

# A Decentralized Model for Automated Multi-attribute Negotiations with Incomplete Information and General Utility Functions

Guoming Lai<sup>†</sup> • Katia Sycara<sup>‡</sup> • Cuihong Li<sup>#</sup>

<sup>†</sup> Tepper School of Business, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA  
15213, USA Tel:412-268-4627, Fax:412-268-5569, Email: guomingl@andrew.cmu.edu

<sup>‡</sup>Robotics Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA

<sup>#</sup>School of Business, University of Connecticut, 2100 Hillside Road, Storrs, CT 06269, USA

## Abstract

This paper presents a decentralized model that allows self-interested agents to reach “*win-win*” agreements in a multi-attribute negotiation. The model is based on an alternating-offer protocol. In each period, the proposing agent is allowed to make a limited number of offers. The responding agent can select the best out of these offers. In the case of rejection, agents exchange their roles and the negotiation proceeds to the next period. To make counteroffers, an agent first uses the heuristic of choosing the offer on an indifference (or “iso-utility”) curve/surface that is closest to the best offer made by the opponent in the previous period, and then taking this offer as the seed, chooses several other offers randomly in a specified neighborhood of this seed offer. Experimental results show that this model induces agents to reach near Pareto optimal agreements in general situations where agents have complex preferences on the attributes and incomplete information. This model does not require the presence of a mediator.

Keywords: Multi-attribute negotiation, Pareto optimality, Win-win, Rational preference, Incomplete information, Self-interested agents

# 1 Introduction

Negotiation is a fundamental and effective way to reach a mutually acceptable agreement among self-interested agents. The seminal work by Nash [7] and Rubinstein [8] initiated a surge of research interest on negotiation. However, most of the existing literature focuses on single attribute negotiations. In reality, multi-attribute negotiations are very common. For example, an employer and a union usually need to simultaneously negotiate wage level, health care and vacations because those issues together determine the utility of the final contract; similarly, a supplier and a buyer usually need to negotiate the price, quality, quantity and delivery time of a procurement contract at the same time. In those situations, the failure of reaching agreement on some issues can lead to the breakdown of the whole negotiation. Besides the necessity, in many situations, people may also be willing to introduce additional issues into their negotiation because all of the parties in the negotiation may benefit from trading off the multiple issues when they have different preferences. For instance, when selling automobiles, dealers can sell the automobiles with a single price, but more often they may also introduce the financing package, insurance package and warranty package into the contract rather than a single price. It is because with some discount on those packages, which may be cheaper for the dealers than to directly lower the price, buyers are more willing to accept the automobile price. On the other hand, buyers may also find it beneficial to negotiate such a contract because the price of buying those packages individually, for example, the insurance plan, may be much higher. Thus, to negotiate multiple issues together may lead them to a “win-win” outcome which otherwise cannot be achieved by negotiating a single issue. Such situations in reality make multi-attribute negotiation important and valuable.

However, a multi-attribute negotiation can be much more difficult than a single-attribute negotiation, due to the following reasons. First, in a multi-attribute negotiation, an agent’s decision depends on her preference over all the issues, which can be very complex. For instance, the preference may need to be characterized by a non-linear utility function with the marginal utility of an issue depending on the values of other issues. Then to reason the opponent’s preference and strategy in the  $n$ -dimensional ( $n > 1$ ) space becomes computationally intractable. Second, with agents negotiating multiple issues simultaneously, it becomes a more complicated decision to make an offer, because for any particular utility level, there may exist a large number of offers in the negotiation space that coincide with that utility level. Which offer to pick becomes a difficult decision especially when an agent does not have complete information about the opponent’s preference in incomplete information environments. This problem can also be viewed as which direction in the negotiation space for an agent to follow in her concession, given the time and negotiation history. This decision, moreover, also impacts the decision how much to concede in each step—an agent can give up less utility by following a concession direction that allows the opponent to retain more utility but at a lower cost of her own utility. Third, in a multi-attribute negotiation, achieving a “win-win” solution is important. Since there exist multiple issues and agents may have different preferences on the issues, given any solution in the negotiation, both agents maybe still can be better off by trading off the values of some issues unless the given solution is already Pareto optimal. *Pareto optimality* is defined as the property that an outcome cannot be further improved (i.e. no agent can get more utility) without sacrificing the other’s utility. However, it is difficult to seek Pareto optimal settlement between self-interested agents when they do not know each other’s preference.

The above issues, however, have not been considered simultaneously in the existing literature on multi-attribute negotiations. First, most of the prior work on multi-attribute negotiations usually assumes that agents have relatively simple (linear) utility functions (e.g. [14, 19, 21]) or binary valued issues (e.g. [12, 16, 24]), which cannot represent the general situations. For example, the utility functions that are widely used in the economics field to represent consumer utility on multiple-goods consumption and also common in real-world situations are usually nonlinear, e.g., Cobb-Douglas utility function, Constant Elasticity of Substitution (CES) utility function and quadratic utility functions [1]. Second, to overcome the negotiating difficulty in the  $n$ -dimensional space, the existing work assumes agents negotiate in an issue-by-issue manner (e.g. [2, 3, 19]), agents are cooperative, or a non-biased mediator is available (e.g. [6, 12, 29]). Thus, the absent in the prior work is a protocol that can not only assist agents efficiently to make offers in an  $n$ -dimensional space with incomplete information but also give self-interested agents sufficient decision flexibility, for instance, the right to select offers to make and the right to accept a given offer. Finally, Pareto optimality is another key aspect that has usually been overlooked.

Complementing the literature, this paper presents a model that addresses the above issues simultaneously. First, the model allows self-interested agents to negotiate multiple attributes under rather mild constraints (see Section 3), and an agent’s information about the other’s utility function and negotiation strategy is incomplete. Second, the model then provides computationally tractable negotiation approaches for agents to make and respond to the offers in each period, which mitigates the difficulty of negotiating in an  $n$ -dimensional space. The model first decomposes agents’ negotiation strategies into three parts: *conceding*, *proposing* and *responding*. The first part determines how much to concede in a step if no

agreement is reached so far. The second part, given the concession generated in the first part, then decides among all those offers within that concession range which are most preferable for the opponent. The last part determines whether to accept one of the offers from the opponent in a given negotiation period. This model allows agents to make multiple (but a limit number of) offers based on a given concession in each step.<sup>1</sup> By allowing an agent to make multiple offers in one round, the model can improve the desirability of an agent's proposal to the opponent without sacrificing her own utility. Moreover, from the multiple offers previously proposed by the opponent, the current proposing agent can get some valuable clue of the opponent's preference and find more desirable offers for the opponent. Third, the model considers Pareto optimality. With the proposing and responding approaches proposed in the model, an agent explicitly seeks the most acceptable offer which is possible for her to find for the opponent in each negotiation round. Although it is intractable to find the exact Pareto optimal offer under incomplete information, this model achieves asymptotic Pareto optimality and the experimental results (in Section 4) show that near Pareto optimality can be achieved with a small number of offers necessary for agents to make in each negotiation step.

Thus, the contribution of this work includes the following aspects: first, while the existing models for multi-attribute negotiations are largely limited to simplified conditions with relatively simple utility functions, complete information or with a non-biased mediator, this work presents an approach that can be applied to more common situations; second, we pro-

---

<sup>1</sup>Such a setting, however, is not rare in applications, for instance, a seller may propose several contracting options with different unit price, delivery time and quality to a buyer, and the latter selects the best one in the menu or rejects all the offers.

vide computationally tractable methods for an agent to make and respond to offers in such complicated situations; third, this work considers Pareto optimality.

The rest of the paper is organized as follows. Section 2 reviews the related work. Section 3 presents the model. Experimental analysis is provided in Section 4. Section 5 concludes.

## 2 Related work

The research work on multi-attribute negotiations has been conducted in the fields of game theory and artificial intelligence (AI). In game theory, the simplest context studied is the one with complete information and cooperative agents. For this context, since agents know the utility functions of each other, it is not hard to compute the Pareto frontier of their negotiation. So, rational agents can reach agreement on this frontier by Nash axioms [7], Kalai-Smorodinsky solution [9], or other solutions. However, these approaches are not applicable in realistic situations due to their strict assumptions. The work considering non-cooperative agents mainly focuses on issue-by-issue negotiation, e.g. [2, 3]. Issue-by-issue negotiation arises because of bounded rationality of agents. It is difficult for an agent with limited computational capability to reason the opponent's strategy in the whole negotiation domain, thus, game theorists propose to negotiate multiple attributes sequentially (i.e. one-by-one). However, an underlying assumption of such a decomposition is that the utility functions of agents are linear additive; Pareto optimality usually cannot be maintained by the models proposed in issue-by-issue negotiation.

In the AI field, the existing work mainly focuses on automated negotiation frameworks and tractable heuristics. For instance, Fatima et al. [19, 20] propose an agenda-based

framework for multi-attribute negotiation. In this framework, agents can propose either a combined offer on multiple issues or a single offer on one issue. Different from the game theoretical research, their work focuses more on tractability. They assume agents adopt time-dependent strategies and can make decisions on the issues independently faced with a combined offer. For example, if there are two issues in a combined offer, say  $x_1$  and  $x_2$ , an agent has two independent strategies  $S_1$  and  $S_2$ , which are used to decide whether to accept  $x_1$  and  $x_2$ . But if one issue is settled, then this issue cannot be negotiated any longer and agents just focus on the remaining issue. Based on such a setting, they show the optimal agendas in different scenarios. Fatima et al. [21] also examine the optimal agendas where agents are allowed to negotiate a package deal or take parallel issue-by-issue negotiation. But their work requires the assumption that agents have linear additive utility functions. Sycara [22, 23] presents a case-based reasoning approach where the automated negotiating agents make offers based on similarity of the multi-attribute negotiation context to previous negotiations. Moreover, the author also uses automatically generated persuasive argumentation as a mechanism for altering the utilities of agents, thus making them more prone to accept a proposal that otherwise they might reject. This is similar as to introduce a mediator into the negotiation. However, Pareto optimality of the system is neglected in the above work.

There exist three pieces of work [6, 12, 29] that introduce a non-biased mediator into the system to help agents negotiate in complex situations. In [6], Ehtamo et al. present a constraint proposal method to generate Pareto-frontier of a multi-attribute negotiation with adopting a non-biased mediator. The mediator generates a constraint in each step and asks agents to find their optimal solution under this constraint. If the feedbacks from agents

coincide, then a Pareto optimal solution of the negotiation is found; otherwise, the mediator updates the constraint based on the feedbacks and the procedure continues. The authors show that their approach can generate the whole Pareto-frontier efficiently. However, in their work, agents have no right to make or accept offers based on their own negotiation strategies. Besides, the assumption that agents can solve multi-criteria-decision-making (MCDM) problems efficiently is essential for their approach. Klein et al. in [12] propose a more tractable and decentralized mediating approach for complex negotiations. The non-biased mediator generates an offer in each period and proposes to both agents. Then agents decide whether to accept the offer based on their own strategies. In their work, they propose two types of negotiation strategies for the agents and examine the equilibrium outcomes as well as the system efficiency. However, a key assumption in their work is that the issues are binary valued. Differently, [29] focuses more on the negotiations with continuously valued issues and provides agents sufficient decision flexibility. In their approach, agents can not only make the acceptance decision, but also choose the reference points based on which the mediator searches for Pareto optimal enhancements. Moreover, a highlight of their work is that the proposed protocol does not necessarily require agents to have a pre-prepared utility function before the negotiation starts but can assist agents to negotiate multiple issues easily with even simple negotiation strategies and also maintain Pareto optimality of the system. However, a common problem of the mediating work is that a non-biased mediator is required to be available. Although the technology achievements in the multi-agent field can make it easier to implement such a mediator by a software agent, there may exist situations where a mediator is not trusted or hard to be implemented. Thus, the work presented in this paper can complement the mediating approach in such domains.



An important issue in multi-attribute negotiation is the tradeoff process between self-interested agents on different issues. Faratin et al. [15] propose a novel idea to make agents trade off on multiple issues. They suggest that agents should apply similarity criteria to trade off the issues, i.e., make an offer on their indifference curve which is most similar to the offer made by the opponent in the last period. In their paper, they define a similarity criterion, based on which they propose the tradeoff algorithm. However, the similarity criterion requires agents have some knowledge about the weights the opponent puts on the issues in the negotiation. A subsequent work [18] proposes a method based on kernel density estimation to learn the weights assigned by the opponent. But the performance might be compromised if agents have no or very little prior information about the real weights the opponent assigns on the issues. Besides, it can become more difficult to define similarity criterion as that suggested in [15] if agents have nonlinear utility functions with interdependent attributes, i.e., the marginal utility of an attribute depends on the values of other attributes. From this point of view, the work presented in this paper shares the common idea with theirs but extends it to more general situations where agents can have nonlinear and interdependent preferences and do not need to have the information of the opponent's preference or strategy. In addition, by allowing an agent to propose multiple offers at a time, the model presented in this paper improves the quality of the solution in terms of Pareto optimality.

Other related work in this field includes: Luo et al. [26] use prioritized fuzzy constraints to represent trade-offs between the different possible values of the negotiation issues and to indicate how concessions should be made when they are necessary; Li and Tesauro [10] propose a method based on combinatorial search and Bayesian updating with the assumption

that agents know the structure of the opponent's utility function; Robu et al. [24] propose an approach based on graph theory and probabilistic influence networks for the negotiations with multiple binary issues with interdependent valuations; etc.

### 3 The model

#### 3.1 The negotiation setting

Consider two self-interested agents  $i \in \{b, s\}$  who need to negotiate a set of attributes  $j \in \{1, 2, \dots, n\}$ . The ranges of the attributes that need to be negotiated are given before the negotiation starts and each of those ranges can be normalized to a continuous range  $\Omega_j = [0, 1]$  with the lower and upper bounds representing the reservation prices of the two agents on this attribute. Without loss of generality, it is assumed that the value that is less than 0 (or more than 1) is not acceptable for agent  $s$  (or  $b$ ). Thus, the negotiation domain can be denoted by  $\Omega = [0, 1]^n$ .

In contrast to the prior work that usually assumes agents have relatively simple preferences on the attributes (e.g. can be characterized by linear utility functions), this paper makes a more common assumption that the preference of each agent is rational, which is widely applied in economics [1].

**Definition** The ordinal preference  $\preceq_i$  of agent  $i$  in the negotiation domain  $\Omega$  is rational if it satisfies the following conditions [1]:

- Strict preference is asymmetric: There is no pair of  $x$  and  $x'$  in  $\Omega$  such that  $x \prec_i x'$  and  $x' \prec_i x$ ;

- Transitivity: For all  $x, x'$  and  $x''$  in  $\Omega$ , if  $x \preceq_i x'$  and  $x' \preceq_i x''$ , then  $x \preceq_i x''$ ;
- Completeness: For all  $x$  and  $x'$  in  $\Omega$ , either  $x \preceq_i x'$  or  $x' \preceq_i x$ ;

where  $\vec{x} \preceq_i \vec{x}'$  (or  $\vec{x} \prec_i \vec{x}'$ ) indicates that the offer  $\vec{x}'$  is at least as good as (or better than)  $\vec{x}$  for agent  $i$ . The first two conditions ensure that agents' preferences are consistent in the negotiation domain and the third condition ensures that any pair of points in the negotiation domain can be compared.

Moreover, this paper assumes that each agent has a utility function  $u_i$  which can represent her ordinal preference, i.e.,  $\vec{x} \preceq_i \vec{x}'$  iff  $u_i(\vec{x}) \leq u_i(\vec{x}')$ . Without loss of generality, the utility range of each agent in the negotiation is normalized to  $[0, 1]$  with the bounds representing the worst/best offers. The paper assumes that  $\mathbf{0}^n/\mathbf{1}^n$  is the best/worst offer for agent  $b$ , i.e.,  $u_b(\mathbf{0}^n) = 1$  and  $u_b(\mathbf{1}^n) = 0$ , and from  $\mathbf{0}^n$  to  $\mathbf{1}^n$  agent  $b$ 's utility is monotonically decreasing, and it is the converse for agent  $s$ .

Finally, the information about the negotiation strategy and the preference of each agent is private.

### 3.2 Alternating-offer protocol

The model adopts Rubinstein's alternating-offer negotiation protocol [8] but it is allowed for agents to make multiple offers each time. In detail, in each period, an agent who behaves as a *proposer* makes a set (with a limited number) of offers to the opponent who behaves as a *responder*. If the responder accepts one of the offers, the negotiation ends; otherwise, agents exchange their roles and the negotiation proceeds to the next period. (The proposing and responding strategies about how many offers to make and how to respond to the offers will

be discussed in Section 3.3.) Such iterations continue until an agreement or the negotiation deadline is reached. This protocol is appropriate when there is a fairly equivalent and symmetric relationship between the two agents. For example, neither a buyer nor a seller has a monopoly in the demand or supply market.

---

### Negotiation Protocol

---

1. Start from the first period and choose one of the agents as the first proposer and the other as the responder;
  2. The proposer proposes a set of offers to the responder;
  3. The responder reacts to the offers;
  4. If the responder accepts one of the offers or the deadline is reached, the negotiation ends; otherwise agents exchange their roles and the negotiation proceeds to the next period.
- 

Figure 1: The alternating-offer protocol

## 3.3 The negotiation strategy

As mentioned in the introduction, an important characteristic of multi-attribute negotiations is that there possibly exist “*win-win*” situations, i.e., for any given solution in the negotiation domain, agents maybe can both be better off by trading off some attributes, except when the solution is already Pareto optimal (see Figure 2 for an example where  $x$  is the tangent point of two indifference curves<sup>2</sup> of the agents). Therefore, the negotiation strategy of a rational agent should take Pareto optimality into account because rational agents should not “leave extra money on the table”.

A negotiation strategy in a multi-attribute negotiation usually includes three parts: conceding, proposing, and responding. The first one is to decide how to concede in the negotiation. In other words, the conceding part decides the reservation utility—the least utility

---

<sup>2</sup>An indifference curve (or surface) of an agent in the negotiation domain consists of the points that are indifferent to the agent (i.e. that have the same utility for the agent).

an agent can accept for each negotiation period. The second part determines the offers that should be proposed to the opponent. Since in a multi-attribute negotiation, for any given utility value, there may exist a number of points in the negotiation domain which have this utility (i.e., the indifference curve/surface of the agent), it becomes essential for an agent to have an effective strategy to select points from this set as proposals offered to the opponent. The responding part suggests whether an agent should accept or reject an offer proposed by the opponent. The detailed description of these three parts in this model is provided in the following.

This model adopts the time-dependent strategy [17] as the conceding strategy (other strategies such as Bayesian learning [30] can also be applied). Time-dependent strategy provides a more tractable approach than the exhaustive reasoning strategies based on game theory, although it cannot promise an optimal outcome in all situations. As the utility range of each agent is given as  $[0,1]$  with the bounds representing the worst and best points, an agent can apply the following equation to concede in a negotiation:

$$s_i(t) = 1 - (1 - ru_i)\left(\frac{t}{T_i}\right)^{\frac{1}{\beta_i}}, \quad (1)$$

where  $s_i(t)$  is the utility that agent  $i$  desires to get in period  $t$ ,  $ru_i$  represents the ultimate reservation utility of agent  $i$  for this negotiation<sup>3</sup>,  $t$  is the current period,  $T_i$  is the deadline of agent  $i$ , and  $\beta_i$  represents the strategy parameter of agent  $i$ .

Based on the utility calculated from Equation (1), agents then can find the corresponding indifference curve/surface in the negotiation domain. (This indifference curve/surface is called the *current indifference curve/surface* in the following description.) However, an

---

<sup>3</sup>The ultimate reservation utility can be zero if an agent prefers an agreement with the worst solution in the negotiation domain to no deal, while it also can be positive if, for instance, an agent has outside options from which she can expect to get a positive utility [11].

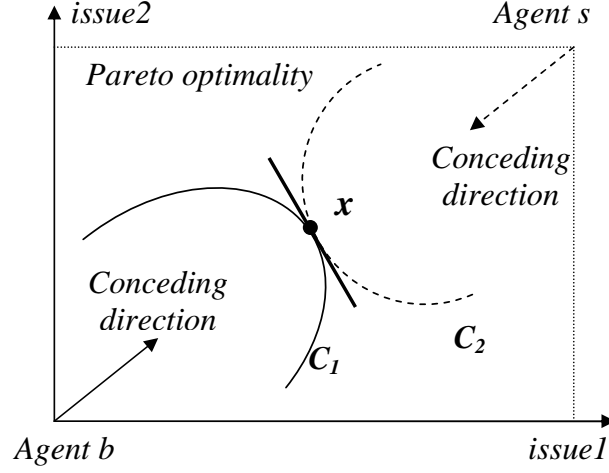


Figure 2: Pareto optimality

important issue that agents need to solve is how to select points from the current indifference curve/surface as the offers to the opponent, i.e., the proposing strategy. An exhaustive proposing approach is to propose the whole curve/surface to the opponent, and then the opponent can find the best points from the curve/surface and make the responding decision.

**Proposition 2.1** *The exhaustive proposing approach leads to a Pareto optimal outcome, with the conceding and responding strategies as given.*

**Proof.** This proposition is straightforward. Given the indifference curve/surface that an agent proposes, the opponent will choose the best offer to make the responding decision. If the opponent accepts the best offer, it is Pareto optimal since there does not exist any better solution on the given indifference curve/surface.  $\square$

However, the exhaustive proposing approach might not be appropriate in some situations, for instance, an agent does not want to let the opponent know her utility function explicitly or to propose a full curve/surface is practically not feasible. But it is usually feasible and reasonable for an agent to make multiple offers of a limited number at one time. For instance, a seller may propose several contracting options with different unit price, de-

livery time, and quality to a buyer each time; an employer may propose several job offers with different position, salary level, job location, and training opportunity to an applicant in each negotiation iteration. The benefit of allowing agents to make multiple offers is that it improves the desirability of the proposal to the opponent without sacrificing an agent's own utility; therefore the negotiation process converges faster to a Pareto optimal solution. Nevertheless, a formal method is still needed to generate the proposing directions. The goal is to find an offer that is as much acceptable to the opponent as possible based on the pace of concession. This paper proposes a heuristic in which an agent first chooses, from her current indifference curve/surface, the offer (assume it is unique for ease of exposition) which has the *shortest distance* to the best offer made by the opponent in the previous period. "The best offer made by the opponent in the previous period" means the offer that provides the agent with the highest utility among all the offers made by the opponent in the previous period. Thus, it implies that the indifference curve/surface of the agent which crosses the best offer is closer than those crossing other offers to the current indifference curve/surface. Then to choose the point on the current indifference curve/surface which has the shortest distance to that best offer may also more likely provide the opponent with the highest utility because such a point might be closer than other points to the opponent's current indifferent curve/surface. But this point might not necessarily be Pareto optimal under the situations with incomplete information and complex utility functions. Then by taking this offer as the seed and choosing a limited number of other offers from the current indifference curve/surface based on it, agents can improve the desirability of the proposal. Note, the exhaustive proposing approach can be viewed as the extreme case of this strategy. The benefit of proposing multiple offers at a time is specifically illustrated with numerical




Figure 3: The *shortest distance* proposing strategy

results in Section 4.2.

Figure 3 presents an example where agents are only allowed to make one offer in each period. The dashed curves are the indifference curves of agent  $s$  and the solid are the indifference curves of agent  $b$ . In period  $t - 4$ , agent  $s$  makes an offer  $x^{t-4}$ , but agent  $b$  rejects it. Then in period  $t - 3$ , agent  $b$  finds an offer  $x^{t-3}$  which has the shortest distance to the point  $x^{t-4}$  on her indifference curve in period  $t - 3$ . Agent  $s$  rejects this offer as well, and she concedes to the second left dashed curve in period  $t - 2$ . Now she finds offer  $x^{t-2}$  which is closest to  $x^{t-3}$  on this curve. Similar iterations continue till an agreement or the deadline is reached.

The proposing strategy then can be formalized as follows. Assume agent  $b$  is the proposer in period  $t$  and she concedes to the utility of  $s_b(t)$  with the corresponding indifference curve/surface  $\mathcal{C}$ . The total number of offers that agent  $b$  plans to make is  $k_b$ . Assume the set of offers agent  $s$  made in period  $t - 1$  is  $\mathbf{X}_{s \rightarrow b}^{t-1} = \{\bar{x}_{s \rightarrow b}^{t-1,1}, \dots, \bar{x}_{s \rightarrow b}^{t-1,k_s}\}$  where  $\bar{x}_{s \rightarrow b}^{t-1,1} \preceq_b \dots \preceq_b \bar{x}_{s \rightarrow b}^{t-1,k_s}$ . (Note, if the best offer is not unique, e.g.,  $\bar{x}_{s \rightarrow b}^{t-1,k_s-1} \sim_b \bar{x}_{s \rightarrow b}^{t-1,k_s}$ , then agent  $b$  can try all of them to derive the corresponding shortest distances by the following



Formula (2) and choose the one that has the smallest shortest distance.) Then agent  $b$  first chooses the offers:

$$\vec{x}_{b \rightarrow s}^{t,l} = \operatorname{argmin}_{\vec{x} \in \mathcal{C}} \|\vec{x} - \vec{x}_{s \rightarrow b}^{t-1,k_s}\|, \quad (2)$$

where  $\|\vec{x} - \vec{y}\|$  represents the distance from point  $\vec{x}$  to  $\vec{y}$  and  $l \in [1, k] \cap \mathbb{Z}$  is the index of the offer where  $k$  is the total number of solutions of Formula (2). If  $k \geq k_b$  (i.e. the number of solutions of Formula (2) is no less than the number of offers that agent  $b$  wants to make), then agent  $b$  randomly chooses  $k_b$  offers from the solutions; otherwise, agent  $b$  takes all of the  $k$  solutions and chooses the other  $(k_b - k)$  offers based on those  $k$  solutions. A simple method can be to choose the other  $(k_b - k)$  offers randomly from the (limited range) neighborhood of the  $k$  solutions:

$$\vec{x}_{b \rightarrow s}^{t,m} = \operatorname{rand}\{\vec{x} | \vec{x} \in \mathcal{C} \text{ and } \|\vec{x} - \vec{x}_{b \rightarrow s}^{t,l}\| \leq \delta(t), l \in [1, k]\}, \quad (3)$$

where  $\delta(t) = \min_{\vec{x} \in \mathcal{C}} \|\vec{x} - \vec{x}_{s \rightarrow b}^{t-1,k_s}\|$  and  $m \in [k + 1, k_b] \cap \mathbb{Z}$  is the index of the offer. Other more complicated methods based on the negotiation history also can be applied.

As discussed in Section 2, Faratin et al. [15] introduce a mechanism to make agents trade off based on the similarity between offers, which is implemented by fuzzy rules. But, first of all, an underlying assumption of the similarity criterion proposed in [15] is that the issues in the negotiation are independent and thus the utility functions are additive. It can become difficult to construct a similarity criterion if the utility functions are nonlinear and the issues are interdependent. Second, in some situations, to define and evaluate suitable fuzzy rules might be difficult. Third, to apply the similarity criterion proposed in their work requires an agent to have some information of the opponent's weights on the issues in order to calculate the similarity value between the offers. Thus, their model either requires

complete information or needs to be provided with a learning mechanism to retrieve the opponent's preference information. But to learn the weights that the opponent puts on the issues precisely is non-trivial (see [18] for a learning method). From this point of view, the approach presented in this paper extends their work to more general situations, since it does not require agents to know any information of the opponent's preference or strategy and can be applied to the negotiations with general utility functions.

The responding strategy applied in this model is to compare the utility of the current best offers made by the opponent with the utility that the agent will concede to in the next period. If the utility of the current best offers is higher, the agent accepts one of the best offers; otherwise, the agent rejects them. Thus, assume agent  $s$  is the responder in period  $t$  and  $\vec{x}_{b \rightarrow s}^{t, best} \in \mathbf{X}_{b \rightarrow s}^t$  represents the current best offer for agent  $s$  made by agent  $b$ , then the responding strategy of agent  $s$  can be formalized as:

$$a_s^t(\mathbf{X}_{b \rightarrow s}^t, u_s) = \begin{cases} \text{accept} & \text{if } s_s(t+1) \leq u_s(\vec{x}_{b \rightarrow s}^{t, best}); \\ \text{reject} & \text{otherwise;} \end{cases} \quad (4)$$

where  $s_s(t+1)$  is the utility agent  $s$  desires in period  $t+1$  which can be calculated by Equation (1). The full negotiation algorithm is described in Figure 4.

## 4 Experimental analysis

This section provides an experimental analysis to evaluate the performance of the model in terms of different utility functions. In the experiments, agents' utility functions are chosen from three types of utility functions: (I) exponential & additive, (II) exponential & inter-dependent, and (III) constant elasticity of substitution (CES) utility functions. All of these

**Inputs:**  $n, u_b, u_s, ru_b, ru_s, T_b, T_s, \beta_b, \beta_s, k_b, k_s$   
**Outputs:**  $\vec{x}$   
/\*assume agent  $b$  is the first proposer\*/  
**Begin:**  
1.  $t = 0, i = b', Rend = false$ ;  
2. **Do while**  $Rend == false$ ;  
3.  $s_i(t) = 1 - (1 - ru_i)(\frac{t}{T_i})^{\frac{1}{\beta_i}}$ ;  
4. **if**  $t == 0$   
5.  $X_i^t = \{\mathbf{0}^n\}$ ;  
6. **else**  
7.  $\vec{x}_{-i}^{t-1, best} = \operatorname{argmax}_{\vec{x} \in X_{-i}^{t-1}} u_i(\vec{x})$ ;  
8.  $X = \operatorname{argmin}_{u_i(\vec{x}) = s_i(t)} \|\vec{x} - \vec{x}_{-i}^{t-1, best}\|$ ;  
/\* $\|\cdot\|$  is the distance between the elements\*/  
9.  $k = |X|$ ;  
/\* $|\cdot|$  is the number of elements in the set\*/  
10. **if**  $k < k_b$   
11.  $\delta(t) = \|X - \vec{x}_{-i}^{t-1, best}\|$ ;  
12.  $S = \{\vec{x} | u_i(\vec{x}) = s_i(t) \text{ and } \|\vec{x} - X\| \leq \delta(t)\}$ ;  
13.  $Y = \operatorname{random}_{(k_b - k)}(S)$ ;  
/\*Randomly select  $k_b - k$  elements from  $S$ \*/  
14.  $X_i^t = X \cup Y$ ;  
15. **else**  
16.  $X_i^t = \operatorname{random}_{k_b}(X)$ ;  
17. **end**  
18. **end**  
19.  $s_{-i}(t+1) = 1 - (1 - ru_{-i})(\frac{t+1}{T_{-i}})^{\frac{1}{\beta_{-i}}}$ ;  
20.  $\vec{x}_i^{t, best} = \operatorname{argmax}_{\vec{x} \in X_i^t} u_{-i}(\vec{x})$ ;  
21. **if**  $s_{-i}(t+1) \leq u_{-i}(\vec{x}_i^{t, best})$   
22.  $\vec{x} = \vec{x}_i^{t, best}, Rend = true$ ;  
23. **elseif**  $t+1 == T_b || t+1 == T_s$   
24.  $\vec{x} = \emptyset, Rend = true$ ;  
25. **else**  
26.  $t = t + 1$ ;  
27.  $i \leftrightarrow -i$ ;  
28. **end**  
29. **end**  
30. return  $\vec{x}$ ;  
**End**

Figure 4: The negotiation algorithm

three types of utility functions are nonlinear. However, they also differ from each other. The first type utility functions are additive and the attributes are not interdependent, for instance,  $u(x_1, \dots, x_n) = 1 - \sum_{j=1}^n w_j x_j^2$ . In the second type, there exist interdependent terms, for instance,  $u(x_1, \dots, x_n) = 1 - [\sum_{j=1}^n w_j x_j^2 + \sum_{k=1}^n \sum_{j=1 \& j \neq k}^n w_{jk} (x_j - x_k)^2]$ . The third type utility functions are widely used by economists to model a broad range of user preferences and they follow the form of  $u(x_1, \dots, x_n) = 1 - [\sum_{j=1}^n w_j x_j^\rho]^{1/\rho}$ . The value range of each attribute in the experiments is set to  $[0, 1]$ , and  $(0, \dots, 0)'$  and  $(1, \dots, 1)'$  are the global optimal solutions for agent  $b$  and  $s$ , respectively. (The detailed utility function settings are described in the following subsections).

Section 4.1 first illustrates the behavior of the model by three specific examples in which agents are only allowed to make one offer in each round, and then an average performance analysis with different utility function settings and different number of offers that agents can make in each period is presented in Section 4.2.

## 4.1 Illustrative examples

**Example 1:** In this example, it is assumed that agent  $b$  and  $s$  need to negotiate three attributes and their utility functions follow type I utility functions:

$$\begin{aligned} u_b(x_1, x_2, x_3) &= 1 - .2x_1^2 - .6x_2^2 - .2x_3^2, \\ u_s(x_1, x_2, x_3) &= 1 - .6(1 - x_1)^2 - .2(1 - x_2)^2 - .2(1 - x_3)^2. \end{aligned}$$

Thus,  $(0,0,0)'$  and  $(1,1,1)'$  are the optima of agent  $b$  and  $s$ , respectively. Both agents follow the time-dependent negotiation strategy:

$$s(t) = 1 - (1 - ru)\left(\frac{t}{T}\right)^{\frac{1}{\beta}},$$

where the ultimate reservation utility  $ru = .2$ , deadline  $T = 20$  and the strategy parameter  $\beta = .8$ .

The negotiation procedure is presented in Figure 5. The dashed curve is the Pareto frontier<sup>4</sup> (Of course, in a negotiation, neither of the agents knows the Pareto-frontier. But it is mathematically calculated here for the comparison purpose of our approach). Agents are only allowed to make one offer each round and a circle in the figure represents a proposal that one of the agents makes in each negotiation period. The square represents the final agreement. Agent  $b$  is the first mover in this example and she proposes her optimal solution  $(0,0,0)'$  as the first offer; agent  $s$  rejects this offer and finds a counteroffer  $(.9244,.8033,.8032)$  which is the point closest to  $(0,0,0)$  on her indifference surface in the second period. The negotiation goes on until the agreement  $(.6777,.2900,.4891)$  is reached in the 8th period, which has a distance **.0795** to the Pareto frontier (Note, this distance represents the shortest distance from the final negotiation agreement to the Pareto frontier).

**Example 2:** In this example, agents negotiate two attributes and their utility functions follow the type II utility functions:

$$u_b(x_1, x_2) = 1 - .1(x_1 - x_2)^2 - .2x_1^2 - .8x_2^2,$$

$$u_s(x_1, x_2) = 1 - .1(x_1 - x_2)^2 - .7(1 - x_1)^2 - .3(1 - x_2)^2.$$

The negotiation strategies agents apply are the same as the one in Example 1. Figure 6 shows the negotiation procedure. Agents make offers alternatingly. Agent  $b$  first proposes her optimal solution  $(0,0)$  in the first period. But agent  $s$  rejects it and counteroffers

---

<sup>4</sup>Pareto frontier is the curve/surface that consists of the Pareto optimal solutions in the negotiation domain.




Figure 5: The negotiation procedure of Example 1

(.8916, .8163) in the second period. Such iterations continue until the final agreement is reached at (.6546, .3593) which has a distance **.0624** to the Pareto frontier. In the figure, the circles below the square are the offers made by agent  $b$ , while the circles above are made by agent  $s$ .

**Example 3:** In this example, agents negotiate two attributes and their utility functions follow CES utility functions:

$$u_b(x_1, x_2) = 1 - [.2x_1^3 + .8x_2^3]^{1/3},$$

$$u_s(x_1, x_2) = 1 - [.7(1 - x_1)^3 + .3(1 - x_2)^3]^{1/3}.$$

The negotiation strategies applied are the same as the one in Example 1. Figure 7 shows the negotiation procedure. Similarly, agent  $b$  first proposes and agent  $s$  responds. In the figure, the circles below the square are the offers made by agent  $b$ , while the circles above are made by agent  $s$ . The negotiation lasts 14 periods and the final agreement is reached at (.6274, .3976) which has a distance **.0410** to the Pareto frontier.




Figure 6: The negotiation procedure of Example 2




Figure 7: The negotiation procedure of Example 3

From the experiments above, we see that, based on this model, even if agents are only allowed to make one offer in each period, the offers are all quite close to the Pareto frontier. Thus, the shortest distance proposing strategy does make the offers approach the Pareto frontier. This result is promising because the utility functions applied in the experiments are much more general than linear additive utility functions that have been usually assumed in the existing literature. Moreover, the model does not require agents to know any information of the opponent's strategy and preference and does not need the existence of a non-biased mediator.

## 4.2 Performance analysis

This subsection evaluates the average performance of the model in terms of Pareto optimality with different utility function settings, different negotiation lengths, and different numbers of offers that agents can make in each period. The experiments are set as follows. First, the utility function types follow the ones used by the three examples in Section 4.1, i.e.,

(I)

$$u_b(x_1, x_2, x_3) = 1 - \sum_{j=1}^3 w_{bj} x_j^2,$$

$$u_s(x_1, x_2, x_3) = 1 - \sum_{j=1}^3 w_{sj} (1 - x_j)^2;$$

(II)

$$u_b(x_1, x_2) = 1 - w_{b1} (x_1 - x_2)^2 - \sum_{j=1}^2 w_{bj+1} x_j^2,$$

$$u_s(x_1, x_2) = 1 - w_{s1} (x_1 - x_2)^2 - \sum_{j=1}^2 w_{sj+1} (1 - x_j)^2;$$



(III)

$$u_b(x_1, x_2) = 1 - [\sum_{j=1}^2 w_{bj} x_j^{\rho_b}]^{1/\rho_b},$$

$$u_s(x_1, x_2) = 1 - [\sum_{j=1}^2 w_{sj} (1 - x_j)^{\rho_s}]^{1/\rho_s};$$

but now the weights  $w_{ij}$  ( $i \in \{b, s\}$ ) in the utility functions are chosen from *uniform distribution*  $U[0, 1]$  and satisfy the condition  $\sum_{j=1}^n w_{ij} = 1$ , and  $\rho_i$  in the third type is chosen from uniform distribution  $U[1, 20]$ . The negotiation strategy follows the time-dependent strategy provided in the above subsection but now the negotiation deadline ( $T$ ) varies from 20 to 100 by 20. The number of offers ( $k$ ) that agents can make in each step also varies from 1 to 5 in the experiments.<sup>5</sup> Then for every type of utility function with each  $T$  and  $k$ , there are 1000 instances run and the average distance from the negotiation agreements to the Pareto frontiers is calculated.

The results of the experiments of the utility types I, II and III are presented in Figure 8, 9 & 10, respectively, with the detailed data shown in Table 1, 2 & 3 (only the experiment data with  $k$  from 1 to 5 is presented in the tables for space concern). In the tables, an element with index  $(i, j)$ —the  $i^{th}$  row and  $j^{th}$  column—represents the average distance between the negotiation agreements and the Pareto frontiers in the experiments, when the number of offers that are allowed to be made in one period is equal to  $i$  and the deadline of the negotiations is equal to  $20 * j$ . For example, the element (1,1) ( $= 0.0826$ ) in Table 1 is the average distance from the negotiation agreements to the Pareto frontiers of the experiments with type I utility functions, one offer allowed to be made in each period, and the deadline equal to 20 periods.

---

<sup>5</sup>The negotiation settings in all the experiments assure that agents can reach agreements.

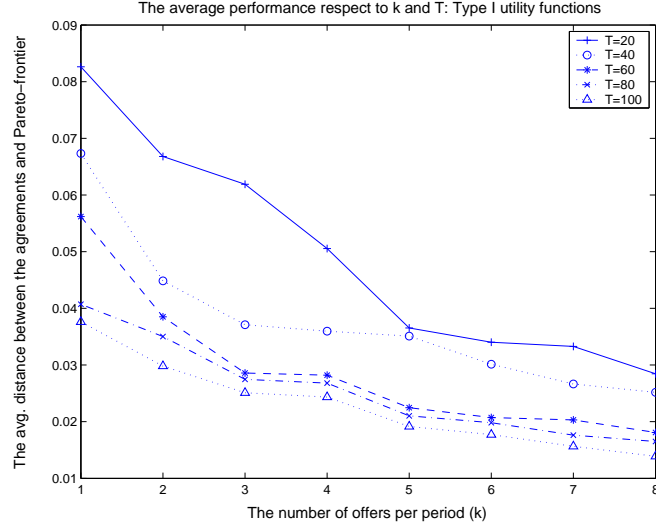


Figure 8: The average performance–Type I utility functions

Table I: The average performance–Type I utility functions

$k \backslash T$	20	40	60	80	100
1	0.0826	0.0673	0.0562	0.0407	0.0376
2	0.0668	0.0448	0.0385	0.0350	0.0298
3	0.0619	0.0371	0.0286	0.0275	0.0251
4	0.0505	0.0360	0.0282	0.0268	0.0243
5	0.0365	0.0350	0.0225	0.0210	0.0191

The results show that the average distance from the final negotiation results to the Pareto frontiers of each type utility function is small even when  $k = 1$ . This indicates that generally the model can provide near Pareto optimal solutions for those utility functions with every limited computational cost. Moreover, as the negotiation deadline increases, the average distance from the negotiation agreements to the Pareto frontiers decreases. This implies that agents concede more slowly when facing a longer negotiation deadline allows the offers to approach the Pareto frontier more closely. Thus, this model works better if the negotiation lasts longer. Furthermore, as  $k$  increases, the average distance between the final agreements and the Pareto frontiers decreases as well. The reason behind this is more




Figure 9: The average performance–Type II utility functions

straightforward. As the number of offers increases, the probability that the best offer that agents choose from the set is Pareto optimal increases. In the extreme case, if agents can exhaust all possible offers, the final agreement is Pareto optimal. Thus, the performance of the model can be improved as the negotiation time or the number of offers that are made in each round increases.

Table II: The average performance–Type II utility functions

$k \backslash T$	20	40	60	80	100
1	0.0315	0.0255	0.0205	0.0182	0.0154
2	0.0243	0.0174	0.0132	0.0107	0.0093
3	0.0216	0.0120	0.0101	0.0080	0.0058
4	0.0145	0.0090	0.0077	0.0053	0.0037
5	0.0128	0.0088	0.0049	0.0038	0.0030

## 5 Conclusion

Multi-attribute negotiation is an important and valuable mechanism in reality for people to reach agreements on multiple issues on which they share the common interests. But a multi-

attribute negotiation is much more complicated than a single-attribute negotiation. Firstly, agents may have complex preferences on continuously valued issues such as nonlinear and interdependent rather than linear additive preferences. Secondly, without any assistance from the negotiation protocol, agents would have to have a more complicated negotiation strategy to negotiate multiple issues simultaneously. Thirdly, in multi-attribute negotiation, rational agents need to take Pareto optimality into consideration. But it is not a trivial goal to maintain Pareto optimality for self-interested agents in incomplete information environments. The existing literature on multi-attribute negotiation has not addressed the above concerns at the same time. It is usually assumed in the prior work that agents have linear additive utility function, attributes are binary, agents are cooperative, or information is complete.

This paper presents a decentralized model for multi-attribute negotiations under incomplete information. The model allows a proposer to make several offers each round. Then, the responder can compare the offers and choose the best one for her to reach a deal or she can reject all of the offers. If the responder rejects the offers, then agents exchange their roles and the negotiation proceeds to the next period. In this model, the decision criterion proposed for an offering agent to make counteroffers is (a) to choose, on an indifference curve/surface, the offer (as the seed offer) that has the shortest distance to the best offer made by the opponent in the previous negotiation period, and then (b) to choose several other counteroffers from a specified neighborhood of the seed offer. This method ensures the offers approach the Pareto frontier as the negotiation proceeds. The numerical experiments show that this model leads agents to reach near Pareto optimal agreement effectively, even in the situations where agents have fairly general utility functions.

This work can be extended in several directions. First, similar as the existing literature,

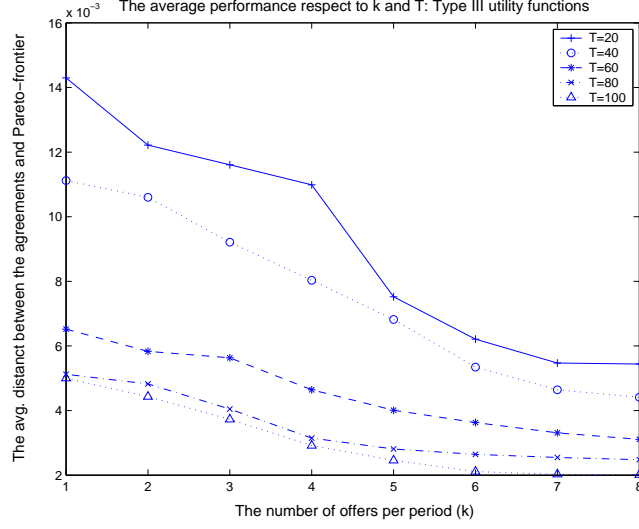


Figure 10: The average performance–Type III utility functions

Table III: The average performance–Type III utility functions

$k \backslash T$	20	40	60	80	100
1	0.0143	0.0111	0.0065	0.0051	0.0050
2	0.0122	0.0106	0.0058	0.0048	0.0044
3	0.0116	0.0092	0.0056	0.0040	0.0037
4	0.0110	0.0080	0.0046	0.0031	0.0029
5	0.0075	0.0068	0.0040	0.0028	0.0024

the current model assumes the utility functions of agents are given. However, there may exist situations in reality that explicit utility functions may not be available for agents before the negotiation starts. Then, it becomes essential for the model to take preference elicitation into consideration. Second, as the goal of this paper is to propose a model for agents to negotiate multiple issues effectively and Pareto optimally, the analysis on strategic behaviors of agents is absent in the paper. In the future work, we shall study more complicated negotiation strategies of an agent, for example, the decision of the number of offers to make in a period, other proposing or responding strategies.

## 6 Acknowledgments

This work was supported by the Office of Naval Research and the Navy Personnel Research, Studies and Technology (NPRST) under Navy contract number N6610-98-D-9501.

## References

- [1] A. Mas-Colell, M. D. Whinston and J. R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
- [2] M. Bac and H. Raff. Issue-by-issue negotiations: the role of information and time preference. *Games and Economic Behavior*, 13:125–134, 1996.
- [3] L.-A. Busch and I. J. Horstmann. Endogenous incomplete contracts: A bargaining approach. *Games and Economic Behavior*, 19:144–148, 1997.
- [4] R. Inderst. Multi-issue bargaining with endogenous agenda. *Games and Economic Behavior*, 30:64–82, 2000.
- [5] K. Lang and R. W. Rosenthal. Bargaining piecemeal or all at once. *The Economic Journal*, 111:526–540, 2001.
- [6] H. Ehtamo, R. Hamalainen, P. Heiskanen, J. Teich, M. Verkama and S. Zions. Generating pareto solutions in a two-party setting: constraint proposal methods. *Management Science*, 45:1697–1709, 1999.
- [7] J. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.

- [8] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97-109, 1982.
- [9] E. Kalai. Proportional solutions to bargaining situations: Intertemporal utility comparisons. *Econometrica*, 45(7):1623–1630, 1977.
- [10] C. Li and G. Tesauro. A strategic decision model for multi-attribute bilateral negotiation with alternating offers. In *4th ACM Conf. on Electronic Commerce*, pages 208–209, 2003.
- [11] C. Li, J.A. Giampapa and K. Sycara. Bilateral negotiation decisions with uncertain dynamic outside options. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Special Issue on Game-theoretic Analysis and Stochastic Simulation of Negotiation Agents*, 36(1):1-13, 2006.
- [12] M. Klein, P. Faratin, H. Sayama and Y. Bar-Yam. Negotiating complex contracts. *Group Decision and Negotiation*, 12(2):111–125, 2003.
- [13] J. Brzostowski and R. Kowalczyk. On possibilistic case-based reasoning for selecting partners for multi-attribute agent negotiation. In *4th Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 273–279, 2005.
- [14] P. Faratin, C. Sierra and N. R. Jennings. Using similarity criteria to make negotiation trade-offs. In *4th Int. Conf. on Multi-Agent Systems*, pages 119–126, 2000.
- [15] P. Faratin, C. Sierra and N. R. Jennings. Using similarity criteria to make issue trade-offs in automated negotiations. *Artificial Intelligence*, 142(2):205–237, 2002.

- [16] Y. Chevaleyre, U. Endriss, J. Lang and N. Maudet. Negotiating over small bundles of resources. In *4th Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 296–302, 2005.
- [17] P. Faratin, C. Sierra and N. R. Jennings. Negotiation decision functions for autonomous agents. *International Journal of Robotics and Autonomous Systems*, 24:159–182, 1998.
- [18] R. M. Coehoorn and N. R. Jennings. Learning an opponent’s preferences to make effective multi-issue negotiation tradeoffs. In *6th Int. Conf. on E-Commerce*, pages 59–68, 2004.
- [19] S. Fatima, M. Wooldridge and N. R. Jennings. Optimal agendas for multi-issue negotiation. In *2nd Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 129–136, 2003.
- [20] S. Fatima, M. J. Wooldridge and N.R. Jennings. An agenda-based framework for multi-issue negotiation. *Artificial Intelligence*, 152:1–45, 2004.
- [21] S. Fatima, M. J. Wooldridge and N.R. Jennings. Optimal negotiation of multiple issues in incomplete information settings. In *3rd Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 1080–1087, 2004.
- [22] K. Sycara. Persuasive argumentation in negotiation. *Theory and Decision*, 28(3):203–242, 1990.
- [23] K. Sycara. Problem restructuring in negotiation. *Management Science*, 37(10):1248–1268, 1991.



- [24] V. Robu, D. Somefun and J. La Poutre. Modeling complex multi-issue negotiations using utility graphs. In *4th Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 280–287, 2005.
- [25] C. Jonker and V. Robu. Automated multi-attribute negotiation with efficient use of incomplete preference information. In *3rd Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 1056–1063, 2004.
- [26] X. Luo, N. R. Jennings, N. Shadbolt, H. Leung and J. H. Lee. A fuzzy constraint based model for bilateral multi-issue negotiations in semi-competitive environments. *Artificial Intelligence Journal*, 148(1-2), 53-102, 2003.
- [27] S. Saha, A. Biswas and S. Sen. Modeling opponent decision in repeated oneshot negotiations. In *4th Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 397–403, 2005.
- [28] G. Lai, C. Li, J. Giampapa and K. Sycara. Literature review on multi-attribute negotiations. *Tech. Report*, Robotics Institute, Carnegie Mellon Univ., Pittsburgh, USA, December, 2004.
- [29] G. Lai, C. Li and K. Sycara. A multi-attribute negotiation model with considering generality, Pareto optimality and tractability. *Tech. Report*, Robotics Institute, Carnegie Mellon Univ., Pittsburgh, USA, December, 2004.
- [30] D. D. Zeng and K. Sycara. Bayesian learning in negotiation. *International Journal of Human-Computer Studies*, 48:125-141, 1998.

- [31] L. Chen and P. Pu. Survey of preference elicitation methods. *EPFL Technical Report IC/2004/67*, Switzerland, 2004.

## 7 Biography

**Guoming Lai** is a Ph.D. candidate in Management of Manufacturing and Automation, a joint program of Tepper School of Business with School of Computer Science at Carnegie Mellon University. He received the B.S. and M.S. degrees in Industrial Engineering from Tsinghua University, China, in 2000 and 2003, respectively. His research interests lie in the interfaces between operations, technology, and economics. His research areas have covered agent mediated e-commerce, multi-agent systems, and supply chain management.

**Katia Sycara** is a Professor in the School of Computer Science at Carnegie Mellon University. She is also the Director of the Laboratory for Agents Technology and Semantic Web Technologies. She holds a B.S in Applied Mathematics from Brown University, M.S. in Electrical Engineering from the University of Wisconsin and PhD in Computer Science from Georgia Institute of Technology. She holds an Honorary Doctorate from the University of the Aegean (2004). She is a member of the Scientific Advisory Board of France Telecom, a member of the Scientific Advisory Board of the Greek National Center of Scientific Research "Demokritos" Information Technology Division. She is a Fellow of the American Association for Artificial Intelligence and the recipient of the 2002 ACM/SIGART Agents Research Award. She is also a fellow of IEEE and a member of the ACM.

Prof. Sycara has given numerous invited talks, has authored more than 250 technical papers dealing with Multiagent Systems, Negotiation, Software Agents, Agent Teams, Web Services, the Semantic Web, and Human-Agent-Robot Teams. She has led multimillion dol-

lar research effort funded by DARPA, NASA, AFOSR, ONR, AFRL, NSF and industry. She is a founding member and member of the Board of Directors of the International Foundation of Multiagent Systems (IFMAS). She is a founding member of the Semantic Web Science Association, and serves as the US co-chair of the Semantic Web Services Initiative. She is a founding Editor-in-Chief of the journal "Autonomous Agents and Multiagent Systems" and is currently serving on the editorial board of 5 additional journals.

**Cuihong Li** is an Assistant Professor of Operations and Information Management at the School of Business, University of Connecticut. She received her Ph.D. in Management of Manufacturing and Automation from Carnegie Mellon University in 2005. Her research interests lie in the interfaces between operations, technology, and economics. Her research areas have covered e-commerce, multi-agent systems, and supply chain management. She is a member of INFORMS.