A Generic Framework for Automated Multi-attribute Negotiation

Guoming Lai

Tepper School of Business

Carnegie Mellon University

5000 Forbes Avenue

Pittsburgh, PA 15213

guomingl@andrew.cmu.edu

Katia Sycara

Robotics Institute

Carnegie Mellon University

5000 Forbes Avenue

Pittsburgh, PA 15213

katia@cs.cmu.edu

Abstract

Automated multi-attribute negotiation provides an important mechanism for distributed decision makers to reach agreements on multiple issues. Agents in a negotiation may have to negotiate multiple issues simultaneously. Moreover, to negotiate multiple issues also provides the opportunity to reach "win-win" solutions. In this paper, we first provide a survey that synthesizes the research on multi-attribute negotiation. We discuss the limitations of the existing research and conclude three key issues: incomplete information, Pareto-optimality and tractability which need to be further studied. We then present a generic framework for automated multi-attribute negotiation with two new mechanisms that incorporate the above issues. Finally, we discuss the challenges and the future work.

Keywords: Automated multi-attribute negotiation, Incomplete information, Pareto-optimality, Mediating, Win-win.

1. Introduction

Automated negotiation provides an important mechanism for distributed decision makers, either human participants or autonomous agents, to reach agreements. With the support of an automated negotiation system, human participants of the negotiations can just input their preferences, requirements, etc. into the system and the representative agents can help negotiate the contents automatically. Such a procedure can not only ease but also accelerate the negotiation processes. Imagine a human coordinator in a large organization who may need to negotiate with many members in the group on different issues at the same day. A face-to-face and one-to-one negotiation approach may make the whole process work slowly and thus impact the operational efficiency of the organization. In the world of autonomous agents, automated negotiation can provide a mechanism for agents to reach agreements on task allocation, resource sharing and surplus division.

To implement automated negotiation requires not only software support systems but also applicable negotiation theories. The existing research on negotiation problems usually can be segmented into two classes on: single-attribute negotiation and multi-attribute negotiation. While the single-attribute negotiation problem has been extensively studied, the research on multi-attribute negotiation is still young. In this paper, we focus on multi-attribute negotiation.

Multi-attribute negotiation is important. First, there exist situations where agents have to negotiate multiple issues at the same time. For example, in the human environment, a supplier and a buyer may need to negotiate the quantity, price and delivery time of a supply contract at the same time; an employer and an applicant may need to negotiate the position, wage and training opportunities simultaneously; in the agent world, it is common that two (or more) agents in an organization need to decide how to allocate multiple tasks or share a set of resources. In those situations, without the agreement on whichever issue may make the whole process halt. Second, besides the necessity, the agents may also benefit from multi-attribute negotiation when they have different preferences over the issues. Being able to trade off one issue for another, the agents may reach an agreement that makes the agents mutually better off. For instance, when selling automobiles, the dealers can sell the automobiles with a single price, but more often they may also introduce the financing package, insurance package, warranty package, spare parts package into the contract. 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, as they may find that the price of buying those packages individually is 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.

As the agents now need to deal with multiple issues simultaneously, multi-attribute negotiation can be much more difficult than single-attribute negotiation, however. The difficulty comes from the following factors. First, in a multi-attribute negotiation, an agent’s utility depends on all the issues. As a result, to make an appropriate offer becomes more complex, since in each step, an agent may find a number of offers that can contribute the same utility level for herself. Which offer to propose is usually nontrivial. This decision impacts the opponent’s utility and then her response decision. If an agent can find an offer that makes the opponent obtain more utility compared to other offers, the opponent may be more willing to accept the offer, and therefore, the agent can concede less and consequently achieve more utility. Second, negotiations in practice often take place in the environments where information is incomplete. The agents may never meet with each other before. While it may be possible for the agents to reason and learn the opponent’s utility function and strategy in a single-attribute negotiation, it becomes much more difficult when there are multiple issues. Moreover, the negotiation context in practice may also vary along with the time. The agents may even not have a complete characterization of their own preferences before a negotiation. The traditional approach to overcome this problem is to apply preference elicitation before a negotiation. However, preference elicitation is a well known difficult and time-consuming procedure (Chen and Pu 2004), especially when the preferences of the agents are complex. Third, in a multi-attribute negotiation, to achieve a Pareto optimal[1] solution is important. Rational

---

[1] 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.

agents shall not leave “extra money” on the table. But to seek a Pareto optimal settlement between self-interested agents in an incomplete information environment is difficult. Therefore, the research on multi-attribute negotiation faces more challenges than that for single-attribute negotiation. A multi-attribute negotiation system need be able to support the agents to negotiate the issues efficiently and robustly, in the domain where the agents may neither know the opponents' preferences nor have a complete characterization of their own preferences.

In this paper, we first review the existing research on multi-attribute negotiation and discuss the gap between the existing work and an applicable automated multi-attribute negotiation system. We point out three key issues which need to be further studied: incomplete information, Pareto-optimality and tractability to support a robust automated multi-attribute negotiation system. We then present a generic framework with two new mechanisms that consider the above issues. Finally, we discuss the future work.

The rest of the paper is organized as follows. Section 2 reviews the related research. Section 3 presents the generic framework. Section 4 outlines the challenges and concludes with the discussion on the future work.

2. Literature review

Although negotiation has been viewed as an important research problem as well as a valuable practical tool for a long time, the research on multi-attribute negotiation problems is still young. We segment the existing work into two categories in the fields of game theory and artificial intelligence (AI).

2.1 Game theoretic models

The goal of the research in game theory is to find optimal negotiation strategies and the corresponding equilibria under different negotiation settings. In game theory, the simplest context studied is the one with complete information and cooperative agents. In this context, since the agents know the utility functions of each other, it is not hard to compute the Pareto-frontier, the collection of Pareto optimal solutions. So, the rational agents can reach agreement on this frontier by Nash axioms, Kalai-Smorodinsky solution, etc. (Nash 1950; Kalai 1977). However, these approaches are not applicable in practice where the agents are usually non-cooperative and the information is incomplete.

The non-cooperative models (e.g. Bac and Raff 1996; Busch and Horstmann 1997; Busch and Horstmann 1999a, 1999b; Lang and Rosenthal 2001, etc.) usually assume that in the negotiation there are two issues and the agents' utility functions are linear and additive on the values of the two issues. The research mainly studies two negotiation protocols: simultaneous negotiation and issue-by-issue negotiation. The research focuses on the questions of: which protocol the agents shall choose to negotiate the issues; which issues should be negotiated first if the agents take issue-by-issue negotiation; what kind of offers the agents shall make; whether the agreement on the two issues will be reached at once in the first period; how the agents divide the two issues in equilibrium; and how the setting of the information and the sizes of the two issues will impact the equilibrium outcomes. For instance, Bac and Raff (1996) based on the Rubinstein's bargaining game, study the case that two self-interested agents need to share two identical pies. The agents can either choose simultaneous negotiation or issue-by-issue negotiation. Bac and Raff show that with complete information the first mover will take simultaneous negotiation by making an offer on both pies and the two agents reach an agreement on the division of the two pies without delay. They also show that if the information is asymmetric then the first mover may choose issue-by-issue negotiation by making an offer on one pie and leave the right on the other pie to the opponent.

A major problem of the game theoretical models is that the research results are difficult to be applied in practice. The agents in real world situations may have much more complex utility functions on the issues rather than linear additive utility functions; the information in the negotiation may be incomplete; there could be more than two issues; and the agents may have limited reasoning capability. In those non-cooperative game models, it is difficult to scale the size of the negotiation up to three issues or to consider nonlinear utility functions under incomplete information. Moreover, Pareto-optimality of the negotiation solutions usually is neglected.

2.2 AI models

In the AI field, the existing work mainly focuses on automated negotiation frameworks and tractable heuristics.

Fatima et al. (2004a, 2004b) propose an agenda-based framework for multi-attribute negotiation. In their framework, the agents can propose either a combined offer on multiple issues or a single offer on one issue. Different from the game theoretical models, their work addresses more on tractability. They assume that the agents adopt time-dependent strategy (first introduced by Faratin, Sierra and Jennings 1998) and the agents may 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 may have two independent strategies S_1 and S_2 which are used to decide whether to accept x_1 and x_2. However, in their work, they make the same assumption as in the game theoretical models that the agents' utility functions are given before the negotiation and they are linear additive. Similarly, Pareto-optimality is not addressed.

Sycara (1990a, 1990b, 1991) uses a case-based reasoning approach for multi-attribute negotiations where the agents make offers based on similarity of the negotiation context (including issues, opponents, and environment) to previous negotiations. Sycara also uses automatically generated persuasive argumentation as a mechanism for altering the utilities of the agents, thus making them more prone to accept a proposal that otherwise they might reject. However, in the research, neither complex utility function nor Pareto-optimality is explicitly considered.

There are also papers that adopt a non-biased mediator in the negotiation. Ehtamo et al. (1999) present a constraint proposal method to generate Pareto-frontier of a multi-attribute negotiation. The mediator generates a constraint in each step and asks the agents to find their optimal solution under this constraint. If the feedbacks from the 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. They show that their approach can generate the whole Pareto-frontier efficiently. But in their work, the negotiation agents do not have any right to make or accept offers based on their own negotiation strategies, which limits its application in the negotiations with self-interested agents. Moreover, their approach relies on the assumption that the agents can solve multi-criteria-decision-making (MCDM) problems efficiently. But it is difficult to require any agent in practice to have such a capability. Klein et al. (2001) propose a mediating approach for negotiating complex contracts with more decision flexibility for the agents. Their approach focuses on the negotiations with binary valued issues. The non-biased mediator generates an offer in each period and proposes to both agents. Then the agents vote whether to accept the offer based on their own strategies. If both agents vote to accept, the mediator mutates the offer (to change the values of some issues in the offer from 0 to 1, or reverse) and repeats the procedure. If at least one agent votes to reject the offer, the mediator mutates the last mutually acceptable offer and repeats the procedure. However, this approach is difficult to be applied in the domains with continuously-valued issues since it is not tractable to mutate the value of an issue if it has a continuous support. Besides, a key assumption they make is that the mediator always can change the contract even if both agents have already voted to accept the contract, which might not be tractable in practice.

An important issue in multi-attribute negotiation is the tradeoff process between self-interested agents on different issues. Faratin, Sierra and Jennings (2002) propose a novel idea to make the agents trade off on multiple issues. They suggest that the 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. However, in this approach, to define and apply the similarity criteria, it is essential that the agents have some knowledge about the weights the opponent puts on the issues in the negotiation. A subsequent work (Coehoorn and Jennings 2004) proposes a method based on kernel density estimation to learn the weights. But the performance still might be compromised if the agents have no or very little prior information about the real weights the opponent assigns on the issues. Moreover, it will be difficult to define and apply the similarity criteria if the agents' utility functions are nonlinear and the issues are interdependent.

Luo et al. (2003) develop a fuzzy constraint based framework for multi-attribute negotiations. In this framework, an agent, say the buyer, first defines a set of fuzzy constraints and submits one of them by priority from the highest to lowest to the opponent, say the seller, during each round. The seller either makes an offer based on the constraints or lets the buyer relax the constraints if a satisfactory offer is not available. The buyer then makes the decision to accept or reject an offer, or to relax some constraints by priority from the lowest to highest, or to declare the failure of the negotiation.

Li and Tesauro (2003) introduce a searching method based on Bayesian rules. In their work, a proposing agent in each negotiation round applies a depth-limited combinatorial search to find a most favorable offer based on her current knowledge about the opponent's type. If the proposal is rejected, this agent updates her knowledge by Bayesian rules about the opponent's type. Then, the agents exchange their roles and the negotiation proceeds following the same pattern. However, their work assumes that the agents know partially about the opponent's utility function as the type of the agent. Moreover, their work does not explicitly address Pareto-optimality of the solution.

Hanson et al. (2003) introduce an automated negotiation system implemented in Java, which can be flexibly used to negotiate multiple issues. Their experiments based on asymmetric negotiation protocols show that near Pareto optimal negotiation solutions can be efficiently achieved for two-attribute Constant Elasticity of Substitution (CES) utility function (for an introduction of CES utility functions, see Mas-Colell, Whinston and Green 1995).

There also exists some research that addresses multi-attribute negotiations on binary issues. For instance, Robu, Somefun and La Poutre (2005) propose an approach based on graph theory and probabilistic influence networks for the negotiations with multiple binary issues; Chevaleyre et al. (2005) address a categorization problem of the agents' utility functions under which the social optimal allocation of a set of indivisible resources (binary issues) is achievable.

We conclude the existing research as follows. First, almost all the models in the existing research are based on the assumption that the agents in a negotiation have explicit utility functions. Some also assume that the agents completely or partially know their opponent's utility function. Second, the existing models either assume a simple utility function (two issues with linear additive utility functions) or focus on binary issues or cooperative negotiations. Finally, Pareto-optimality and tractability have not been considered simultaneously in the models.

3. A generic framework for automated multi-attribute negotiation

In this section, we propose a generic framework that incorporates three key issues: incomplete information, Pareto-optimality and tractability to bridge the gap between the existing theories and the application requirements. In Subsection 3.1, we describe the assumptions of the framework. In Subsection 3.2, we describe the negotiation protocol. In Subsection 3.3, we describe the negotiation strategy.

3.1 The negotiation setting

In this framework, we consider the case with two self-interested agents $i \in \{b, s\}$ who need to negotiate a set of issues $j \in \{1, 2, ..., n\}$ in T periods. The range of each issue j 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 issue. Without loss of generality, we assume 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 that agents have relatively simple preferences on the issues (e.g. can be characterized by linear utility functions), we make a more general assumption that the preference of each agent is rational and strictly convex, which is widely applied in economics (Mas-Colell, Whinston and Green 1995).

Definition 1: The ordinal preference $\lesssim_i$ of agent i in the negotiation domain is rational and strictly convex if it satisfies the following conditions:

- Strict preference is asymmetric: There is no pair of $\boldsymbol{x}$ and $\boldsymbol{x}'$ in Ω such that $\boldsymbol{x} \prec_i \boldsymbol{x}'$ and $\boldsymbol{x}' \prec_i \boldsymbol{x}$;
- Transitivity: For all $\boldsymbol{x}$, $\boldsymbol{x}'$ and $\boldsymbol{x}''$ in Ω, if $\boldsymbol{x} \lesssim_i \boldsymbol{x}'$ and $\boldsymbol{x}' \lesssim_i \boldsymbol{x}''$, then $\boldsymbol{x} \lesssim_i \boldsymbol{x}''$;
- Completeness: For all $\boldsymbol{x}$ and $\boldsymbol{x}'$ in Ω, either $\boldsymbol{x} \lesssim_i \boldsymbol{x}'$ or $\boldsymbol{x}' \lesssim_i \boldsymbol{x}$;
- Strict convexity: For any solution $\boldsymbol{x}$, the set of solutions that an agent prefers to $\boldsymbol{x}$ is strictly convex;

where $\boldsymbol{x} \lesssim_i \boldsymbol{x}'$ (or $\boldsymbol{x} \prec_i \boldsymbol{x}'$) indicates that the offer $\boldsymbol{x}'$ is at least as good as (or better than) $\boldsymbol{x}$ for agent i.

The first two conditions ensure that the agents' preferences are consistent in the negotiation domain; the third condition ensures that any pair of points in the negotiation domain can be compared; the last condition ensures that agents' preferences on each issue are monotone if the values of the other issues are fixed, i.e., if the value of an issue increases, when the values of the other issues are fixed, the utility of an agent is monotonically increasing or decreasing. This last condition implies that each Pareto optimal solution of a multi-attribute negotiation is on a joint tangent hyperplane of a pair of indifference curves (or surfaces)[2] of the two agents and the Pareto frontier is a continuous curve. This condition makes it tractable to find near Pareto optimal solutions when the issues are continuously-valued and the information is incomplete. Although the last condition is much stricter compared to the other three conditions, it is still quite general and holds in many real negotiation environments. The linear, CES and quadratic utility functions all satisfy these four conditions.

Based on the above conditions, we normalize the utility range of each agent to [0,1] with the bounds representing the worst/best offers in the negotiation space. We assume 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*. We assume that the agents may or may not have her utility function elicited, but given a limited number of offers, an agent can judge the utility level of the offers and find the best. This assumption is indeed the basic requirement of any preference elicitation procedures.

3.2 The alternating-offer protocol

We adopt Rubinstein's alternating-offer game (Rubinstein 1982) but allow agents to make multiple offers each time. In detail, in each period, an agent who acts as a proposer

[2] An indifference curve (surface) of an agent consists of the points that are indifferent to the agent.

makes one (or multiple with a limited number) offer to the opponent who acts 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. Such iterations continue until an agreement or the negotiation deadline is reached.

3.3 The negotiation strategy

We divide the negotiation strategy of an agent into three components: *conceding*, *responding*, and *proposing*. The conceding strategy is to decide how to concede in the negotiation. In other words, it decides the reservation utility—the least utility an agent desires in each negotiation period; the responding strategy determines whether an agent should accept or reject an offer proposed by the opponent; the proposing strategy is one which determines the offers that should be proposed to the opponent. We describe these three components one-by-one in the following.

3.3.1 The conceding strategy

We have assumed that each agent's preference can be characterized by a utility function, and the range of the utility function is [0,1], i.e., it is one-dimensional. We therefore can apply the existing conceding strategies developed for single-attribute negotiations to determine how much utility to concede in each period. Particularly, we propose that the agents can adopt the time-dependent strategy (Faratin, Sierra and Jennings 1998) due to its tractability for applications. The time-dependent strategy can be characterized by:

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

where $s_i(t)$ is the reservation utility of agent *i* in period *t* (i.e., the opponent's offer in this period needs to provide at least this utility level for agent *i* to accept it); ru_i is the ultimate reservation utility of agent *i* for this negotiation (i.e., the least utility level that agent *i* is willing to accept); and $\beta_i > 0$ represents the strategy parameter of agent *i*. Following this strategy, an agent desires high reservation utility levels at the beginning since there is still plenty of time left; an agent concedes gradually depending on the time past if there is no agreement reached yet; and at the end of the negotiation an agent concedes to the ultimate reservation utility for this negotiation. β_i controls the concession characteristic. If $\beta_i < 1$, agent *i* concedes slowly at the beginning but fast when the time approaches the deadline; if $\beta_i > 1$, the agent concedes fast at the beginning but slowly when the time approaches the deadline; and if $\beta_i = 1$, the agent concedes evenly during the whole negotiation. ru_i may be fixed during the whole negotiation or may be adjusted by the agent time-by-time if the agent receives outside options during the negotiation (see Li, Giampapa and Sycara 2006).

3.3.2 The responding strategy

The responding strategy directly depends on the conceding strategy. We propose that an agent can compare the utility of the current best offer made by the opponent with the utility that the agent will concede to in the next period. If the utility of the current best offer is higher, the agent accepts the offer; otherwise, the agent rejects it. Thus, suppose agent *b* is the responder and ***x*** is the best offer made by agent *s* in period *t*, then agent *b*'s responding strategy follows

$$a_b(x) = \begin{cases} accept, & \text{if } s_b(t+1) \le U_b(x); \\ reject, & \text{otherwise.} \end{cases} \tag{2}$$

where $a_i(\boldsymbol{x})$ represents the reaction function of agent i to the offer $\boldsymbol{x}$.

3.3.3 The proposing strategy

Among these three component strategies, the proposing strategy is the most complicated one, since in a multi-attribute negotiation, for any given utility level that an agent concedes to, there may exist a number of different points in the negotiation space which can contribute this utility level (i.e., all the points on the indifference curve/surface with this utility level). It then becomes essential for an agent to have an effective approach to find and select desirable points from this set as proposals offered to the opponent.

In particular, we design two new mechanisms that can be used as the proposing component in different situations. The first negation mechanism (in Subsection 3.3.3.1) is called shortest-distance proposing mechanism, which can be applied where the agents do not know their opponent's utility function but do have a utility function of their own. The second negotiation mechanism (in Subsection 3.3.3.2) is called Pareto-optimal mediating mechanism, which can be applied where the agents neither know the opponent's utility function nor have their own utility functions elicited. The second mechanism requires a mediating third party which can be implemented by a software agent.

3.3.3.1 The shortest-distance proposing mechanism

For this strategy, we assume that each agent does have an elicited utility function before the negotiation, but the utility function is private information.

As an agent has an explicit utility function, the agent can find the corresponding indifference curve/surface in the negotiation space with the utility level equal to the reservation utility calculated in Equation (1) for period *t*. The simplest approach is to propose the whole indifference curve/surface to the opponent and the opponent looks for the best point on the curve/surface and makes the responding decision. We call this approach the exhaustive proposing approach. 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. In such situations, the agents need to find an alternative approach.

Therefore, we propose a mechanism called *shortest-distance proposing mechanism*. In this mechanism, an agent first chooses, from her current indifference curve/surface, the offer 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. To choose the point on the current indifference curve/surface of the agent which has the shortest distance to that best offer may more likely provide the opponent with the highest utility because such a point might be closer than the other points to the opponent's current indifferent curve/surface.

Figure 1 presents an example of the shortest-distance proposing mechanism. In this example, 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 $\boldsymbol{x}^{t-4}$, but agent *b* rejects it. Then in period *t-3*, agent *b* finds an offer $\boldsymbol{x}^{t-3}$ on her indifference curve in period *t-3* which has the shortest distance to the offer $\boldsymbol{x}^{t-4}$. Agent *s* rejects this offer as well, and she concedes to the second left dashed curve in period *t-2*. Now she finds offer $\boldsymbol{x}^{t-2}$ on this curve which is closest to $\boldsymbol{x}^{t-3}$. Similar iterations continue until an agreement or the deadline is reached.

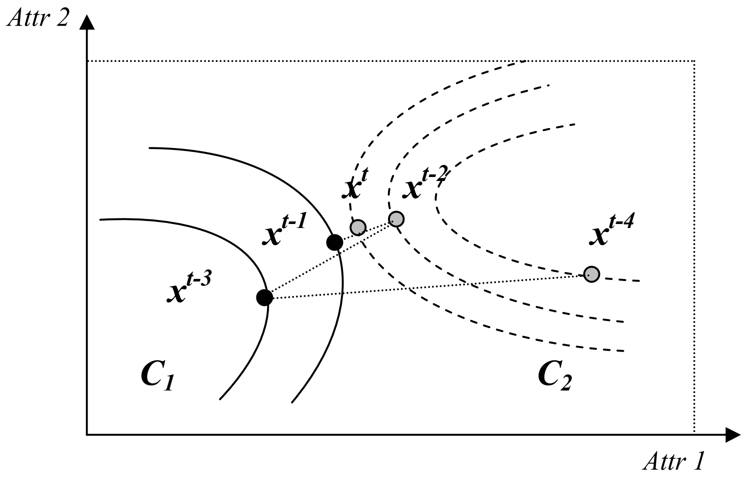

Figure 1: The shortest distance proposing protocol

However, the point chosen by the shortest distance protocol might not necessarily be Pareto optimal, especially if the information is incomplete and the utility functions are complex. Therefore, we propose that the agent can take this point as the seed offer, and based on it, choose a limited number of other offers from the current indifference curve/surface. The agent proposes these offers together to the opponent. By doing so, the agent can improve the desirability of the proposals. Although to propose the whole curve/surface may not be tractable, to make a limited number offers (e.g., two or three) in each period usually can be feasible and reasonable. For instance, a seller may propose several contracting options with different unit price, delivery time and quality to a buyer at each time; an employer may propose several job offers with different position, salary level, job location and training opportunity to an applicant.

Therefore, the shortest-distance proposing mechanism can be formalized as follows. Assume agent *b* is the proposer in period *t* and she concedes to the reservation utility level s_b(t) with the corresponding indifference curve/surface *C*. The total number of offers that agent *b* plans to make is k_b. Assume $\boldsymbol{x}_\mathbf{s}^{t-1}$ is the best offer for agent *b* among all the offers proposed by agent *s* in period *t-1*, then agent *b* first chooses the offer which has the shortest distance to $\boldsymbol{x}_\mathbf{s}^{t-1}$ by

$$x_b^{t,1} = \arg\min_{x \in C} \left\| x - x_s^{t-1} \right\| \tag{3}$$

where ||***x-y***|| represents the distance from point ***x*** to ***y***. (Note that if there are several offers among the offers made by the opponent in the previous period that can contribute the same highest utility level, agent *b* can try all of them and find all the corresponding points following Formula (3). The agent then chooses the one which has the shortest distances.) If k_b >1, agent *b* will want to make multiple offers. We propose that agent *b* can take $\boldsymbol{x}_\mathrm{b}^{t,1}$ as the seed and randomly choose other k_b-1 offers from the neighborhood of $\boldsymbol{x}_\mathrm{b}^{t,1}$ by:

$$x_b^{t,m} = rand\{x \mid x \in C \text{ and } \left\| x - x_b^{t,1} \right\| \le \delta(t)\} \tag{4}$$

where $\delta(t) = \left\| x_b^{t,1} - x_s^{t-1} \right\|$, and 2≤m≤ k_b indexes an offer.

This approach can be very easily applied in an automated negotiation system. Certainly, other more complicated methods based on the negotiation history also can be designed. In the following, we provide a numerical example.

Numerical Example 1: There are two agents who negotiate two issues within maximal 20 periods in this example. The agents have the CES utility function as

$$U_b(x) = 1 - \left[0.2 x_1^3 + 0.8 x_2^3 \right]^{\frac{1}{3}};$$

$$U_s(x) = 1 - \left[0.7(1-x_1)^3 + 0.3(1-x_2)^3 \right]^{\frac{1}{3}};$$

where $\boldsymbol{x}$=($\boldsymbol{x}_1$, $\boldsymbol{x}_2$) denotes an offer with the value of the first (second) issue equal to $\boldsymbol{x}_1$ ($\boldsymbol{x}_2$).

The agents both follow a conceding strategy of $s_i(t) = 1-(1-0.2)\left(\frac{t}{20}\right)^{\frac{1}{0.8}}$ and the proposing strategy proposed above. The agents in this example only make one offer in each period. Figure 2 shows the negotiation procedure and the offers made by the two agents. The dashed curve is the Pareto-frontier. The square is the final agreement, the circles below the square are the offers made by agent b, and the circles above the square are made by agent s. The negotiation lasts 14 periods and the final agreement is reached at (0.6274, 0.3976) which has a distance 0.0410 to the Pareto frontier. From the figure, we can see that all the offers are very close to the Pareto-frontier.

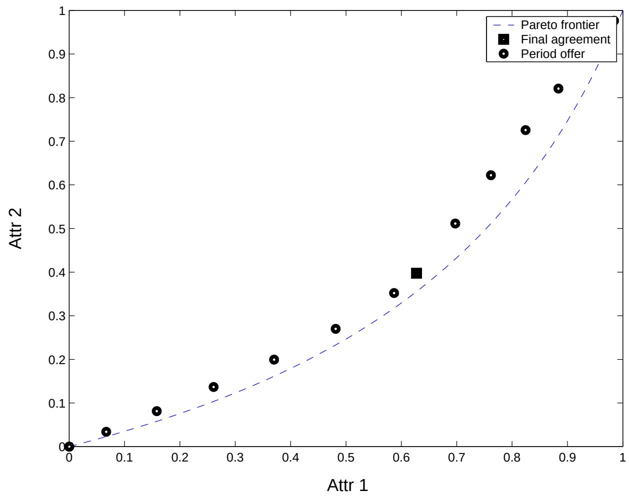

Figure 2: The negotiation procedure of example 1

3.3.3.2 The Pareto-optimal mediating mechanism

In the above subsection, we assume that the agents know their own utility functions. However, there are situations where the agents may not have an elicited utility function. A traditional approach to solve this problem is to first apply a preference elicitation procedure to characterize the agents' preferences and then they negotiate based on the utility functions. However, preference elicitation is a well-known difficult and time-consuming process, especially when the agents have complex preferences. In this subsection, we propose a *Pareto-optimal mediating mechanism*. Under this mechanism, the agents can negotiate multiple issues even if they do not have an explicit utility function on the issues before the negotiation. The only assumption we need is that the agents can judge the utility of a limited set of points. In the following, we first provide the structure of the mechanism, and then describe the mediating procedure in detail.

I. The structure of the mechanism. We adopt a non-biased mediator who can be implemented by an autonomous agent. This protocol first decomposes the n-dimensional negotiation space into a sequence of negotiation base lines. In each period, the mediator provides a (linear) negotiation base line and the proposing agent is required to propose a *base offer* on this line. Although the agents may not have explicit utility functions, the agents still can apply Equation (1) to determine their reservation utility in each period. Based on this reservation utility, an agent then can find a point on the negotiation base line which the agent think can contribute the utility equal to the reservation utility. This process is not difficult for an agent since the negotiation base line is a linear line and the utilities of the points on this line are monotonically increasing or decreasing based on the assumptions we have made. It is indeed similar as a step of a preference elicitation procedure.

Based on the base offer, the mediator works with the two agents to find a (near) Pareto optimal point that is mutually better than the base offer. This point if found is returned as the offer in this period to the responding agent. If the offer is rejected, the base line is updated and the negotiation proceeds to the next period. Such a procedure iterates until an agreement or the deadline is reached. The mechanism can be formally described as:

Step 0: We connect the two best offers (i.e, $\mathbf{0}^n$ and $\mathbf{1}^n$) for the agents by a line which acts as the first negotiation base line. Then, one agent is chosen (randomly or by some rule) to be the first mover in this negotiation.

Step 1: The proposer in the current period makes a proposal on the present negotiation base line. This proposal becomes the base offer for the mediator to find a Pareto optimal point in Step 2.

Step 2: The mediator works with the two agents to find a Pareto optimal point based on the base offer from Step 1 (see "**II. The mediating procedure**" for detail). The procedure goes to Step 3 when such a point is reached.

Step 3: The responder makes her decision in this period. If she accepts the offer found in Step 2, the procedure ends with this solution; else if she rejects it but the deadline is reached, the procedure ends with the negotiation breakdown; otherwise the mediator connects the point found in Step2 with the offer in the previous period by a line which becomes the new negotiation base line for the next period, the two agents exchange their roles and the process goes back to Step 1.

Figure 3 presents a depiction of the mechanism. In the figure, agent *b* makes a base offer $\boldsymbol{x}$ on the first negotiation base line *L1* which connects (0,0) and (1,1). Then, based on $\boldsymbol{x}$, the mediator works with the two agents and finds an offer $\boldsymbol{y}$ which is located on the Pareto-frontier (the black bold curve in the figure). But agent *s* rejects this offer, and then the negotiation base line is updated to the line *L2* that connects $\boldsymbol{y}$ and (1,1). Agent s makes a based offer ***p*** on L2 and the mediator suggests ***q***. The procedure repeats until an agreement is reached or the negotiation breaks down when the time deadline is reached.

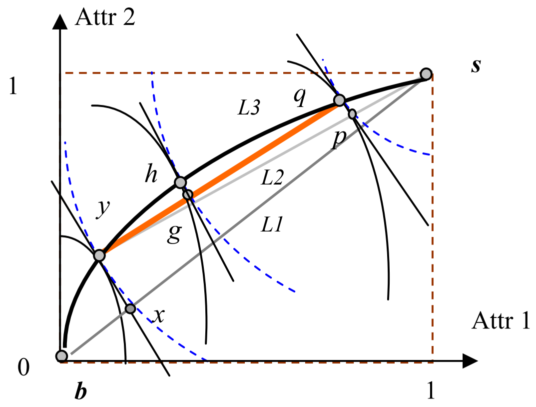

Figure 3: A depiction of the mediating negotiation protocol

II. The mediating procedure. The mediator's role is to find a Pareto optimal enhancement for a base offer in each negotiation period. For clarity of presentation, we describe the mediating procedure with a two-issue case. For a point ***x*** in the two-dimensional space, we use $x[1]$ to represent the value of issue 1 and $x[2]$ to represent the value of issue 2. Since to search for an exact Pareto optimal point is computationally intractable given that the information is incomplete, we apply an asymptotic approach with the following concepts:

Definition 2 Given a point $\boldsymbol{x}=(x[1],x[2])$ in the two-dimensional space, we call the range from the second issue value $x[2]-\varepsilon$ to $x[2]+\varepsilon$ (see range *A* in Figure 4) as the ε-range of point *x*.

Definition 3 A point ***x*** is an *ε-satisfying Pareto optimal solution* if one of the following two properties is satisfied: (1) there does not exist any point that is mutually better than ***x***

for both agents; (2) all the points mutually better than $\boldsymbol{x}$, if existing, are located in the ε-range of $\boldsymbol{x}$.

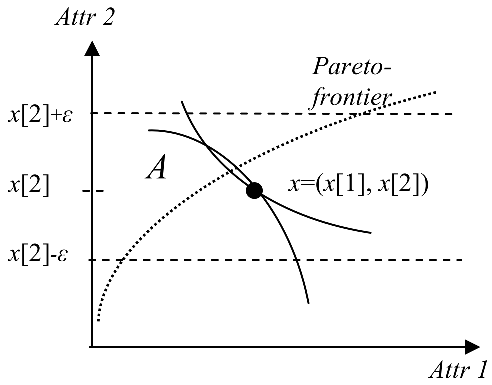

Figure 4: The definition of ε-satisfaction

From Definition 3, it is direct to see that when $\varepsilon \to 0$, an ε-satisfying solution $\boldsymbol{x}$ is Pareto optimal. This is because: when $\varepsilon \to 0$, $A \to \emptyset$ (see range A in Figure 4); we know that if $A = \emptyset$ (there is no point mutually better than $\boldsymbol{x}$), then $\boldsymbol{x}$ is Pareto optimal.

Definition 4 For a point $\boldsymbol{x}$, the range in the negotiation space that still needs to be searched for an ε-satisfying Pareto optimal solution, is called the *necessary range (NR) of* $\boldsymbol{x}$.

In the following, we use the value of the second issue to characterize the necessary range. At the beginning of the mediating procedure in each negotiation period, the lower bound of the necessary range (NRL) is 0 and the upper bound (NRU) is 1.

With the above definitions, the mediating procedure can be described as:

Step 0: The mediator sets the value of ε.

Step 1: Given the latest point $\boldsymbol{x}_n=(x_n[1],x_n[2])$ in the searching history[3], the mediator first checks whether it is ε-satisfying. To do this, the mediator can first apply a query to let each agent report a point where the value of its second issue is equal to $x_n[2]+\varepsilon$ and the agent is indifferent between this point and $\boldsymbol{x}_n$. Assume the mediator receives x_ε^{1} from agent *b* and x_ε^{2} from agent *s*. Then there can be two scenarios:

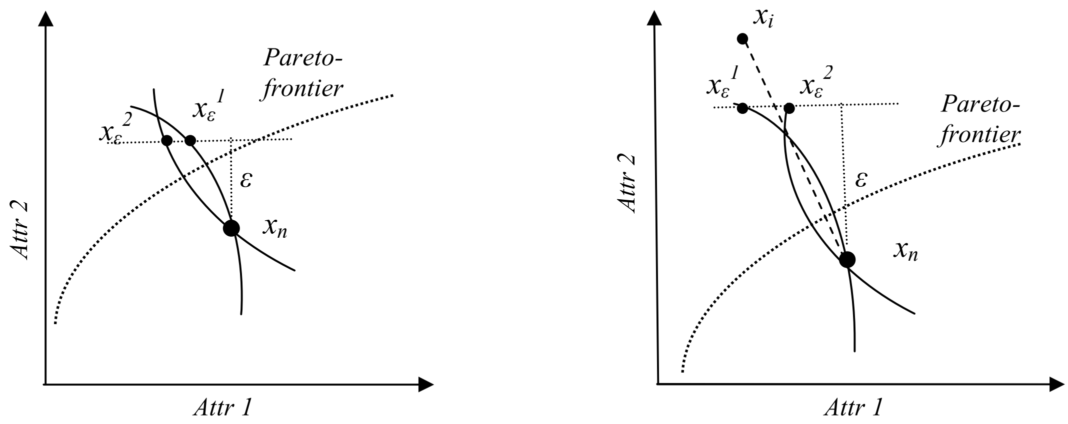

Figure 5. The scenarios of the results for a query

Scenario 1: $x_\varepsilon^{1}[1]>x_\varepsilon^{2}[1]$ (See the left subplot of Figure 5).

This scenario indicates that there do exist solutions mutually better than $\boldsymbol{x}_{\mathrm{n}}$ with the value of the second issue equal to $x_n[2]+\varepsilon$. Then, the procedure returns 'up' as the search direction from $\boldsymbol{x}_{\mathrm{n}}$ and goes to Step 2.

Scenario 2: $x_\varepsilon^{1}[1]\leq x_\varepsilon^{2}[1]$ (See the right subplot of Figure 5).

This scenario indicates that there does not exist a point such that the value of its second issue is no smaller than $x_n[2]+\varepsilon$ and it is mutually better than $\boldsymbol{x}_{\mathbf{n}}$. Then, the mediator needs to do a similar query to check the other direction, i.e. to let each agent report the point

[3] The searching history stores the series of points $(x_0,x_1,\ldots,x_{\mathrm{n}})$ that the mediator has found in the current negotiation period, where $\boldsymbol{x}_0$ is the base offer and $x_0 \prec x_1 \prec x_2 \ldots \prec x_{\mathrm{n}}$ for both agents. The sign '$\prec$' is defined as follows: $x_0 \prec x_1$, means that x_0 is less preferable than x_1 for both agents.

where the value of its second issue is equal to $x_n[2]-\varepsilon$ and the agent is indifferent between this point and $\boldsymbol{x}_n$. Then, the mediator checks whether there is a point mutually better than $\boldsymbol{x}_n$ and the value of its second issue is equal to $x_n[2]-\varepsilon$. If there is, the procedure returns 'down' as the search direction and goes to Step 2. If there is not such a point either, x_n is ε-satisfying by Definition 3. Thus, the mediator returns $\boldsymbol{x}_n$ as the offer for this negotiation period.

Step 2: Given the point $\boldsymbol{x}_n$ and the direction from Step 1, the mediator can apply Algorithm 1 (see Figure 6) to get the next point x_{n+1} that is mutually better than $\boldsymbol{x}_n$ for both agents. In this algorithm, the mediator first needs to get the necessary range, which is characterized by the lower bound, *NRL,* and upper bound, *NRU,* of the necessary range. *NRL* and *NRU* are global variables and are updated after each query the mediator processes in this negotiation period. The function '*rankquery*($\boldsymbol{x}_n$,*yDim*)' in this algorithm is to let agents report the indifferent point of x_n and the value of its second issue equal to *yDim*. After the mediator gets $\boldsymbol{x}_{n+1}$, the procedure goes back to Step 1 and repeats until an ε-satisfying point is reached. This ε-satisfying point is returned as the offer of this period, and the responder makes her decision whether to accept it or not.

Such a mediating procedure indeed mimics a preference elicitation process. However, this procedure is processed along with the negotiation and only the points in some local areas in the negotiation space are queried. The procedure saves a lot of computational effort compared to a traditional preference elicitation process, since it avoids characterizing the agents' preferences on the whole negotiation space.

For the cases where there are more than two issues in the negotiation, the mediator can process the above procedure in a sequence of two-dimensional spaces. For instance, if there are three issues: *attr* 1, *attr* 2 and *attr* 3, given a reference point $\boldsymbol{x}_0$, the mediator can first process the above procedure in the (*attr* 1×*attr* 2)-space with the value of *attr* 3 fixed at $x_0[3]$. After a point, say $\boldsymbol{x}_1$, is found which is ε-satisfying in the (*attr* 1×*attr* 2)-space, the mediator begins to do the same procedure in the (*attr* 2×*attr* 3)-with the value of *attr* 1 fixed at $x_1[1]$, and later in the (*attr* 3×*attr* 1) -space. The mediator can repeat this process until the point reached is ε-satisfying in each two-dimensional space. In the following, we provide a numerical example.

```
Inputs: x_n, dir, NRL, NRU
Outputs: x_{n+1}, NRL, NRU
Begin:
    if (dir= ='up')
       yDim=NRU;
       NRL=x_n[2];
    elseif (dir= ='down')
       yDim=NRL;
       NRU=x_n[2];
    end
    yDim=(yDim+x_n[2])/2;
    x^1_indf=rankquery_1(x_n,yDim);
    x^2_indf =rankquery_2(x_n,yDim);
    while (x^1_indf[1]≤ x^2_indf [1])
      if (dir= ='up')
         NRU=yDim;
      elseif (dir= ='down')
         NRL=yDim;
      end
      yDim=(yDim+x_n[2])/2;
      x^1_indf=rankquery_1(x_n,yDim);
      x^2_indf =rankquery_2(x_n,yDim);
    end
    x_{n+1}=((x^1_indf [1]+x^2_indf [1])/2, yDim);
    return x_{n+1}, NRL, NRU;
End
```

Figure 6. Algorithm 1

Numerical Example 2: There are two agents who negotiate three issues within maximal 20 periods in this example. The agents' preferences can be represented by the following quadratic utility functions[4]:

$$U_b(x) = 1 - \left[0.2x_1^{\,2} + 0.6x_2^{\,2} + 0.2x_3^{\,2}\right];$$
$$U_s(x) = 1 - \left[0.6(1-x_1)^2 + 0.2(1-x_2)^2 + 0.2(1-x_3)^2\right].$$

Agent *b* follows a conceding strategy: $s_b(t) = 1-(1-0.2)\left(\frac{t}{20}\right)^{\frac{1}{0.4}}$ and agent s follows a conceding strategy: $s_s(t) = 1-(1-0.3)\left(\frac{t}{20}\right)^{\frac{1}{2}}$. The accuracy parameter ε is set to 0.01. The results are shown in Table 1 and Figure 8. From Table 1, we see the maximal number of queries processed in a negotiation period is 70, but it decreases quickly, as the base line is turning closer to the Pareto-frontier (due to the update mechanism processed in every period). During the last negotiation period, only 11 queries are processed. The offers made in each period are very close to the Pareto-frontier which are ε-satisfying Pareto optimal.

For more numerical experiments related to the above two proposing mechanisms, we refer the readers to Lai et al. (2006a, 2006b).

[4] By saying agents' preferences can be characterized by some utility functions, we do not mean that agents know their utility functions explicitly on the whole negotiation space but mean that given a limited number of points, agents can compare them and say which one they prefer.

Periods	Proposals	Period offers	Queries
1	(0,0,0)	(0,0,0)	0
2	(0.604,0.604,0.604)	(0.848,0.383,0.651)	70
3	(0.059,0.027,0.045)	(0.077,0.009,0.027)	27
4	(0.591,0.259,0.443)	(0.674,0.187,0.408)	50
5	(0.214,0.050,0.115)	(0.233,0.033,0.092)	35
6	(0.562,0.148,0.328)	(0.580,0.133,0.315)	24
7	(0.377,0.074,0.184)	(0.386,0.065,0.173)	26
8	(0.506,0.107,0.261)	**(0.510,0.104,0.257)**	11

Table 1. The negotiation data of Experiment 2

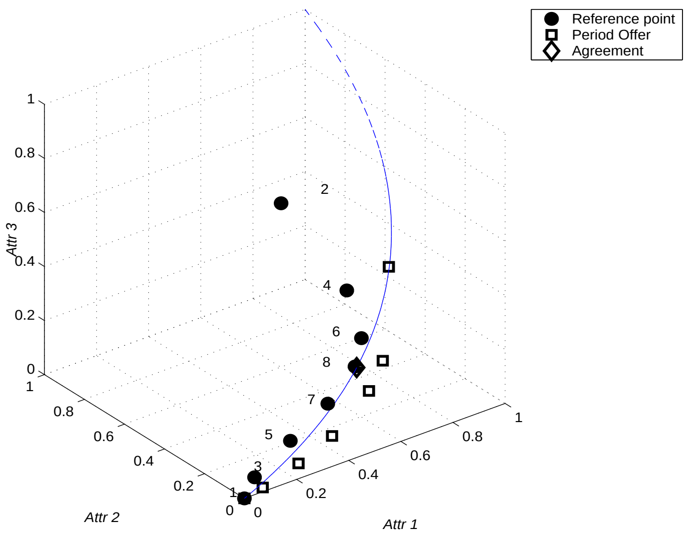

Figure 8. The negotiation procedure of Experiment 2

4. Conclusions and future work

Automated negotiation is an important and useful mechanism for both the negotiations between human participants and those in the world of autonomous agents. To implement automated negotiation, applicable negotiation theories are essential. The existing research has well studied single-attribute negotiation, but not much addresses multi-attribute negotiation.

This paper reviews the existing research on multi-attribute negotiation and points out three key issues: incomplete information, Pareto-optimality and tractability, which need to be considered to bridge the gap between the existing work and applicable automated multi-attribute negotiation systems. This paper then presents a modeling framework to incorporate these issues. This framework divides a negotiation strategy into three parts: conceding, responding and proposing. Based on this structure, the agents can apply the existing conceding and responding strategies developed for single-attribute negotiation problems. This framework applies two generic proposing mechanisms for different situations. When the agents have their own utility functions but not their opponent's, the agents can use the shortest-distance proposing mechanism to select offers in each period. When the agents do not have their own utility functions elicited, the agents can apply the Pareto-optimal mediating mechanism to negotiate without the need to do preference elicitation exhaustively. The mediating mechanism also simplifies the agents' decisions by decomposing the original *n*-dimensional space into a series of linear negotiation base lines.

However, there are still some broader issues that need to be studied further. First, the existing research including ours usually focuses on the negotiations with either continuously-valued or binary-valued issues. But in practice, negotiations can include both types of issues. The existing theories need to be extended further to consider such situations. Second, the existing research has seldom looked at the situations where the agents even do not have their own utility functions elicited before the negotiation. Although we address this problem, we do need a non-biased mediator available. In practice, there may exist situations where it is not easy to include a mediator. Therefore, some further work is needed to develop completely decentralized mechanisms that can be applied in such situations. Finally, most of the conceding strategies developed in the existing research assume that the agents have fixed ultimate reservation utilities for the negotiation. However, in practice, there may exist situations where the agents receive outside options during a negotiation. To adopt outside options in the multi-attribute negotiation framework is valuable to be studied.

Acknowledgement

This research has been sponsored in part by AFOSR Grant F49620-01-1-0542 and in part by an IBM subcontract to CMU of ARL ITA Grant FA8650-06-C-7606.

References

Bac, M. and H. Raff. (1996). "Issue-by-Issue Negotiations: the Role of Information and Time Preference," Games and Economic Behavior 13, 125-134.

Blum, A., J. Jackson, T. Sandholm, and M. Zinkevich. (2004). "Preference Elicitation and Query Learning," Journal of Machine Learning Research 5, 649-667.

Busch, L.-A. and I. J. Horstmann. (1997). "A Comment on Issue-by-Issue Negotiations," Games and Economic Behavior, 19:144-148.

Busch, L.-A. and I. J. Horstmann. (1999a). "Endogenous Incomplete Contracts: A Bargaining Approach," Canadian Journal of Economics 32, 956-975.

Busch, L.-A. and I. J. Horstmann. (1999b). "Signaling via an Agenda in Multi-issue Bargaining with Incomplete Information," Economic Theory 13, 561-575.

Busch, L.-A. and I. J. Horstmann. (2002). "The Game of Negotiations: Ordering Issues and Implementing Agreements," Games and Economic Behavior 41, 169-191.

Busch, L.-A. and I. J. Horstmann. (1997). "Bargaining Frictions, Bargaining Procedures and Implied Costs in Multiple-Issue Bargaining," Economica 64, 669-680.

Chen, L. and Pu, P. (2004). *Survey of Preference Elicitation Methods.* EPFL Technical Report IC/2004/67, Switzerland.

Chevaleyre, Y., U. Endriss, J. Lang and N. Maudet. (2005). "Negotiating over Small Bundles of Resources," In *the proceeding of the 4th Int. Conf. on Autonomous Agents and Multi-Agent Systems*, 296-302.

Coehoorn, R. M. and N. R. Jennings. (2004). "Learning an opponent's preferences to make effective multi-issue negotiation tradeoffs," In *the proceeding of the 6th Int. Conf. on E-Commerce*, 59-68.

Ehtamo, H., R. P. Hamalainen, P. Heiskanen, J. Teich, M. Verkama, and S. Zionts. (1999). "Generating Pareto Solutions in a Two-Party Setting: Constraint Proposal Methods," Management Science 45, 1697-1709.

Faratin, P., C. Sierra, and N. R. Jennings. (1998). "Negotiation decision functions for autonomous agents," International Journal of Robotics and Autonomous Systems 24, 159-182.

Faratin, P., C. Sierra, and N. R. Jennings. (2000). "Using similarity Criteria to Make Negotiation Trade-Offs," In *the Proceeding of the 4th International Conference on Multi-Agent Systems,* 119-126.

Faratin, P., C. Sierra, and N. R. Jennings. (2002) "Using Similarity Criteria to Make Issue Trade-offs in Automated Negotiations," Artificial Intelligence 142, 205-237.

Fatima, S., M. J. Wooldridge, and N. R. Jennings. (2004a). "An Agenda-Based Framework for Multi-Issue Negotiation," Artificial Intelligence 152, 1-45.

Fatima, S., M. J. Wooldridge, and N. R. Jennings. (2004b). "Optimal negotiation of multiple issues in incomplete information settings," In *the Proceeding of the 3rd International Conference on Autonomous Agents and Multi-Agent Systems*, 1080-1087.

Fershtman, C. (1990). "The Importance of the Agenda in Bargaining," Games and Economic Behavior 2, 224-238.

Fershtman, C. (2000). "A Note on Multi-issue Two-Sided Bargaining: Bilateral Procedures," Games and Economic Behavior 30, 216-227.

Hanson, J., G. Tesauro, J. Kephart, E. Sniblesome. (2003). "Multi-Agent Implementation of Asymmetric Protocol for Bilateral Negotiations," In *the Proceeding of the ACM Conference on Electronic Commerce 2003*, 224-225.

Hudson, B. and T. Sandholm. (2004). "Effectiveness of Query Types and Policies for Preference Elicitation in Combinatorial Auctions," In *the Proceeding of the 3rd International Conference on Autonomous Agents and Multi-Agent Systems*, 386-393.

In, Y. and R. Serrano. (2003). "Agenda Restrictions in Multi-Issue Bargaining (II): Unrestricted Agendas," Economics Letters 79, 325-331.

In, Y. and R. Serrano. (2004). "Agenda Restrictions in Multi-Issue Bargaining," Journal of Economic Behavior and Organization 53, 385-399.

Inderst, R. (2000). "Multi-issue Bargaining with Endogenous Agenda," Games and Economic Behavior 30, 64-82.

Kalai, E. (1977). "Proportional Solutions to Bargaining Situations: Intertemporal Utility Comparisons," Econometrica 45, 1623-1630.

Klein, M., P. Faratin, H. Sayama, and Y. Bar-Yam. (2003). "Negotiating Complex Contracts," Group Decision and Negotiation 12, 111-125.

Lai, G., C. Li, K. Sycara, and J. Giampapa. (2004). *Literature review of multi-attribute negotiations,* Technical Report, CMU-RI-TR-04-66, Carnegie Mellon University, Pittsburgh, USA.

Lai, G., C. Li and K. Sycara. (2006a). "Efficient Multi-attribute Negotiation with Incomplete Information," Group Decision and Negotiation 15, 511-528.

Lai, G., K. Sycara and C. Li. (2006b). "A Decentralized Model for Multi-attribute Negotiations," In *the Proceedings of the Eighth International Conference on Electronic Commerce (ICEC)*, 3-10.

Lang, K. and R. W. Rosenthal. (2001). "Bargaining Piecemeal or All at Once," The Economic Journal 111, 526-540.

Li, C., and G. Tesauro (2003). "A strategic Decision Model for Multi-attribute Bilateral Negotiation with Alternating Offers," In *the Proceeding of the ACM Conference on Electronic Commerce 2003*, 208-209.

Li, C., J. Giampapa, and K. Sycara. (2006). "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), 31-44.

Luo, X., N. R. Jennings, N. Shadbolt, H. Leung and J. H. Lee. (2003). "A Fuzzy Constraint Based Model for Bilateral Multi-issue Negotiations in Semi-competitive Environments," Artificial Intelligence 148, 53-102.

Mas-Colell, A., M. D. Whinston and J. R. Green. (1995). *Microeconomic Theory*. Oxford University Press, New York.

Nash, J. (1950). "The Bargaining Problem," Econometrica 18, 155-162.

Robu, V., D. Somefun and J. La Poutre. (2005). “Modeling Complex Multi-issue Negotiations Using Utility Graphs,” In *the proceeding of the 4th Int. Conf. on Autonomous Agents and Multi-Agent Systems*, 280-287.

Rubinstein, A. (1982). “Perfect Equilibrium in a Bargaining Model,” Econometrica 50, 97-109.

Santi, P., V. Conitzer, and T. Sandholm. (2004). “Towards a Characterization of Polynomial Preference Elicitation with Value Queries in Combinatorial Auctions,” In *the Proceedings of the Conference on Computational Learning Theory,* 1-16.

Sycara, K. (1990a). “Persuasive Argumentation in Negotiation,” Theory and Decision 28, 203-242.

Sycara, K. (1990b). “Negotiation Planning: An AI Approach,” European Journal of Operational Research 46, 216-234.

Sycara, K. (1991). “Problem Restructuring in Negotiation,” Management Science 37, 1248-1268.

Weinberger, C. J. (2000). “Selective Acceptance and Inefficiency in a Two-Issue Complete Information Bargaining Game,” Games and Economic Behavior 31, 262-293.

Zeng, D. and K. Sycara. (1998). “Bayesian learning in negotiation,” International Journal of Human-Computer Studies 48, 125-141.