Negotiation and Argumentation in Multi-agent Systems

Fundamentals, Theories, Systems and Applications

Editor:
Fernando Lopes
LNEG–National Research Institute
Lisbon, Portugal

Co-Editor:
Helder Coelho
University of Lisbon
Lisbon, Portugal

Bentham e Books

Send Orders for Reprints to reprints@benthamscience.net
Negotiation and Argumentation in Multi-agent Systems, 2014, 233-253

CHAPTER 8

RANA: a Relationship-aware Negotiation Agent

Carles Sierra[1,*] and John Debenham[2]

[1] IIIA, CSIC, Campus UAB, 08193 Bellaterra, Catalonia, Spain
Email: `sierra@iiia.csic.es`
[2] QCIS, UTS, Broadway, NSW 2007, Australia
Email: `john.debenham@uts.edu.au`

Abstract. Much has been written on the use of rhetorical argumentation to alter the beliefs of a partner agent within a particular negotiation. The problem addressed in this chapter is the *measurement* of the long-term value of rhetorical argumentation in repeated interactions between a pair of agents, and of the *management* of such argumentation to achieve strategic aims concerning the strength of the agents' *relationships*. RANA is a relationship-aware negotiation agent in the context of information-based agents [1] that have embedded tools from information theory enabling them to measure and manage strategic information.

Keywords: Argumentation, Negotiation, Bargaining, Social relationships, Information theory, Rhetorics, Trust, Software Agent, Agent architecture, Multi-agent System.

8.1 Introduction

Human agents generally place great value on their relationships with others particularly in the conduct of business [2, 3]. Business relationships develop as a subtle byproduct of interaction. Our premiss is that if artificial agents are to conduct business automatically then they too will need to understand the value of business relationships, and will need tools to build and manage

*Corresponding author: Carles Sierra
Telephone: +34-93-580 95 70 *Fax*: +34-93-580 96 61

Fernando Lopes and Helder Coelho (Eds.)

them. An agent's *relationships* is a model that somehow summarises its full interaction history with the intention of enabling the agent to exhibit a strategic social sense. Summary measures such as trust [6] may appear explicitly in the relationship model—there may be many such measures in this model. This chapter describes a framework for representing relationships, and describes strategies for strengthening and weakening them using rhetorical argumentation.

The term *argumentation* is commonly used to refer to both classical argumentation and rhetorical argumentation. The term *classical argumentation* is commonly used to refer to the generation of arguments, for example as logical proofs, that support or reject particular courses of action that may be taken during a decision making process. The term *rhetorical argumentation* is commonly used to refer to the use of rhetoric particles with the aim of altering an agent's beliefs. One well-known set of rhetoric particles is: *informs*, *appeals*, *rewards* and *threats* [7]. These four particles may be used to alter an agent's beliefs so as to make proposed deals, or contracts, either more or less attractive. In this chapter we use rhetoric to strengthen or to weaken the relationship between two agents. Our focus is on three rhetoric particles: *informs*, *opinions*, and *advice*. This work is based on the idea that relationships may be strengthened by exchanging reliable and relevant information, opinions and advice.

This work is based on our information-based agents [1]. Information-based agents are endowed with utilitarian machinery that enables them to pursue their goals in a conventional manner. This machinery is augmented with tools drawn from information theory that enable the agents to value the contents of every utterance received. The intuition being that when an agent utters the utterance gives away information. When an utterance is received it may be valued in terms of the information gain in the receiving agent's world model. These valuations may be summarised in various ways to model the ongoing value of information sources. These summaries may take account of the type of illocutionary particle, and the source from which it came. However, the contents of an utterance, no matter what it is, may prove to be incorrect or worthless for some reason, and the receiving agent may have no way of determining its true worth for some time after receipt. When the true value is known, the information-based machinery does two things: first it uses this knowledge to update its model of the speaker as an information giver in terms of how freely it passes on information, and second, it maintains an ongoing model the overall integrity of the speaker for utterances of each type. By applying this idea to various aspects of agent dialogue, this machinery is used in this chapter to build models of relationships that an agent has with other agents. These relationship models then prove valuable, for example, when deciding which agent to do business with for a given type of transaction.

Section 8.2 describes the rhetoric particles and the communication language with an emphasis on *informs*, *opinions*, and *advice*. Our LOGIC framework [16] is used to describe and formalise the characteristics of relationships between agents in Section 8.3. This model contains four components. In addition, two models, for trust and integrity, are required to support the selection of interaction partners. They are described in Section 8.4. Section 8.5 finally draws the work together in a discussion of strategies.

8.2 Communication: Rhetoric Particles and Language

In this section we detail the language that agents use to build relationships with a particular emphasis on the three rhetoric particles: `inform`, `opinion` and `advise`. These three particles are used as follows:

- An `inform` communicative act informs the listener that a proposition is true. That is, the speaker intends that the listener believes that the proposition may be verified.
- "An `opinion` communicative act is a speaker's evaluation of a particular aspect of a thing in context, where the *context* is the set of all things that the thing is being, explicitly or implicitly, evaluated with or against. The set of valuations of all things in the context calibrates the valuation space." [4] For example, "Sydney is more fun than Melbourne". The context can loosely specified: "today's weather could not have been worse". Dealing with opinions is predicated on an understanding of the speaker's intended context.
- "An `advise` communicative act 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. It is a directive in Searle's classification of speech acts." [4] The speaker of an `advise` statement may intend that the listener consider some action to be desirable in some sense, for example, "If I were you I would spend at least two weeks in Barcelona". Alternatively, the speaker of an `advise` statement may intend that the listener modifies its beliefs, for example, "If I were you I would spend more time in Sydney on your holiday than in Melbourne as you plan".

These three communicative acts are, in a sense, in a loose hierarchy. An `inform` act is concerned with a statement that may be verified by observing world states, an `opinion` act with the speaker's personal evaluation, and an `advice` act with the speaker's goals on the listener's future states.

The FIPA semantics [8] is expressed in terms of the *rational effect* that illocutionary particles have. The rational effect of the three particles just described may be cast in two different ways: first, the effect that they have on the future internal state or future actions of the listener; and second, the effect that they have on the relationship between speaker and listener when the integrity of the uttered particle has been evaluated. For example, for the first, if John informs Carles that Ferran Adrià is cooking at a local restaurant then Carles may immediately go for lunch, and for the second, if when Carles arrives at the restaurant it is closed for renovations then Carles may use John's communication of inaccurate information to decrease his belief in John's information-giving ability and so to reduce the strength of his relationship with John. This second sense of rational effect is of particular relevance to this chapter.

8.2.1 Rhetoric Particles

The `inform` communicative act is widely used. In the following, agent i informs agent j of proposition p. Its meaning is commonly taken to be as specified in [8] that is extended below to include a dual rational effect:

$$
\begin{array}{ll}
\multicolumn{2}{l}{\texttt{<}i\texttt{, inform(}j\texttt{, }p\texttt{)>}} \\
\texttt{FP:} & \mathrm{B}_i p \wedge \neg \mathrm{B}_i(\mathrm{Bif}_j p \vee \mathrm{Uif}_j p) \wedge \\
 & \mathrm{B}_i \mathrm{I}_j \ \mathrm{Done}(\texttt{<}j\texttt{, eval(}p\texttt{, }x\texttt{)>}, \phi) \\
\texttt{RE1:} & \mathrm{B}_j p \\
\texttt{RE2:} & \mathrm{Done}(\texttt{<}j\texttt{, eval(}p\texttt{, }x\texttt{)>}, \phi)
\end{array}
$$

where in FIPA notation FP is the *feasibility precondition*, RE is the *rational effect*, and

$\mathrm{B}_i p$ means that agent i believes proposition p,

$\mathrm{Bif}_i p \equiv \mathrm{B}_i p \vee \mathrm{B}_i \neg p$,

$\mathrm{Uif}_i p \equiv \mathrm{U}_i p \vee \mathrm{U}_i \neg p$, where $\mathrm{U}_i p$ means that agent i is uncertain about p.

$\mathrm{I}_i a$ means that agent i intends to perform action a,

In $\mathrm{Done}(\texttt{<}j\texttt{, eval(}p\texttt{, }x\texttt{)>}, \phi)$, the construct $\texttt{<}j\texttt{, eval(}p\texttt{, }x\texttt{)>}$ means that after the integrity of proposition p is known to agent j, that agent rates the integrity as x, where proposition ϕ is true when the integrity of p is known to j, and $\mathrm{Done}(a, \phi)$ means that a has just taken place and that ϕ was true just before that.

In `RE2`: the evaluation is done when agent j has had the opportunity to exploit the contents of the inform and has a view the value of its integrity. This evaluation is performed on a fuzzy scale, `eval` $\in [0, 1]$, where 0 (means "is totally useless"), and 1 (means "is most valuable").

The FIPA specification does not include the `opinion` and `advise` communicative acts. These are defined following in the FIPA notation.

Following the definition of the `opinion` communicative act given above, its representation will contain:

- "the thing that is the *subject* of the opinion" [4],
- "the *aspect*, or attribute, of the thing that is being evaluated" [op.cit.],
- "a distribution over some evaluation space representing the *rating* of the aspect of the thing in the context" [op.cit], and
- "optionally the *context* in which the evaluation is made, and a *reason* supporting the opinion" [op.cit.].

For example, "The visual appearance (i.e. the aspect) of the Maserati Corsa (i.e. the subject) is superb (i.e. the rating)" where the context may be "Italian cars". An `opinion` act may also contain a `reason` that is intended to justify the rating; for example "The New York Times rated the Maserati Corsa second their list of Top Sports Cars of the 2000s".

The *rating* in an opinion is performed over a generally understood *rating space*. In the above Maserati example the rating space could have been {superb, ok, hideous}. In general the rating will be expressed as a probability distribution over the rating space. In the example the rating is $< 1, 0, 0 >$ over this rating space. If a rating is expressed as a single probability, for example "I rate the Maserati Corsa as 'superb' with confidence 0.8", then this is taken to be equivalent to a distribution with the vacant slots filled with the maximum entropy distribution, for example $< 0.8, 0.1, 0.1 >$ [18].

If a speaker performs an `opinion` action then this suggests that the speaker:

- believes that she knows a particular intention of the listener,
- believes that an opinion she holds is relevant to that intention, and
- believes that her own opinion may influence that intention.

Suppose that, agent i, informs agent j that i's rating of an aspect, s, of a thing, t, is e in (the optional) context c for the (optional) reason, r. It would be very convenient if:

```
<i, opinion(j,s,t,e[,c,r])>
```

could be defined as:

$<i,\ \mathtt{inform}(j, \mathtt{Rates}(i,s,t,e[,c,r]))>$

where $\mathtt{Rates}(i,s,t,e[,c,r])$ is a proposition meaning agent i's rating of aspect s of thing t is e, but this is rather weak with rational effect:

$\mathtt{RE1{:}B}_j\ \mathtt{Rates}(i,s,t,e[,c,r])$

it would be unreasonable to require the rational effect to be that j now holds the same opinion as i:

$\mathtt{RE1{:}B}_j\ \mathtt{Rates}(j,s,t,e[,c,r])$

but it is perfectly reasonable to expect that j's rating has been positively affected by the utterance:

$<i,\ \mathtt{opinion}(j,s,t,e[,c,r])>$

FP: $\mathrm{B}_j\ \mathtt{Rates}(j,s,t,e'[,c,r])\ [\wedge\ \mathrm{B}_i\mathrm{I}_j c\ \wedge\ \mathrm{B}_i r]\ \wedge$
$\mathrm{B}_i\mathrm{I}_j\ \mathtt{Done}(<j,\ \mathtt{eval}(s,t,e,x[,c,r])>,\phi)$

RE1: $\mathrm{B}_j\ \mathtt{Rates}(j,s,t,e''[,c,r])\ \wedge \mathtt{Closer}(e'',e,e')$

RE2: $\mathtt{Done}(<j,\ \mathtt{eval}(s,t,e,x[,c,r])>,\phi)$

where $\mathtt{Closer}(x,y,z)$ means that rating y is closer to rating x than rating z.

That is, having uttered an opinion about t, i believes that j's rating of t has moved closer to i's rating than it was prior to the utterance bing made. We assume some suitable distance measure and that `eval` is redefined in line with it.

We consider dual rational effects for the `advise` communicative act as for `inform` and `opinion` above. A speaker may utter an `advise` communicative act if she believes she knows the listener's intentions or the listener's plans. For example, "I advise you to go to Sydney for your holiday." may imply that the speaker believes the listener intends to holiday elsewhere. Another example, "I advise you to book your Sydney holiday now" may imply that the speaker believes he knows that the listener has an active plan to do otherwise. Only the first case is considered here.

As we have noted, an `advise` communication may contain advice *either* to the listener to utter, *or* that the listener should modify his beliefs. Advice to the listener to utter is called an `advise` *action*. Advice that the listener should modify his beliefs is called an `advise` *belief change*. Additionally, the speaker of an `advise` belief change will be interested to know whether the listener has taken his advice. In this sense, an `advise` belief change is a strong `inform`. These two cases are detailed below. In addition, such advice may suggest that the listener modifies his goals, his intentions or his plans—these three cases are not discussed here.

By performing an `advise` *action* the speaker indicates that:

- he "believes he knows that the listener holds a particular intention" [4],
- "his knowledge of facts concerning the listener's intention is better than the listener's knowledge of them" [op.cit.],
- "he intends the listener to believe that the advised action is in the listener's interests" [op.cit.], and
- "the listener may act otherwise" [op.cit.].

Consider now agent i advises agent j that she should perform action a subject to the condition that agent j intends to achieve goal c. Two feasibility preconditions are given; they represent agent i's belief that her knowledge is superior in some way to that of agent j. Two rational effects represent two possible motives for agent i making the utterance.

$$
\begin{array}{ll}
\multicolumn{2}{l}{<i,\ \mathtt{advise}(j,a,c)>} \\
\text{FP:} & B_iI_jc \ \wedge\ B_i(\ W_i(c) \vdash\ W_{j\backslash i}(c)\ )\wedge \\
& \neg B_iI_j\ \mathrm{Done}(<j,a>)\ \wedge \\
& B_iI_j\ \mathrm{Done}(<j,\ \mathrm{eval}(a,c,x)>,\phi) \\
\text{or:} & B_iI_jc \ \wedge\ B_i\,(\mathbb{H}(W_i(c)) < \mathbb{H}(W_{j\backslash i}(c)))\wedge \\
& \neg B_iI_j\ \mathrm{Done}(a)\ \wedge \\
& B_iI_j\ \mathrm{Done}(<j,\ \mathrm{eval}(a,c,x)>,\phi) \\
\text{RE1:} & \mathrm{Done}(<j,a>) \\
\text{RE2:} & \mathrm{Done}(<j,\ \mathrm{eval}(a,c,x)>,\phi)
\end{array}
$$

where:

$\mathrm{eval}(a,c,x)$ is the action of evaluating action a as x in context c, and the proposition ϕ is true when this evaluation is performed;
$W_i(c)$ denotes that part of i's world model concerning c;
$W_{j\backslash i}(c)$ denotes i's beliefs concerning j's beliefs concerning c;
$W_i(c) \vdash W_{j\backslash i}(c)$ denotes that $W_{j\backslash i}(c)$ is implied by a subset of $W_i(c)$;
$\mathbb{H}(S)$ denotes the overall uncertainty—possibly as entropy.

Second, an `advise` *belief change*, the two feasibility preconditions are alternative representations of i's beliefs of the greater value of her knowledge, and the two rational effects represent two possible motives for acting:

$$
\begin{array}{ll}
\multicolumn{2}{l}{<i,\ \mathtt{advise}(j,p,c)>} \\
\text{FP:} & B_i\ I_j\ c\ \wedge\ B_i(\ W_i(c) \rightarrow\ W_{j\backslash i}(c)\ )\wedge\ B_i\neg B_j\ p\wedge \\
& B_iI_j\ \mathrm{Done}(<j,\ \mathrm{eval}(p,c,x)>,\phi) \\
\text{or:} & B_i\ I_j\ c\ \wedge\ B_i\,(\mathbb{H}(W_i(c)) < \mathbb{H}(W_{j\backslash i}(c)))\wedge \\
& B_i\neg B_j\ p\ \wedge B_iI_j\ \mathrm{Done}(<j,\ \mathrm{eval}(p,c,x)>,\phi) \\
\text{RE1:} & B_iB_j\ p \\
\text{RE2:} & \mathrm{Done}(<j,\ \mathrm{eval}(p,c,x)>,\phi)
\end{array}
$$

where `eval(p,c,x)`, $\mathbb{W}_i(c)$, $\mathbb{W}_{j\setminus i}(c)$, $(\mathbb{W}_i(c) \rightarrow \mathbb{W}_{j\setminus i}(c))$ and $\mathbb{H}(S)$ are as above. We note the difference between the `RE1` above and as in the `inform`.

The agents use a communication language, U, that is discussed in detail in [4]. It contains three fundamental primitives: $\text{Commit}(\alpha, \beta, \varphi)$ meaning that α commits to β that φ will occur within some future world state, $\text{Observe}(\alpha, \varphi)$ meaning that a world state is observed in which φ occurs, and $\text{Done}(u)$ meaning that action u has been performed.

8.3 The LOGIC Framework for Agent Relationships

All that an agent knows is represented in its full interaction history. This history is typically very large and agents may contain tools for summarising it in various ways [12]. The LOGIC framework is expressed in terms of summary measures that may be applied to represent aspects of the relationships that an agent has with other agents. This framework is described in detail in [5]. The LOGIC framework for categorising information is illustrated in Figure 8.1. The LOGIC framework is closely related to traditional agent conceptual components:

- $L = \{B(\alpha, \varphi)\}$, that is a set of *beliefs*.
- $O = \{\text{Plan}(\langle\alpha_1, \text{Do}(p_1)\rangle, \ldots, \langle\alpha_n, \text{Do}(p_n)\rangle\}$, 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*.

The *relationship model*, $\mathcal{R}^t_{\alpha\beta}$, that agent α has of agent β contains four components and uses the ideas of the LOGIC framework:

- an *intimacy model*, $J^t_{\alpha\beta}$, that represents how much α knows of β's private information [13]—this summarises the information gain in utterances passed from β to α,
- a *reliability model*, $R^t_{\alpha\beta}$, that represents *how reliable* $J^t_{\alpha\beta}$ is,
- a *reflection model*, $K^t_{\alpha\beta}$, in which α represents how much she believes β knows of her private information, and
- a *balance model*, $B^t_{\alpha\beta}$, that represents the difference growth rates of $J^t_{\alpha\beta}$ and $K^t_{\alpha\beta}$.

A *categorising function* is used to place the contents of all utterances received into the five LOGIC categories:

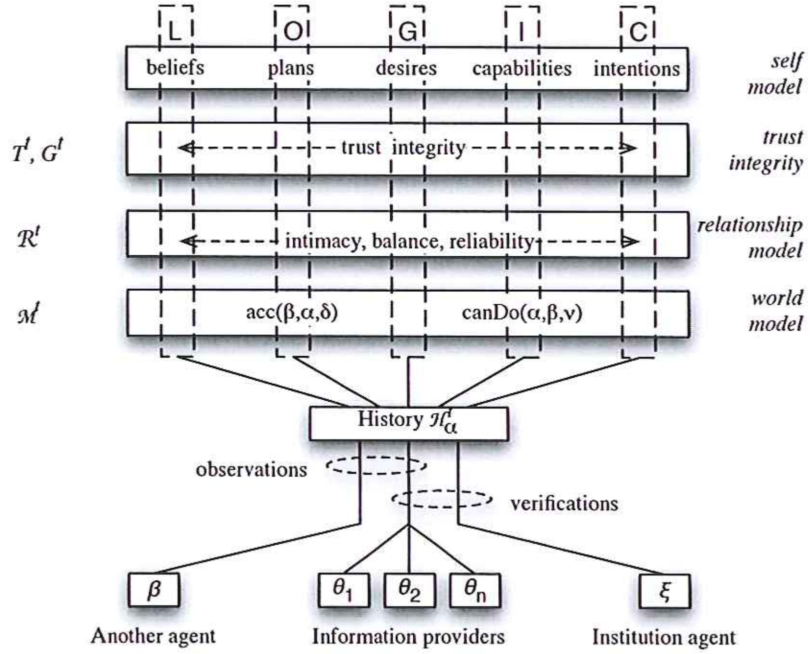

Fig. 8.1 The five categories of the LOGIC framework into which information is categorised in the relationship model.

- *legitimacy* is information concerning commitments that have been made,
- *options* is information concerning commitments that an agent *may be* prepared to make,
- *goals* is information concerning an agent's goals,
- *independence* is information about other agents who may be capable of satisfying a given agent's needs, and
- *commitments* is information about all commitments that an agent has.

To illustrate the process, suppose for example that John makes the utterance u = "Protos is fine wine", and suppose that as a result of a subsequent evaluation $\texttt{eval}(u) = 0.9$. Then α will update its estimate of $R_{\alpha\text{John}(\mathsf{L},\text{wine})}$ where L∈LOGIC is Legitimacy. Epistemic probabilities that represent the reliability of forecasts that an event will occur may be combined as follows. If the prior probability is ω and x and y represent the reliability of two forecasts then they may be combined using the standard method for prior ω: $\text{comb}(x, y, \omega) = \frac{x \times y \times (1-\omega)}{(x \times y \times (1-\omega)) + (1-x) \times (1-y) \times \omega}$. Suppose $r = R^{t-1}_{\alpha\text{John}(\mathsf{L},\text{wine})}$.

To combine the accumulated value r with the single observation $\texttt{eval}(u)$ it is reasonable to moderate the relative significance of $\texttt{eval}(u)$ by:

$$e = (\rho \times \omega) + (1 - \rho) \times \texttt{eval}(u) \tag{8.1}$$

and to define: $R^t_{\alpha\text{John}(\text{L},\text{wine})} = \text{comb}(r, e, \omega)$, where ρ is the learning rate. This approach is generalised below to deal with concepts that are close to, but not identical to, terms in the utterance, for example, $R_{\alpha\text{John}(\text{L},\text{beverages})}$.

Given an utterance u and an evaluation $\texttt{eval}(u)$, suppose that f is the LOGIC category of u, $f \in v(u)$, where v is the categorising function described in Section 8.3. For any category $c \in \mathcal{O}$, define e using Equation 8.1. To define $R^t_{\alpha\beta(f,c)}$ given u we moderate the value of e to an extent determined by $\text{Close}(c, u)$. Note that if $e > \omega$ then $\texttt{eval}(u) > \omega$, and vice versa. Consider the two cases: $e > \omega$ and $e < \omega$. If $e > \omega$ then we moderate e to:

$$e' = e \times \text{Close}(c, u)$$

provided that $e' > \omega$; if $e' < \omega$ then u and c are too far removed from each other and no update occurs. Similarly, if $e < \omega$ then we moderate e to:

$$e'' = (\text{Close}(c, u) \times (e - 1)) + 1$$

provided that $e'' < \omega$, and if $e'' > \omega$ then no update occurs. Let $r = R^{t-1}_{\alpha\beta(f,c)}$ and define:

$$R^t_{\alpha\beta(f,c)} = \begin{cases} \text{comb}(r, e', \omega) & \text{if } e > \omega \text{ and } e' > \omega, \\ \text{comb}(r, e'', \omega) & \text{if } e < \omega \text{ and } e'' < \omega, \\ r & \text{otherwise.} \end{cases} \tag{8.2}$$

The definition of comb assumes that successive reliability observations are statistically independent. This assumption is perhaps unreasonable, by moderating $\texttt{eval}(u)$ to e using Equation 8.1 with the learning rate ρ may be seen to compensate for it. This assumption avoids the estimation of the degree of dependence of these repeated observations. The conditions in Equation 8.2 ensures that the update is only applied when the degree of similarity, Sim, is reasonably large [10]. When $\text{Sim} = 1$, $e' = e'' = e$. As the value of Sim decreases to 0: e' decreases to 0, and e'' increases to 1. Precisely, this restricts the update to those values of e' and e'' that are "on the same side of" 0.5 as e.

Now to illustrate the procedure to update the various models with a simple example. Consider the utterance, u, in which agent *John* sends his opinion to agent j that the quality of agent `Carles`'s olives is good with probability 0.8:
<*John*, `opinion(`
`j, olive-quality, Carles,` $\mathbb{P}$`(eval=good)=0.8)`>
where the understood evaluation space is (*good*, *bad*, *indifferent*).

The two rational effects of this utterance are that:

- agent j's estimate of the quality of Carles' olives, that will be represented as a random variable X, will be such that $\mathbb{P}(X = good)$ is closer to 0.8 than it was prior to the utterance being made, and
- in due course, agent j will evaluate John's advice-giving ability in the region of Carles and olives.

The utterance u will effect j's world model. This effect will be determined first by j's choice of random variables, X_i, that make up its world model, second by its update functions, U_{X_i}, that give meaning to utterances, and third by its prior estimate of the reliability of the utterance. For example, if agent j had a random variable representing its beliefs concerning the quality of Carles' olives then that random variable may have update functions that lead to an alteration in the probability distribution associated with it as a result of the receipt of the utterance.

It is reasonable to suggest that agent j would categorise u as an Option (in that it conveys an opinion that relates to possible contracts to purchase olives) and as Independence (in that it conveys an opinion concerning Carles' ability to satisfy j's needs concerning olives). In which case, agent j updates both $J^t_{j\text{Carles}(\text{O},c)}$ and $J^t_{j\text{Carles}(\text{I},c)}$ for any c semantically close to "olives". These two updates would immediately be reflected in the corresponding balance models.

Then at some time later, perhaps when agent j has sampled Carles' olives, agent j updates his reliability models of John: $R^t_{j\text{John}(\text{O},\text{Carles})}$, $R^t_{j\text{John}(\text{I},\text{Carles})}$, $R^t_{j\text{John}(\text{O},\text{olives})}$ and $R^t_{j\text{John}(\text{I},\text{olives})}$.

8.4 Trust and Integrity

The relationship model described in Section 8.3 is revised as each utterance is received, and is used to manage information revelation through argumentation.

In Section 8.5.1 we consider the problem of choosing a negotiation partner with the aim of satisfying a given need. Two measures support this process; both of them are measures of complete dialogues. First *trust* that measures the difference between commitments made and enactment observed [15]. Second *integrity* that uses the eval function to measures the difference between initial expectation and final evaluation of the entire dialogue sequence. Trust and integrity are fundamentally different as is illustrated by a dialogue leading to a contract to purchase a case of red wine. The way in which the supplier enacts his contractual commitments, such as whether the correct wine is delivered on time, determines the contribution that that dialogue makes to the buyer's estimate of the supplier's trust for red wine. When the buyer has consumed the wine, and has had the opportunity to assess the condition of the wine, he then determines the contribution that that dialogue makes to the integrity of the supplier for red wine.

8.4.1 Trust Model

If β makes contractual commitments during an interaction, for example by committing to the truth of information exchanged or by committing to act, α will observe the difference, if any, between β's commitment, φ, and the true enactment, φ', as advised by the institution agent ξ. The relationship between commitment and enactment is denoted by, "$\mathbb{P}^t(\text{Observe}(\varphi')|\text{Commit}(\varphi))$ simply as $\mathbb{P}^t(\varphi'|\varphi) \in \mathcal{M}^t$" [1].

"In the absence of in-coming messages the conditional probabilities, $\mathbb{P}^t(\varphi'|\varphi)$, should tend to ignorance as represented by the *decay limit distribution*" [1]. $\Phi = \{\varphi_1, \varphi_2, \ldots, \varphi_m\}$ denotes the set of all possible enactments with prior distribution $\mathbb{P}^{t-1}(\varphi'|\varphi)$. Suppose that an observation is received from ξ, we estimate the posterior $\mathbb{P}^t(\varphi'|\varphi)$ as follows.

For any commitment ϕ, suppose φ_k is observed—this observation will be used to update the estimate $\mathbb{P}^{t-1}(\varphi'|\varphi)$ that summarises the previous observations made. We moderate the significance of this single observation using the approach in Equation 8.1 and define:

$$e = \frac{\rho}{m} + (1 - \rho)$$

where m is the number of possible enactments in Φ, ρ is the learning rate and the prior is the maximum entropy distribution. We moderate the effect of the single observation (φ_k, ϕ) on the distribution $\mathbb{P}^t(\varphi'|\varphi)$ using a semantic similarity function Sim [10]:

$$e' = e \times \text{Sim}(\phi, \varphi)$$
$$e'' = (\text{Sim}(\phi, \varphi) \times (e - 1)) + 1$$

The prior probability is: $r = \mathbb{P}^{t-1}(\varphi'|\varphi)$, and as previously:

$$\mathbb{P}^t(\varphi'|\varphi) = \begin{cases} \text{comb}(r, e', \frac{1}{m}) & \text{if } e > \frac{1}{m} \text{ and } e' > \frac{1}{m}, \\ \text{comb}(r, e'', \frac{1}{m}) & \text{if } e < \frac{1}{m} \text{ and } e'' < \frac{1}{m}, \\ r & \text{otherwise.} \end{cases} \quad (8.3)$$

$\mathbb{P}^t(\varphi'|\varphi)$ is defined over all possible commitments φ, and is at too fine a granularity to be useful. Measures that summarise $\mathbb{P}^t(\varphi'|\varphi)$ are now described. These measures generalise what are commonly called *trust* of contractual enactments and *reliability* of information into a single computational framework. Tools from information theory are employed to construct these measures.

If an agent has a clear idea of the "ideal enactment" of a contract by the negotiation partner then by representing this as a probability distribution over enactment space, trust can be measured as the relative entropy of this ideal distribution with respect to the expected enactment also represented as a probability distribution over enactment space.

Perhaps a more realistic approach is based on the preferences over enactments. Representing the probability that the agent will, at the time of enactment, prefer any enactment to that which is specified in the contract, then use the distribution of expected enactments to generate a trust measure as "expected preferability".

Finally, another sense of trust is consistency of behaviour. This captures the notion of "you can trust John to sell poor wine". This may be measured as the entropy of the distribution of expected enactments. That is the degree of uncertainty in the behaviour of the negotiation partner.

For a formal treatment of these three measures see [17].

8.4.2 Integrity Model: $G^t_{\alpha\beta}$

α's estimate of β's *integrity* is the strength of α's belief that when enacting its contractual commitments β will do so taking account of α's interests—as opposed to executing the contract exactly as specified despite any subsequent observations. For example, "I did not purchase your mushrooms from the market because the only ones available were badly damaged by the recent rain.". Integrity is measured on a finite, fuzzy scale, `eval` $\in [0, 1]$. The evaluation space E must contain 0 (meaning "is of absolutely no use"), and must contain 1 (meaning "valued most highly"). We extend the `eval` function that was used for individual utterances in Section 8.2.1 to evaluate dialogues in this section.

In some situations the evaluation stage (described in Section 8.3) may take place a considerable time after the enactment stage; for example, "John advised me to use Spanish tiles in our kitchen—six years later they are still greatly admired" that implicitly rates the quality of John's advice. This illustrates why business relationships may take time to develop.

The integrity model, $G^t_{\alpha\beta}$, is required to do the following. Suppose α has a particular need ν, the integrity model $G^t_{\alpha\beta}$ aims to estimate the integrity of each agent in satisfying ν. This estimate is extracted from past commitment dialogues recorded in $\mathcal{H}^t_\alpha$. Suppose that α has held past commitment dialogues with β then for each such dialogue extract $C^t_{\alpha\beta}$ consisting of: an abstraction, using the is-a, $\leqslant$, relation the ontology, of the need that triggered the dialogue, the current contextual information and the final evaluation of the dialogue [9].

$G^t_{\alpha\beta}$ aims to form beliefs on the evaluation of future commitment dialogues with agent β based on $C^t_{\alpha\beta}$ by treating the evaluations, `eval`, as values of the dependent variable. Given a need, ν and context Θ^t agent α will form an *expectation*, $\epsilon(\nu, \Theta^t)$, of agent β's behaviour in satisfying that need given the context. The relationship between expectation ϵ and evaluation ϵ' is represented using conditional probabilities, $\mathbb{P}^{t'}_{\alpha\beta}(\epsilon'|\epsilon)$.

Any attempt to estimate $\mathbb{P}^{t'}_{\alpha\beta}(\epsilon'|\epsilon)$ has to deal with the unbounded variation in context Θ^t. It is reasonable to assume that the set of 'essentially different' contexts, Γ, is finite. We estimate $\mathbb{P}^t_{\alpha\beta}(\epsilon'|\epsilon(\nu,\gamma))$ for $\gamma \in \Gamma$. Suppose that $(e_i, (\nu', \gamma'))$ is observed where $e_i \in E = \{e_i\}_{i=1}^m$ the finite evaluation space—this observation will be used to update the estimate $\mathbb{P}^{t-1}(\epsilon'|\epsilon)$ that summarises the previous observations made. As in Section 8.4.1 we moderate the significance of this single observation using:

$$e = \frac{\rho}{m} + (1 - \rho)$$

where ρ is the learning rate and the prior is the maximum entropy distribution. Following Section 8.4.1 we moderate the affect of the single observation $(e_i, (\nu', \gamma'))$ on the distribution $\mathbb{P}^t(\epsilon' = e_i|\epsilon(\nu,\gamma))$ using the Sim function:

$$e' = e \times \text{Sim}(\nu, \nu') \times \text{Sim}(\gamma, \gamma')$$
$$e'' = (\text{Sim}(\nu, \nu') \times \text{Sim}(\gamma, \gamma') \times (e - 1)) + 1$$

The prior probability of observing $(e_i, \epsilon(\nu, \gamma))$ is: $r = \mathbb{P}^{t-1}(\epsilon' = e_i|\epsilon(\nu,\gamma)))$ and as previously:

$$\mathbb{P}^t(\epsilon' = e_i|\epsilon(\nu,\gamma))) = \begin{cases} \text{comb}(r, e', \frac{1}{m}) & \text{if } e > \frac{1}{m} \text{ and } e' > \frac{1}{m}, \\ \text{comb}(r, e'', \frac{1}{m}) & \text{if } e < \frac{1}{m} \text{ and } e'' < \frac{1}{m}, \\ r & \text{otherwise.} \end{cases} \quad (8.4)$$

This estimate for $\mathbb{P}^t_{\alpha\beta}(\epsilon'|\epsilon(\nu,\gamma))$ enables α to construct a probability distribution of the evaluation that will be observed if α selects β to satisfy need ν in context γ. It may be convenient to summarise β's expected integrity given particular circumstances. Summary measures may be constructed as for trust in the previous section. One approach is to define an distribution of what α considers to be ideal: $\mathbb{P}^t_I(\epsilon'|\epsilon(\nu,\gamma))$. Then to define integrity as the relative entropy between this ideal distribution and the estimated distribution:

$$G(\alpha,\beta,\nu,\gamma) = 1 - \sum_{\epsilon} \mathbb{P}^t_I(\epsilon'|\epsilon(\nu,\gamma)) \log \frac{\mathbb{P}^t_I(\epsilon'|\epsilon(\nu,\gamma))}{\mathbb{P}^t_{\alpha\beta}(\epsilon'|\epsilon(\nu,\gamma))}$$

As the evaluation space is metricated and totally ordered it is simpler to define integrity as expectation: $G(\alpha,\beta,\nu,\gamma) = \sum_i e_i \times \mathbb{P}^{t-1}(\epsilon' = e_i|\epsilon(\nu,\gamma)))$.

8.5 'Relationship-aware' Negotiation Strategies

An *argumentation strategy* determines the utterance that an agent should make in a dialogue given the agent's history. Why should an agent wish to have an argumentation strategy that is 'relationship-aware'? In real negotiations:

- the extent to which a partner agent will meet its commitments is uncertain (as is just about everything else), and
- each agent knows more about his private information than any other agent (unless the agent is very stupid)—this issue is called the *asymmetry of information*.

The first of these is closely associated with the concept of *trust*. The pervasive *asymmetry of information* raises the question of whether an agent will take advantage of its necessarily superior private information [11]. One way in which human agents manage the asymmetry of information is by building relationships with their negotiation partners.

Human agents consider building relationships in both personal and corporate negotiations. For example, when buying wine some humans will read reviews and price lists, develop some level of expertise and then decide which to buy. Other agents will visit their trusted wine merchant and ask for advice. On the corporate level, in the 1970s IBM offered substantial discounts to customers who only used IBM equipment and thus encouraged the development of a close relationship with their customers. Many organisations negotiate corporate deals for items such as laptop computers that include some degree of personal service with their employees. Humans rely on relationships in an attempt to be "looked after", or "not taken advantage of", in their dealings.

For software agents a relationship exists between any pair of agents that have interacted. An agent may desire to build relationships with other agents for essentially two separate purposes. First, to attempt to address the information asymmetry between himself and his negotiation partners in contractual negotiations. Second, to share strategic information; for example, "John usually gives me a 10% discount.". The relationship model aims to support both of these purposes.

Any negotiation strategy has to address two issues. First to choose who to negotiate with, and second, how to manage a negotiation.

8.5.1 Selecting an Interaction Partner

We assume that α has a set of high-level *aims* that includes its strategic aims and its principles, and knows what its future *needs* are. The selection of an interaction partner is considered in two steps. First, a set of partners are chosen for each need, ν, this set is called the *pool* for ν. Second, given that a need has triggered a partner is chosen from the pool for that need.

The pool for each need is chosen from agents who are known to α using the trust model and the integrity model, and from unknown agents using a socially derived reputation model—reputation lies outside the present discussion [19]. The integrity of all models decays in time to their decay limit unless the model is refreshed with new observations. That is, the uncertainty of information extracted from the model will increase under the same conditions. This means that the agents should attempt to manage their models by sustaining a sufficient refresh rate to ensure that $\mathbb{H}(M^t_{\alpha\beta})$ does not decrease. An obvious exception being that α has determined that she no longer wished to deal with β. This means that such an agent should limit the number of potential interaction partners taking account of the frequency with which their models are refreshed.

To illustrate the need to control the size of a need's pool, suppose that α selects a partner on the basis of a random variable, X, that is defined over an evaluation space of k distinct terms—that could for example be `excellent`, `mediocre` and `shocking`. Let h denote the entropy of X, $\mathbb{H}(X)$. The maximum entropy of such a variable is: $\overline{h} = -\log_2(\frac{1}{k})$, and the minimum is: $\underline{h} = -\log(1) = 0$; for example, if $k = 4$ then $\overline{h} = 2$. If the integrity of X decays then: $\frac{dh}{dt} = \mu \times (\overline{h} - h)$ and decay from $\underline{h}$ to $\overline{h}$ is given by: $h = \overline{h} - (\overline{h} - \underline{h}) \times e^{-\mu t}$ where t is time and μ is the integrity decay constant. The integrity of X decays by 10% in each time step if $\mu = 0.1055$. Suppose the size of the pool is n and that partners are selected in rotation at each time step. Then each partner will be selected when its X estimate has decayed for n time steps, and if $n = 3$ then the entropy of its X estimate will have decayed by 27% by the time each partner is chosen, and if $n = 10$ then the decay will be 65%. This illustrates the obvious intuition that the smaller the pool the more accurate are the performance estimates for agents in the pool.

α maintains a strategic view on its desired form of the entire relationship model for each agent in the pool. This applies particularly to its desired form of the intimacy model. This is α's *relationship target* $T^t_{\alpha\beta}$ model for agent β. This pair of intimacy models is represented in the {L,O,G,I,C} framework: $T^t_{\alpha\beta(f,c)} = \left(TJ^t_{\alpha\beta(f,c)}, TK^t_{\alpha\beta(f,c)}\right)$ where $TJ_{\alpha\beta}$ and $TK_{\alpha\beta}$ are respectively the targets for $J_{\alpha\beta}$ and $K_{\alpha\beta}$. The relationship target is intended to capture the sense in which human agents express their relationship aspirations; for example "we really need to find a red wine supplier whose advice we can rely on". In Section 8.5.2 we will describe how α proactively moulds a relationship progressively towards its aspirational target form—given the cooperation of its negotiation partner.

When an interaction takes place the intimacy models, $\left(J^t_{\alpha\beta(f,c)}, K^t_{\alpha\beta(f,c)}\right)$, will change. Prior to a negotiation, α expresses the extent that it wishes these models to be moved by the negotiation towards the target intimacy by articulating a *negotiation target*, $N^t_{\alpha\beta(f,c)} = \left(NJ^t_{\alpha\beta(f,c)}, NK^t_{\alpha\beta(f,c)}\right)$, that is α's aspirations at time t for the intimacy to be achieved when the negotiation is complete. At a finer level of granularity, the *interaction target*, $A^t_{\alpha\beta(f,c)} = \left(AJ^t_{\alpha\beta(f,c)}, AK^t_{\alpha\beta(f,c)}\right)$, expresses α's aspirations for the intimacy to be achieved during the next exchange of utterances. Given the uncertainty in behaviour in any negotiation, these targets represent the agent's aspirations concerning the resulting action of other agents rather than expectations of what those actions will be. They are rough indications only of what α hopes may occur. We will see that the specification of these targets constrains the individual utterances that α makes.

Given a need ν and a pool for ν it remains to select an interaction partner from the pool. For each pair of potential partners, β and β', we assume that α is able to estimate, using the relationship, trust and integrity measures described above, $\mathbb{P}^t(\beta > \beta')|\nu$, the probability that a negotiation with β to satisfy ν will lead to a better outcome than with β' [14]. From these pairwise comparisons an approximate estimate may be derived* of $\mathbb{P}^t(\beta \gg)|\nu$, the "probability that a negotiation with β to satisfy ν will lead to a better outcome than with any of the other potential partners" [1]. Then select a negotiation partner using the stochastic strategy: $\mathbb{P}^t(\text{Select } \beta_i|\nu) = \mathbb{P}^t(\beta_i \gg)|\nu$. That is, the probability of selecting β_i is equal to the probability that β_i is the best choice for ν.

Having selected the interaction partner, and having set the relationship target, α now manages the interaction itself. We think of individual utterances as consisting of illocutions with contractual implications, such as `accept`, and argumentative illocutions such as `opinion`. We use the term *strategy* to refer to the agent's method of determining its contractual illocutions, and *tactics* to refer to the method for determining its argumentation. Negotiation strategies have been widely discussed elsewhere. In the following section we discuss tactics.

8.5.2 The RANA Agent Architecture

The RANA agent architecture is an information-based agent for relationship-aware negotiation. We first describe the relationship between RANA and other information-based agents, and then discuss the tactics that RANA employs.

Figure **8.2** shows the information-based negotiation agent, argumentation agent and the RANA agent. These three agent architectures are associated with an increasingly rich communication language, and with additional models all of which are summaries of $\mathcal{H}^t_\alpha$. The honour model measures the veracity of arguments and is discussed in [19]. Reputation is the social construction of shared opinion and is discussed in [19].

α maintains a pool of potential partners for each need as described in Section 8.5.1. Suppose that some *event* occurs that triggers one of α's needs, the *relationship targets*, *trust* and *integrity* estimates and the current state of the *relationship model* together determine a negotiation *partner*, β, with whom α will attempt to satisfy its need, and, at the same time, to alter the state of the relationship model in line with a *negotiation target*. When interaction commences the negotiation target, that may alter during the negotiation, determines the *interaction target*. α's negotiation strategy then determines

*See for example [20]

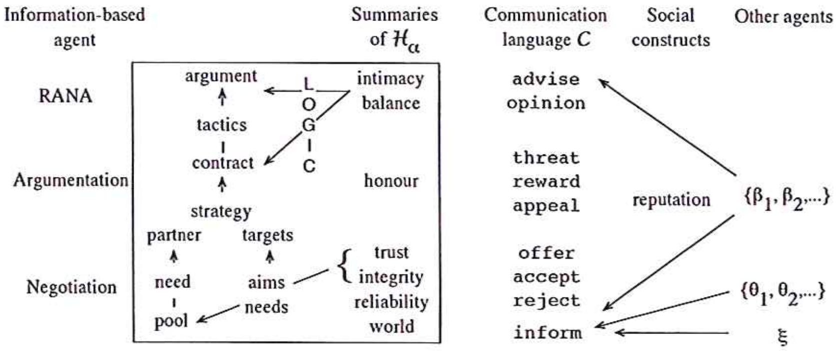

Fig. 8.2 Three information-based architectures: negotiation, argumentation and RANA. These three architectures are related to an increasingly rich communication language, and to the models required all of which are summaries of $\mathcal{H}^t_\alpha$—see the discussion in Section 8.5.2. Reputation is the social construction of shared opinion and is discussed in [19]. The honour model is concerned with the veracity of arguments, it is described in [17].

the utilitarian illocutionary content of its next utterance. Finally its tactics (see below) wrap the utilitarian content in argumentation.

In Section 8.5.1 we introduced the notion of a *target* for intimacy that may be defined for interactions, dialogues and relationships for both for an agent and tentatively for its interaction partner. Information-based agents manage the information aspects of interaction as well as the utilitarian aspects, and the RANA agent has both information-based targets and utilitarian targets. These targets evolve in time as described above, they substantially constrains the utterances that α can make.

For example, when constructing a simple offer, the interaction targets will constrain: the value of the offer to the speaker, the estimated value to the listener, the private information loss to the speaker and the expected information gain to the listener. Similarly when making an argumentative utterance, such as "In my opinion Sierra Enterprises is an excellent company for investment."—the act of uttering this statement has both information-based and utilitarian implications.

α's *tactics* wraps a possibly empty set of utilitarian particles in argumentative particles to form a complete *utterance*.

One general strategy for managing this process in the context of an interaction is the "equitable information revelation strategy" that aims to reveal information of similar value to the interaction partner as has been observed in recent utterances from her [21]. The difference in this discussion is that we now have the full structure of the {L,O,G,I,C} framework across which we may, if we wish, balance information value in the various {L,O,G,I,C} categories. To manage argumentation, including `inform`, `threaten`, `reward`, `appeal`, `opinion` and `advice`, the tactics considers the short-term, medium-term and long-term effect that the argumentation is intended to have on β's future states and actions, particularly if those actions may then alter α's states. The effective use of argumentation needs to consider more than information revelation. One key strength of information-based agents [1] is that they are able to manage both the utilitarian-based and the information-based aspects of interaction. Finally we recall our former remark that, unlike utilitarian-based valuations, information-based valuations can always be computed,

Acknowledgments

Research supported by the Agreement Technologies CONSOLIDER project under contract CSD2007-0022 and INGENIO 2010, by the Spanish Ministry of Science through the CBIT project TIN2010-16306, by the ERA-NET project ACE, and by the Agreement Technologies COST Action, IC0801.

Conflict of Interest

The authors confirm that this article content has no conflict of interest.

References

1. C. Sierra and J. Debenham. Information-based agency. In *Proceedings of Twentieth International Joint Conference on Artificial Intelligence IJCAI-07*, Hyderabad, India, January 2007, pp. 1513–1518.
2. W. Ulaga and A. Eggert. Relationship value in business markets: The construct and its dimensions. *Journal of Business To Business Marketing*. Vol. 12, no. 1, pp. 73 – 99, February 2005.
3. P. Rauyruena and K. E. Miller. Relationship quality as a predictor of B2B customer loyalty. *Journal of Business Research*. Vol. 60, no. 1, pp. 21–31, January 2007.
4. C. Sierra and J. Debenham. Agent Argumentation with Opinions and Advice. In *Proceedings of Thirtieth International Conference on Research and Development in Intelligent Systems AI-2010*, Cambridge, UK, 14-16 December 2010, pp. 21–34.
5. C. Sierra and J. Debenham. When Trust Is Not Enough. In *Proceedings of Twelth International Conference on Conference E-Commerce and Web Technologies ECWeb-2011*, Toulouse, France, 29 August - 2 September 2011, pp. 246–257.
6. J. Sabater and C. Sierra. Review on computational trust and reputation models. *Artificial Intelligence Review*. Vol. 24, no. 1, pp. 33–60, September 2005.

7. C. Sierra, N. Jennings, P. Noriega, and S. Parsons, *Proceedings of the 4th International Workshop on Intelligent Agents IV, Agent Theories, Architectures, and Languages,* London, UK: Springer-Verlag, 1997, ch. A Framework for Argumentation-Based Negotiation, pp. 177–192.
8. FIPA, "Communicative act library specification," Foundation for Intelligent Physical Agents, Geneva, Switzerland, Tech. Rep. SC00037J, 2002.
9. Y. Kalfoglou and M. Schorlemmer. IF-Map: An ontology-mapping method based on information-flow theory. In *Journal on Data Semantics I*, Lecture Notes in Computer Science, S. Spaccapietra, S. March, and K. Aberer, Eds., Springer-Verlag: Heidelberg, Germany, 2003, Vol. 2800, pp. 98–127.
10. Y. Li, Z. A. Bandar, and D. McLean. An approach for measuring semantic similarity between words using multiple information sources. *IEEE Transactions on Knowledge and Data Engineering*. Vol. 15, no. 4, pp. 871 – 882, July / August 2003.
11. J. S. Adams, "Inequity in social exchange. In *Advances in experimental social psychology*, L. Berkowitz, Ed., New York: Academic Press, 1965, Vol. 2.
12. H. Sondak, M. A. Neale, and R. Pinkley. The negotiated allocations of benefits and burdens: The impact of outcome valence, contribution, and relationship. *Organizational Behaviour and Human Decision Processes*. 3, pp. 249–260, December 1995.
13. K. L. Valley, M. A. Neale, and E. A. Mannix. Friends, lovers, colleagues, strangers: The effects of relationships on the process and outcome of negotiations. In *Research in Negotiation in Organizations*, R. Bies, R. Lewicki, and B. Sheppard, Eds., JAI Press, 1995, vol. 5, pp. 65–94.
14. M. H. Bazerman, G. F. Loewenstein, and S. B. White. Reversal of preference in allocation decisions: judging an alternative versus choosing among alternatives. *Administration Science Quarterly*. 37, pp. 220–240, 1992.
15. D. Artz and Y. Gil. A survey of trust in computer science and the semantic web. *Web Semantics: Science, Services and Agents on the World Wide Web*. Vol. 5, no. 2, pp. 58–71, June 2007.
16. 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, pp. 1026–1033.
17. ——, Trust and honour in information-based agency. In *Proceedings Fifth International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2006*, P. Stone and G. Weiss, Eds., Hakodate, Japan: ACM Press, New York, May 2006, pp. 1225 – 1232.
18. D. MacKay, *Information Theory, Inference and Learning Algorithms*, Cambridge University Press, 2003.
19. C. Sierra and J. Debenham. Information-based reputation. In *First International Conference on Reputation: Theory and Technology (ICORE'09)*, M. Paolucci, Ed., Gargonza, Italy, 2009, pp. 5–19.
20. ——, Information-based deliberation. In *Proceedings Seventh International Conference on Autonomous Agents and Multi Agent Systems AAMAS-2008*, L. Padgham, D. Parkes, J. Müller, and S. Parsons, Eds., Estoril, Portugal: ACM Press, New York, May 2008.
21. P. Faratin, C. Sierra, and N. Jennings. Using similarity criteria to make issue trade-offs in automated negotiation. *Journal of Artificial Intelligence*. Vol. 142, no. 2, pp. 205–237, 2003.