Argumentation with Advice

John Debenham[1] and Carles Sierra[2]

[1] QCIS, UTS, Broadway, NSW 2007, Australia
[2] IIIA, CSIC, Campus UAB, 08193 Bellaterra, Catalonia, Spain

Abstract. This paper is concerned with *rhetorical argumentation* that aims to alter the beliefs of the listener, and so to influence his future actions, as opposed to *classical argumentation* that is concerned with the generation of arguments, usually as logical proofs, for and against a given course of action. Rhetorical argumentation includes rhetoric moves such as *Threat*, *Reward* and *Appeal*. Rhetorical argumentative utterances generated by an agent contribute to the strength of its relationship with the listener. This paper examines advice and the rhetoric particle *"I advise you . . . "* that may be used to strengthen such relationships.

1 Introduction

The study of argumentation is in two camps: first, *classical argumentation* that is concerned with the generation of arguments, usually as logical proofs, for and against a given course of action that support decision making processes; and second, *rhetorical argumentation* that aims to alter the beliefs of the listener, and is the focus of this paper. The seminal work [1] builds on the notion of one argument "attacking" another; we are more interested in how to *counter* the effect of the partner agent's arguments rhetorically, and how to lead a dialogue towards some desired outcome. Rhetorical argumentation includes moves such as *Threat*, *Reward* and *Appeal*; although no formal model of the meaning of these speech acts has been proposed yet. *Argumentation* in this sense is concerned with building (business) *relationships* through shaping another agent's reasoning, beliefs and expectations [2].

Agents may attempt to counter their partner's arguments with *Inform* statements. The subject of an inform may be factual, e.g. "today is Tuesday", or non-factual, e.g. "this movie is exciting". *Opinions* are non-factual informative speech acts, they are the speaker's evaluation of a particular aspect of a thing in context, and may be used in an attempt to build relationships with the listener. *Advice* is opinion that is uttered with the aim of either changing the listeners beliefs or influencing the listener's future actions, e.g. "if I were you I would buy the Nikon". We give the semantics of advice utterances and describe their strategic use advice in argumentative dialogue [3].

In this paper an agent's rationality is based on two basic suppositions: *everything* in the world is constantly changing and not *all* facts can be known by an agent. An agent will have its model of: the world, of the other agents and of itself evolving at all time, and does not have, for instance, a *fixed* set of preferences. As it continually receives information from the environment, i.e. it is situated in it, its beliefs change. In particular, an agent changes its models both to *manage* its future dialogue and *because* of what has already been said.

C. Huemer and T. Setzer (Eds.): EC-Web 2011, LNBIP 85, pp. 136–147, 2011.

When agents engage in argumentative dialogue they may attempt to discover the objectives, needs or preferences of the other agent. This has the direct consequence of updating the model of the other agent and so enabling the conversation to progress. Section 2 discusses the communication language and advice illocutions. Section 3 proposes a rational agent architecture that contains the necessary components to give (higher-order) semantics to these illocutions in Section 4. Section 5 concludes.

2 Communication Framework

The communication language we consider, U, contains three fundamental primitives:[1] $\text{Commit}(\alpha, \beta, \varphi)$ 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, $\text{Observe}(\alpha, \varphi)$ to represent that a certain state of the world, φ, is observed, and $\text{Done}(u)$ to represent the event that a certain action u[2] has taken place. In our language, norms, contracts, and information chunks will be represented as instances of $\text{Commit}(\cdot)$ where α and β can be individual agents or institutions, U is the set of expressions u defined as:

$$
\begin{aligned}
u &::= \mathit{illoc}(\alpha, \beta, \varphi, t) \mid u; u \mid \textbf{Let } \mathit{context} \textbf{ In } u \textbf{ End} \\
\varphi &::= \mathit{term} \mid \text{Done}(u) \mid \text{Commit}(\alpha, \beta, \varphi) \mid \text{Observe}(\alpha, \varphi) \mid \varphi \wedge \varphi \mid \\
&\qquad \varphi \vee \varphi \mid \neg\varphi \mid \forall v.\varphi_v \mid \exists v.\varphi_v \\
\mathit{context} &::= \varphi \mid id = \varphi \mid \mathit{prolog\_clause} \mid \mathit{context}; \mathit{context}
\end{aligned}
$$

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 u, and term represents logical predicates. t represents a point in time.[3] We will note by Φ the set of expressions φ used as the propositional content of illocutions.

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:

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

or, "If I tell you who I buy my tomatoes from then would you keep that information confidential?" as:

$$
\begin{aligned}
&\text{Offer}(\alpha, \beta, \exists\delta. (\text{Commit}(\alpha,\beta,\text{Done(Inform}(\alpha,\beta,\text{provider}(\delta,\alpha,\text{tomato})))) \wedge \\
&\qquad\qquad \forall\gamma.\ \forall\ \text{t. Commit}(\beta,\alpha,\neg\text{Done(Inform}(\beta,\gamma,\text{provider}(\delta,\alpha,\text{tomato}), \text{t}))))
\end{aligned}
$$

[1] We will not detail this language as our focus is on new illocutionary moves requiring higher-order semantics.

[2] Without loss of generality we will assume that all actions are dialogical.

[3] Usually dropped in the examples to simplify notation.

In order to define the *terms* of the language introduced above (e.g. $pay(\beta, \alpha, y)$ or $discount(y, 10\%)$) 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)).[4]

We model ontologies following an algebraic approach [4] as: An ontology is a tuple $\mathcal{O} = (C, R, \leq, \sigma)$ where:

1. C is a finite set of concept symbols (including basic data types);
2. R is a finite set of relation symbols;
3. $\leq$ is a reflexive, transitive and anti-symmetric relation on C (a partial order)
4. $\sigma : R \to C^+$ is the function assigning to each relation symbol its arity

where $\leq$ 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 they are in the structure defined by the $\leq$ relation. Semantic distance plays a fundamental role in strategies for information-based agency. How signed contracts, $Commit(\cdot)$ about objects in a particular semantic region, and their execution $Observe(\cdot)$, *affect* our decision making process about signing future contracts on nearby semantic regions is crucial to modelling the common sense that human beings apply in managing trading relationships. A measure [5] bases the *semantic similarity* between two concepts on the path length induced by $\leq$ (more 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:

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

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.

Agents give advice when they perceive that the listener has less experience in an area. Advice is thus a rhetorical move that uses the asymmetry of information between two agents. It is a genuine ecological move as it makes full sense in the context of a dialogue where both sides are revealing their positions and thus its meaning can only be determined in the context of the agents' mutual evolving models of each other.

In the context of negotiation advice makes sense before the signing of the contract — warning the other agent about potential consequences, "I advise you not buy a reflex camera for your grand mother, they are too bulky", or afterwards to justify a contract violation, "if I were you I would be happy with receiving bottles from the 2008 vintage instead of 2007, they are much better". They are naturally composed of a comparison between contracts or options and a justification.

[4] Axioms defined over the concepts and relations are omitted here.

3 Argumentation Agent Architecture

This Section describes how argumentative interactions are managed by our agent using the LOGIC illocutionary framework [6] that was originally proposed for agents whose sense of distributive justice spanned equity, equality and need. [6] focussed heavily on the *prelude stage* of a negotiation where agents prepare using the five LOGIC dimensions [7]. The five LOGIC dimensions are quite general:

- Legitimacy concerns *information* that may be part of or relevant to contracts signed.
- Options concerns *contracts* where a contract is a set of commitments one for each agent in the contract.
- Goals are the *objectives* of the agents.
- Independence concerns the agent's *outside options* — i.e. the set of agents are capable of satisfying the agent's needs.
- Commitments are the *commitments* that an agent may have.

and are used in this paper to manage all incoming communications including the exchange of "I advise you..." argumentative illocutions. A more formal representation model for LOGIC is:

- $L = \{B(\alpha, \varphi)\}$, that is a set of *beliefs*.
- $O = \{\text{Plan}(\langle\alpha_1, \text{Do}(p_1)\rangle, \dots, \langle\alpha_n, \text{Do}(p_n)\rangle\}$, that is a set of *joint plans*
- $G = \{D(\alpha, \varphi)\}$, that is a set of *desires*.
- $I = \{\text{Can}(\alpha, Do(p))\}$, that is a set of *capabilities*.
- $C = \{I(\alpha, \text{Do}(p))\} \cup \{\text{Commit}(\alpha, \text{Do}(p))\}$, that is a set of *commitments* and *intentions*.

Our description is from the point of view of agent α in a *multiagent system* with a finite number of other agents $\mathcal{B} = \{\beta_1, \beta_2, \dots\}$, and a finite number of *information providing agents* $\Theta = \{\theta_1, \theta_2, \dots\}$ that provide the *context* for all events in the system — Θ^t denotes the state of these agents at time t. The only thing that α 'knows for certain' is its *history* of past communication that it retains in the repository $\mathcal{H}_\alpha^t$. Each *utterance* in the history contains: an illocutionary statement, the sending agent, the receiving agent, the time that the utterance was sent or received. Utterances are organised into dialogues, where a *dialogue* is a finite sequence of related utterances.

α acts to satisfy a *need*, ν, that are considered in context (ν, Θ^t), and does so by communicating an utterance, (μ, β), containing an illocutionary statement, $\mu \in U$, to another agent, $\beta \in \mathcal{B}$. If an utterance is part of a complete dialogue, d, that aimed to satisfy a need then the dialogue is tagged with: the triggering need, ν, the prevailing context, Θ^t, and an *ex post* rating $r \in R$ of how satisfactorily the dialogue satisfied the need. Such a *rated dialogue* has the form: $\boldsymbol{d} = (d, \nu, \Theta^t, r) \in \mathcal{H}_\alpha^t$.

Agent α observes the actions of another agent β in the context Θ^t. Observations are of little value unless they can be verified. α may not posses a sufficient variety of sensory input devices. Sensory inadequacy is dealt with by invoking a truthful *institution agent*, ξ, that promptly reports what it sees. So if β commits to delivering twelve sardines at 6:00pm, or states that "it will rain tomorrow" and is committed to the truth of that prediction, then α will eventually verify those commitments when ξ advises what occurs. If β passes an "I advise you..." message to α, or even a simple Inform(...) message, we assume that β is committed to the validity of the contents.

All communication is recorded in $\mathcal{H}_\alpha^t$ that in time may contain a large amount of data. To make this data useful to α's strategies it is summarised and categorised using the LOGIC framework. To achieve this α requires a categorising function $v : U \rightarrow \mathcal{P}(\{\mathsf{L,O,G,I,C}\})$ where U is the set of utterances. The power set, $\mathcal{P}(\{\mathsf{L,O,G,I,C}\}$, is required as some utterances belong to multiple categories. For example, "I will not pay more for Protos[5] than the price that John charges" is categorised as both Option *and* Independence.

World Model. α's world model, $\mathcal{M}^t$, is the first way in which $\mathcal{H}_\alpha^t$ is summarised. α's proactive reasoning machinery identifies the aspects of the world that α is interested in. They are represented in $\mathcal{M}^t$ as probability distributions, (X_i), in first-order probabilistic logic $\mathcal{L}$. Each of α's plans, s, contains constructors for a set of distributions $\{X_i\} \in \mathcal{M}^t$ together with associated *update functions*, $K_s(\cdot)$, such that $K_s^{X_i}(\mu)$ is a set of linear constraints on the posterior distribution for X_i. $\mathcal{M}^t$ is then maintained from utterances received using *update functions* that transform utterances into constraints on $\mathcal{M}^t$.

Proactive reasoning is described in [8]. For example, in a simple multi-issue contract negotiation α may estimate $\mathbb{P}^t(\text{acc}(\beta,\alpha,\delta))$, the probability that β would accept contract δ, by observing β's responses. The distribution $\mathbb{P}^t(\text{acc}(\beta,\alpha,\delta))$ is classified as an Option in LOGIC. Using shorthand notation, if β sends the message $\text{Offer}(\delta_1)$ then α may derive the constraint: $K^{\text{acc}(\beta,\alpha,\delta)}(\text{Offer}(\delta_1)) = \{\mathbb{P}^t(\text{acc}(\beta,\alpha,\delta_1)) = 1\}$, and if this is a counter offer to a former offer of α's, δ_0, then: $K^{\text{acc}(\beta,\alpha,\delta)}(\text{Offer}(\delta_1)) = \{\mathbb{P}^t(\text{acc}(\beta,\alpha,\delta_0)) = 0\}$. 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" [9]. As another example, the predicate $\text{canDo}(\alpha,\beta,\nu)$ meaning β is able to satisfy α's need ν — this predicate is classified as Independence in LOGIC.

Updating $\mathcal{M}^t$ is complicated when the integrity of utterances received are questionable — it would certainly be foolish for α to believe completely every utterance received. For completeness the procedure for doing this, and for attaching an *a priori* belief to utterances (see Equation 7), is summarised in Section 3.1. If at time t, α receives such an utterance u that may alter this world model then the (Shannon) *information* in u with respect to the distributions in $\mathcal{M}^t$ is: $\mathbb{I}^t(u) = \mathbb{H}(\mathcal{M}^t) - \mathbb{H}(\mathcal{M}^{t+1})$. 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: $\mathbb{H}(\mathcal{N}^t) - \mathbb{H}(\mathcal{N}^{t+1})$. $\mathcal{M}^t$ may contain distributions in any of the five LOGIC categories, where $\mathbb{H}$ is Shannon entropy.

Intimacy and Balance Model. The *intimacy* and *balance* model is the second way in which $\mathcal{H}_\alpha^t$ is summarised. *Intimacy* is degree of closeness, and *balance* is degree of fairness. Informally, *intimacy* measures how much one agent knows about another agent's private information, and *balance* measures the extent to which information revelation

[5] A fine wine from the 'Ribera del Duero' region, Spain.

between the agents is 'fair'. The *intimacy* and *balance* model is structured using the LOGIC illocutionary framework and the ontology $\mathcal{O}$[6]. For example, the communication $\text{Accept}(\beta, \alpha, \delta)$ meaning that agent β accepts agent α's previously offered deal δ is classified as an Option, and $\text{Inform}(\beta, \alpha, \textit{info})$ meaning that agent β informs α about *info* and commits to the truth of it is classified as Legitimacy. The *intimacy* and *balance* model contains two components per agent: first α's model of β's private information, and second, α's model of the private information that β has about α.

The *intimacy* of α's relationship with β_i, I_i^t, is the amount that α knows about β_i's private information and is represented as real numeric values over $\{\mathsf{L,O,G,I,C}\} \times \mathcal{O}$. Suppose α receives utterance u from β_i and that category $f \in v(u)$. For any concept $c \in \mathcal{O}$, define $\Theta(u, c) = \max_{c' \in u} \theta(c', c)$. Denote the value of I_i^t in position (f, c) by $I_{i(f,c)}^t$ then: $I_{i(f,c)}^t = \rho \times I_{i(f,c)}^{t-1} + (1 - \rho) \times \mathbb{I}^t(u) \times \Theta(u, c)$ for any c, where ρ is the discount rate and $\mathbb{I}^t(u)$ is as defined above. α's estimate of β_i's intimacy on α, J_i^t, is constructed similarly. The *balance* of α's relationship with β_i, B_i^t, is the element by element numeric difference of I_i^t and J_i^t.

Trust, Reliability and Honour. The third way in which α summarises $\mathcal{H}_\alpha^t$ is with trust, reliability and honour measures. These concepts are all concerned with the relationship between commitment and enactment. Trust is concerned with the relationship between a signed contract (the commitment) and the execution of the contract (the enactment). Reliability is is concerned with the relationship between information (where the truth of the information is the commitment) and its subsequent verification (the enactment). Honour is similarly concerned with arguments.

We represent the relationship between commitment and enactment using conditional probabilities, $\mathbb{P}(u'|u)$. If u is a commitment and u' the corresponding subsequent observation then $\mathbb{P}(u'|\, u)$ is the probability that u' will be observed given that u had been promised. For example, if u is an "I advise you…" message from agent β then the conditional probability, $\mathbb{P}(u'|u)$, is an estimate of α's expectation of what will eventually be observed, and the uncertainty in the validity of β's communication is the entropy $\mathbb{H}(u'|u)$.

[10] describes three aspects of the relationship between commitment and enactment:

1. as the difference between our expectation $\mathbb{P}(u'|\, u)$ and a distribution that describes what we would ideally like to observe $\mathbb{P}_I(u'|\, u)$:

$$1 - \sum_{u'} \mathbb{P}_I^t(u'|u) \log \frac{\mathbb{P}_I^t(u'|u)}{\mathbb{P}_\beta^t(u'|u)}$$

2. as expected preferability of the enactment compared to the commitment:

$$\sum_{u'} \mathbb{P}^t(\text{Prefer}(u', u)) \mathbb{P}_\beta^t(u'|u)$$

[6] Only a subset of the ontology is required. The idea is simply to capture "How much has Carles told me about wine", or "how much do I know about his commitments (possibly with other agents) concerning cheese".

3. as predictability of those enactments that are preferable to the commitment:

$$1 + \frac{1}{B^*} \cdot \sum_{u' \in \Phi_+(u,v,\kappa)} \mathbb{P}^t_+(u'|u) \log \mathbb{P}^t_+(u'|u)$$

where if $u \leq v$ in the ontology let: $\Phi_+(u,v,\kappa) = \{u' \mid \mathbb{P}^t(\text{Prefer}(u',u,v)) > \kappa\}$ for some constant κ, and $\mathbb{P}^t_+(u'|u)$ is the normalisation of $\mathbb{P}^t_\beta(u'|u)$ for $u' \in \Phi_+(u,v,\kappa)$,

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

There is no neat function mapping the concepts of trust, reliability and honour into the five LOGIC categories. For example, the relationship between contractual commitment and contractual enactment is concerned with both Options *and* Commitment. Alternatively, the relationship between the commitment and enactment of an argument is concerned with Legitimacy *and* what ever else the argument is about. However the five LOGIC categories together provide a complete framework for representing these concepts.

Self Model. Finally, α's *self model* is not directly related to communication. It represents the LOGIC relationships between the agent's components and the various summaries of the communications received.

3.1 Updating $\mathcal{M}^t$

α's world model, $\mathcal{M}^t$, at time t is a set of random variables, $\mathcal{M}^t = \{X_i, \ldots, X_n\}$ each representing an aspect of the world that α is interested in. In the absence of in-coming messages the integrity of $\mathcal{M}^t$ decays. α may have background knowledge concerning the expected integrity as $t \to \infty$. Such background knowledge is represented as a *decay limit distribution*. One possibility is to assume that the decay limit distribution has maximum entropy whilst being consistent with observations. Given a distribution, $\mathbb{P}(X_i)$, and a decay limit distribution $\mathbb{D}(X_i)$, $\mathbb{P}(X_i)$ decays by:

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

where Δ_i is the *decay function* for the X_i satisfying the property that $\lim_{t\to\infty} \mathbb{P}^t(X_i) = \mathbb{D}(X_i)$. 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^t_i and $\mathbb{D}^t(X_i)$.

The following procedure updates $\mathcal{M}^t$ for all utterances $u \in U$. Suppose that α receives a message u from agent β at time t. Suppose that this message states "I advise you that something is so" with probability z, and suppose that α attaches an epistemic belief $\mathbb{R}^t(\alpha,\beta,u)$ to u — a method for estimating $\mathbb{R}^t(\alpha,\beta,u)$ is given below. 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*[7], $K_s(\cdot)$, such that $K^{X_i}_s(u)$ is a set of linear

[7] A sample update function for the distribution $\mathbb{P}^t(acc(\beta,\alpha,\delta))$ is given above.

constraints on the posterior distribution for X_i. Denote the prior distribution $\mathbb{P}^t(X_i)$ by $\boldsymbol{p}$, and let $\boldsymbol{p}_{(u)}$ be the distribution with minimum relative entropy[8] with respect to $\boldsymbol{p}$: $\boldsymbol{p}_{(u)} = \arg\min_{\boldsymbol{r}} \sum_j r_j \log \frac{r_j}{p_j}$ that satisfies the constraints $K_s^{X_i}(u)$. Then let $\boldsymbol{q}_{(u)}$ be the distribution:

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

and then let:

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

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

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

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

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

$\mathbb{R}^t(\alpha, \beta, u)$ 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 observation. Suppose that u is received from agent β at time t and is verified by the institution agent, ξ, as u' at some later time t'. Denote the prior $\mathbb{P}^u(X_i)$ by $\boldsymbol{p}$. Let $\boldsymbol{p}_{(u)}$ be the posterior minimum relative entropy distribution subject to the constraints $K_s^{X_i}(u)$, and let $\boldsymbol{p}_{(u')}$ be that distribution subject to $K_s^{X_i}(u')$. We now estimate what $\mathbb{R}^u(\alpha, \beta, u)$ should have been in the light of knowing *now*, at time t', that u should have been u'.

The idea of Equation 2, is that $\mathbb{R}^t(\alpha, \beta, u)$ should be such that, *on average* across $\mathcal{M}^t$, $\boldsymbol{q}_{(u)}$ will predict $\boldsymbol{p}_{(u')}$ — 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 u. The *observed reliability* for u and distribution X_i, $\mathbb{R}^t_{X_i}(\alpha, \beta, u)|u'$, on the basis of the verification of u with u', is the value of k that minimises the Kullback-Leibler distance:

$$\mathbb{R}^t_{X_i}(\alpha, \beta, u)|u' = \arg\min_k \mathbb{K}(k \cdot \boldsymbol{p}_{(u)} + (1 - k) \cdot \boldsymbol{p} \parallel \boldsymbol{p}_{(u')})$$

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

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

[8] Given a probability distribution $\boldsymbol{p}$, the *minimum relative entropy distribution* $\boldsymbol{q} = (q_1, \ldots, q_I)$ subject to a set of n linear constraints $\boldsymbol{g} = \{g_j(\boldsymbol{p}) = \boldsymbol{a_j} \cdot \boldsymbol{q} - c_j = 0\}, j = 1, \ldots, n$ (that must include the constraint $\sum_i q_i - 1 = 0$) is: $\boldsymbol{q} = \arg\min_{\boldsymbol{r}} \sum_j r_j \log \frac{r_j}{p_j}$. This may be calculated by introducing Lagrange multipliers $\boldsymbol{\lambda}$: $L(\boldsymbol{q}, \boldsymbol{\lambda}) = \sum_j q_j \log \frac{q_j}{p_j} + \boldsymbol{\lambda} \cdot \boldsymbol{g}$. Minimising L, $\{\frac{\partial L}{\partial \lambda_j} = g_j(\boldsymbol{p}) = 0\}, j = 1, \ldots, n$ is the set of given constraints $\boldsymbol{g}$, and a solution to $\frac{\partial L}{\partial q_i} = 0, i = 1, \ldots, I$ leads eventually to $\boldsymbol{q}$. Entropy-based inference is a form of Bayesian inference that is convenient when the data is sparse [11] and encapsulates common-sense reasoning [12].

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 in $\mathcal{M}^t$ that u affects, then the *observed reliability* of β on the basis of the verification of u with u' is:

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

For any concept $c \in \mathcal{O}$, $\mathbb{R}^t(\alpha,\beta,c)$ is α's estimate of the reliability of information from β concerning c. In the absence of incoming communications the integrity of this estimate will decay in time by: $\mathbb{R}^t(\alpha,\beta,c) = \chi \times \mathbb{R}^{t-1}(\alpha,\beta,c)$ for decay constant $\chi < 1$ and close to 1. On receipt of communication u is subsequently verified as u':

$$\mathbb{R}^t(\alpha,\beta,c) = \mu \times \mathbb{R}^{t-1}(\alpha,\beta,c) + (1-\mu) \times \mathbb{R}^t(\alpha,\beta,u)|u' \tag{7}$$

where μ is the learning rate, that estimates the reliability of β's advice on any concept c. If $\mathbf{X}(u)$ are independent the predicted *information* in u is:

$$\mathbb{I}^t(u) = \sum_{X_i \in \mathbf{X}(u)} \mathbb{I}^t_{X_i}(\alpha,\beta,u) \tag{8}$$

Suppose α sends message u to β where u is α's private information, then assuming that β's reasoning apparatus mirrors α's, α can estimate $\mathbb{I}^t(\beta,\alpha,u)$. This completes the the update process for $\mathcal{M}^t$.

4 Advice Interaction

An *opinion* is a speaker's evaluation of a particular aspect of a thing in context. *Advice* is a speaker's evaluation of a particular aspect of a thing in the context of the speaker's beliefs of the listener's context. An "I advise you..." illocution is a form of advice [13]. It is a directive in Searle's classification of speech acts. This illocution gives advice to the listener to take some action, for example, "I advise you I would buy that Ferrari." It is not an assertive. Such advice will only be considered seriously by the listener if he believes that the speaker's beliefs about him are accurate. In terms of this work, this is indicated by a degree of intimacy in the appropriate section of the LOGIC framework.

An agent may be motivated to issue an "I advise you..." illocution *either* to develop a reputation for giving good advice — in the LOGIC framework this develops intimacy particularly in the L dimension — *or* to directly influence the listener's actions possibly to the benefit of the speaker "If I were you I would accept the offer I made you yesterday". The rational effect of these two examples are different. In the first example, whether the listener follows the advice is not important, what matters is whether he believes at some time that the advice was good, in the second example, the intention is that the listener will follow the advice.

"I advise you..." illocutions may be issued with varying degrees of knowledge of the state of the listener. For example, "I advise you to buy the Ferrari." assumes that the speaker has beliefs about the listener's intentions — such as he intends to buy a car. Another example, "If I were you I would offer them €100 now" assumes that the speaker has beliefs about both the listener's intentions *and* the state of his active plans. For simplicity we restrict these beliefs to the listener's intentions.

In common usage, an "I advise you..." illocution may contain advice *either* to act (i.e. advice that the listener should utter) as described above, *or* that the listener should modify his mental attitudes "I advise you to count on tomorrow being fine". The first of these is an *"I advise you" action*, and the second, that advises the agent to modify his beliefs, is an *"I advise you" belief change*[9]. In addition, such advice may advise the listener to modify his goals, his intentions or his plans — these three cases are omitted for brevity. A definition of an *"I advise you" action* is given in Table 1. The definition of an *"I advise you" belief change* is not presented here.

5 Discussion

In this paper we have argued that a a rich model of rationality is required to properly model agents in a changing world. Particularly important is the need to model dialogical moves that refer to the agent's internal models (beliefs, or goals) that are updated as a dialogue develops. Traditional constructivist approaches share a more static view of the world. Dialogues may influence internal models along a number of dimensions. In this paper we have followed a simplified version of the approach of [6] classifying them as beliefs, plans, desires, capabilities and intentions. This model is very flexible and clear in representing and classifying the evolving pieces of information that an agent's memory requires in order to correctly interpret and generate illocutionary moves. We have given a precise model of how this evolution of the memory can be implemented using concepts drawn from information-theory. Finally, a formal description of a prototypical dialogical move, "I advise you ...", is given. We have argued, that if agents are to be situated in a changing world, they need to incorporate an ecological mind that among other things requires a higher order interpretation of communication languages. This is so, because self-reference and the reference to whole dialogues is unavoidable in argumentative information exchanges.

As future lines of work, we plan to extend this approach to further advice-giving illocutions, and to revisit other classical dialogical moves such as those found in negotiation dialogues (e.g. propose, accept, reject). The *evolution* of our information-theoretic agents is being further examined in the development of negotiation agents in the Diplomacy game: we plan to use a Diplomacy testbed (www.dipgame.org) to obtain experimental results from agents interacting with human beings using rich languages that have illocutionary moves similar to the one modelled here.

[9] In line with the remarks at the beginning of this section the assertive "Tomorrow will be fine" may be treated as an Inform; when that statement is verified by the listener he will feed that into his estimate of the speaker's reliability as in Section 3.1. The directive "I advise you to count on tomorrow being fine." is a richer statement. It relies on the speaker's weather forecasting ability *and* on the accuracy of his beliefs of the listener. In particular, it relies on the accuracy of the speaker's beliefs concerning the significance of tomorrow's weather to what ever the listener is doing. That is it relies on a level of intimacy. The subsequent evaluation of this piece of advice will then effect the speaker's intimacy represented in the LOGIC model.

Table 1. Advice actions in FIPA-style format. The two feasibility preconditions are alternative representations of i's beliefs of the superiority of his knowledge, and the two rational effects represent two possible motives for uttering the illocution.

Summary	The sender (for example, i) informs the receiver (for example, j) that the sender believes the receiver should perform some action (for example, a) if the receiver's intentions includes some goal (for example, c)
Message Content	A tuple consisting of an action expression denoting the action that is advised, and an intention that the receiver may hold.
Description	I_Advise_You indicates that the sending agent: • believes he knows the receiving agent holds a particular intention • believes his knowledge of facts concerning the receiving agent's intention is better than the receiving agent's knowledge of it • intends the receiving agent to believe that the action is in his interests • believes that the receiving agent may act otherwise
Formal Model	$<i,\ \mathtt{i\_advise\_you}(j,a,c)>$ FP1: $B_i\ I_j\ c \wedge B_i\ W_i(c) \rightarrow W_{j\setminus i}(c) \wedge B_i\ \mathrm{Agent}(j,\ a) \wedge \neg B_i\ I_j\ \mathrm{Done}(a)$ FP2: $B_i\ I_j\ c \wedge B_i\ (\mathbb{H}(W_i(c)) < \mathbb{H}(W_{j\setminus i}(c))) \wedge B_i\ \mathrm{Agent}(j,\ a)\ \neg B_i\ I_j\ \mathrm{Done}(a)$ RE1: $B_j\ I_i\ \mathrm{Done}(<j,\ \mathrm{rates}(a,\ x)>, \phi)$ where $\mathrm{rates}(a, x)$ is the action of rating action a as x, and ϕ is true when the rating is performed RE2: $\mathrm{Done}\ (a)$ $W_i(c)$ denotes all of i's beliefs concerning c — i.e. that part of i's world model $W_{j\setminus i}(c)$ denotes i's beliefs concerning all of j's beliefs concerning c $W_i(c) \rightarrow W_{j\setminus i}(c)$ denotes that everything in $W_{j\setminus i}(c)$ can be derived from a subset of $W_i(c)$ $\mathbb{H}(S)$ denotes the overall uncertainty of the set of beliefs S — possibly as entropy

Examples	Agent i advises agent j that from his understanding of agent j's intentions agent j should accept an offer from agent k to sell a Nikon camera to agent j.
	`(i_advise_you`
	` :sender (agent-identifier :name i)`
	` :receiver (set (agent-identifier :name j))`
	` :content`
	`  "((advise-action (agent-identifier :name j)`
	`   (accept-proposal`
	`    :sender (agent-identifier :name j)`
	`    :receiver (set (agent-identifier :name k))`
	`    :content`
	`     "accept the Nikon"`
	`  ("want camera"))"`
	` :language fipa+if_I_were_you+advise-action)`
	where `advise-action` is an action that the receiver is advised to perform

Acknowledgements. Research supported by the Agreement Technologies CONSOLIDER project under contract CSD2007-0022 and INGENIO 2010 and by the Agreement Technologies COST Action, IC0801.

References

1. Dung, P.M.: On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. Artificial Intelligence 77, 321–358 (1995)
2. Greene, K., Derlega, V.J., Mathews, A.: Self-disclosure in personal relationships. In: Vangelisti, A., Perlman, D. (eds.) Cambridge Handbook of Personal Relationships, pp. 409–427. Cambridge University Press, Cambridge (2006)
3. Krause, A., Horvitz, E.: A utility-theoretic approach to privacy and personalization. In: Fox, D., Gomes, C.P. (eds.) Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence, AAAI 2008, Chicago, Illinois, USA, pp. 1181–1188. AAAI Press, Menlo Park (2008)
4. Kalfoglou, Y., Schorlemmer, W.M.: IF-map: An ontology-mapping method based on information-flow theory. In: Spaccapietra, S., March, S., Aberer, K. (eds.) Journal on Data Semantics I. LNCS, vol. 2800, pp. 98–127. Springer, Heidelberg (2003)
5. Li, Y., Bandar, Z.A., McLean, D.: An approach for measuring semantic similarity between words using multiple information sources. IEEE Transactions on Knowledge and Data Engineering 15, 871–882 (2003)
6. Sierra, C., Debenham, J.: The LOGIC Negotiation Model. In: Proceedings Sixth International Conference on Autonomous Agents and Multi Agent Systems AAMAS 2007, Honolulu, Hawai'i, pp. 1026–1033 (2007)
7. Fischer, R., Ury, W., Patton, B.: Getting to Yes: Negotiating agreements without giving in. Penguin Books (1995)
8. Sierra, C., Debenham, J.: Information-based deliberation. In: Padgham, L., Parkes, D., Müller, J., Parsons, S. (eds.) Proceedings Seventh International Conference on Autonomous Agents and Multi Agent Systems AAMAS 2008, Estoril, Portugal, ACM Press, New York (2008)
9. Sierra, C., Debenham, J.: Information-based agency. In: Proceedings of Twentieth International Joint Conference on Artificial Intelligence, IJCAI 2007, Hyderabad, India, pp. 1513–1518 (2007)
10. Sierra, C., Debenham, J.: Trust and honour in information-based agency. In: Stone, P., Weiss, G. (eds.) Proceedings Fifth International Conference on Autonomous Agents and Multi Agent Systems AAMAS 2006, Hakodate, Japan, pp. 1225–1232. ACM Press, New York (2006)
11. Cheeseman, P., Stutz, J.: On the relationship between bayesian and maximum entropy inference. In: Fischer, R., Preuss, R., von Toussaint, U. (eds.) Bayesian Inference and Maximum Entropy Methods in Science and Engineering, pp. 445–461. American Institute of Physics, Melville (2004)
12. Paris, J.: Common sense and maximum entropy. Synthese 117, 75–93 (1999)
13. Decapua, A., Huber, L.: 'If I were you . . . ': Advice in American English. Multilingua 14, 117–132 (1995)