

An Agent Architecture for an Uncertain World

John Debenham

University of Technology, Sydney
Faculty of Engineering and IT
PO Box 123, Broadway 2007, Australia
debenham@it.uts.edu.au

Carles Sierra

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

Abstract

Successful negotiators look beyond a purely utilitarian view. We propose a new agent architecture that was inspired by the observation that “Everything that an agent says gives away (valuable) information.” It is intended for agents who are uncertain about their environment. Information-based agency uses tools from information theory, and includes techniques for managing information exchange including: the acceptance of contracts, the estimation of trust, reliability, honour and confidence.

1 Introduction

An agency framework is grounded in information-based concepts [4]; it aims to provide a basis for managing the uncertainties that pervade real-world negotiation such as eProcurement contract negotiation. In so far as uncertainty is the problem, information is the solution — “*It’s what you know that matters*”. Information-based agency provides a range of tools and techniques that enable agents to manage their information resources, including the proactive acquisition of information, in the presence of pervasive uncertainty.

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 [1]. 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). Agents have an *internal language* \mathcal{L} used to build a probabilistic *world model*.

The world model, \mathcal{M}^t consists of probability distributions that represent the agent’s (uncertain) beliefs about the world. These distributions are in effect summaries of the observations that the agent has made. Information-based agents also construct coarse-grained summary measures (e.g. trust, reputation, and reliability [6]) 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.

Agent α receives all utterances expressed in \mathcal{C} in an inbox \mathcal{X} where they are time-stamped and sourced-stamped. An utterance u 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 $\mathbb{R}^t(\alpha, \beta, u)$ that normally decays with time. α acts in response to an utterance 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 2, other messages are dealt with by α ’s reactive reasoning described in Section 3.

2 Proactive Reasoning

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.

The distributions in \mathcal{M}^t are determined by α ’s plans that are in turn 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 sign a proposed contract leading to a probability of signing $\mathbb{P}^t(\text{sign}(\alpha, \beta, \chi, \delta))$. This is discussed further in Section ??.

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 , $\mathbb{P}^t(\text{confide}(\alpha, \beta, o) = f_j)$ that α can trust β with confidential information in context o , and any other distribution as required.

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*, $\mathbb{D}(X_i)$, for the random variable X_i . 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. In practice, the decay limit distribution may be known. For example, if X_i represents α 's beliefs about tomorrow's weather then the published averages of the weather for that time of year may be used to generate $\mathbb{D}(X_i)$.

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 Reactive Reasoning

In the absence of in-coming utterances the integrity of \mathcal{M}^t decays by Equation 1. The following procedure updates \mathcal{M}^t for all percepts expressed in language \mathcal{C} . Suppose that α receives a message u 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 $\mathbb{R}^t(\alpha, \beta, u)$ to u — this probability is α 's estimate of β 's *confidence* that reflects α 's level of personal *caution*. Each of α 's active plans, s , contains constructors for a set of distributions $\{X_i\} \subseteq \mathcal{M}^t$ together with associated *update functions*, $J_s(\cdot)$, such that $J_s^{X_i}(u)$ is a set of linear constraints on the posterior distribution for X_i . Examples of these update functions are given in Section ??.

Denote the prior distribution $\mathbb{P}^t(X_i)$ by \vec{p} , and let $\vec{p}_{(u)}$ be the distribution with minimum relative entropy¹ with respect to \vec{p} : $\vec{p}_{(u)} = \arg \min_{\vec{r}} \sum_j r_j \log \frac{r_j}{p_j}$ that satisfies the constraints $J_s^{X_i}(u)$. Then let $\vec{q}_{(u)}$ be the distribution:

¹ Given 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 [2] and encapsulates common-sense reasoning [5].

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

and then let:

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

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

Finally, merging Equation 3 and Equation 1, we obtain the method for updating a distribution X_i on receipt of an utterance u :

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

This procedure deals with integrity decay, and with two probabilities: first, the probability z in the percept u , and second the belief $\mathbb{R}^t(\alpha, \beta, u)$ that α attached to u . Equation 3 is intended to prevent weak information from decreasing the certainty of $\mathbb{P}^{t+1}(X_i)$.

3.1 Confidence

The confidence estimate, $\mathbb{R}^t(\alpha, \beta, u)$, mentioned in the previous section is an epistemic probability that takes account of α 's personal caution. An empirical estimate of $\mathbb{R}^t(\alpha, \beta, u)$ may be obtained by measuring the 'difference' between commitment and enactment. Suppose that utterance u is received from agent β at time t entailing a commitment $\text{Commit}(\beta, \alpha, \varphi)$, e.g. $u = \text{Accept}(\beta, \alpha, \varphi, t)$ (and is verified by ξ as $u' = \text{Inform}(\xi, \alpha, \text{Observe}(\xi, \varphi'), t')$ at some later time t' . Denote the prior $\mathbb{P}^t(X_i)$ by \vec{p} . Let $\vec{p}_{(u)}$ be the posterior minimum relative entropy distribution subject to the constraints $J_s^{X_i}(u)$, and let $\vec{p}_{(u')}$ be that distribution subject to $J_s^{X_i}(u')$. We now estimate what $\mathbb{R}^t(\alpha, \beta, u)$ should have been in the light of knowing *now*, at time t' , that φ should have been φ' .

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

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

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

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

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

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

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

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

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

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

For each commitment $\text{Commit}(\beta, \alpha, \varphi)$ established at time t and derived from an utterance u that has been verified with as $\text{Observe}(\alpha, \varphi')$ derived from u' , the *observed confidence* that α has for agent β in φ is:

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

where Sim measures the semantic distance between two sections of the ontology, and ν is the learning rate. Over time, α notes the context of the various $\text{Commit}(\beta, \alpha, \varphi)$ derived from utterances u received from β , and over the various contexts calculates the relative frequency, $\mathbb{P}^t(\text{Commit}(\beta, \alpha, \varphi))$.

4 Uncertain Information

Agent α 's world model, \mathcal{M}^t , at time t is a set of random variables, $\mathcal{M}^t = \{X_i, \dots, X_n\}$ each representing an aspect of the world that α is interested in. In a multiagent system it is natural to measure the uncertainty of a random variable in terms of the cost, in some sense, of communicating the true value of it from one agent to another. One such sense is the lower bound on the expected number of binary questions that are always guaranteed to discover the true value of a random variable, X ; this is given by the *entropy*, $\mathbb{H}(X) = \sum_i -p_i \log p_i$, where the p_i are values of the probability mass function for X , [4].

When a negotiation terminates it is normal for agents to review what the cost of the negotiation *ex post*; for example, "I got him to sign up, but had to tell him about our plans to close our office in Girona". It is not feasible to define the value of information in terms of the value derived from subsequent enactments that may have benefited from that information unless those subsequent enactments are known in advance. Without knowing what use the recipient will

make of the "Girona information", it is not possible to relate the value of this act of information revelation to the value of outcomes and so to preferences.

If β passes an utterance to α , α evaluates this act in two ways. First, it is valued for the strategic significance of the information that it contains, precisely it is measured as the expected increase in utility that α expects to enjoy given that it has the information. Second, it is valued because the sending agent *was prepared to divulge* the information in the utterance, precisely it is measured as the decrease in uncertainty that the receiving agent has over the sending agent's private information — this is the *information measure*. All utterances received are qualified by α with a confidence belief probability as described in Section 3. From α 's point of view, β 's *private information* is everything that β knows and that α does not know with certainty. Due to the persistent effect of Equation 1, this will include much of what β knows.

α 's world model, \mathcal{M}^t , is a set of probability distributions. If at time t , α receives an utterance u that may alter this world model (as described in Section 3) then the (Shannon) *information* in u with respect to the distributions in \mathcal{M}^t is: $\mathbb{I}^{t+1}(u) = \mathbb{H}(\mathcal{M}^t) - \mathbb{H}(\mathcal{M}^{t+1})$. Let $\mathcal{N}^t \subseteq \mathcal{M}^t$ be α 's model of agent β . If β sends the utterance u to α then the *information* about β within u is: $\mathbb{H}(\mathcal{N}^t) - \mathbb{H}(\mathcal{N}^{t+1})$.

5 Uncertainty in Enactment

If an agent signs a contract it is then committed to enact its commitments in that contract. If an agent transmits information it is implicitly committed to the world being in some state when conditions hold — e.g. "If you open that bottle of wine Carles will be furious." If an agent makes a threat or offers a reward — e.g. "If you purchase all of your eggs from me I will give you a 20% discount." — then this too commits the speaker to ensuring that the world is in some state under certain conditions. In summary, all information that an agent communicates will one way or another be a commitment or a promise that the world or the agents in it will be, are, or have been in certain states possibly subject to associated conditions being true. If agent α receives a commitment from β , α will be interested in any variation between β 's commitment, φ , and what is actually observed (or is advised by the institution agent ξ), as the enactment, φ' . We denote the relationship between commitment and enactment, $\mathbb{P}^t(\text{Observe}(\alpha, \varphi')|\text{Commit}(\beta, \alpha, \varphi))$ simply as $\mathbb{P}^t(\varphi'|\varphi) \in \mathcal{M}^t$.

In the absence of in-coming utterances the conditional probabilities, $\mathbb{P}^t(\varphi'|\varphi)$, should tend to ignorance as represented by the *decay limit distribution* and Equation 1. We now show how Equation 4 may be used to revise $\mathbb{P}^t(\varphi'|\varphi)$ as observations are made. Let the set of possible enactments be $\{\varphi_1, \varphi_2, \dots, \varphi_m\} \subseteq \Phi$ with prior distribution

$\vec{p} = \mathbb{P}^t(\varphi'|\varphi)$. Suppose that utterance u is received, we estimate the posterior $\vec{p}_{(u)} = (p_{(u)i})_{i=1}^m = \mathbb{P}^{t+1}(\varphi'|\varphi)$, as follows.

First, if $u \in \mathcal{C}$ entails that φ_k is observed then α may use this observation to estimate $p_{(\varphi_k)k}$ as having some value d at time $t+1$. We then estimate the distribution $\vec{p}_{(\varphi_k)}$ by applying the principle of minimum relative entropy as in Equation 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 by appealing to the structure of the ontology using the $\text{Sim}(\cdot)$ function. Given $\text{Observe}(\alpha, \phi')$ after $\text{Commit}(\beta, \alpha, \phi)$ define the vector \vec{s} by

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

for $i = 1, \dots, m$. \vec{s} 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{s} .

6 Uncertainty in Preferences

Agent α 's *preferences* is a relation defined over a set, S , where $s_1 \prec_\alpha s_2, s_1, s_2 \in S$ denotes “ α prefers s_2 to s_1 ”. When discussing preferences in negotiation, the set S is often the *outcome space*. Elements in an outcome space may be described either by the world being in a certain state or by a concept in the ontology having a certain value. If an agent knows its preferences then it may use results from game theory or decision theory to achieve a preferred outcome in some sense. For example, an agent may prefer the concept of price (from the ontology) to have lower values than higher, or to purchase wine when it is advertised at a discount (a world state).

In practice, the articulation of a preference relation may not be simple, particularly across a large multi-issue space. Consider the problem of specifying a preference relation for a collection of fifty cameras with different features, from different makers, with different prices, both new and second-hand. This is a multi-issue evaluation problem. It is realistic to suggest that “a normal intelligent human being” may not be able to place the fifty cameras in a preference ordering with certainty, or even to construct a meaningful probability distribution to describe it. The complexity of articulating preferences over real negotiation spaces poses a practical limitation on the application of preference-based strategies.

In contract negotiation the outcome of the negotiation, (a', b') , is the enactment of the commitments, (a, b) , in that

contract, where a is α 's commitment and b is β 's. Some of the great disasters in market design [3], for example the Australian Foxtel fiasco, could have been avoided if the designers had considered how the agents were expected to enact (a', b') their commitments (a, b) after the contract is signed.

Consider a contract (a, b) , and let $(\mathbb{P}_\alpha^t(a'|a), \mathbb{P}_\alpha^t(b'|b))$ denote α 's estimate of what will be enacted by α and β respectively if (a, b) is signed. Further assume that the pair of distributions $\mathbb{P}_\alpha^t(a'|a)$ and $\mathbb{P}_\alpha^t(b'|b)$ are independent², and that α is able to estimate $\mathbb{P}_\alpha^t(a'|a)$ with confidence. α will only be confident in her estimate of $\mathbb{P}_\alpha^t(b'|b)$ if β 's actions are constrained by norms, or if α has established a high degree of trust in β . If α is unable to estimate $\mathbb{P}_\alpha^t(b'|b)$ with reasonable certainty then put simply: she won't know what she is signing. For a utilitarian α , $(a_1, b_1) \prec_\alpha (a_2, b_2)$ if she prefers $(\mathbb{P}_\alpha^t(a'_2|a_2), \mathbb{P}_\alpha^t(b'_2|b_2))$ to $(\mathbb{P}_\alpha^t(a'_1|a_1), \mathbb{P}_\alpha^t(b'_1|b_1))$ in some sense.

7 Summary

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.

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] 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.
- [3] P. Klemperer. What really matters in auction design. *The Journal of Economic Perspectives*, 16(1):169–189, 2002.
- [4] D. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003.
- [5] J. Paris. Common sense and maximum entropy. *Synthese*, 117(1):75 – 93, 1999.
- [6] C. Sierra and J. Debenham. Trust and honour in information-based agency. In P. Stone and G. Weiss, editors, *Proceedings Fifth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2006*, pages 1225 – 1232, Hakodate, Japan, May 2006. ACM Press, New York.

²That is we assume that while α is executing commitment a she is oblivious to how β is executing commitment b and *vice versa*.