

An eNegotiation Framework

John Debenham

University of Technology, Sydney,
Faculty of Information Technology,
PO Box 123, NSW 2007, Australia
debenham@it.uts.edu.au

Abstract. Negotiation between two trading agents is a two-stage process. First, the agents exchange offers whilst acquiring and exchanging information. Second, they attempt to reach a mutually satisfactory agreement in the light of that information. This process aims to reach informed decisions in eMarket bargaining by integrating the exchange of offers and the acquisition and exchange of information drawn from the environment. Negotiation proceeds by a loose alternating offers protocol that is intended to converge when the agents believe that they are fully informed. The eNegotiation framework is a prototype for developing agents that exploit an information-rich environment in one-to-one negotiation. This work is part of a program that is investigating deeper issues of market evolution and business network development in an immersive, virtual worlds eMarket environment.

1. Introduction

The economic (game theoretic) and multiagent interest in negotiation became fused in “Agent Mediated Electronic Commerce” which has attracted increasing attention since the mid-1990s. Much of this work has studied economically rational strategies and mechanisms in an eMarket context; for example, the use of auction mechanisms for multi-unit and multi-attribute trade [1]. From an artificial intelligence perspective humans do not always strive to be utility optimisers. Sometimes they do so, for example, traders taking short-term positions in exchanges act to make short-term profits. And sometimes they do not, for example when buying a house or an automobile an intelligent human may have intelligent, justifiable but completely irrational motives. The study of negotiation between artificially intelligent agents should not be bound by economic rationality; see also [2].

Here, two agents reach a mutually beneficial agreement on a set of issues because they have chosen to become well informed. The aim is to reach *well-informed* decisions rather than *economically rational* decisions. “Good negotiators, therefore, undertake integrated processes of knowledge acquisition combining sources of knowledge obtained at and away from the negotiation table. They learn in order to plan and plan in order to learn” [3]. The work described here attempts to encapsulate this intimate link between the negotiation itself and the knowledge generated both by and because of it. During a negotiation, an agent may actively acquire information that it may, or may not, choose to place on the negotiation table. Information is a strategic weapon in competitive interaction [4].

The eNegotiation framework is a prototype for developing negotiation agents that exploit an information-rich environment in one-to-one negotiation. The

negotiation agents are responsible for conducting the negotiation on behalf of another. They are purely self-interested agents. Eventually, the negotiation architecture will consist of two further generic agent types: mediation agents and observer agents. *Mediation agents* are impartial observers of negotiations and develop expertise in how to lead multi-attribute negotiations to a successful conclusion [5]. *Observer agents* analyse unsuccessful negotiations. A failed negotiation is a missed business opportunity and so may be a clue to the discovery of an innovative, evolutionary way of doing business.

The eNegotiation framework is being built by members of the eMarkets research group in the Faculty of IT at UTS in Sydney: <http://research.it.uts.edu.au/emarkets/> where the project is named the “Curious Negotiator” [6]. Minghui Li, a PhD student in the group is building the eNegotiation framework. To date a version has been roughed up in Java. Woralak Kongdenfha, another PhD student is developing negotiation strategies for the framework. Bots that access the World Wide Web have been designed and built by Dr Debbie Zhang, a Senior Researcher in the group. A fully integrated demonstrable version of the framework will be operational in July when it will be trialed on an exemplar application described below and made publicly available. Plans are to then re-engineer the framework using the Australian Jack agent building-environment and to apply it to a variety of richer negotiation situations.

The following *exemplar application* will be used in initial trials of the framework. An agent has owned a digital camera for three years and wants to upgrade to a better model. This agent has set some money aside, and has a pretty good idea of the features that she wants in her next digital camera. She is in no hurry to buy. She has let it be known that she will consider buying a second-hand camera as long as it is in as-new condition. On receiving an opening offer from a seller to buy a camera she then negotiates with that seller—hopefully leading to a mutually satisfactory deal. This problem is simple in that the actual negotiation is single-attribute—namely the price of the camera. It involves an assessment of the extent to which the camera offered is suited to her requirements—this involves multi-attribute reasoning. It also involves an assessment of the market value of the camera. Both of these two issues are resolved by reference to information on the World-Wide-Web that is accessed by information gathering bots. The negotiation uses a “loose” alternating offers protocol in which each agent may make an offer at any time but is under no contractual obligation until a deal is made [7]. During a negotiation the agent is expected to be able to feed information to its opponent such as “that model does not have the features that I am looking for” and “your price is higher than I could buy it for on the Web”. This exemplar application has been chosen to enable us to make the whole negotiation apparatus work in a satisfactory way on a moderately simple problem. Future plans are to use the augmented negotiation framework in the group’s two long-term projects that are in eMarket evolution and in the development of business networks in an eMarket environment.

Game theoretic analyses of *bargaining* are founded on the notion of agents as utility optimisers in the presence of complete and incomplete information about their opponents [8]. To deal with incomplete information, agents are typically assumed to know the probability distribution of that which is not known completely [9]. Likewise, game theoretic analyses of auction mechanisms in *markets* are typically based on assumptions concerning the probability distributions of the utilities of the agents in both the private and public value cases [10]. Further, game theoretic analyses of *exchanges* focus on price formation, strategies and mechanisms for agents

with complete and incomplete information. Game theory tells us what to do, and what outcome to expect, in many well-known negotiation situations, but these strategies and expectations are derived from assumptions. For example, in auction theory, the most restrictive of these are that the players are utility optimisers with SIPV (symmetric, independent, private valuations) [10].

Impressive though the achievements in game theory are, the significance of that work in describing how humans negotiate—and so too how artificially intelligent systems should negotiate—is questionable. When confronted with the SIPV formula for an optimal bid in a first-price sealed-bid auction a colleague once remarked: “Look, I may simply want to buy it”. The sense in this remark is that utility optimisation is not always the motivation for a negotiation.

The work here is based on the notion that when an intelligent agent buys a hat, a car, a house or a company she does so because she *feels comfortable* with the general terms of the deal. This “feeling of comfort” is achieved as a result of information acquisition and validation. This information includes both that which the agent is given and information that which the agent chooses to acquire. Here negotiation is as much of an information exchange process as it is an offer exchange process. The exchange of information is equally as important as the exchange of offers—one feeds off the other. In so far as game theory aims to reach a *rational* solution, the aim here is to reach an *informed* solution.

2. Managing an Information-Rich Environment

The form of negotiation considered is between two agents in an information-rich environment. That is, the agents have access to general information that may be of use to them. They also exchange information as part of the negotiation process. The two agents are called “me” and my opponent ω . The environment here is the Internet, in particular the World Wide Web, from which information is extracted on demand using special purpose ‘bots’. So my agent, “me”, may receive information either from ω or from one of these bots. In a ‘real life’ negotiation, the sort of information that is tabled during a negotiation includes statements such as “this is the last bottle available”, “you won’t get a better price than this”, and so on. To avoid the issue of natural language understanding, and other more general semantic issues, the interface between each of the negotiating agents and these information sources is represented using the language of first-order, typed predicate logic, and a set of pre-agreed, pre-specified predicates. All information passed to “me” is expressed in this way.

As well as being unambiguous, the use of first-order typed predicate logic has the advantage of admitting metrics that describe, for example, how “close” two statements are. These metrics are useful in enabling the agent to manage the information extracted from the environment in a strategic way. The terms predicate, term, function, variable and constant are used in their well-accepted sense. In typed logic a term, constant or variable appearing as part of a proposition belongs to a certain domain determined by the argument in which it occurs in the predicate of that proposition. If functions can be avoided then these metrics are simpler. Functions are not used in this discussion, although their inclusion does not introduce any technical problems beyond making the Herbrand universe unbounded. The notation: $\langle \text{variable} \rangle / \langle \text{term} \rangle$ denotes the substitution in which the named variable, the *subject*, is replaced by the term wherever that variable occurs in an expression.

Using the usual notation and terminology for Horn clauses, given a set of facts—ie: unit clauses with just one positive predicate—the following defines a partial ordering of those facts:

$P(\underline{x}) \geq_s P(\underline{y})$ if there exists set of substitutions J whose subjects are variables in \underline{x} such that: $\underline{y} = \underline{x} / J$

where \underline{x} and \underline{y} are complete sets of arguments of P . If this holds then $P(\underline{x})$ *subsumes* $P(\underline{y})$. This ordering between two literals captures the notion of one being *more general* than the other.

Given two positive propositions, p_1 and p_2 , both of which have the same predicate symbol, then the *unification* of those two propositions is denoted by $p_1 \cap p_2$, and the *anti-unification*—using the terminology of Reynolds—by $p_1 \cup p_2$. [The anti-unification of two positive literals is the unique literal that subsumes both of them and of all such literals is minimal with respect to the partial order \geq_s .]

Suppose that G is a finite set of ground literals that is closed with respect to unification and anti-unification. G is a non-modular lattice where unification and anti-unification are the meet and join respectively. Given a literal c , the function $\lambda_G : \{\text{the set of all literals}\} \rightarrow \{\text{the positive integers}\}$ is defined as $\lambda_G(c) =$ the number of literals in G with which c is resolvable. $\lambda_G(c)$ is a measure of the “generality” of c with respect to unification over G . It is a measure of how “general” c is over G . Given an isotone valuation v on the vertices of G , the function:

$$\delta_1(c_1, c_2) = v[c_1] - v[c_2]$$

is a measure of how much “more general” c_1 is compared with c_2 . Further

$$\delta_2(c_1, c_2) = v[c_1 \cup c_2] - v[c_1 \cap c_2]$$

is a measure of the “distance” between c_1 and c_2 . There is a problem with δ_2 when the meaning of the partial ordering is such that: if $c_1 \leq c_2$ then knowing c_2 means that knowing c_1 brings no new information—this is the case with the partial ordering \geq_s . In this case the function:

$$\delta_3(c_1, c_2) = v[c_1 \cup c_2] - v[c_2]$$

is a measure of how much more “novel” c_1 is compared with c_2 . [Another possibility is $v[c_1] - v[c_1 \cap c_2]$ but that is not so useful when the intersection is likely to be empty—ie: at the “bottom” of the lattice—if $c_1 \leq c_2$ then $\delta_3(c_1, c_2) = 0$.

Suppose that J is any subset of G , given a literal c , then:

$$\delta_i(c, J) = \min_{s \in J} \delta_i(c, s) \quad \text{for } i = 1, 2, 3.$$

are generalisations of the measures between two vertices to measures between a vertex and a set of vertices. Suppose that S is the set of literals that agent “me” knows, let H be set of all possible literals obtained by instantiating the predicates in $\{c\} \cup S$ over the Herbrand universe for $\{c\} \cup S$, then the above two measures are:

$$\delta_1(c, S) = \min_{s \in S} [\lambda_H[c] - \lambda_H[s]]$$

$$\delta_2(c, S) = \min_{s \in S} [\lambda_H[c \cup s] - \lambda_H[c \cap s]]$$

$$\delta_3(c, S) = \min_{s \in S} [\lambda_H[c \cup s] - \lambda_H[s]]$$

If there are functions present then the Herbrand universe will be non-finite. In this case H may be defined as the set of such instantiations up to a certain chosen function nesting depth. As long as that chosen nesting depth is greater than that occurring in the literals in $\{c\} \cup S$ the resulting measures are generally useful.

This leads to a general approach to managing information in negotiation:

- develop a predicate logic representation
- introduce a partial ordering on that representation that captures the meaning of “generalisation” in a way that is relevant to the application
- define lattice meet and join operations
- define an isotone valuation on the lattice that is consistent with the ordering—this valuation need not necessarily treat each argument (domain) in the same way (as in the example above)—it may be useful to introduce weights that give some arguments more significance than others.

and then apply the ideas above. The introduction of two additional measures, of “cost” of acquiring the information and of “belief” in its validity, complete the machinery required to manage the acquisition and maintenance of an information base in a strategic way.

To illustrate this, in the exemplar system one job that the information bots do is to find the cheapest camera that has a certain set of features. This is achieved by scraping the sites of various information providers $\{P_i\}$. My agent knows what sites are available to it. To my agent, the bot that scrapes site P_j is represented as:

Cheapest($P_j: \{P_i\}, f_1:F_1, \dots, f_n:F_n; c:C, p:\$$) [1]

where f_i is the i 'th feature (such as, whether the camera has an on-board, electronic flash), F_i is the domain (ie: all possible values) for each of these features, c is the cheapest camera with those features (according to source P_j), and p is its price (according to source P_j). The “;” separates the “inputs” from the “outputs”. It is assumed that each domain F_i contains the features dc_i meaning “I don't care what f_i is”. So to my agent the bots appear as a set of as yet uninstantiated literals. If my agent chooses to activate the j 'th bot then it does so by setting the P_j and all of the f_i to particular values and the result of the scraping exercise is then, for example:

Cheapest($P_2, \text{“no flash”}, \dots, dc_n; \text{Kodak123}, \54)

meaning that “according to site P_2 the Kodak123 camera is the cheapest camera available with the set of features specified and it costs \$54”. In so far as that information is valid, it can only be assumed to be valid at the time that bot is activated. Here it is assumed that information is valid for the period of the negotiation. There is also the matter of how reliable the information is—with what degree of belief should the agent believe it to be true? This is achieved by attaching a certainty factor to each retrieved tuple that is based solely on an assessment of the general reliability of the information provider P_j .

The “don't care” value, dc , may be a value for any or all of the $\{f_i\}$. [Sites vary in their accuracy and completeness—some do not contain information for every such combination.] At some time after activation a bot, information will be received by my agent also in form [1] but with the “output” values (ie: values for c_x and p_x) specified. This completes the first step—the predicate logic representation. Introduce the partial order defined by the values of the first $(n+1)$ arguments only. If the first $(n+1)$ arguments of expressions e_1 and e_2 are the same except that e_1 contains dc values where e_2 does not then $e_1 \geq_D e_2$. \geq_D is the partial ordering. If e_1 and e_2 can be “unified” by setting dc input values in one expression to the corresponding value in the other then define $e_1 \cap e_2$ to be the expression with that “unified” set of input values, [output values remaining undefined as they play no part in this ordering], otherwise $e_1 \cap e_2$ is the empty expression. In defining this ordering the constant dc is treated rather like a logic variable in conventional unification, but *it is a logical*

constant. The expression $e_1 \cup e_2$ is defined similarly—ie: to reflect the sense of “generality” in the “Cheapest” predicate. Let S be the set of ground literals that the agent believes to be true, then define H as above by treating dc as a “variable”, and λ_H as above.

After all of this the measures of “how general” statements are [11], and how “close” statements are, may be used to strategically gather information. In the exemplar system this amounts to managing the “information lattice” in a fairly common sense manner. It enables the agent to explore “nearby” alternatives, and to obtain a “more general feeling” for price. More importantly, it has been shown that this approach may be used for first-order resolution logic—it is proposed that this approach may be applied to any application that can be expressed in first-order logic, including those of greater complexity, to strategically manage the acquisition of information. This approach relies on the hand crafting of the lattice structure and the lattice ordering for each predicate involved.

3. Accepting an Offer

A mechanism decides whether to accept an offer or to cease the negotiation—neither of these may occur in which case either one of these two agents has to “think again”. This mechanism is the first step towards automatic negotiation in which agents aim to reach comfortable agreements through managing both the exchange of offers and requests for information in an integrated way. The goal here is to design something that works. The solution described does work, but is not claimed to be the best, or even a particularly good, way representing “comfort”. The following section embeds this mechanism into a negotiation framework.

Suppose that a negotiation between two agents, my opponent ω and me, stands alone—that it is not part of a related sequence of negotiations. Each agent aims to reach an agreement on a deal. A *deal* Δ is a commitment for me to do something, τ (my “terms”), subject to the other agent agreeing to do something, υ , $\Delta = (\tau, \upsilon)$. Suppose my negotiation with ω is conducted in the light of information, ι , (ie: ι is the information available to me, but not necessarily to ω) then:

$\text{Acc}(\omega, \tau, \upsilon, \iota)$

is a proposition in which the predicate “Acc” means “taking account of information ι , if I agree to do τ , subject to ω agreeing to do υ , then I will feel comfortable with the deal”. This raises the two questions of defining what it means to say that I feel comfortable with a deal, and determining the point in time at which that determination should be made. It makes acceptable sense to the author to state “looking back on it, I made the right decision at the time”—those two questions are not addressed further here. The information, ι , consists of any information tabled by either party as part of the negotiation process, and any information extracted by me from the environment as described in the previous section. The commitments, τ and υ , may be quite complex and may include staggered payments or related “sweeteners”—in general they will contain multiple attributes [12]. A deal is normally associated with a time, or time frame, when the respective commitments of the two agents are to be performed. In addition, the “Acc” predicate could contain an assessment of the guarantees, assurances, escrow services and “after sales service”, γ —but this has been omitted for simplicity. In addition to the Acc predicate, a “Cea” predicate aims to detect when a negotiation should cease—possibly because it appears to be a “waste of time”. The Cea predicate is not described here.

A probability is attached to the “Acc” proposition so that it may be used as the basis for making a decision on whether or not to accept the terms τ . This probability aims to represent the likelihood that the proposition will turn out to be true in due course—see the comment above concerning the point in time at which “I may feel comfortable”. A fuzzy set is introduced to define “Accept” to some level of confidence χ . Its membership function $\mu_A: [0, 1] \rightarrow [0, 1]$ is monotonic increasing. μ_A is applied to




Figure 1. Fuzzy Accept

$P(\text{Acc}(\omega, \tau, \nu, \iota))$. This function is a matter for each agent to specify, and in general may be a function of the terms, τ , and maybe ν , and so there is no claim that this work is leading to context-independent super-negotiator. See Fig. 1 for an example of a piece-wise linear fuzzy membership function μ_A . The μ_A function just scales the values—it accepts offers when $P(\text{Acc}) \geq \mu_A^{-1}(\chi)$. It is included so that the framework will support experiments with cooperative agents seeking to reach a compromise deal with equivalent certainty—the μ_A function enables agents to have different “types” analogous to Kasbah’s “anxious”, “cool-headed”, and “frugal” [13]. The value χ_{\min} is the greatest value of confidence χ for which both:

$P(\text{Acc}(\omega, \tau, \nu, \iota)) \in \text{Accept}$, and
 $P(\text{Cea}(\omega, \tau, \nu, \iota)) \in \text{Cease}$

are true. χ_{\min} is the confidence value at which the agent will be equally disposed to accept the offer *and* to cease negotiation—values of χ at or below χ_{\min} are of no interest here. The level of confidence χ is vaguely related to the notion of “risk averseness” in game theory. If, to confidence χ ,

$P(\text{Acc}(\omega, \tau, \nu, \iota)) \notin \text{Accept}$, and
 $P(\text{Cea}(\omega, \tau, \nu, \iota)) \notin \text{Cease}$

then the agent continues to negotiate as described in the following section.

The probability attached to the “Acc” proposition—and similarly for the “Cea” proposition—is derived from probabilities attached to four other propositions:

$\text{Suited}(\iota, \nu)$, $\text{Good}(\iota, \omega)$, $\text{Fair}(\iota, \tau, \nu)$, and $\text{Me}(\iota, \tau, \nu)$

meaning respectively: “taking account of information ι , terms ν are perfectly *suited* to my needs”, “taking account of information ι , ω will be a *good* agent for me to be doing business with”, “taking account of information ι , τ are generally considered to be *fair* terms in exchange for ω agreeing to do ν given that ω is impeccable”, and “independent of what everybody else believes, on strictly subjective grounds, taking account of information ι , τ are fair terms in exchange for ω agreeing to do ν given that ω is impeccable”. The last two of these four explicitly ignore the suitability of ν and factors out the appropriateness of the opponent ω . If ω ’s reputation is doubtful then the mechanism will compensate for these doubts by looking for better terms than would be comfortable for me if I had dealt with an impeccable opponent. The difference in meaning between the third and fourth proposition is that the third captures the concept of “a fair market deal” and the fourth a strictly subjective “what ν is worth to me”. The “Me” proposition is closely related to the concept of private valuations in game theory.

To deal with “Suited” depends on what the terms ν are. In the exemplar system the terms will be a second-hand camera if I am buying, and money if I am selling.

There are a large number of sites on the World Wide Web that may be used to estimate the extent to which a digital camera will suit me. For example:

<http://www.dpreview.com/reviews/compare.asp>

Those sites are used in the exemplar system to attach a level of confidence in the “Suited” proposition. This is described in detail in the following section. In the digital camera market area there are a variety of such advisory sites $\{S_1, \dots, S_n\}$. These sites differ by the questions that they ask. If the agent’s preferences are known sufficiently well to answer some of the questions on one of these sites then it is a matter of seeing whether the camera, v , is amongst the recommended choices. So each of the sites $\{S_i\}$ may be used to generate evidence for “Suited”. These sites are designed to help a buyer choose a camera—they are used here for a rather different task: to check that a given camera is suitable for me. A Bayesian revises the *a priori* probability attached to “Suited” in the light of evidence obtained from the $\{S_i\}$. For any seller it is assumed here that ω is a commitment to deliver money, and that the probability of money being “suitable” is 1.0. That is, $P(\text{Suited}(\text{“money”})) = 1.0$. Another approach to dealing with the suitability of a deal is used in some B2B e-procurement reverse auctions [14]. There the buyer specifies various packages and attaches a factor to them as a real number ≥ 1.0 . The sellers bid on the various packages. The bids are then multiplied by the factor associated with the relevant package to determine the winner—ie: the lowest bid.

Attaching a probability to “Good” involves an assessment of the reliability of the opposing agent. For some retailers (sellers), information—of possibly questionable reliability—may be extracted from sites that rate them. For individuals, this may be done either through assessing their reputation established during prior trades in a multi-agent environment [15], or through the use of some intermediate escrow service that in turn is deemed to be “reliable” in which case the assessment is of the individual *and* the escrow service. In the exemplar system this factor has been ignored as the camera is assumed to be offered for sale by an individual. That is, $P(\text{Good}(\omega)) = 1.0$.

Attaching a probability to “Fair” is achieved by reference to an appropriate market. When a deal is to purchase retail goods for money this may be achieved by using bots to scrape a set of sites. In the exemplar system the second-hand camera, v , is assumed to be in “as new” condition, and so its fair market value, α_v , is obtained by discounting the best price, β_v , that is found for that item in new condition, say by 30%. The sites scraped are those merchant sites that are linked from the sites used for “Suited” above. These are shown as $\{\bar{S}_1, \dots, \bar{S}_n\}$. There is no causal link between the $\{S_i\}$ and the $\{\bar{S}_i\}$. If the terms τ specify that money is to be exchanged for the camera then the set $\{\tau : \text{Fair}(\tau, \omega)\}$ is defined by a fuzzy membership function such as that shown in Figure 2.

Attaching a probability to “Me” is a purely subjective matter. For a single-attribute, real-valued τ it is dealt with in a similar way to “Fair” above using a fuzzy membership function that is zero at my upper limit, $\bar{\sigma}_v$, and 1.0 at the greatest price that “I would be delighted to pay”, σ_v . Figure 3 shows a possible fuzzy membership function for “Me” for such a τ . That function aims to encapsulate the notion of “ τ




Figure 2. Fuzzy Fair

are personally fair terms to me in exchange for ω doing v ". My upper limit, $\bar{\sigma}_v$, is my "valuation" in game theoretic terms—it is a link between this informed-decision approach and economically rational approaches.

The whole "accept an offer" apparatus is shown in Figure 4. There is no causal relationship between the four propositions, with the possible exception of the third and fourth. So to link the probabilities associated with the five propositions, the probabilities are treated as epistemic probabilities and form the nodes of a simple Bayesian net. The weights on the three arcs of the resulting Bayesian net are meaningful in practical terms. They are a subjective, high-level representation of what "comfortable" means to an agent. More important, the resulting net divides the problem of assessing "Accept" into four more convenient sub-problems.

Figure 4 shows the structure of the hybrid net for the general case—ie: including the machinery for "Good". The top half of the hybrid net combines data obtained from the Internet through the "fuzzy filters" represented by \bullet 's on that figure. There "Good" is derived as a result of scraping a set of sites $\{M_1, \dots, M_m\}$ as would most likely be the case in a 'real' electronic business application. The product of these "fuzzy filters" are treated as probabilities and attached to the bottom half of the hybrid net that is a conventional Bayesian belief network. The probability attached to "Acc" is then passed through the fuzzy filter μ_A also denoted by a \bullet on that figure. The conditionals on the Bayesian network are subjective—they are easy to specify because 12 of them are zero—for the cases in which I believe that either Fair or Me is "false". With fairly optimistic priors of 0.9 on each of the four evidence nodes, and the conditionals set to:

$P(\text{Acc} \mid \text{Suited}, \text{Good}, \text{Fair}, \text{Me}) = 1.0$, $P(\text{Acc} \mid \sim\text{Suited}, \text{Good}, \text{Fair}, \text{Me}) = 0.7$,
 $P(\text{Acc} \mid \text{Suited}, \sim\text{Good}, \text{Fair}, \text{Me}) = 0.8$, $P(\text{Acc} \mid \sim\text{Suited}, \sim\text{Good}, \text{Fair}, \text{Me}) = 0.6$
the probability $P(\text{Acc}) = 0.77$, and so the membership function, μ_A , for the fuzzy set "Accept" and the confidence value χ should be defined so that $\mu_A(0.77) < \chi$. It then remains to access the information from the available sources to, hopefully, increase $P(\text{Acc})$ so that the deal is acceptable.

4. The Negotiation Process

The ideas in the previous two sections are summarised in Figure 5. My agent is a hybrid agent. Its beliefs are derived first from incoming messages in the "Information received" box, expressed in pre-determined predicates, from ω (and from "John")—these beliefs activate the deliberative machinery. They are time-stamped on arrival as they include offers from ω . Second, beliefs are also derived from information that arrives because a plan has activated a macro tool—these beliefs




Figure 3. Fuzzy Me

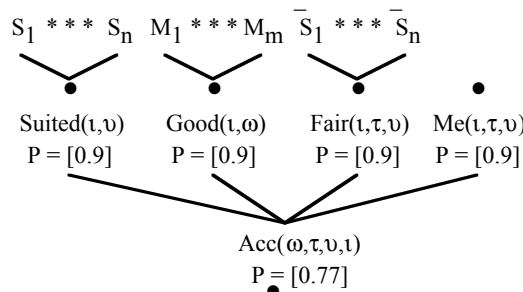


Figure 4. Hybrid net to accept an offer

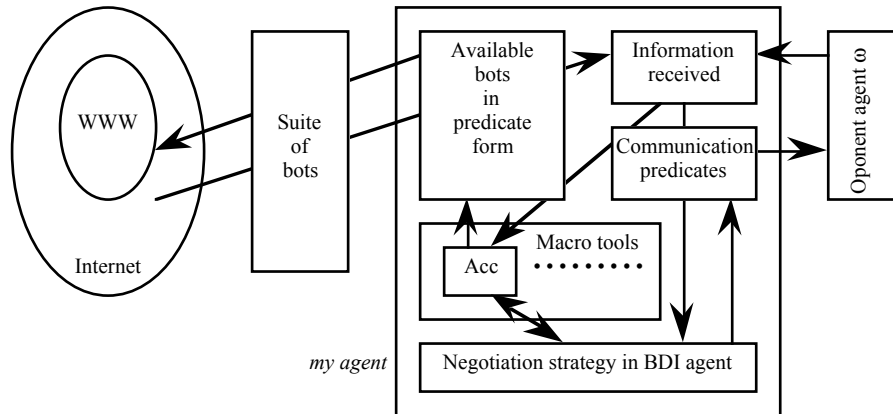


Figure 5. The negotiation framework

trigger the reactive logic. My agents' negotiation strategy—which has yet to be described—is embedded in its plans. It is guided by the lattice structure that is superimposed on each predicate. Those structures are built by hand and are intended to capture something of the meaning of “generality” for each predicate.

The background of information in which a negotiation takes place will in general be continually changing. This dynamic nature of information greatly complicates its management—dealing with belief revision is one consequence of this. For example, most people who contemplate upgrading their personal computing equipment are confronted with the problem “shall I buy now or wait to see what becomes available next month”. In the markets for stocks and bonds the background information is continually changing, and only experienced traders are able to detect deep patterns in the information shift on which to build their trading strategies. To simplify the discussion here we assume that the background information remains constant during the period of each negotiation. Specifically we assume that any information derived from information sources during a negotiation remains valid for the period of the negotiation.

A negotiation here is a sequence of offers tabled together with attendant information. The offers in this sequence may originate from the two agents in turn—the alternating offers model—or the agents may table offers at random. As time progresses during a negotiation my agent will have access to an increasing amount of information—shown in the lower half of Figure 6. The source of that information will be the information bots, information tabled by my opponent ω , and the deals themselves. Of this information, only the terms offered by my agent, τ , are certain. The validity of the information extracted by the bots will be uncertain, as will the information tabled by ω , and even ω 's terms, ν , may be uncertain in that if I accept an offer, (τ, ν) , ω may not meet her commitments—she may “renege” on the deal.




Figure 6. The negotiation process

The μ_A value of acceptable offers will be above my confidence level χ —five offers are shown in the top half of Figure 6.

The negotiation process is in two stages—first, the exchange of offers and information, and, second, the end-game. As the information is assumed to be static, there will come a time during the exchange of offers and information when the agents believe that have all the information that they require. They are then “fully informed”, and should be in a position to take a final position. The development of the offer sequence, the construction of counter-offers [16], and the tabling of information are all part of the negotiation game. This can be a game of bluff and counter-bluff in which an agent may not even intend to close the deal if one should be reached. Strategies for playing the “negotiation game” in an information rich environment are the subject of on-going research. The aim here is to present the machinery that those strategies will have at their disposal.

4.1 Stage 1: Offer, counter-offer, re-offer and information exchange

The current version of the exemplar application has two pre-defined bot predicates in addition to the “Cheapest” predicate (described above):

Features($c:C, P_j: \{P_i\}; f_1:F_1, \dots, f_n:F_n, r:*s, p:\$$)

Cameras($P_j: \{P_i\}, f_1:F_1, \dots, f_n:F_n, r:*s, p:\$; \{c_k:C\}$)

where the meaning of Features and Cameras is self-evident—in Features “ $r:*s$ ” is a rating in $*s$ (typically one to five). The lattice structure for Cameras is defined similarly to that for Cheapest. The structure for Features is null—it could be given a trivial structure by including “ dc ” in C but this serves no useful purpose.

There are four pre-defined interaction predicate templates:

Message(“John”, “me”, IWant($f_1:F_1, \dots, f_n:F_n; c:C$))

Message($a_1:A, a_2:A, Offer(\tau: \{ \$, C \}, v: \{ \$, C \})$)

Message($a_1:A, a_2:A, Accept \tau: \{ \$, C \}, v: \{ \$, C \})$)

Message($a_1:A, a_2:A, Unacceptable((\tau: \{ \$, C \}, v: \{ \$, C \}), <reason>$)

where the meaning of IWant is self-evident, “ A ” is the domain of agent names (eg: me or ω), and $<reason>$ is derived here by identifying the evidence node with lowest probability in the Accept Bayesian net. A “Message” is from agent a_1 to agent a_2 and here contains either an “Offer”, an “Accept”, or a reason for an offer being “Unacceptable”. At present my agent does not understand received “Unacceptable” messages but it sends them if wishes to do so. It should not do so as a matter of course as that would then give ω a sure way to discover my threshold. My agent is presently unable to receive tabled information from ω .

If an agent’s confidence in accepting an offer is close to her threshold then she may decide to seek further information—either from the bots or from ω —in the hope that she will be able to accept it. An issue here is what information to request, particularly when costs are involved.

In ‘real’ negotiation it is common for an agent to table information in an attempt to encourage the opposition to see her point of view: “I really wanted a blue one”. Advising an agent that its terms are not quite what I wanted is not difficult, but we can do more subtle things than that. In the exemplar system, when assessing an offer my agent may ask its information bots “to find if there is a camera whose features nearly match those of the one on offer but which sells new for less than the offer price”. If such a camera exists then it is used as fodder for the Unacceptable predicate.

In the exemplar application my agent's offer strategy is to make a, possibly empty, sequence of counter offers. The price in the initial counter offer, if any, is the greatest such that both $P(\text{Fair}) = P(\text{Me}) = 1.0$ and the confidence in $\text{Accept} > \chi$. Then the subsequent prices are chosen so that the confidence in Accept tends asymptotically to χ so as to more or less reach χ in 5 steps. This sequence will be empty if $P(\text{Suited})$ and $P(\text{Good})$ are such that confidence in Accept does not exceed χ for any positive price. If ω were to discover this simple strategy then it should take advantage of it. This is not seen as a weakness in these experiments that aim primarily at trialing the information management machinery.

More important, it is proposed that more complex applications may be supported in the same way as long as the information in the application is expressed in first-order predicates with lattice orderings for each predicate that capture "generalisation". For example, the same machinery could in principle handle a negotiation to purchase a house where features could include the number of bedrooms, the size of the house, the "aspect" (sunny or otherwise), the extent to which there are views from the property, the prime construction method, number of car spaces and so on.

4.2 Stage 2: Fully informed agents and the end-game

If my agent believes that it has all the information it requires to make an informed decision—with the exception of the position of the opposing agent—and no matter what the opposing agent's position is it will make no difference to my agent's position then my agent is *fully informed*. An agent may believe that it is fully informed prior to any negotiation taking place, or it may come to believe that it is fully informed during the exchange of offers and information.

If both agents in a negotiation believe that they are fully informed then they should be prepared to enter a final, single round of the negotiation process that determines a final, definitive outcome of either "no deal" or "a deal will be conducted under these terms..." [17].

If:

$$\{ (\tau, \nu) \mid P(\text{Acc}_{\text{me}}(\omega, \tau, \nu, \tau_{\text{me}})) \in \text{Accept}_{\text{me}} \} \cap \\ \{ (\tau, \nu) \mid P(\text{Acc}_{\omega}(\text{me}, \nu, \tau, \tau_{\omega})) \in \text{Accept}_{\omega} \} = \text{PosDeals}$$

is non-empty then a deal is possible. There will typically be a large or unbounded number of such possible deals. The structure of the admissible deals in this intersection may be quite complex if the negotiation involves multiple attributes. The comfort level of any acceptable deal will be no less than my level of confidence χ_{me} , and clearly no greater than 1.0. It will be less than 1.0 as the deal must also be acceptable to my opponent ω . The best confidence level that I may hope for, β_{me} , for will be determined by ω 's confidence threshold χ_{ω} , and *vice versa*. This is illustrated in Figure 7. The $[\chi, \beta]$ intervals are vaguely related to the idea of players having different "types" in the game-theoretic analysis of bargaining [9].

There are no ready answers to the following questions. Could some analogue of the Nash Bargaining Solution of 1950 [9] be used here? In so far as the Nash solution yields a fair distribution of the surplus utility, perhaps the basis for a final definitive outcome could be based on a fair distribution of "comfort"? Perhaps some mechanism that split each agent's range of feasible confidence in fair proportions would satisfy an analogue of the Myerson-Satterthwaite [1983] result [19] for the linear equilibrium in "split the difference" bargaining? The linear equilibrium in the

raw form of “split the difference” bargaining yields the highest expected payoff of any mechanism *ex ante*, but it is not truth-revealing—there are associated truth-revealing mechanisms—this mechanism is not optimal *ex post*. These questions require some serious analysis, and have yet to be addressed. Two possible ways in which a final, single round of negotiation could be managed are now described.

First, two fully informed agents table their final, proposed deals simultaneously. If either the terms that I propose are unsatisfactory to ω , or *vice versa*, then there is no deal. Otherwise a trusted intermediary sets about constructing a compromise deal based on the agents’ two proposals. If the difference between the deals concerns one continuous attribute—such as an amount of money—then the intermediary may simply “split the difference”—although this will not lead to truth-revealing behaviour by the agents. If there are multi-attributes, or attributes with discrete values then this intermediation will be more complex, although the second method following may be employed.

Second, two fully informed agents make their whole reasoning apparatus available to a trusted intermediary who determines whether the set PosDeals is empty. If it is not then the intermediary attempts to find a deal that divides the two agent’s $[\beta, \chi]$ intervals into equal proportions $\pi:1$. If this proves impossible, as may be the case if one of the deal attributes is non-continuous, then a deal is selected at random from the pair of deals that most closely match with one deal’s μ_A value above and one below the $\pi:1$ division.




Figure 7.

Expected confidence range

5. Conclusion

The eNegotiation framework is the first step in the design of agents that can exploit information-rich environments in one-to-one negotiation. Such agents should be adept at handling both the exchange of information and the exchange of offers. The negotiation protocol used is a loose exchange of offers and information in an initial stage, followed by an end-game mechanism when the agents believe that they have become fully informed. This protocol assumes that sufficient information is available and that the currency of the information is sustained throughout the negotiation process. Semantic issues [18] are avoided by introducing pre-defined, pre-agreed predicates. The information is managed by representing it in these predicates and then by overlaying a lattice structure with isotone valuations on each predicate. It is proposed that this method may be applied to a wide range of negotiation situations, and if so then general negotiation strategies may be developed to manage both the offers and the information. To date this work has raised more problems than it has solved. Initial indications are that it appears to work and so it is expected to lead to the design of sound negotiation agents that strive to make informed decisions rather than economically rational decisions.

References

- [1] Martin Bichler. The Future of E-Markets: Multi Dimensional Market Mechanisms. Cambridge University Press, 2001.

- [2] Felix Brandt and Gerhard Weiß. Antisocial Agents and Vickrey Auctions. in Intelligent Agents VIII proceedings ATAL 2001. Springer Verlag, 2002.
- [3] Watkins, M. Breakthrough Business Negotiation—A Toolbox for Managers. Jossey-Bass, 2002.
- [4] S.S. Fatima, M. Wooldridge and N. R. Jennings. The Influence of Information on Negotiation Equilibrium. In: Agent Mediated Electronic Commerce IV, Springer-Verlag, 2002.
- [5] Howard Raiffa. Negotiation Analysis: The Science and Art of Collaborative Decision Making. Harvard, 2002.
- [6] Rob Saunders. Curious Design Agents and Artificial Creativity: A Synthetic Approach to the Study of Creative Behaviour. Ph.D. Dissertation. University of Sydney, Department of Architectural and Design Science, 2001.
- [7] Sarit Kraus. Strategic Negotiation in Multiagent Environments. MIT Press, 2001.
- [8] Muthoo, A. Bargaining Theory with Applications. Cambridge UP, 1999.
- [9] Martin J. Osborne and Ariel Rubinstein. Bargaining and Markets. Academic Press, 1990.
- [10] Elmer Wolfstetter. Topics in Microeconomics. Cambridge UP, 1999.
- [11] Hilderman, R.J. and Hamilton, H.J. "Evaluation of Interestingness Measures for Ranking Discovered Knowledge." In Cheung, D., Williams, G.J., and Li, Q. (eds.), Proceedings of the 5th Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD'01), Hong Kong, April, 2001, pp. 247-259.
- [12] Gerding, E.H., van Bragt, D.D.B. and La Poutre, J.A. Multi-issue negotiation processes by evolutionary simulation: validation and social extensions. In proceedings Workshop on Complex Behavior in Economics. Aix-en-Provence, France, May 4-6, 2000.
- [13] Anthony Chavez, Daniel Dreilinger, Robert Guttman and Pattie Maes. A Real-Life Experiment in Creating an Agent Marketplace. in proceedings of the Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM97). London, UK, April 1997.
- [14] Birgit Burmeister, Tobias Ihde, Thomas Kittisteiner, Benny Moldavanu and Jörg Nikutta.. A Practical Approach to Multi-Attribute Auctions. in proceedings Third International Workshop on Negotiations in Electronic Markets, Aix-en-Provence, France, 2002.
- [15] S. Ramchurn, N. Jennings, C. Sierra and L. Godo. "A Computational Trust Model for Multi-Agent Interactions based on Confidence and Reputation". in proceedings workshop on "Proc. 5th Int. Workshop on Deception, Fraud and Trust in Agent Societies" AAMAS-2003, Melbourne, Australia, July 2003.
- [16] Peyman Faratin, Carles Sierra and Nick Jennings. Using Similarity Criteria to Make Issue Trade-Offs in Automated Negotiation. Journal of Artificial Intelligence. 142 (2) 205-237, 2003.
- [17] Jeffrey S. Rosenschein and Gilad Zlotkin. Rules of Encounter. MIT Press, 1998.
- [18] Steven Wilmott, Monique Calisti and Emma Rollon. Challenges in Large-Scale Open Agent Mediated Economies. In: Agent Mediated Electronic Commerce IV, Springer-Verlag, 2002.
- [19] R. Myerson and M. Satterthwaite. Efficient Mechanisms for Bilateral Trading. Journal of Economic Theory, 29, 1–21, April 1983.