Information-Based Planning and Strategies

John Debenham

Abstract The foundations of information-based agency are described, and the principal architectural components are introduced. The agent's deliberative planning mechanism manages interaction using plans and strategies in the context of the relationships the agent has with other agents, and is the means by which those relationships develop. Finally strategies are described that employ the deliberative mechanism and manage argumentative dialogues with the aim of achieving the agent's goals.

1 Introduction

This paper is in the area labelled: *information-based agency* [9]. Information-based agency is founded on two premises. First, everything in its world model is uncertain [2]. Second, everything that an agent communicates gives away valuable information. Information, including arguments, may have no particular utilitarian value [6], and so may not readily be accommodated by an agent's utilitarian machinery.

An information-based agent has an identity, values, needs, plans and strategies all of which are expressed using a fixed ontology in probabilistic logic for internal representation and in an illocutionary language [8] for communication. All of the forgoing is represented in the agent's deliberative machinery. We assume that such an agent resides in a electronic institution [1] and is aware of the prevailing norms and interaction protocols. In line with our "Information Principle" [8], an information-based agent makes no *a priori* assumptions about the states of the world or the other agents in it — these are represented in a world model, $\mathcal{M}^I$, that is inferred solely from the messages that it receives.

The world model, $\mathcal{M}^I$, is a set of probability distributions for a set of random variables each of which represents the agent's expectations about some point of

John Debenham

University of Technology, Sydney, Australia e-mail: debenham@it.uts.edu.au

Please use the following format when citing this chapter:

Debenham, J., 2008, in IFIP International Federation for Information Processing, Volume 276; *Artificial Intelligence and Practice II*; Max Bramer; (Boston: Springer), pp. 4554.

interest about the world or the other agents in it. We build a history of interaction by noting each commitment made (commitments to act, commitments to the truth of information or to the validity of an argument), and by relating each of them to subsequent observations of what occurs. Tools from information theory are then used to summarise these historic (commitment, observation) pairs — in this way we have defined models of *trust*, *honour*, *reliability* and *reputation* [8]. Further we have defined the *intimacy* and *balance* of both dialogues and relationships [10] in terms of our ‘LOGIC’ illocutionary framework. All of these notions make no presumption that our agents will align themselves with any particular strategy.

In related papers we have focussed on argumentation strategies, trust and honour, and have simply assumed that the agent has a kernel deliberative system. In this paper we describe the deliberative system for an information-base agent.

2 Plans

A plan p is $p(a_p, s_p, t_p, u_p, c_p, g_p)$ where:

- a_p is a conditional action sequence — i.e. it is conditional on future states of the world, and on the future actions of other agents. We think of plans as probabilistic statecharts in the normal way where the arcs from a state are labelled with “event / condition / action” leading into a P symbol that represents the lottery, s_p , that determines the next state as described following:
- $s_p : S \rightarrow \mathbb{P}(S_p = s) \equiv \mathbf{s}$ where S is the set of states and S_p is a random variable denoting the state of the world when a_p terminates¹.
- $t_p : S \rightarrow \mathbb{P}(T_p = t) \equiv \mathbf{t}$ where T_p is a random variable denoting the time that a_p takes to execute and terminate for some finite set of positive time interval values for t .
- $u_p : S \rightarrow \mathbb{P}(U_p = u) \equiv \mathbf{u}$ where U_p is a random variable denoting the gross utility gain, excluding the cost of the execution of a_p for some finite set of utility values for u .
- $c_p : S \rightarrow \mathbb{P}(C_p = c) \equiv \mathbf{c}$ where C_p is a random variable denoting the cost of the execution of a_p for some finite set of cost values for c .
- $g_p : S \rightarrow \mathbb{P}(G_p = g) \equiv \mathbf{g}$ where G_p is a random variable denoting the expected information gain to α and to β of the dialogue that takes place during the execution of the plan each expressed in $\mathcal{G} = \mathcal{F} \times \mathcal{O}$.

The distributions above are estimated by observing the performance of the plans as we now describe.² In the absence of any observations the probability mass functions for S_p , T_p , U_p , C_p and G_p all decay at each and every time step by:

¹ For convenience we assume that all action sequences have a “time out” and so will halt after some finite time.

² An obvious simplification would be to use point estimates for t_p , u_p , c_p and each element of g_p , but that is too weak a model to enable comparison.

$$\mathbb{P}^{t+1}(X_i) = \lambda \times \mathbb{D}(X_i) + (1 - \lambda) \times \mathbb{P}^t(X_i) \quad (1)$$

for some constant $\lambda : 0 < \lambda < 1$, where λ is the *decay rate*.

The implementation of a_p does not concern us. We do assume that the way in which the plans are implemented enables the identification of common algorithms and maybe common methods within different plans. Given two plans p and q , the function $\text{Sim}(p, q) \in [0, 1]$ measures the similarity of their action sequences a_p and a_q in the sense that their performance parameters are expected to be correlated to some degree.

Estimating S_p . Denote the prior estimate by $\mathbf{s}^t$. When a plan terminates, or is terminated, the world will be in one of p 's end states. Call that state z . Then the observed distribution for $\mathbf{s}^{t+\delta t}$ will have the value 1 in position z . On the basis of this observation the agent may be inclined to fix its estimate for s_z^{t+1} at γ where $s_z^t \leq \gamma \leq 1$. The posterior distribution $\mathbf{s}^{t+1}$ is defined as the distribution with minimum relative entropy with respect to $\mathbf{s}^t$: $\mathbf{s}^{t+1} = \arg \min_{\mathbf{r}} \sum_j r_j \log \frac{r_j}{s_j^t}$ that satisfies the constraint $s_z^{t+1} = \gamma$. If $\gamma = s_z^t$ then the posterior is the same as the prior. If $\gamma = 1$ then the posterior is certain with $\mathbb{H}(\mathbf{s}^{t+1}) = 0$. One neat way to calibrate γ is in terms of the resulting information gain; that is to measure γ in terms of the resulting *learning rate* μ :

$$\mathbb{H}(\mathbf{s}^{t+1}) = (1 - \mu) \times \mathbb{H}(\mathbf{s}^t) \quad (2)$$

where $\mu : 0 < \mu < 1$.

Estimating T_p , U_p , C_p and G_p . Just as for estimating S_p , when the plan terminates α will have observations for the values of these variables, and as a result may wish to increase the corresponding frequency in the posterior to some new value. Using the method described above for estimating S_p , the posterior distribution is the distribution with minimum relative entropy with respect to the prior subject to the constraint that the frequency corresponding to the observation is increased accordingly.

Further, for these four variables we use the $\text{Sim}(\cdot, \cdot)$ function to revise the estimates for ‘nearby’ plans. In [9] two methods for using a $\text{Sim}(\cdot, \cdot)$ function to revise estimates are described — the situation here is rather simpler. Consider the variable C_p . Applying the method in the paragraph ‘Estimating S_p .’, suppose a value had been observed for C_p and as a result of which c_j^{t+1} had been constrained to be γ . Consider any plan q for which $\text{Sim}(p, q) > 0$. Denote $\mathbb{P}(C_q = c)$ by $\mathbf{d}$. The posterior distribution $\mathbf{d}^{t+1}$ is defined as the distribution with minimum relative entropy with respect to $\mathbf{d}^t$: $\mathbf{d}^{t+1} = \arg \min_{\mathbf{r}} \sum_j r_j \log \frac{r_j}{d_j^t}$ that satisfies the constraint: $d_j^{t+1} = \gamma'$ where γ' is such that:

$$\mathbb{H}(\mathbf{d}^{t+1}) = (1 - \mu \times \text{Sim}(p, q)) \times \mathbb{H}(\mathbf{d}^t) \quad (3)$$

where $0 \leq \text{Sim}(p, q) \leq 1$ with higher values indicating greater similarity.

3 Planning

If an agent's needs could potentially be satisfied by more than one plan then a mechanism is required to select which plan to use. As the execution of plans incurs a cost we assume that α won't simply fire off every plan that may prove to be useful. A random variable, V_p , derived from the expectations of S_p , T_p , U_p , C_p , G_p and other estimates in $\mathcal{M}^t$ represents the agent's expectations of each plan's overall *performance*. V_p is expressed over some finite, numerical valuation space with higher values being preferred.

The mechanisms that we describe all operate by selecting plans stochastically. We assume that there is a set of P candidate plans $\{p_i\}$ with corresponding random variables V_{p_i} representing performance, and plan p_j is chosen with probability q_j where $\sum_k q_k = 1$. Let $\mathcal{N}^t = \{V_{p_k}^t\}_{k=1}^P$. The integrity of the performance estimates for random variable V_{p_i} are maintained using the method "Estimating S_p " in Section 2. If p_i is selected at time t then when it terminates the observed performance, $v_{p_i, \text{ob}}^t$, is fed into that method.

First, consider the naïve mechanism that selects plan p_j by: $q_j = 1$ for $j = \arg \max_i \mathbb{E}(V_{p_i})$. This mechanism is well-suited to a one-off situation. But if the agent has continuing need of a set of plans then choosing the plan with highest expected payoff may mean that some plans will not be selected for a while by which time their performance estimates will have decayed by Equation 1 to such a extent that may never be chosen. An agent faces the following dilemma: the only way to preserve a reasonably accurate estimate of plans is to select them sufficiently often — even if they they don't perform well today perhaps one day they will shine.

The simple method: $q_i = \frac{1}{P}$ selects all plans with equal probability. The following method attempts to prevent the uncertainty of estimates from decaying above a threshold, τ , by setting $q_j = 1$ where:

if $\exists i \cdot \mathbb{H}(V_{p_i}) > \tau$ **then** let $j = \arg \max_k \mathbb{H}(V_{p_k})$
else let $j = \arg \max_k \mathbb{E}(V_{p_k})$

this method may deliver poor performance from the '**then**' and good performance from the '**else**', but at least it attempts to maintain some level of integrity of the performance estimates, even if it does so in an elementary way.

A strategy is reported in [4] on how to place all of one's wealth as win-bets indefinitely on successive horse races so as to maximise the rate of growth; this is achieved by proportional gambling, i.e. by betting a proportion of one's wealth on each horse equal to the probability that that horse will win. This result is interesting as the strategy is independent of the betting odds. Whether it will make money will depend on the punter's ability to estimate the probabilities better than the bookmaker. The situation that we have is not equivalent to the horse race, but it is tempting to suggest the strategies:

$$q_i = \frac{\mathbb{E}(V_{p_i})}{\sum_k \mathbb{E}(V_{p_k})} \quad (4)$$

$$q_i = \mathbb{P}(V_{p_i} > V_{p_j}), \forall V_{p_j} \in \mathcal{N}, j \neq i \quad (5)$$

For the second strategy: q_i is the probability that p_i 's performance is the better than that of all the other plans. With this definition it is clear that $\sum_i q_i = 1$. Both strategies will favour those plans with a better performance history. Whether they will prevent the integrity of the estimates for plans with a poor history from decaying to a meaningless level will depend on the value of λ in Equation 1, the value of μ in Equation 2, and on the frequency with which plans are activated. As the estimates for plans that perform well, and plans that perform badly, all decay to the maximum entropy decay limit $\mathbb{D}(V_{p_i})$ if they are not invoked, both of these strategies indirectly take account of the level of certainty in the various performance estimates.

We consider now the stability of the integrity of the performance estimates in time. If plan p_j is *not* executed the information loss in X_j^t for one time step due to the effect of Equation 1 is: $\lambda \times \mathbb{H}(X_j^t)$. If no plans in $\mathcal{N}$ are executed during one time step then the total information loss in $\mathcal{N}$ is: $\lambda \times \sum_k \mathbb{H}(X_k^t)$. If plan p_j is executed the information gain in X_j^t due to the effect of Equation 2 is: $\mu \times \mathbb{H}(X_j^t)$, but this observation may effect the other variables in $\mathcal{N}^t$ due to Equation 3, and the total information gain in $\mathcal{N}$ is: $\mu \times \sum_k \text{Sim}(p_j, p_k) \times \mathbb{H}(X_k^t)$. Assuming that at most one plan in $\mathcal{N}^t$ is executed during any time step, and that the probability of one plan being executed in any time step is χ ; the expected net information gain of $\mathcal{N}^{t+1}$ compared with $\mathcal{N}^t$ is:

$$\chi \cdot \mu \cdot \sum_j q_j \cdot \sum_k \text{Sim}(p_j, p_k) \cdot \mathbb{H}(X_k^t) - \lambda \cdot \sum_k \mathbb{H}(X_k^t) \quad (6)$$

If this quantity is negative then the agent may decide to take additional steps to gain performance measurements so as to avoid the integrity of these estimates from consistently declining.

We now consider the parameters λ and μ to be used with the strategy in Equation 4. The effect of Equation 1 on variable V_i after t units of time is:

$$(1 - (1 - \lambda)^t) \times \mathbb{D}(V_{p_i}) + (1 - \lambda)^t \times V_{p_i}^{t_0}$$

The probability that plan p_i will be activated at any particular time is:

$$\chi \times \frac{\mathbb{E}(V_{p_i})}{\sum_k \mathbb{E}(V_{p_k})}$$

and the mean of these probabilities for all plans is: $\frac{\chi}{\bar{p}}$. So the mean number of time units between each plan's activation is: $\frac{N}{\chi}$. In the absence of any intuitive value for λ , a convenient way to calibrate λ is in terms of the expected total decay towards $\mathbb{D}(V_{p_i})$ between each activation — this is expressed as some constant ϕ , where $0 < \phi < 1$. For example, $\phi = \frac{1}{2}$ means that we expect a 50% decay between activations. The value of λ that will achieve this is: $\lambda = 1 - (1 - \phi)^{\chi \div N}$. Then the value for μ is chosen so that the expression (6) is non-negative. Using these values should ensure that the probability distributions for the random variables V_i remain within reasonable bounds, and so remain reasonably discriminating.

It would be nice to derive a method that was optimal in some sense, but this is unrealistic if the only data available is historic data such as the V_{p_i} . In real situations the past may predict the future to some degree, but can not be expected to predict performance outcomes that are a result of interactions with other autonomous agents in a changing environment. As a compromise, we propose to use (5) with values for λ and μ determined as above. (5) works with the whole distribution rather than (4) that works only with point estimates, but is algebraically simpler. These methods are proposed on the basis that the historic observations are all that α has.

4 Preferences

Agent α 's *preferences* is a relation defined over an *outcome space*, where $s_1 \prec_\alpha s_2$ denotes “ α prefers s_2 to s_1 ”. Elements in the 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.

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 [5], for example the Australian Foxtel fiasco, could have been avoided if the designers had considered how the agents were expected to deviate (a', b') from 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 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 independent [3]³ 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.

³ That is we assume that while α is executing commitment a she is oblivious to how β is executing commitment b and *vice versa*.

One way to manage contract acceptance when the agent's preferences are unknown is to found the acceptance criterion instead on the simpler question: "how certain am I that (a, b) is a good contract to sign?" — under realistic conditions this is easy to estimate⁴.

So far we have not considered the management of information exchange. When a negotiation terminates it is normal for agents to review what the negotiation has cost *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 attach an intrinsic value to information that is related to the value derived from enactments. 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 outcomes and so to preferences.

While this negotiation is taking place how is the agent to decide whether to reveal the "Girona information"? He won't know then whether the negotiation will terminate with a signed contract, or what use the recipient may be able to make of the information in future, or how any such use might affect him. In general it is unfeasible to form an expectation over these things. So we argue that the decision of whether to reveal a piece of information should *not* be founded on anticipated negotiation outcomes, and so this decision should not be seen in relation to the agent's preferences. The difficulty here is that value is derived from information in a fundamentally different way to the realisation of value from owning a commodity, for example⁵.

A preference-based strategy may call upon powerful ideas from game theory. For example, to consider equilibria α will require estimates of $\mathbb{P}'_{\beta}(a'|a)$ and $\mathbb{P}'_{\beta}(b'|b)$ in addition to $\mathbb{P}'_{\alpha}(a'|a)$ and $\mathbb{P}'_{\alpha}(b'|b)$ — these estimates may well be even more speculative than those in the previous paragraph. In addition she will require knowledge about β 's utility function. In simple situations this information may be known, but in general it will not.

5 Information-based strategies

An information-based agent's deliberative logic consists of:

1. The agent's *raison d'être* — its mission — this may not be represented in the agent's code, and may be implicit in the agent's design.

⁴ In multi-issue negotiation an agent's preferences over each individual issue may be known with certainty. Eg: she may prefer to pay less than pay more, she may prefer to have some feature to not having it. In such a case, if some deals are known to be unacceptable with certainty, some are known to be acceptable with certainty, and, perhaps some known to be acceptable to some degree of certainty then maximum entropy logic may be applied to construct a complete distribution representing 'certainty of acceptability' over the complete deal space. This unique distribution will be *consistent* with what is known, and *maximally noncommittal* with respect to what is *not* known.

⁵ If a dialogue is not concerned with the exchange of anything with utilitarian value, then the two agents may feel comfortable to balance the information exchanged using the methods in [10].

2. A set of *values*, Π , — high level principles — and a fuzzy function $v : (S \times A \times \Pi) \rightarrow \text{fuz}$, that estimates, when the world is in state $s \in S$, whether the agent performing action $a \in A$ supports or violates a value $\pi \in \Pi$.
3. A *strategy* that provides an overarching context within which the plans are executed — see Section 5. The strategy is responsible for the evolution of the relationships between the agents, and for ensuring that plans take account of the state of those relationships.
4. A hierarchy⁶ of *needs*, N , and a function $\sigma : N \rightarrow \mathcal{P}(S)$ where $\sigma(n)$ is the set of states that satisfy need $n \in N$. Needs turn ‘on’ spontaneously, and in response to *triggers*, T ; they turn ‘off’ because the agent believes they are satisfied.
5. A set of *plans*, P — Section 2.

In this model an agent knows with certainty those states that will satisfy a need, but does *not* know with certainty what state the world is in.

We now describe the strategic reasoning of an information-based agent. This takes account of the, sometimes conflicting, utilitarian and information measures of utterances in dialogues and relationships. This general definition may be instantiated by specifying functions for the ψ_i in the following.

The following notation is used below. R_i^t denotes the relationship (i.e. the set of all dialogues) between α and β_i at time t . *Intimacy* is a summary measure of a relationship or a dialogue and is represented in $\mathcal{G}$. We write I_i^t to denote the intimacy of that relationship, and $I(d)$ to denote the intimacy of dialogue d . Likewise B_i^t and $B(d)$ denotes balance.

The Needs Model. α is driven by its needs. When a need fires, a plan is chosen to satisfy that need using the method in Section 3. If α is to contemplate the future she will need some idea of her future needs — this is represented in her *needs model*: $v : \mathcal{T} \rightarrow \times^n[0, 1]$ where $\mathcal{T}$ is time, and: $v(t) = (n_1^t, \dots, n_N^t)$ where $n_i^t = \mathbb{P}(\text{need } i \text{ fires at time } t)$.

Setting Relationship Targets. On completion of each dialogue of which α is a part, she revises her aspirations concerning her intimacy with all the other agents. These aspirations are represented as a *relationship target*, T_i^t , for each β_i , that is represented in $\mathcal{G}$. Let $\mathbf{I}^t = (I_1^t, \dots, I_o^t)$, $\mathbf{B}^t = (B_1^t, \dots, B_o^t)$ and $\mathbf{T}^t = (T_1^t, \dots, T_o^t)$, then $\mathbf{T}^t = \psi_1(v, \mathbf{I}^t, \mathbf{B}^t)$ — this function takes account of all β_i and aims to encapsulate an answer to the question: “Given the state of my relationships with my trading partners, what is a realistic set of relationships to aim for in satisfaction of my needs?”.

Activating Plans. If at time t , some of α ’s active needs, N_{active}^t , are not adequately⁷ being catered for, N_{neglect}^t , by existing active plans, P_{active}^t , then select P_{active}^{t+1} to take account of those needs:

⁶ In the sense of the well-known Maslow hierarchy [7], where the satisfaction of needs that are lower in the hierarchy take precedence over the satisfaction of needs that are higher.

⁷ For each need n , $\sigma(n)$ is the set of states that will satisfy n . For each active plan p , $\mathbb{P}(S_p = s)$ is probability distribution over the possible terminal states for p . During p ’s execution this initial estimation of the terminal state is revised by taking account of the known terminal states of executed sub-plans and $\mathbb{P}(S_{p'} = s)$ for currently active sub-plans p' chosen by p to satisfy sub-goals. In this way we continually revise the probability that P_{active}^{t-1} will satisfy α ’s active needs.

$$P_{\text{active}}^{t+1} = \psi_2(P_{\text{active}}^t, N_{\text{neglect}}^t, N_{\text{active}}^t, \mathbf{I}^t, \mathbf{T}^t)$$

The idea being that α will wish select P_{active}^{t+1} so as to move each observed intimacy I_i^t towards its relationship target intimacy T_i^t . Having selected a plan p , $\mathbb{E}(U_p)$ and $\mathbb{E}(G_p)$ assist α to set the *dialogue target*, D_i^t , for the current dialogue [10]. In Section 3 we based the plan selection process on a random variable V_p that estimates the plan's performance in some way. If α is preference-aware then V_p may be defined in terms of its preferences.

Deactivating Plans. If at time t , a subset of α 's active plans, $P_{\text{sub}}^t \subset P_{\text{active}}^t$, adequately caters for α 's active needs, N_{active}^t , then:

$$P_{\text{active}}^{t+1} = \psi_3(P_{\text{active}}^t, N_{\text{active}}^t, \mathbf{I}^t, \mathbf{T}^t)$$

is a minimal set of plans that adequately cater for N_{active}^t in the sense described above. The idea here is that P_{active}^{t+1} will be chosen to best move the observed intimacy I_i^t towards the relationship target intimacy T_i^t as in the previous paragraph.

The work so far describes the selection of plans. Once selected a plan will determine the actions that α makes where an action is to transmit an utterance to some agent determined by that plan. Plans may be bound by interaction protocols specified by the host institution.

Executing a Plan — Options. [10] distinguishes between a *strategy* that determines an agent's Options from which a single kernel action, a , is selected; and *tactics* that wrap that action in argumentation, a^+ — that distinction is retained below. Suppose that α has adopted plan p that aims to satisfy need n , and that a dialogue d has commenced, and that α wishes to transmit some utterance, u , to some agent β_i . In a multi-issue negotiation, a plan p will, in general, determine a set of Options, $A_p^t(d)$ — if α is preference aware [Section 4] then this set could be chosen so that these options have similar utility. Select a from $A_p^t(d)$ by:

$$a = \psi_4(A_p^t(d), \Pi, D_i^t, I(d), B(d))$$

that is the action selected from $A_p^t(d)$ will be determined by α 's set of values, Π , and the contribution a makes to the development of intimacy.

If d is a bilateral, multi-issue negotiation we note four ways that information may be used to select a from $A_p^t(d)$. (1) α may select a so that it gives β_i similar information gain as β_i 's previous utterance gave to α . (2) If a is to be the opening utterance in d then α should avoid making excessive information revelation due to ignorance of β_i 's position and should say as little as possible. (3) If a requires some response (e.g. a may be an offer for β_i to accept or reject) then α may select a to give her greatest expected information gain about β_i 's private information from that response, where the information gain is either measured overall or restricted to some area of interest in $\mathcal{M}^t$. (4) If a is in response to an utterance a' from β_i (such as an offer) then α may use entropy-based inference to estimate the probability that she should accept the terms in a' using nearby offers for which she knows their acceptability with certainty [9].

Executing a Plan — Tactics. The previous paragraph determined a kernel action, a . Tactics are concerned with wrapping that kernel action in argumentation, a^+ . To achieve this we look beyond the current action to the role that the dialogue plays in the development of the relationship:

$$a^+ = \psi_5(a, T_i^t, I_i^t, I(d), B_i^t, B(d))$$

In [10] *stance* is meant as random noise applied to the action sequence to prevent other agent's from decrypting α 's plans. Stance is important to the argumentation process but is not discussed here.

6 Conclusion

In this paper we have presented a number of measures to value information including a new model of confidentiality. We have introduced a planning framework based on the kernel components of an information-based agent architecture (i.e. decay, semantic similarity, entropy and expectations). We have defined the notion of strategy as a control level over the needs, values, plans and world model of an agent. Finally, the paper overall offers a model of negotiating agents that integrates previous work on information-based agency and that overcomes some limitations of utility-based architectures (e.g. preference elicitation or valuing information).

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. J. Halpern. *Reasoning about Uncertainty*. MIT Press, 2003.
3. M. Jaeger. Representation independence of nonmonotonic inference relations. In *Proceedings of KR'96*, pages 461–472. Morgan Kaufmann, 1996.
4. J. J. Kelly. A new interpretation of information rate. *IEEE Transactions on Information Theory*, 2(3):185–189, September 1956.
5. P. Klemperer. What really matters in auction design. *The Journal of Economic Perspectives*, 16(1):169–189, 2002.
6. D. Lawrence. *The Economic Value of Information*. Springer-Verlag, 1999.
7. A. H. Maslow. A theory of human motivation. *Psychological Review*, 50:370–396, 1943.
8. 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.
9. C. Sierra and J. Debenham. Information-based agency. In *Proceedings of Twentieth International Joint Conference on Artificial Intelligence IJCAI-07*, pages 1513–1518, Hyderabad, India, January 2007.
10. C. Sierra and J. Debenham. The LOGIC Negotiation Model. In *Proceedings Sixth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2007*, Honolulu, Hawai'i, May 2007.