a}_j \cdot \vec{p} - c_j = 0\}, j = 1, \dots, J$ (that must include the constraint $\sum_i p_i - 1 = 0$) is: $\vec{p} = \arg \min_{\vec{r}} \sum_j r_j \log \frac{r_j}{q_j}$.

This may be calculated by introducing Lagrange multipliers $\vec{\lambda}$: $L(\vec{p}, \vec{\lambda}) = \sum_j p_j \log \frac{p_j}{q_j} + \vec{\lambda} \cdot \vec{g}$. Minimising L , $\{\frac{\partial L}{\partial \lambda_j} = g_j(\vec{p}) = 0\}, j = 1, \dots, J$ is the set of given constraints \vec{g} , and a solution to $\frac{\partial L}{\partial p_i} = 0, i = 1, \dots, I$ leads eventually to \vec{p} . Entropy-based inference is a form of Bayesian inference that is convenient when the data is sparse [Cheeseman and Stutz, 2004] and encapsulates common-sense reasoning [Paris, 1999].

satisfies the constraints $J_s^{X_i}(\mu)$. Then let $\vec{q}(\mu)$ be the distribution:

$$\vec{q}(\mu) = \mathbb{R}^t(\alpha, \beta, \mu) \times \vec{p}(\mu) + (1 - \mathbb{R}^t(\alpha, \beta, \mu)) \times \vec{p} \quad (2)$$

and then let:

$$\mathbb{P}^t(X_{i(\mu)}) = \begin{cases} \vec{q}(\mu) & \text{if } \vec{q}(\mu) \text{ is more interesting than } \vec{p} \\ \vec{p} & \text{otherwise} \end{cases} \quad (3)$$

A general measure of whether $\vec{q}(\mu)$ is more interesting than \vec{p} is: $\mathbb{K}(\vec{q}(\mu) \parallel \mathbb{D}(X_i)) > \mathbb{K}(\vec{p} \parallel \mathbb{D}(X_i))$, where $\mathbb{K}(\vec{x} \parallel \vec{y}) = \sum_j x_j \ln \frac{x_j}{y_j}$ is the Kullback-Leibler distance between two probability distributions \vec{x} and \vec{y} .

Finally merging Eqn. 3 and Eqn. 1 we obtain the method for updating a distribution X_i on receipt of a message μ :

$$\mathbb{P}^{t+1}(X_i) = \Delta_i(\mathbb{D}(X_i), \mathbb{P}^t(X_{i(\mu)})) \quad (4)$$

This procedure deals with integrity decay, and with two probabilities: first, the probability z in the percept μ , and second the belief $\mathbb{R}^t(\alpha, \beta, \mu)$ that α attached to μ .

In a simple multi-issue contract negotiation α may estimate $\mathbb{P}^t(\text{acc}(\beta, \alpha, \delta))$, the probability that β would accept δ , by observing β 's responses. Using shorthand notation, if β sends the message $\text{Offer}(\delta_1)$ then α may derive the constraint: $J^{\text{acc}(\beta, \alpha, \delta)}(\text{Offer}(\delta_1)) = \{\mathbb{P}^t(\text{acc}(\beta, \alpha, \delta_1)) = 1\}$, and if this is a counter offer to a former offer of α 's, δ_0 , then: $J^{\text{acc}(\beta, \alpha, \delta)}(\text{Offer}(\delta_1)) = \{\mathbb{P}^t(\text{acc}(\beta, \alpha, \delta_0)) = 0\}$. In the not-atypical special case of multi-issue bargaining where the agents' preferences over the individual issues *only* are known and are complementary to each other's, maximum entropy reasoning can be applied to estimate the probability that any multi-issue δ will be acceptable to β by enumerating the possible worlds that represent β 's "limit of acceptability" [Debenham, 2004].

3.3 Reliability

$\mathbb{R}^t(\alpha, \beta, \mu)$ is an epistemic probability that takes account of α 's personal caution. An empirical estimate of $\mathbb{R}^t(\alpha, \beta, \mu)$ may be obtained by measuring the 'difference' between commitment and enactment. Suppose that μ is received from agent β at time u and is verified by ξ as μ' at some later time t . Denote the prior $\mathbb{P}^u(X_i)$ by \vec{p} . Let $\vec{p}(\mu)$ be the posterior minimum relative entropy distribution subject to the constraints $J_s^{X_i}(\mu)$, and let $\vec{p}(\mu')$ be that distribution subject to $J_s^{X_i}(\mu')$. We now estimate what $\mathbb{R}^u(\alpha, \beta, \mu)$ should have been in the light of knowing *now*, at time t , that μ should have been μ' .

The idea of Eqn. 2, is that $\mathbb{R}^t(\alpha, \beta, \mu)$ should be such that, *on average* across \mathcal{M}^t , $\vec{q}(\mu)$ will predict $\vec{p}(\mu')$ — no matter whether or not μ was used to update the distribution for X_i , as determined by the condition in Eqn. 3 at time u . The *observed reliability* for μ and distribution X_i , $\mathbb{R}_{X_i}^t(\alpha, \beta, \mu) | \mu'$, on the basis of the verification of μ with μ' , is the value of k that minimises the Kullback-Leibler distance:

$$\mathbb{R}_{X_i}^t(\alpha, \beta, \mu) | \mu' = \arg \min_k \mathbb{K}(k \cdot \vec{p}(\mu) + (1 - k) \cdot \vec{p} \parallel \vec{p}(\mu'))$$

The predicted *information* in the enactment of μ with respect to X_i is:

$$\mathbb{H}_{X_i}^t(\alpha, \beta, \mu) = \mathbb{H}^t(X_i) - \mathbb{H}^t(X_{i(\mu)}) \quad (5)$$

that is the reduction in uncertainty in X_i where $\mathbb{H}(\cdot)$ is Shannon entropy. Eqn. 5 takes account of the value of $\mathbb{R}^t(\alpha, \beta, \mu)$.

If $\mathbf{X}(\mu)$ is the set of distributions that μ affects, then the *observed reliability* of β on the basis of the verification of μ with μ' is:

$$\mathbb{R}^t(\alpha, \beta, \mu)|\mu' = \frac{1}{|\mathbf{X}(\mu)|} \sum_i \mathbb{R}_{X_i}^t(\alpha, \beta, \mu)|\mu' \quad (6)$$

If $\mathbf{X}(\mu)$ are independent the predicted *information* in μ is:

$$\mathbb{I}^t(\alpha, \beta, \mu) = \sum_{X_i \in \mathbf{X}(\mu)} \mathbb{I}_{X_i}^t(\alpha, \beta, \mu) \quad (7)$$

Suppose α sends message μ to β where μ is α 's private information, then assuming that β 's reasoning apparatus mirrors α 's, α can estimate $\mathbb{I}^t(\beta, \alpha, \mu)$.

For each formula φ at time t when μ has been verified with μ' , the *observed reliability* that α has for agent β in φ is:

$$\begin{aligned} \mathbb{R}^{t+1}(\alpha, \beta, \varphi) = & (1 - \nu) \times \mathbb{R}^t(\alpha, \beta, \varphi) + \\ & \nu \times \mathbb{R}^t(\alpha, \beta, \mu)|\mu' \times \text{Sim}(\varphi, \mu) \end{aligned}$$

where Sim measures the semantic distance between two sections of the ontology as introduced in Section 2.1, and ν is the learning rate. Over time, α notes the context of the various μ received from β , and over the various contexts calculates the relative frequency, $\mathbb{P}^t(\mu)$. This leads to an overall expectation of the *reliability* that agent α has for agent β :

$$\mathbb{R}^t(\alpha, \beta) = \sum_{\mu} \mathbb{P}^t(\mu) \times \mathbb{R}^t(\alpha, \beta, \mu)$$

3.4 Commitment and Enactment

The interaction between agents α and β will involve β making contractual commitments and (perhaps implicitly) committing to the truth of information exchanged. No matter what these commitments are, α will be interested in any variation between β 's commitment, φ , and what is actually observed (as advised by the institution agent ξ), as the enactment, φ' . We denote the relationship between commitment and enactment, $\mathbb{P}^t(\text{Observe}(\varphi')|\text{Commit}(\varphi))$ simply as $\mathbb{P}^t(\varphi'|\varphi) \in \mathcal{M}^t$.

In the absence of in-coming messages the conditional probabilities, $\mathbb{P}^t(\varphi'|\varphi)$, should tend to ignorance as represented by the *decay limit distribution* and Eqn. 1. We now show how Eqn. 4 may be used to revise $\mathbb{P}^t(\varphi'|\varphi)$ as observations are made. Let the set of possible enactments be $\Phi = \{\varphi_1, \varphi_2, \dots, \varphi_m\}$ with prior distribution $\vec{p} = \mathbb{P}^t(\varphi'|\varphi)$. Suppose that message μ is received, we estimate the posterior $\vec{p}_{(\mu)} = (p_{(\mu)_i})_{i=1}^m = \mathbb{P}^{t+1}(\varphi'|\varphi)$.

First, if $\mu = (\varphi_k, \varphi)$ is observed then α may use this observation to estimate $p_{(\varphi_k)_k}$ as some value d at time $t + 1$. We estimate the distribution $\vec{p}_{(\varphi_k)}$ by applying the principle of minimum relative entropy as in Eqn. 4 with prior \vec{p} , and the posterior $\vec{p}_{(\varphi_k)} = (p_{(\varphi_k)_j})_{j=1}^m$ satisfying the single constraint: $J^{(\varphi'|\varphi)}(\varphi_k) = \{p_{(\varphi_k)_k} = d\}$.

Second, we consider the effect that the enactment ϕ' of another commitment ϕ , also by agent β , has on \vec{p} . This is achieved in two ways, first by appealing to the structure of the ontology using the $\text{Sim}(\cdot)$ function, and second by introducing a valuation function.

The $\text{Sim}(\cdot)$ method. Given the observation $\mu = (\phi', \phi)$, define the vector \vec{t} by

$$t_i = \mathbb{P}^t(\varphi_i|\varphi) + (1 - |\text{Sim}(\phi', \phi) - \text{Sim}(\varphi_i, \varphi)|) \cdot \text{Sim}(\varphi', \phi)$$

for $i = 1, \dots, m$. \vec{t} is not a probability distribution. The multiplying factor $\text{Sim}(\varphi', \phi)$ limits the variation of probability to those formulae whose ontological context is not too far away from the observation. The posterior $\vec{p}_{(\phi', \phi)}$ is defined to be the normalisation of \vec{t} .

The valuation method. α may wish to value φ in some sense. This value will depend on the future use that α makes of it. So α estimates the value of φ using a probability distribution (p_1, \dots, p_n) over some *evaluation space* $E = (e_1, \dots, e_n)$. $p_i = w_i(\varphi)$ is the probability that e_i is the correct evaluation of the enactment φ , and $\vec{w} : \mathcal{L} \times \mathcal{L} \rightarrow [0, 1]^n$ is the *evaluation function*.

For a given φ_k , $(\mathbb{P}^t(\varphi_1|\varphi_k), \dots, \mathbb{P}^t(\varphi_m|\varphi_k))$ is the prior distribution of α 's estimate of what will be observed if β committed to φ_k , and $\vec{w}(\varphi_k) = (w_1(\varphi_k), \dots, w_n(\varphi_k))$ is α 's evaluation over E of β 's commitment φ_k . α 's expected evaluation of what will be observed that β has committed to φ_k is $\vec{w}^{\text{exp}}(\varphi_k)$:

$$w^{\text{exp}}(\varphi_k)_i = \sum_{j=1}^m \mathbb{P}^t(\varphi_j|\varphi_k) \cdot w_i(\varphi_j)$$

for $i = 1, \dots, n$. Now suppose that α observes the enactment ϕ' of another commitment ϕ also by agent β . Eg: α may buy wine and cheese from the same supplier. α may wish to revise the prior estimate of the expected valuation $\vec{w}^{\text{exp}}(\varphi_k)$ in the light of the observation (ϕ', ϕ) to:

$$(\vec{w}^{\text{rev}}(\varphi_k) | (\phi'|\phi)) =$$

$$\vec{g}(\vec{w}^{\text{exp}}(\varphi_k), \text{Sim}(\phi', \phi), \text{Sim}(\varphi, \phi), \vec{w}(\varphi), \vec{w}(\phi), \vec{w}(\phi')))$$

for some function \vec{g} — the idea being, for example, that if the commitment, ϕ , a contract for cheese was not kept then α 's expectation that the commitment, φ , a contract for wine will not be kept should increase. We estimate the posterior $\vec{p}_{(\phi', \phi)}$ by applying the principle of minimum relative entropy as in Eqn. 4 with prior \vec{p} and $\vec{p}_{(\phi', \phi)} = (p_{(\phi', \phi)_j})_{j=1}^m$ satisfying the n constraints:

$$J^{(\varphi'|\varphi)}((\phi', \phi)) = \left\{ \sum_{j=1}^m p_{(\phi', \phi)_j} \cdot w_i(\varphi_j) = \right.$$

$$\left. g_i(\vec{w}^{\text{exp}}(\varphi_k), \text{Sim}(\phi', \phi), \text{Sim}(\varphi, \phi), \vec{w}(\varphi), \vec{w}(\phi), \vec{w}(\phi')) \right\}_{i=1}^n$$

This is a set of n linear equations in m unknowns, and so the calculation of the minimum relative entropy distribution may be impossible if $n > m$. In this case, we take only the m equations for which the change from the prior to the posterior value is greatest. That is, we attempt to select the most significant factors.

4 Summary Measures

The measures here generalise what are commonly called *trust*, *reliability* and *reputation* measures into a single computational framework. It they are applied to the execution of

contracts they become trust measures, to the validation of information they become reliability measures, and to socially transmitted behaviour they become reputation measures.

Ideal enactments. Consider a distribution of enactments that represent α 's "ideal" in the sense that it is the best that α could reasonably expect to happen. This distribution will be a function of α 's context with β denoted by e , and is $\mathbb{P}_I^t(\varphi'|\varphi, e)$. Here we measure the relative entropy between this ideal distribution, $\mathbb{P}_I^t(\varphi'|\varphi, e)$, and the distribution of expected enactments, $\mathbb{P}^t(\varphi'|\varphi)$. That is:

$$M(\alpha, \beta, \varphi) = 1 - \sum_{\varphi'} \mathbb{P}_I^t(\varphi'|\varphi, e) \log \frac{\mathbb{P}_I^t(\varphi'|\varphi, e)}{\mathbb{P}^t(\varphi'|\varphi)} \quad (8)$$

where the "1" is an arbitrarily chosen constant being the maximum value that this measure may have. This equation measures one, single commitment φ . It makes sense to aggregate these values over a class of commitments, say over those φ that are in the context o , that is $\varphi \leq o$:

$$M(\alpha, \beta, o) = 1 - \frac{\sum_{\varphi: \varphi \leq o} \mathbb{P}_\beta^t(\varphi) [1 - M(\alpha, \beta, \varphi)]}{\sum_{\varphi: \varphi \leq o} \mathbb{P}_\beta^t(\varphi)}$$

where $\mathbb{P}_\beta^t(\varphi)$ is a probability distribution over the space of commitments that the next commitment β will make to α is φ . Similarly, for an overall estimate of β 's reputation to α :

$$M(\alpha, \beta) = 1 - \sum_{\varphi} \mathbb{P}_\beta^t(\varphi) [1 - M(\alpha, \beta, \varphi)]$$

Preferred enactments. The previous measure requires that an ideal distribution, $\mathbb{P}_I^t(\varphi'|\varphi, e)$, has to be specified for each φ . Here we measure the extent to which the enactment φ' is preferable to the commitment φ . Given a predicate $\text{Prefer}(c_1, c_2, e)$ meaning that α prefers c_1 to c_2 in environment e . An evaluation of $\mathbb{P}^t(\text{Prefer}(c_1, c_2, e))$ may be defined using $\text{Sim}(\cdot)$ and the evaluation function $\vec{w}(\cdot)$ — but we do not detail it here. Then if $\varphi \leq o$:

$$M(\alpha, \beta, \varphi) = \sum_{\varphi'} \mathbb{P}^t(\text{Prefer}(\varphi', \varphi, o)) \mathbb{P}^t(\varphi' | \varphi)$$

and:

$$M(\alpha, \beta, o) = \frac{\sum_{\varphi: \varphi \leq o} \mathbb{P}_\beta^t(\varphi) M(\alpha, \beta, \varphi)}{\sum_{\varphi: \varphi \leq o} \mathbb{P}_\beta^t(\varphi)}$$

Certainty in enactment. Here we measure the consistency in expected acceptable enactment of commitments, or "the lack of expected uncertainty in those possible enactments that are better than the commitment as specified". If $\varphi \leq o$ let: $\Phi_+(\varphi, o, \kappa) = \{\varphi' | \mathbb{P}^t(\text{Prefer}(\varphi', \varphi, o)) > \kappa\}$ for some constant κ , and:

$$M(\alpha, \beta, \varphi) = 1 + \frac{1}{B^*} \cdot \sum_{\varphi' \in \Phi_+(\varphi, o, \kappa)} \mathbb{P}_+^t(\varphi'|\varphi) \log \mathbb{P}_+^t(\varphi'|\varphi)$$

where $\mathbb{P}_+^t(\varphi'|\varphi)$ is the normalisation of $\mathbb{P}^t(\varphi'|\varphi)$ for $\varphi' \in \Phi_+(\varphi, o, \kappa)$,

$$B^* = \begin{cases} 1 & \text{if } |\Phi_+(\varphi, o, \kappa)| = 1 \\ \log |\Phi_+(\varphi, o, \kappa)| & \text{otherwise} \end{cases}$$

As above we aggregate this measure for those commitments in a particular context o , and measure reputation as before.

Computational Note. The various measures given above involve extensive calculations. For example, Eqn. 8 contains $\sum_{\varphi'}$ that sums over *all* possible enactments φ' . We obtain a more computationally friendly measure by appealing to the structure of the ontology described in Section 2.1, and the right-hand side of Eqn. 8 may be approximated to:

$$1 - \sum_{\varphi': \text{Sim}(\varphi', \varphi) \geq \eta} \mathbb{P}_{\eta, I}^t(\varphi'|\varphi, e) \log \frac{\mathbb{P}_{\eta, I}^t(\varphi'|\varphi, e)}{\mathbb{P}_\eta^t(\varphi'|\varphi)}$$

where $\mathbb{P}_{\eta, I}^t(\varphi'|\varphi, e)$ is the normalisation of $\mathbb{P}_I^t(\varphi'|\varphi, e)$ for $\text{Sim}(\varphi', \varphi) \geq \eta$, and similarly for $\mathbb{P}_\eta^t(\varphi'|\varphi)$. The extent of this calculation is controlled by the parameter η . An even tighter restriction may be obtained with: $\text{Sim}(\varphi', \varphi) \geq \eta$ and $\varphi' \leq \psi$ for some ψ .

5 Strategies

An approach to issue-tradeoffs is described in [Faratin *et al.*, 2003]. That strategy attempts to make an acceptable offer by "walking round" the iso-curve of α 's previous offer (that has, say, an acceptability of c) towards β 's subsequent counter offer. In terms of the machinery described here:

$$\arg \max_{\delta} \{ \mathbb{P}^t(\text{acc}(\beta, \alpha, \delta)) | \mathbb{P}^t(\text{acc}(\alpha, \beta, \chi, \delta)) \approx c \} \quad (9)$$

In multi-issue negotiation the number of potential contracts that satisfy, or nearly satisfy, a utilitarian strategy such as this will be large — there is considerable scope for refining the choice of response.

Everything that an agent communicates gives away information. So even for purely self-interested agents interaction is a semi-cooperative process. α 's machinery includes measures of expected information gain in messages either sent to, or received from, β . For example, α can refine a strategy such as Eqn. 9 by replying to a message μ from β with μ' that attempts to give β equitable information gain: $\mathbb{I}^{t-1}(\alpha, \beta, \mu) \approx \mathbb{I}^t(\beta, \alpha, \mu')$. This refinement aims to be *fair* in sharing α 's private information. Unlike the quasi-utilitarian measures, both the measure of information gain and the semantic measure (see below) apply to *all* messages.

The structure of the ontology is provided by the $\text{Sim}(\cdot)$ function. Given two contracts δ and δ' containing concepts $\{o_1, \dots, o_i\}$ and $\{o'_1, \dots, o'_j\}$ respectively, the (non-symmetric) distance of δ' from δ is the vector $\vec{\Gamma}(\delta, \delta') = (d_k : o'_k)_{k=1}^j$ where $d_k = \min_x \{\text{Sim}(o_k, o'_x) | x = 1, \dots, j\}$, $o'_k = \sup(\arg \min_x \{\text{Sim}(o_k, x) | x = o'_1, \dots, o'_j\}, o_k)$ and the function $\sup(\cdot, \cdot)$ is the supremum of two concepts in the ontology. $\vec{\Gamma}(\delta, \delta')$ quantifies how different δ' is to δ and enables α to "work around" or "move away from" a contract under consideration.

More generally, suppose that α has selected a plan s that is attempting to select the 'best' action in some sense. Then s will construct a distribution $D = (d_i) \in \mathcal{M}^t$ where d_i is the probability that z_i is the best action that α can take, and $\{z_1, \dots, z_r\}$ are the available actions. s reduces $\mathbb{H}(D)$ to an acceptable level before committing itself to act — here α is an

uncertainty minimising agent. s reduces the entropy of \mathcal{M} by acquiring information as we now describe.

Distribution D will typically be defined in terms of other distributions. For example, the ‘best’ action may be defined in terms of the trustworthiness of β (Section 4), quasi-utilitarian distributions such as those described in Section 3.1, measures of information gain in Section 3.3, and measures of semantic context in Section 2.1. Suppose that D is defined in terms of $\{X_i\}$. The integrity of each X_i will develop in time by Eqn. 4. The update functions for X_i , J^{X_i} , identify potential ‘entropy reducing’ information that may perhaps be obtained from the information sources $\{\theta_i\}_{i=1}^t$. If α requests information φ from θ_j that is delivered as φ' it is managed as described in Section 3.4. That section also showed how nearby observations, (ϕ', ϕ) , may be used to update distributions, so identifying additional information that may reduce entropy.

Let $M_{\alpha\beta}^t$ be the set of time-stamped messages that α has sent to β , and $M_{\beta\alpha}^t$ likewise both at time t . The next action that α takes in this dialogue will be to send β a message μ that is determined by its strategy. α 's *strategy* is a function $f(\cdot)$ of the following form:

$$\begin{aligned} \mu &= \arg \max_z f(\mathbb{I}^t(\beta, \alpha, z), \\ &\{\mathbb{I}^s(\beta, \alpha, x^s) \mid x^s \in M_{\alpha\beta}^t\}, \{\mathbb{I}^s(\alpha, \beta, x^s) \mid x^s \in M_{\beta\alpha}^t\}, \\ &\{\vec{\Gamma}(z, x^s) \mid x^s \in M_{\alpha\beta}^t \cup M_{\beta\alpha}^t\}, \\ &(\mathbb{P}^t(\text{acc}(\beta, \alpha, z), \mathbb{P}^t(\alpha, \beta, \chi, z)) \mid z \text{ is a proposal}) \end{aligned}$$

It takes account of the acceptability (expected utility) of proposals, the expected information gain in each communication and the semantic relationship between proposals, both as seen by itself and by β .

6 Conclusion

Information-based agency integrates in an homogeneous agent architecture the exchange of information and the negotiation of norms and contracts using a rich communication language. These agents' world models are manipulated by the agent's proactive and reactive machinery using ideas from information theory. These agents are aware of the value of the information in every utterance, the semantic relationship between utterances as well as their utilitarian value. Summary measures that generalise trust, reliability and reputation illustrate the strength of the architecture. These agents' negotiation strategies are based on three dimensions: informational, rational and semantic.

Acknowledgments Carles Sierra has been supported by the EU funded STREP project OpenKnowledge and the Spanish MEC project Webi-2. Both authors are supported by the Australian ARC Discovery Project ‘The Curious Negotiator’.

References

[Arcos *et al.*, 2005] Josep Lluís Arcos, Marc Esteva, Pablo Noriega, Juan Antonio Rodríguez, and Carles Sierra. Environment engineering for multiagent systems. *Journal on Engineering Applications of Artificial Intelligence*, 18, 2005.

[Cheeseman and Stutz, 2004] P. Cheeseman and J. Stutz. *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, chapter On The Relationship between Bayesian and Maximum Entropy Inference, pages 445 – 461. American Institute of Physics, Melville, NY, USA, 2004.

[Debenham, 2004] J. Debenham. Bargaining with information. In N.R. 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 Press, New York, July 2004.

[Faratin *et al.*, 2003] 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.

[Jennings *et al.*, 2001] Nick Jennings, Peyman Faratin, Alessandro Lomuscio, Simon Parsons, Carles Sierra, and Mike Wooldridge. Automated negotiation: Prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, 10(2):199–215, 2001.

[Kalfoglou and Schorlemmer, 2003] Yannis Kalfoglou and Marco Schorlemmer. IF-Map: An ontology-mapping method based on information-flow theory. In Stefano Spaccapietra, Sal March, and Karl Aberer, editors, *Journal on Data Semantics I*, volume 2800 of *Lecture Notes in Computer Science*, pages 98–127. Springer-Verlag: Heidelberg, Germany, 2003.

[Kraus, 1997] S. Kraus. Negotiation and cooperation in multi-agent environments. *Artificial Intelligence*, 94(1–2):79–97, 1997.

[Li *et al.*, 2003] Yuhua Li, Zuhair A. Bandar, and David McLean. An approach for measuring semantic similarity between words using multiple information sources. *IEEE Transactions on Knowledge and Data Engineering*, 15(4):871 – 882, July / August 2003.

[MacKay, 2003] D. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003.

[Paris, 1999] J. Paris. Common sense and maximum entropy. *Synthese*, 117(1):75 – 93, 1999.

[Rosenschein and Zlotkin, 1994] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. The MIT Press, Cambridge, USA, 1994.

[Sierra and Debenham, 2006] C. Sierra and J. Debenham. Trust and honour in information-based agency. In *Proceedings Fifth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2006*, Hakodate, Japan, May 2006. ACM Press, New York.