Modeling Mental States in Agent Negotiation

Walid S. Saba

Knowledge & Language Engineering Group
School of Computer Science
University of Windsor
401 Sunset Ave., Windsor, ON N9B-3P4 Canada
saba@cs.uwindsor.ca
http://santana.cs.uwindsor.ca

Abstract

We describe a mental state model for agents negotiating in a virtual marketplace. Buying and selling agents enter the marketplace with an 'attitude' formulated as a complex function of prior experience(s), market conditions, product information as well as personal characteristics such as importance of time *vs.* importance of price and the commitment level to the purchase/sale of the product. While buying and selling agents are generally assumed to have opposing interests, the mental state model can be extended to other types of situations where agents are deliberating in a cooperative manner.

Introduction

We are in the process of building DALIA, an environment for distributed, artificial and linguistically competent intelligent agents. While there is no consensus on what exactly is an *intelligent agent*, there clearly is a core meaning that the agent community agrees upon, which, at a minimum, includes the following: (*i*) an agent is an autonomous module/system that is expected to be an expert at performing a specific task; (*ii*) agents are situated, i.e., they operate in a *dynamic* environment of which they must be aware; and (*iii*) agents are expected to be capable of performing commonsense reasoning and to exhibit flexible problem solving behavior. Other characteristics include learning, mobility and communication (see Bradshaw (1997) and Jennings and Wooldridge (1998)).

In our view, intelligent agents must also have a certain level of linguistic competency and must be able to perform commonsense reasoning in a highly dynamic and uncertain environment. There are clearly several challenges in formalizing this type of commonsense reasoning where various temporal and modal aspects must be taken into consideration in a highly dynamic and uncertain environment. Our long-term objective is motivated by the following scenario, involving a certain buyer B:

$$
\begin{aligned}
&((\textit{very likely PC prices will keep going down for a while})\\
&\wedge (B \textit{ can wait for another few months to buy a PC}))\\
&\quad \supset (\neg (B \textit{ stumbles on a very good deal}) \supset\\
&\qquad\qquad B \textit{ should probably wait a few months to buy a PC})
\end{aligned}
$$

We will not be concerned with this aspect of the work here[1]. Instead, our focus in this paper will be on the modeling of an agent's mental state (attitude) in agent negotiations. As autonomous agents might potentially be engaged in a number of dialogue types, such as (competitive) negotiation, inquiry, persuasion, or deliberation, to name a few (see Walton and Krabbe, 1995), the nature as well as the role of mental states might differ substantially in different types of dialogues. With respect to the agents' mental states, different types of dialogues can arguably be broadly classified into two types: *competitive* and *cooperative* dialogues. In a cooperative setting (such as in a deliberation dialogue), agents typically *cooperate* to reach a consensus on some course of action. In this framework there is no initial fixed goal and the object of the (cooperative) negotiation is to reach such a common goal. In a competitive setting (such as negotiation between buying and selling agents), however, agents are assumed to have *competing interests*. In the former case, therefore, there are no hidden goals, while in the latter agents must indeed have a hidden agenda.

Cooperative vs. Competitive Dialogues

In a competitive environment (e.g., negotiation between buying and selling agents), participants have a fixed goal and a well-defined utility function that participants use to measure their progress towards achieving their goal. In negotiating for a certain resource, for example, agents typically compute a function that measures the utility of one or more attributes (e.g., cost). These attributes are typically

[1] See (Saba and Corriveau, 1997; Saba and Corriveau, 2001) for our concurrent work on reasoning and language understanding.

assumed to take on values in a partially ordered domain. In such a framework, participants have a fixed goal (say a price range), and progress (of the negotiation) is measured by how far (or how close) the participants are from their respective goals. Thus progress and termination criteria can be easily formulated in such a model.

In a deliberation (e.g., see Hitchcock *et al.* (2001)), which is a form of cooperative negotiation, participant agents have no fixed initial commitment to any potential course of action, and the outcome of the deliberation is generally a consensus on an optimal course of action. This necessarily means that the outcome of a negotiation in a cooperative environment is in principle sub-optimal from the perspective of any individual agent. As argued in (Saba, 2001*a*), given the absence of individual goals and the absence of any initial goal it is quite difficult to formulate progress and termination criteria in such a model. In other words, since individual participants have no fixed goal against which they can measure a certain proposal, it is not clear how progress can be measure, and, consequently, it is not all that clear how termination can be guaranteed. We argue however that even in a cooperative setting, such as a deliberation dialogue, participants must have some differing goals, which must be reflected in their individual mental states. What makes a cooperative setting different from that of a competitive setting is that in the former participants must share at least one aspect of their goal, while this constraint cannot be expected in the latter case.

In this paper we thus argue that regardless of the dialogue type agent dialogues are a function of the participant's mental states (i.e., the participant's goals and attitudes). Moreover, we argue that a single mental state model for both the competitive and the cooperative models can be developed. Assuming that a mental state (reflecting an agent's attitude) is given as an n-tuple $\langle a_1, \ldots, a_n \rangle$ where $a_1, \ldots, a_n$ are attributes from a partially ordered domain, a cooperative model is one where all the participants agree on at least one of these attributes. In the competitive setting no such constraint can be assumed.

Negotiating with an Attitude

As part of our ongoing effort on DALIA - an environment for distributed, artificial and linguistically competent intelligent agents - we have designed a virtual marketplace environment where buying and selling agents negotiate on behalf of their clients. Our negotiation model has several common features with a number of existing models (e.g., Chavez and Maes, 1996; Esmahi and Dini, 1999; Kumar and Feldman, 1998). However, in these models the notion of a mental state of an agent, which as will be argued below plays an important role in the negotiation process, has received little attention. Specifically, we argue that an agent's mental state defines an overall *context of negotiation*, which consequently defines an agent's strategy and ultimately determines how a negotiation proceeds. In the current model an agent's mental state is taken to be an agent's level of commitment, an agent's prior experience, and the degree to which time and price are important to the buyer and/or seller. It must be noted here that in exploring the effect an agent's mental state has on the overall negotiation process we were not primarily concerned with how comprehensive our mental state model is, but in building the framework of a negotiation model where prior experiences and the current mental state of an agent play a important role. Finally, and although some form of learning from experience has been previously suggested (e.g., Wong *et al.*, 2000), the interaction between an agent's mental state and an agent's prior experience in formulating *new* negotiation strategies is novel in our model.

Basic Components of the Model

The virtual marketplace we have developed is an environment that at any point in time has a list of buyers and sellers. Buyers and sellers in the environment are also assumed to have access to two knowledge sources: (*i*) an ontology for domain-specific product information; and (*ii*) a set of general commonsense rules. Agents are assumed to learn from experience using case-based reasoning (CBR) techniques: negotiations along with product information, market conditions and the outcome of the negotiation are stored as experiences for use in future negotiations[1]. A high-level view of the model is shown in figure 1 below.

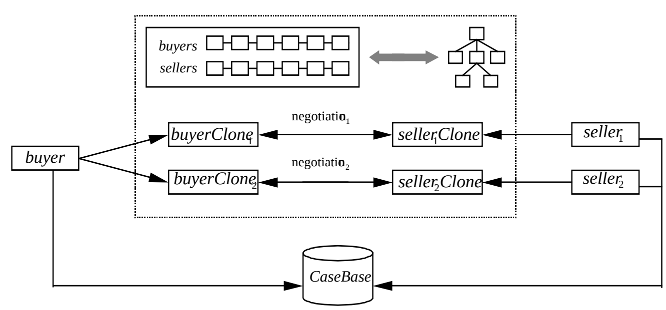

Figure 1. Basic components of the virtual marketplace.

[1] Market conditions, which in the current model are simply the ratio of supply to demand, are clearly very simplistic. A more realistic model must eventually assume a certain probability distribution of buyers and sellers and must incorporate profit projections based on these distributions. Kephart and Greenwald (1999) is a good example of such a recent study.

The process starts when clients (users) create buyer and seller agents that are sent to the virtual marketplace. Buyers and sellers are registered in the marketplace where a list of buyers and sellers is maintained. In the current model it is buyers that are assumed to be proactive, i.e., it is buyers that look for and initiate a negotiation with sellers.

Here's an overview of the process from a buyer's perspective:

- A buyer *b* is created by some client.
- *b* enters the marketplace (*b* is added to list of buyers)
- *b* retrieves a publicly available price range for the product in question
- *b* computes its own price range as a function of its attitude and the publicly available price
- *b* retrieves a list of sellers *S*, selling the same product.
- *b* sends an (asynchronous) message to each seller $s \in S$ requesting a negotiation
- Sellers return negotiation clones or decline to negotiate
- For each seller clone *sc*, *b* creates a buyer clone *bc*
- A negotiation starts between each pair of *sc* and *bc*
- Buyers start bidding with the minimum of their price range, while sellers start bidding with their maximum (price ranges of agents are hidden)
- A deal is made when the buyer's maximum reaches the seller's minimum
- No deal is made if the buyer's maximum falls short of the seller's minimum
- Agents might potentially save the experience for future use

We explore this process in some detail below. First some definitions are in order.

Definition 2.1 *An Agent's Attitude* (*aat*) is a hidden mental state comprised of a triple representing the importance of time, the importance of price and the commitment level of an agent, respectively: $aat = \langle x_1, x_2, x_3 \rangle$, where $x_i \in [0,1]$ [1]. For an agent with *attitude* = ⟨1.0,0.5,0.8⟩, for example, time is a priority, and the commitment level is rather high although it is somewhat indifferent to the price (an ideal buyer attitude from the perspective of a seller).

If we take the absolute Euclidean distance as an equivalence operator, as done in (Faratin *et al.*, 2000), we can compute a measure of similarity between two attitudes, $0 \le AS(attitude_1, attitude_2) \le 1$, as:

[1] Of course, one can imagine other attributes that could also be added to the mental state.

$$AS(\langle x_1, x_2, x_3 \rangle, \langle y_1, y_2, y_3 \rangle)$$
$$= (1 - |x_1 - y_1|) \wedge (1 - |x_2 - y_2|) \wedge (1 - |x_3 - y_3|)$$

Although the two t-norm functions commonly used for conjunction (*minimum* and *product*) seem equally plausible here, we have (admittedly) arbitrarily chosen to use *product* in the current model.

Definition 2.2 *Public Price Range*: All agents in the marketplace are assumed to have access to a publicly available price range that can be obtained from a product-specific ontology: $PPR(prod) = [pmin, pmax]$. A measure of similarity between two public price ranges $0 \le RS(range_1, range_2) \le 1$ is computed as follows:

$$RS([min_1, max_1], [min_2, max_2])$$
$$= \left(1 - \frac{|min_1 - min_2|}{min_1 + min_2}\right) \wedge \left(1 - \frac{|max_1 - max_2|}{max_1 + max_2}\right)$$

Definition 2.3 *A Negotiation* is a process that can be in any of the following states: (*i*) $\texttt{DONE}^+$; (*ii*) $\texttt{DONE}^-$; or (*iii*) $\texttt{DONE}^0$, corresponding respectively to a negotiation that completed successfully (deal was made), a negotiation that completed unsuccessfully (no deal is possible), and to a negotiation that is still in progress. A negotiation *record* is an (ordered) list *L* of offers and counter offers, and an outcome: $L = \langle \{\langle Offer, CounterOffer \rangle\}, Outcome \rangle$; where $Outcome \in \{\texttt{DONE}^+, \texttt{DONE}^-\}$. In addition to the outcome, an important feature of a negotiation is the time it took, which could be crudely approximated as $length(L)$, i.e., as the length of the negotiation record.

The specific offers and counter-offers are also important features of a negotiation record. Since the length of the negotiation has already been taken into account, one only needs to consider the average offer and counter-offer of a negotiation. That is,

$$avgOffer(< L, outcome >) = \left(\frac{1}{length(L)}\right)\left(\sum_{\langle offer, counterOffer \rangle \in L} offer\right)$$
$$avgCounterOffer(< L, outcome >)$$
$$= \left(\frac{1}{length(L)}\right)\left(\sum_{\langle offer, counterOffer \rangle \in L} counterOffer\right)$$

Taken the average offer and counter offer of a negotiation, the time a negotiation takes, as well as the negotiation outcome, a similarity measure between two negotiations $0 \le NS(\langle L_1, Outcome_1 \rangle, \langle L_2, Outcome_2 \rangle) \le 1$ can now be defined as follows:

$$NS(\langle L_1, outcome_1 \rangle, \langle L_2, outcome_2 \rangle)$$
$$= \begin{cases} \text{MAX}(0, Sim(L_1, L_2) - \varepsilon) & \text{if } outcome_1 \ne outcome_2 \\ Sim(L_1, L_2) & otherwise \end{cases}$$

where ε is a bias against the difference in the outcome and

$$Sim(L_1, L_2) = \left(1 - \frac{|length(L_1) - length(L_2)|}{length(L_1) + length(L_2)}\right) \wedge \left(1 - |avgOffer(L_1) - avgOffer(L_2)|\right) \wedge \left(1 - |avgCounterOffer(L_1) - avgCounterOffer(L_2)|\right)$$

Note that since there is always at least one offer and counter offer $length(L_1) + length(L_2) \neq 0$.

Definition 2.4 A new *Agent Experience* results after every negotiation. In addition to the negotiation record, an experience record contains information about the agent's current (mental state) attitude, the public price range, the agent's price range, as well as market conditions. This is an example experience:

```
ProductCategory    ⇒  PersonalComputer
PublicPriceRange     ⇒  [1000,3000]
AgentPriceRange    ⇒  [1000,2000]
Attitude           ⇒  [1.0,0.5,0.8]
SupplyDemandRatio  ⇒  ⟨3,1⟩
Negotiation        ⇒
  ⟨{⟨1000,1500⟩,⟨1100,1400⟩,⟨1200,1200⟩},DONE+⟩
```

This represents an agent's experience in buying a personal computer, when the supply-to-demand ratio was 3 to 1, the agent was highly committed to buying, the price was not much of a factor, but time was crucial. Under those circumstances, the negotiation was successfully completed after three offers and counter offers.

Definition 2.5 An *Agent's Price Range* is a function of its attitude (definition 2.1) and the public price range (definition 2.2). There are several functions that are plausible here. However note that the lower the minimum price range of a buyer the more time a negotiation will take. Thus, the buyer's minimum price range must be higher than the public's minimum if the importance of time is high for that buyer. Moreover, when the commitment of a buyer is high the agent should be as close to the public's maximum as possible, since a high commitment indicates an urgent need for the product. However, the importance of the price should balance the degree of commitment (in a sense, a high commitment is a form of desperation, and the importance of price provides a sobering effect). A seller reasons in exactly the opposite direction. The following two functions are one suggestion on how such reasoning might plausibly be captured, where $\langle m,n,k\rangle$ represents an agent's attitude[1]:

$$APR^{seller}([pmin, pmax], \langle m, n, k\rangle) = [pmin + (pmin)(n)(1-k)/10,\ pmax - (m)(pmax - pmin)/10]$$

$$APR^{buyer}([pmin, pmax], \langle m, n, k\rangle) = [pmin + (m)(pmax - pmin)/10,\ pmax - (pmax)(n)(1-k)/10]$$

Definition 2.6 *Supply/Demand Ratio*: Part of an agent's experience is the relation of the price to the supply/demand ratio, which is simply $|S|/|B|$ where S and B are the current list of sellers and buyers of a certain product, respectively.

[1] A similar strategy has been suggested in (Mudgal and Vassileva, 2000), although only the price was considered a factor in formulating a bidding strategy.

The Negotiation Process

A buyer b and a seller s enter the marketplace with attitudes $\langle bm,bn,bk\rangle$ and $\langle sm,sn,sk\rangle$, respectively. Consequently,

- b computes its price range: $[bp_{max}, bp_{min}] \leftarrow APR^{buyer}([pmin,pmax],\langle bm,bn,bk\rangle)$
- s computes its price range: $[sp_{max}, sp_{min}] \leftarrow APR^{seller}([pmin,pmax],\langle sm,sn,sk\rangle)$
- b hides its maximum, bp_{max}, and starts its bidding with $bbid \leftarrow bp_{min}$
- s hides its minimum, sp_{min}, and starts its bidding with $sbid \leftarrow sp_{max}$
- With each successive offer and counter offer buyers and sellers update their respective biddings as follows: $bbid \leftarrow bbid + \alpha$ and $sbid \leftarrow sbid - \beta$ where α and β are the buyer's step increment and the seller's step decrement, respectively. In the current model α and β are initially $(pmax - pmin)/100$
- A negotiation is always in one of these states: (*i*) $\texttt{DONE}^{+}$ if $(bbid \geq sbid)$; (*ii*) $\texttt{DONE}^{-}$ if $(bp_{max} < sp_{min})$; and (*iii*) $\texttt{DONE}^{0}$ if neither of the previous is true; or, equivalently, if $((bbid < sbid) \wedge (bp_{max} \geq sp_{min}))$

It must also be noted that in general a buyer has more than one negotiation thread running concurrently (one thread with every potential seller). The buyer waits for each (negotiation) thread to return a result $r \in \{(\texttt{DONE}^{+}, price), (\texttt{DONE}^{-}, price)\}$. If time is more important than price, the buyer exits the marketplace (terminating all its clones) as soon as a result $r = (\texttt{DONE}^{+}, price)$ is received, otherwise, the buyer waits for all (negotiation) threads to terminate and selects the one that found the best deal (if any).

The Ontology

Agents in the virtual marketplace have access to a domain-specific ontology of product information. Currently, the ontology and the domain knowledge are quite limited, and the main functionality of the ontology is to provide a public price range for specific product categories. Using the notion of semantic distance in a semantic network (Sussna, 1993), we also compute a simple measure of conceptual similarity between two products as follows[1]:

$$
\begin{aligned}
&CS(prod_1, prod_2) = \frac{1}{d_1 + d_2} \\
&where \\
&l = \mathrm{LUB}(prod_1, prod_2) \\
&d_1 = dist(prod_1, l) \\
&d_2 = dist(prod_2, l)
\end{aligned}
$$

where LUB is the least upper bound of two concepts in the ontology, and the distance between two concepts, $dist(c_1,c_2)$, is the number of 'isa' links from c_1 to c_2. (note that $dist(c_1,c_2) \neq 0$).

Computing a conceptual similarity between two products based solely on the product category is not sufficient for our purposes. The reason for this is the following: when buying a scanner, one might recall their experience in buying a printer. In this case the conceptual similarity of the product categories seems to be sufficient. However, this is a very simplistic view, since one would hardly recall their experience in buying a (computer) mouse when one is buying a (computer) monitor, although both are "computer products". Clearly, the price range is also crucial. That is, our experience in buying big items with similar price ranges might be similar even though the product categories might be different. The similarity between two products pd_1 (with price pr_1) and pd_2 (with price pr_2) can be computed as follows:

$$
PS(pr_1, pr_2, pd_1, pd_2) = CS(pd_1, pd_2) \wedge \left(1 - \frac{|pr_1 - pr_2|}{pr_1 + pr_2}\right)
$$

Learning as Improving One's Attitude

Regardless of the outcome of a negotiation, both buyers and sellers occasionally save their experience in a *case base* for future use. When dealing with a case base one has to carefully consider a strategy for (*i*) case representation; (*ii*) case indexing and retrieval; (*iii*) case matching; and (*iv*) case adaptation (see [10]). We consider these very briefly here. A case (experience) in our model has the following structure:

```
〈 ProductCategory: prod,        (e.g., PersonalComputer)
  ProductName: pname,           (e.g., Intell PIII)
  PublicPriceRange: ppr,        [pmin,pmax]
  AgentPriceRange: apr,         [bmin,bmax]
  Attitude: att,                (e.g., 〈1.0,0.5,0.8〉)
  SupplyDemandRatio: sdr,       (e.g., 〈3,1〉)
  Negotiation: neg
  (〈{〈1000,1500〉,〈1100,1400〉,〈1200,1200〉},DONE+〉)〉
```

[1] To support recommender agents (Glover *et al.*, 1999), the ontology must be considerably extended. Saba (2001*b*) reports on such work.

Cases are indexed in the case base by the product category and the average price, computed as $(pmin+pmax)/2$. When searching for 'relevant' cases (or experiences), a perfect match cannot be expected, instead the search is conducted as follows: two lists of all cases, corresponding to failed and successful experiences, are generated. Cases included in these lists are those that match the search criteria within a certain threshold. When searching for a relevant experience, cases are matched as follows:

$$
\begin{aligned}
&Match(c_1, c_2) \\
&= \frac{1}{3}\begin{pmatrix} PS(prod(c_1), prod(c_2)) \\ + AS(att(c_1), att(c_2)) + RS(ppr(c_1), ppr(c_2)) \end{pmatrix}
\end{aligned}
$$

Agents learn from experience in our model by using prior experiences to adjust their attitude and the increment they make during negotiation (recall α and β from section 2.2) prior to computing their $[bmin, bmax]$ price range. Since the attitude, its corresponding price range and the bidding increment affect the entire negotiation process, agents will overtime tend to minimize the negotiation time and the maximum bidding price. Moreover, the proportion of successful negotiations will increase over time since agents also learn from negative experiences (those that completed unsuccessfully). Using a matching experience, an agent's attitude and bidding increment are updated as follows:

$$
\begin{aligned}
&Adjust(\langle prod, pname, ppr, apr, aat, sdr, neg\rangle, \langle \alpha, aat\rangle) \\
&= \begin{cases} \langle \alpha', aat'\rangle \leftarrow f_{inc}(\langle \alpha, aat\rangle) & \text{if } outcome(neg) = \mathrm{DONE}^- \\ \langle \alpha', aat'\rangle \leftarrow f_{dec}(\langle \alpha, aat\rangle) & \text{if } outcome(neg) = \mathrm{DONE}^+ \end{cases}
\end{aligned}
$$

The functions f_{inc} and f_{idec} update the bidding increment and the attitude based on previous results as follows: f_{idec} hardens the bidding increment and the attitude (for every occurrence of a previous success), while f_{inc} loosens the bidding increment and the attitude (for every occurrence of a previous failure). The reasoning behind this process is to let agents find an optimal attitude/bidding-practice that maximizes the number of successes. The exact threshold by which the attitudes and the bidding increments (decrements) are updated based on previous successes (failures) is still being investigated, although one plausible approach is to start with a percentage of the initial values.

A final point should be made about case updates. When a negotiation terminates, a search and match is done on the case base. When searching for updating or adding a new experience a match between two cases is done as follows:

$$Match(c_1, c_2) = \frac{1}{4}\begin{pmatrix} w_1 \times PS(prod(c_1), prod(c_2)) + w_2 \times AS(att(c_1), att(c_2)) + \\ w_3 \times RS(ppr(c_1), ppr(c_2)) + w_4 \times NS(neg(c_1), neg(c_2)) \end{pmatrix}$$

Currently we assign equal weights, w_i, to all attributes of a case, although we plan to test various weighting schemes, perhaps using a machine learning experiment. If a strong match is not found, the new case represents a 'novel' experience and is added to the case base. When a strong match occurs, the two cases are 'merged' resulting in a modification of an existing experience:

$$Merge\begin{pmatrix} \langle p_1, pn_1, ppr_1, apr_1, att_1, sdr_1, neg_1 \rangle, \\ \langle p_2, pn_2, ppr_2, apr_2, att_2, sdr_2, neg_2 \rangle \end{pmatrix} = \left\langle \begin{matrix} \text{LUB}(p_1, p_2), (pn_1 = pn_2), (ppr_1 = ppr_2), avg(apr_1, apr_2), \\ avg(att_1, att_2), avg(sdr_1, sdr_2), \text{MIN}(neg_1, neg_2) \end{matrix} \right\rangle$$

An Interim Recap

We have briefly described a virtual marketplace where buying and selling agents that learn from experience autonomously negotiate on behalf of their clients. In such a competitive setting agents are assumed to have a hidden mental state, which agents tend to optimize over time using their (private) prior experiences[1]. As we mentioned in section 1, this model could be easily extended to situations where agent negotiation is cooperative rather than competitive. We discuss this extension below.

Mental States in Cooperative Settings

Unlike the competitive negotiation model described above, agents in a cooperative setting (such as those engaged in a deliberation) must share at least one attribute value from the attributes of their mental state.

One can imagine various scenarios where agents that have slightly different goals might nevertheless negotiate in a cooperative manner. Since different negotiation tasks would require a different set of relevant attributes for the mental state, we could simply re-state the previous problem of the competing buying and selling agents. We argue that if the buyers and sellers agreed on an essential attribute of the mental state 3-tuple, then the resulting negotiation would be more cooperative than competitive. Let us recall that in the simple negotiation model described in section 2 buying and selling agents were assumed to have a mental state which is triple $\langle x_1, x_2, x_3 \rangle$ representing the importance of time, the importance of price and the commitment level of an agent, respectively, where $x_i \in [0,1]$. Suppose that for some two agents b and s we have the following:

$aat(b) = \langle x_1, x_2, 1.0 \rangle$, where $x_1 \in [0,1]$ and $x_2 \in [0,1]$
$aat(s) = \langle x_1, x_2, 1.0 \rangle$, where $x_1 \in [0,1]$ and $x_2 \in [0,1]$

That is, both b and s agree on the commitment level: b is rather desperate to buy and s is rather desperate to sell. We argue that in such a situation the negotiation between b and s would be more cooperative than competitive. Now here's what the agreement on one attribute results in:

Claim. The agreement on the commitment level attribute of the buyer's and seller's attitudes makes the negotiation process more cooperative than competitive.

Proof. Recall that the buyer and seller compute their private ranges as a function of the public price range and their attitude as follows:

$$APR^{seller}([pmin, pmax], \langle m, n, k \rangle) = [pmin + (pmin)(n)(1-k)/10, pmax - (m)(pmax - pmin)/10]$$

$$APR^{buyer}([pmin, pmax], \langle m, n, k \rangle) = [pmin + (m)(pmax - pmin)/10, pmax - (pmax)(n)(1-k)/10]$$

Clearly, the higher the commitment level, the higher is the maximum of the buyer and the lower is the minimum of the seller. If this is the case, then according to the process described in section 2.2, the negotiation in such a case would most successfully converge faster than otherwise, which could be interpreted as though the agents cooperated more so than competed. Clearly, the more attributes b and s agree on the more cooperative the negotiation is. Intuitively, therefore, cooperative negotiations are those where agents tend to agree more so than disagree on their hidden goals (mental states).

From our perspective, therefore, the mental state model defines the boundaries of such negotiation tasks and it subsumes both competitive as well as cooperative dialogue types. Note also that according to the learning strategy described above, when agents agree on one or more attribute, experiences become more or less public, and consequently, the buyer's increment and the seller's decrement tend to converge. This, in turn, would further move the

---

[1] We are currently in revamping the implementation of the current prototype, and will report in various aspects of the implementation in (Saba, forthcoming). Interested readers can visit `http://sanatan.cs.uwindsor.ca` for updates and to request further information.

negotiation process from a competing to a cooperating process.

Concluding Remarks

In this paper we presented an approach to agent negotiation based on the mental state model. We argued that a mental state model subsumes both competitive as well as cooperative dialogue types. We first described a prototype of a virtual marketplace environment where buying and selling agents that learn from experience autonomously competitively negotiate on behalf of their clients in a highly dynamic and changing (uncertain) environment. It was consequently argued that a cooperative model is simply a model where agents agree more so than disagree on the relevant attributes of the mental state. The current model is still too simplistic, and to date the main focus has been on building the appropriate infrastructure that allows us to experiment with various reasoning models and strategies. Nevertheless we have incorporated a powerful learning strategy and a mental state model that seem to provide the agents with a novel problem solving behavior. Currently we are focusing on incorporating various commonsense reasoning strategies to account for the various temporal and modal aspects of a highly dynamic and changing environment. Moreover, we believe that some form of uncertainty reasoning must be incorporated into the model since market conditions and agent□s beliefs are rarely crisply defined.

Acknowledgements

This research is supported by the National Science and Engineering Research Council of Canada (NSERC) and the University of Windsor. Thanks are also due to two anonymous reviewers.

References

Bradshaw, J. (Ed.). 1997. *Software Agents*, AAAI Press/MIT Press.

Chavez, A. and Maes, P. 1996. Kasbah: An Agent Marketplace for Buying and Selling Goods, In *Proc. of the 1st Int. Conf. on the Practical Application of Intelligent Agents and Multi-Agent Technology.*

Esmahi, L. and Dini, P. 1999. Toward an Open Virtual Market Place for Mobile Agents, In *Proc. of IEEE 8th Int. Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises.*

Faratin, P., Sierra, C. and Jennings, N. R. 2000. Using Similarity Criteria to Make Negotiation Trade-Offs, In *Proc. of 4th International Conference of Multi-Agent Systems (ICMAS-2000)*, Boston, MA, pp. 119-126.

Faratin, P., Jennings, N. R., and Sierra, C. 1999. Automated Negotiation for Provisioning Virtual Private Networks using FIPA-Compliant Agents, In *Proceedings of of AAAI Workshop on Negotiation.*

Glover, E., *et. al.* 1999. Recommending Web Documents Based on User Profiles, In *Proceedings of ACM SIGIR'99 Workshop on Recommender Systems.*

Hitchcock, D., McBurney, P. and Parsons, S. 2001. A Framework for Deliberation Dialogues, In H. V. Hansen, C. W. Tindale, J. A. Blair and R. H. Johnson (Editors): *Argumentation and its Applications. Proc. of the Fourth Biennial Conference of the Ontario Society for the Study of Argumentation (OSSA 2001)*

Jennings, N. R. and Wooldridge, M. (Eds.) 1998. *Agent Technology: Foundations, Applications, and Markets*, Springer-Verlag.

Kephart, J. O., Greenwald, A. R. 1999. Shopbot Economics, *Lecture Notes in Artificial Intelligence*, Volume 1638, Springer-Verlag, pp. 208-220.

Kolonder, J. 1993. *Case-Based Reasoning*, Morgan-Kaufman Publishers.

Kumar, M. and Feldman S. I. 1998. Business Negotiation on the Internet, *IBM Institute for Advanced Commerce*, IBM T. J. Watson Research Center Report, Yorktown Heights, NY.

Mudgal, C. and Vassileva, J. 2000. An Influence Diagram Model for Multi-Agent Negotiation, In *Proceedings of 4th International Conference of Multi-Agent Systems (ICMAS-2000)*, Boston, MA, pp. 7-12.

Saba, W. S. forthcoming. A Framework for Competitive and Cooperative Agent Negotiation, to be submitted to *Knowledge and Information Systems: An International Journal.*

Saba, W. S. 2001*a*. Commentary on Hitchcock *et al.* (2001), appeared in OSSA-2001.

Saba, W. S. 2001*b.* The Ontology of Commonsense Knowledge, In *Proc. of Workshop on Ontologies (Joint German/Austrian Conference on AI, KI-2001*)

Saba, W. and Corriveau, J.-P. 2001. Plausible Reasoning and the Resolution of Quantifier Scope Ambiguities, *Studia Logica*, 67:271-289.

Saba, W. and Corriveau, J.-P. 1997. A Pragmatic Treatment of Quantification in Natural Language, In *Proceedings of the 14th National Conference on Artificial Intelligence (AAAI-97)*, pp. 610-615.

Sussna, M. 1993. Word sense disambiguation for free-text indexing using a massive semantic network, In *Proceedings of the 2nd International Conference on Information and Knowledge Management*, pp. 1-5.

Walton, D. N. and Krabbe, E. W. 1995. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning.* Albany, NY: SUNY Press.

Wong, W. Y., Zhang, D. M. and Kara-Ail, M. 2000. Negotiating with Experience, In *Proceedings of the AAAI-2000 Workshop on Knowledge-Based Electronic Markets - KBEM'00.*