

2019-05

Multi-Criteria Multi-Participant Automated Negotiation: Belief Propagation-based Proposal Preparation and Real Time Opponent Learning

Eshragh, Faezeh

Eshragh, F. (2019). Multi-Criteria Multi-Participant Automated Negotiation: Belief Propagation-based Proposal Preparation and Real Time Opponent Learning (Doctoral thesis, University of Calgary, Calgary, Canada). Retrieved from <https://prism.ucalgary.ca>.
<http://hdl.handle.net/1880/110438>

Downloaded from PRISM Repository, University of Calgary

UNIVERSITY OF CALGARY

Multi-Criteria Multi-Participant Automated Negotiation:

Belief Propagation-based Proposal Preparation and Real Time Opponent Learning

by

Faezeh Eshragh

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE
DEGREE OF DOCTOR OF PHILOSOPHY

GRADUATE PROGRAM IN GEOMATICS ENGINEERING

CALGARY, ALBERTA

MAY, 2019

© Faezeh Eshragh 2019

Abstract

Automated negotiation has received considerable attention in the past few decades as a computer tool for modeling human negotiations. The aim of automated negotiation is capturing the model of interactions during the negotiation process and improving the efficiency and quality of real-world negotiations. Due to the complexity of negotiations, there are many challenges in modeling different aspects of this process. One of the important types of negotiation, which is the focus of this thesis, is multi-issue multi-participant argumentation-based negotiation. In such negotiations, several participants with different viewpoints and perspectives negotiate over several criteria. The involved parties in these negotiations exchange proposals (a set of values assigned to negotiation issues) and receive their opponents' evaluation of the offered proposals as well as possible arguments. The primary goal of the negotiation is finding a solution that can satisfy all the involved parties. However, in multi-issue multi-participant negotiations, finding such a solution can be quite challenging because: 1- participants have different and, sometimes, conflicting preferences about the negotiation issues; and 2- these preferences are not usually revealed to others. The higher the number of negotiation issues (i.e., the dimensions of the search space for a satisfactory solution), the higher the number of unknown preferences and therefore, the harder to reach an agreement. Therefore, the negotiation process can take a long time before approaching a possible agreement.

The current thesis studies two critical aspects of automated negotiation: proposal preparation and opponent modeling. The order of the offered proposals in consecutive rounds of the negotiation directly impacts the pace of reaching an agreement. Therefore, selecting the right proposal for each round based on the interactions in the previous rounds is the key to effective negotiation. In this thesis, a novel proposal-preparation solution is proposed. It represents the negotiation issues and participants' preferences via a graphical model and applies belief propagation to optimize this graph, the output of which is a proposal to offer to the participants. The thesis also discusses the problem of unknown preferences of the participants in this negotiation context. A re-

cursive Bayesian filtering algorithm is proposed to learn/estimate the preferences of the opponents only through the limited information they exchange as the negotiation proceeds. The proposed approaches are then applied to two case studies to investigate their impact on the efficiency of the negotiation process. The experimental results show that using the presented proposal preparation and opponent modeling techniques, the efficiency of the negotiation process is increased by up to 85% in both case studies.

Acknowledgements

I would like to thank Dr. Behrouz Far for his guidance while working towards this dissertation. His profound knowledge and professional attitude always motivated me. I also would like to express my sincere gratitude to my co-supervisor Dr. Mozhdeh Shahbazi for her invaluable assistance, helpful advises, and calm reassurances. This research would not have been where it is without her insightful ideas and constant support.

Furthermore, I would like to thank my supervisory committee members Dr. Scott Heckbert and Dr. Steve Liang for their valuable feedback and their kind support.

I would like to offer special thanks to Dr. Danielle Marceau, who, although no longer with us, continues to inspire me by her wisdom, experience, and dedication. I am also grateful for the financial support as this thesis was partially supported by the funds available to the Schulich Chair in GeoSpatial Information Systems and Environmental Modeling held by Dr. Danielle Marceau. This study was also funded by a Discovery grant awarded by the Natural Sciences and Engineering Research Council of Canada (NSERC).

Finally, yet most importantly, heartfelt thanks to my husband for his tremendous support and patience, my parents for their unconditional trust and encouragement and my little son for always cheering me up even at the toughest moments of my research.

Table of Contents

| | |
|---|------|
| Abstract | ii |
| Acknowledgements | iv |
| Table of Contents | v |
| List of Tables | vii |
| List of Figures | viii |
| List of Symbols | x |
| 1 Introduction | 1 |
| 1.1 Automated Negotiation | 1 |
| 1.1.1 Applications and Challenges | 2 |
| 1.1.2 Problem statement | 5 |
| 1.1.3 Objectives | 7 |
| 1.1.4 Main Contributions | 8 |
| 1.1.5 Assumptions and limitations | 9 |
| 1.1.6 Thesis structure | 9 |
| 2 Belief Propagation-based Proposal Preparation in Automated Negotiation | 10 |
| 2.1 Abstract | 13 |
| 2.2 Introduction | 13 |
| 2.2.1 Related work | 15 |
| 2.2.2 Objectives and contributions | 18 |
| 2.3 Automated Negotiation | 19 |
| 2.3.1 Agent-based model | 21 |
| 2.3.2 Problem statement | 21 |
| 2.3.3 Belief propagation-based proposal preparation (BPPP) | 22 |
| 2.3.4 Z-scoring approach for selecting among the prepared proposals | 25 |
| 2.3.5 Argument handling model | 27 |
| 2.3.6 Simple frequency-based argument processing | 28 |
| 2.4 Experiments | 29 |
| 2.4.1 Case study A: energy-system planning in Alberta | 30 |
| 2.4.2 Data preparation and implementation | 31 |
| 2.4.3 Results and discussions | 35 |
| 2.4.4 Case study B: King County House Sales | 39 |
| 2.5 Conclusion and Future Work | 42 |
| 3 Comprehensive Review of Machine Learning Approaches in Automated Negotiation: Examples in Environmental Resource Management | 44 |
| 3.1 Abstract | 46 |
| 3.2 Introduction | 46 |
| 3.3 Automated negotiation | 49 |
| 3.4 Machine learning techniques in automated negotiation | 53 |
| 3.4.1 Bayesian learning | 54 |
| 3.4.2 Reinforcement learning | 56 |
| 3.4.3 Evolutionary learning | 57 |
| 3.4.4 Artificial neural network | 59 |

| | | |
|-----|---|-----|
| 3.5 | Automated negotiation in environmental resource management | 60 |
| 3.6 | Conclusion | 66 |
| 4 | Real-time Opponent Learning in Automated Negotiation using Recursive Bayesian Filtering | 69 |
| 4.1 | Abstract | 72 |
| 4.2 | Introduction | 72 |
| | 4.2.1 Objective and Contributions | 78 |
| | 4.2.2 Assumptions | 79 |
| 4.3 | Related Work | 79 |
| 4.4 | Methodology | 85 |
| | 4.4.1 Multi-Agent System | 85 |
| | 4.4.2 Problem Statement | 86 |
| | 4.4.3 Proposal Evaluation Model | 86 |
| | 4.4.4 Unscented Particle Filtering | 90 |
| 4.5 | Experiments | 101 |
| | 4.5.1 Case study A: Energy-System Planning in Alberta | 102 |
| | 4.5.2 Case study B: King County House Sales | 112 |
| 4.6 | Conclusion and Future Work | 122 |
| 5 | General Conclusions and Discussions | 124 |
| 5.1 | Summary and Conclusions | 124 |
| 5.2 | Research Perspectives and Future Work | 127 |
| 6 | Appendix | 129 |

List of Tables

| | | |
|-----|---|-----|
| 2.1 | Significant stakeholders in the project on energy-system planning | 31 |
| 3.1 | Learning techniques applied in automated negotiation | 61 |
| 3.2 | Applications of automated negotiation in environmental resource management studies | 66 |
| 3.3 | Suitability of learning approaches in automated negotiation for environmental resource management | 66 |
| 4.1 | Significant stakeholders in the project on energy-system planning; Source: (Eshragh et al., 2018) | 103 |

List of Figures and Illustrations

| | | |
|------|---|-----|
| 2.1 | Proposal preparation flowchart | 11 |
| 2.2 | Negotiation Protocol | 20 |
| 2.3 | A schematic representation of the proposed negotiation system | 22 |
| 2.4 | The method of calculating the z-score for each alternative proposal | 26 |
| 2.5 | The study area located near the Slave River and Forth Smith city at the border of Alberta and the Northwest Territories | 32 |
| 2.6 | Selected routes in the data preparation phase | 34 |
| 2.7 | Effect of number of Ranges on rounds of negotiation | 35 |
| 2.8 | Effect of parameter η on rounds of negotiation | 36 |
| 2.9 | The attribute values of the offered proposals in two different settings- the case of 15 states; (a) First-nation values; (b) wildlife values; (c) forest values; (d) wetland values | 37 |
| 2.10 | Disagreement distances in negotiation over the first case study with and without using the BPPP approach. The dashed lines represent the end of negotiation. | 38 |
| 2.11 | Disagreement distances for a sequence of offered proposals in the Utility-based approach vs. the BPPP approach (a) With arguments (b) Without arguments. The dashed lines represent the end of negotiation. | 40 |
| 2.12 | Disagreement distances in negotiation over the second case study with and without AH model. The dashed lines represent the end of negotiation. | 41 |
| 2.13 | Disagreement distances for a sequence of offered proposals in the Utility-based approach vs. the BPPP approach. The dashed lines represent the end of negotiation. | 42 |
| 4.1 | Opponent modeling and learning flowchart | 70 |
| 4.2 | Negotiation Process | 77 |
| 4.3 | The system flowchart | 85 |
| 4.4 | Objective function for a) less-is-better attributes b) more-is-better attributes | 87 |
| 4.5 | Attribute and proposal score functions a) Score function of the Year-Built attribute b) Score function of the Price attribute c) Score function of the proposal | 88 |
| 4.6 | UPF process in each negotiation round | 92 |
| 4.7 | Update and correct a particle using stakeholders' feedback | 92 |
| 4.8 | Transition function f pseudo-code | 95 |
| 4.9 | Mapping uniformly selected index SI to the cumulative distribution domain | 99 |
| 4.10 | Resampling algorithm pseudo-code | 100 |
| 4.11 | The study area located near the Slave River and Forth Smith city at the border of Alberta and the Northwest Territories; Source: (Eshragh et al., 2018) | 104 |
| 4.12 | Selected routes in the data preparation phase; Source: (Eshragh et al., 2018) | 105 |
| 4.13 | Effect of λ_p on the negotiation efficiency | 107 |
| 4.14 | Number of negotiation rounds using UPF and frequency-based approaches | 109 |
| 4.15 | Estimated upper bound preference limit for FN attribute using UPF and frequency-based approaches | 111 |
| 4.16 | Number of negotiation rounds using UPF and frequency-based approaches with and without arguments | 111 |

| | | |
|------|---|-----|
| 4.17 | The study area-King County, Washington, US | 113 |
| 4.18 | Houses for sale in King County | 114 |
| 4.19 | Developed website for collecting users' preference profiles; (a) Acquiring user's criteria; (b) Offering a property; (c) Representing attributes of the proposed property; (d) Acquiring user's feedback about the offered property | 115 |
| 4.20 | Average improvement using UPF approach with arguments over Frequency-based approach with arguments | 116 |
| 4.21 | Average improvement using UPF approach without arguments over Frequency-based approach without arguments | 116 |
| 4.22 | Quality of the final agreement using UPF approach and Frequency-based approach | 117 |
| 4.23 | Average improvement with different initial state estimations for users U1, U4, U6, and U14 | 118 |
| 4.24 | Estimated upper preference limit for price attribute using UPF approach and Frequency-based Approach | 119 |
| 4.25 | Proposed values for a) Price b)Year-Built attributes using the UPF estimations of stakeholder's criteria | 120 |
| 4.26 | Proposed values for a) Price b)Year-Built attributes using the Frequency-based estimations of stakeholder's criteria | 121 |
| 4.27 | Negotiation results with UPF approach with and without BP proposal preparation . | 121 |

List of Abbreviations

| Abbreviation | Definition |
|--------------|---|
| ABM | Agent-Based Model |
| ABMI | Alberta Biodiversity Monitoring Institute |
| ABN | Argumentation Based Negotiation |
| AEP | Alberta Environment and Parks |
| AH | Argument Handling |
| AI | Artificial Intelligence |
| ANN | Artificial Neural Network |
| BL | Bayesian Learning |
| BP | Belief Propagation |
| BPPP | Belief Propagation-based Proposal Preparation |
| CP-Net | Preference Network |
| EL | Evolutionary Learning |
| EP | Evolutionary Programming |
| ES | Evolution Strategies |
| FN | First-Nations |
| FSM | Finite State Machine |
| GA | Genetic Algorithm |
| GIS | Geographic Information System |
| LCP | Least Cost Path |
| MAS | Multi-Agent System |
| MCMC | Markov chain Monte Carlo |
| ML | Machine Learning |
| MRF | Markov Random Field |

| | |
|-----|-----------------------------------|
| PDF | Probability Distribution Function |
| RBF | Recursive Bayesian Filtering |
| RL | Reinforcement Learning |
| SIR | Sampling-Importance Resampling |
| SIS | Sequential Importance Sampling |
| UPF | Unscented Particle Filtering |

Chapter 1

Introduction

1.1 Automated Negotiation

Negotiation is one of the most common means of resolving conflicts in social interactions (Van Kleef et al., 2006). It can be defined as a discussion between two or more parties with conflicting interests aiming to reach an agreement (Pruitt and Carnevale, 1993). The involved participants may be individuals or groups of people who negotiate over single or multiple issues simultaneously. The agreement, which might be a mutually acceptable deal, allocation of resources or rules of behavior, has to satisfy all the participants to some extent. In the past few decades, negotiation has been studied from different perspectives such as psychology (Pruitt and Carnevale, 1993), economy (Kreps, 1990), and computer science (Jennings et al., 2001). The main aim of such studies is to understand the complicated nature of the negotiation process and make it more efficient, where negotiation efficiency is primarily measured in terms of the number of rounds of negotiation before reaching an agreement.

One of the critical scientific approaches in this domain, which has its roots in Artificial Intelligence (AI), is automated negotiation. It is a distributed search in the space of potential agreements, facilitated by an agent-based model (ABM). An ABM consists of a set of intelligent elements, called agents, designed to mimic human behaviors. Agents are autonomous entities with specific goals and. They can make decisions, based on some predefined or learned rules, and act towards achieving their objectives. Agents may interact with each other and with their environment using observation through sensors and consequent actuators. Through these interactions agents can improve their knowledge of the environment and use that knowledge to achieve their goal.

Agent-based models are very similar to multi-agent systems (MAS) in the way they are defined. However, their objectives and usages are quite different. While ABM is mainly used for studying

the collective behavior of agents, specifically in natural research area, the MAS objective is more about designing the agents with the aim of solving specific practical or engineering problems.

Automated negotiation consists of three main components: i) negotiation protocol, ii) negotiation object, and iii) negotiation strategy (Lomuscio et al., 2001). The *negotiation protocol* defines the set of rules governing the interactions between the agents. The *negotiation object* corresponds to the range of issues over which the negotiation happens and finally, the *agents' strategy*, is used by agents to act according to the negotiation protocol to reach a satisfactory agreement. Chapter 3 discusses these components in more details.

Three main approaches have been employed in automated negotiation: i) game theory, ii) heuristic approach and iii) argumentation-based approach (Jennings et al., 2001). The details of these approaches can be found in section 3.3 of the third chapter. Among these approaches the argumentation-based approach is followed by this research as it is a better fit to the problem under study.

1.1.1 Applications and Challenges

Automated negotiation has been first employed in the context of distributed artificial intelligence in the 1980s (Davis and Smith, 1983; Malone et al., 1983) where agents interact and negotiate to solve problems in a distributed way. It, then, has received considerable attention in other domains, such as supply chain management (Fink, 2006; Van Der Putten et al., 2006), political studies (Aragonès and Dellunde, 2008), e-commerce (Dworman et al., 1996; Zeng and Sycara, 1998; Faratin et al., 2002; Ramchurn et al., 2007; McBurney et al., 2002; Huang et al., 2010; Cao et al., 2015), policy disputes over natural resources (Thoyer et al., 2008; Carraro et al., 2007; Eshragh et al., 2017; Rodriguez-Fernandez et al., 2019), business process management (Jennings et al., 2000; Norman et al., 1996), grid computing (Kraus, 2001b; Wolski et al., 2003), load balancing (Fatima and Wooldridge, 2001), M-services (Paurobally et al., 2003), data, task, and resource allocation (Kraus, 2001b; Bigham and Du, 2003; Lin et al., 2006a), event scheduling (Modi et al., 2005; Hossain, 2012), crowdsourcing (Azaria et al., 2012) and energy exchange in remote communities (Alam

et al., 2013).

The negotiation scenario may vary based on the application context. In this thesis, our focus is on multi-issue, multi-participant negotiations where a one-to-many type of negotiation occurs between several negotiators. That is, one of the participants plays the proposer role and prepares and offers proposals to others. The rest of the participants are stakeholders who evaluate the received proposals based on their preferences and send their evaluation results (decisions) to the proposer. In this thesis, the argumentation-based approach of automated negotiation is used. Thus, the stakeholders are able to provide arguments along with their decisions. The proposer agent readjusts its understanding of the stakeholders using the received evaluations/arguments and prepares the next proposal accordingly. The new proposal is, then, sent to the stakeholders, and the process continues till all the parties agree over the offered proposal.

Automating the explained negotiation scenario requires two principal modules: i) Proposal preparation, and ii) Opponent modeling.

Proposal-preparation module

This module is in charge of selecting a proposal (i.e., a point in the space of possible solutions) at each round of negotiation. The order of the offered proposals during the negotiation impacts the satisfaction of the involved parties as well as the pace of the negotiation. It also affects the outcome of the negotiation. It is, therefore, crucial for the proposer agent to take every participant into account and come up with a proposal that is beneficial to itself while it has a high likelihood to be accepted by others. In multi-issue negotiations, the dimensions and size of the space of possible solutions can be considerably large. Thus, it is quite challenging to find such a mutually acceptable proposal at a reasonable time.

Several studies about proposal-preparation strategies have been conducted in the literature. These approaches range from time-dependent concession making strategies to more advanced ones based on the dynamic behavior of the opponents. For example, in (Faratin et al., 1998) the authors presented a set of tactics to determine the value of a proposal attribute using criteria such as time,

resources and previous offers and counter-offers. The tactics can be time-dependent, imitative or a combination of both. In the time-dependent tactics, the value of the negotiation attribute is initialized with a constant, and it is increased/decreased (based on the attribute type) as the time passes. For example, in “Boulware” tactics (Faratin et al., 1998), an exponential or polynomial function with a small concession-making rate is used for readjusting the values of the proposal attributes. This tactic provides values more in favor of the proposer, and therefore, the negotiation can proceed slowly. Another time-dependent tactic discussed in (Faratin et al., 1998) is the “conceder” tactics, in which the concession-making rate is significant. Therefore, the proposer concedes a lot in each round no matter what the opponent preferences are. Behavior-dependent tactics, on the other hand, determine the values of proposal attributes based on the counter-offers the agent receives from its opponents. Relative Tit-for-Tat, random absolute Tit-for-Tat and averaged Tit-for-Tat are some of these behavior-dependent tactics that work based on percentage or average of percentages of changes in opponent’s counter-offers (Faratin et al., 1998). These tactics are therefore inapplicable in negotiation scenarios in which one agent is in charge of offering proposals, and the rest can only evaluate those proposals. The other problem with this approach is that the negotiation doesn’t proceed if one of the parties stops making concessions.

Cheng et. al. (Cheng et al., 2006) proposed a fuzzy inference system in which several suppliers and buyers (the negotiating agents) submit offers and a set of matching agents matches supplies to buyers by finding the most similar proposals to each negotiator. The matching process is based on similarity measures, which are defined as a function of issue similarities and the weight of the issues. That is, the negotiating agents need to reveal their preferences on weights and even their utility functions to the matching agents. Having all this information, the matching agents try to find the closest suppliers and buyers and the negotiation will continue among the participants. For generating a new offer, each agent tries to find a proposal that is most beneficial to itself while it is at the same time the most similar to the opponent’s last offer. The problem is, thus, reduced to a multi-objective decision-making problem. The problem with this approach is the amount of shared

information. Also, it cannot be applied to the argumentation-based negotiation context. A further discussion on existing techniques for proposal preparation is presented in section 2.2.1.

Opponent-modeling module

As explained in Section 1.1, the agents' strategies remain private during the negotiations; that is, negotiators reveal only a part of their preferences through their decisions, and the rest remains hidden (Niemann and Lang, 2009). For example, when negotiating over the price of a product, the buyer may give the seller some clues about his/her available budget, but he/she usually does not let the seller know his/her exact price preferences (Zeng and Sycara, 1998). In such cases, the proposer needs to learn about the other participants using their decisions about the offered proposals and, if available, the supplementary information such as their arguments. The learning/estimation process happens in the opponent modeling component. The knowledge gained about other parties through opponent modeling is then fed back to the proposal-preparation module to help the proposer find and offer more agreeable proposals; thus, increasing the probability of reaching an agreement in a timely manner.

There is a vast area of literature discussing opponent modeling that will be presented in details in Chapter 4. However, the proposed approaches are not applicable in our targeted negotiation context due to several reasons. For instance, most approaches mine the data from previous similar negotiations to find the model of negotiators based on their behavior (Robu and La Poutre, 2006) and others determine the opponents' class based on some predefined categories and treat them according to their classes (Hindriks and Tykhonov, 2008). The above-mentioned approaches are not applicable in our target negotiation context since there is not any historical data or classification information available.

1.1.2 Problem statement

The two explained modules, proposal preparation and opponent modeling, are fundamental components of automated negotiation and this study focuses on designing and developing them. For

preparing proposals, three main challenges need to be addressed:

1. Considering multi-issue negotiations, how to select a proposal from the large, multi-dimensional space of possible solutions so that it a) has a high probability of being accepted by the participants; b) maximizes the proposer's benefit; and c) considers the interrelations and correlations among the negotiation issues/attributes.
2. Considering multi-participant negotiations, when offering a proposal, how to take into account the preferences of every participant without compromising the others'.
3. Considering argumentation-based negotiations, how to use the arguments received from the opponents to a) learn their evaluation strategies and their agreement space; and b) improve the probability of their agreement with the proposal.

To address these issues, the state-of-the-art techniques are either using pre-known information about opponents or follow a proposal-based approach using concession-making techniques. As discussed in Section 1.1.1 (The proposal preparation module), these techniques are not applicable to the context of argumentation-based negotiations in which prior information about the participants is not available. Preparing proposals with high agreement probability for more than one opponent (multi-participant negotiation) in a multi-issue argumentation-based negotiation is, therefore, a challenge that is not addressed in the literature. The main challenges for opponent modeling are:

1. The decision-making (proposal-evaluation) models of the opponents need to be mathematically defined for computer agents. The closer the model to real human behaviors, the more complicated the model; i.e., the model involves a more significant number of parameters to estimate.
2. No prior knowledge about the negotiators is available. Therefore, at the beginning of negotiation, the agents have no pre-classification information about their opponents.

3. No data about previous similar negotiations with similar participants is available. Therefore, data-mining approaches, supervised and unsupervised machine-learning approaches, are not applicable.
4. The goal is reaching an agreement in fewer rounds of negotiation even if there are many negotiation issues involved. Therefore, learning approaches whose time-complexity dramatically increases by the number of parameters to be learned, such as reinforcement learning, cannot be applied due to time constraints.

1.1.3 Objectives

The main objective of this research is to develop an automated multi-issue multi-participant negotiation approach that mimics real-world negotiations while improving the efficiency of the negotiation process. This objective is achieved by developing systematic approaches proposed for modeling participants' behavior, learning their preferences and then adjusting the offered proposal based on the learned criteria with the aim of reaching a win-win agreement. This research has, therefore, three specific objectives:

1. Handling the vagueness of the evaluation outcomes received from the stakeholders by representing them in the form of numerical information processable by an agent-based model
2. Providing the proposer agent with a real-time learning/estimation approach so that it can learn other participants' preferences in a reasonable number of negotiation rounds using solely the proposal evaluation outcomes and arguments it receives from them and no extra information.
3. Developing a proposal-preparation approach so that the offered proposal has a high probability of being accepted by all the opponents while being beneficial to the proposer regardless of the number of negotiation issues and/or participants.

4. Assessing the applicability of the proposed approaches using relevant case studies.

The developed negotiation approach is extensively assessed via two case studies of energy-system planning (explained in Section 2.4.1) and real estate (explained in Section 2.4.2), both with a large space of possible solutions, one involving several contradicting participants and the other involving a large number of negotiation issues. The first case-study is an example of negotiations in the context of environmental resource management. This type of negotiations is an important case of multi-issue multi-participant negotiations, and, due to their crucial role construction and development projects, they are an inseparable part of our daily life. The second case study is about the negotiations that happen in the context of buying/selling a real estate. This case is used to evaluate the scalability and performance of the proposed approach when the negotiation involves a larger number of issues.

1.1.4 Main Contributions

Based on the stated objectives, the main contributions of this study are:

1. Developing an estimation approach based on unscented particle filtering so that the proposer agent can learn its opponents' preferences in a few rounds of negotiation. The stakeholders' evaluation models are defined using fuzzy sets. The only required information is the fuzzy scores the stakeholders assign to the offered proposals as well as any argument they may provide in support of their decisions. The efficiency of this learning process is high even in multi-issue, multi-participant negotiations.
2. Proposing a graphical model to represent the negotiation issues and to map their relations with one another and with participants' preferences. This graphical model is optimized using belief propagation to find the appropriate values for the proposal attributes. This model is able to handle a large number of attributes (issues) and is augmented with a z-scoring approach to handle the conflicting preferences of multiple participants.

3. Applying the proposed approaches to the case studies of energy-system and real estate and proving the efficiency of the proposed negotiation model using extensive experiments.

1.1.5 Assumptions and limitations

The following assumptions are made throughout this thesis:

1. Opponents are logical and do not act randomly; i.e., their strategy does not change dramatically during the negotiation.
2. The set of the negotiation issues (attributes) is defined before the negotiation starts and cannot be modified during the negotiation process

1.1.6 Thesis structure

The remainder of this thesis is outlined as follows. Chapter 2 presents the proposed probabilistic graphical model and the belief propagation-based approach that infers this model to find a proper proposal in every round of negotiation. Chapter 3 discusses the history of using machine learning techniques in automated negotiation and particularly in the context of environmental resource management as a very well-known example of multi-issue multi-participant negotiations. Chapter 4 presents the proposed opponent learning approach. Finally, Chapter 5 presents a summary of the research outcomes. This chapter ends with conclusions and possible directions for future research.

Chapter 2

Belief Propagation-based Proposal Preparation in Automated Negotiation

Article Presentation

Background

This chapter is part of the thesis related to the third and fourth objectives, namely developing a proposal preparation approach and assessing it using different case studies. This chapter is originally published as a book chapter "Using Belief Propagation-based Proposal Preparation for Automated Negotiation over Environmental Issues" in "Theory and Application of Reuse, Integration, and Data Science" by Springer in 2018.

To address the challenges of proposal preparation in automated negotiation, this chapter presents a novel approach that models the proposal-selection problem as a Markov Random Field (MRF) and optimizes it using min-sum loopy belief propagation. This approach takes into account all the participants' preferences and, therefore, reduces the disagreements over the offered proposals. This proposal-preparation component works closely with argument-handling and opponent-learning components to improve the results as the negotiation proceeds. However, as the main focus of this chapter is proposal preparation, a conventional technique of opponent learning (i.e., frequency-based approach) is used in this chapter. This way, we can make sure that the improvements in the negotiation efficiency are all made by the suggested proposal-preparation approach. The opponent-learning component is further addressed in Chapter 4.

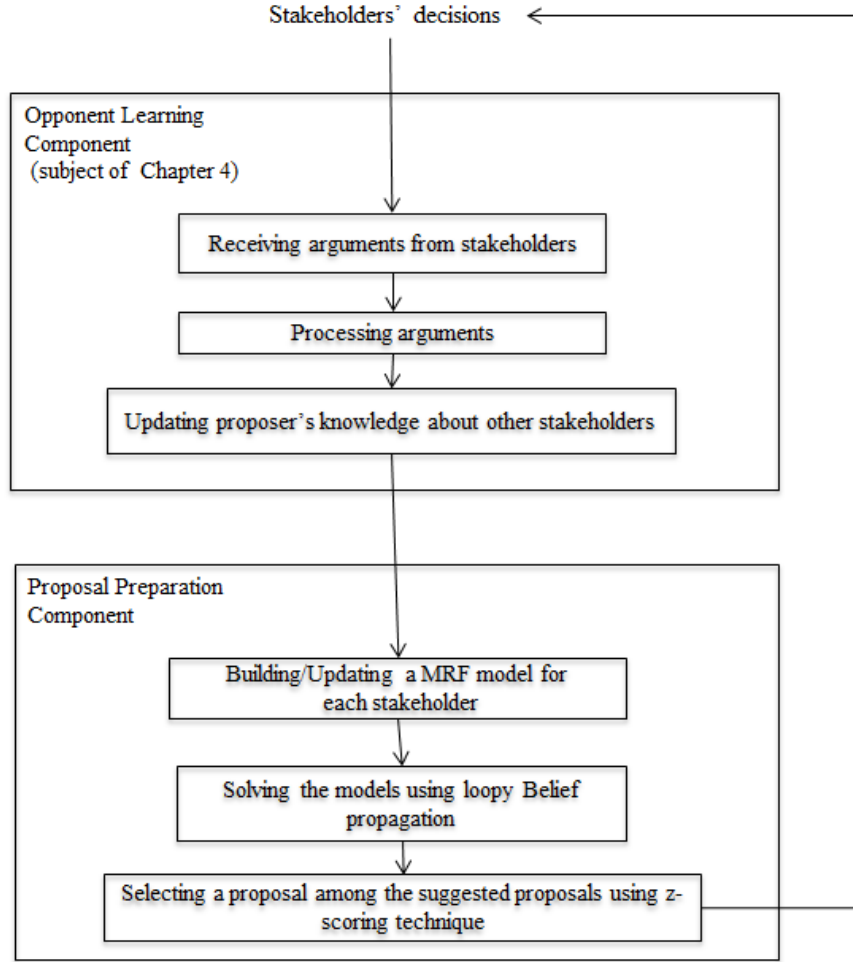


Figure 2.1: Proposal preparation flowchart

General Methodology

The negotiation starts by the proposer with sending a proposal to stakeholder agents and receiving their evaluations (accept/reject + arguments). The argument-handling component then processes the arguments and revises the proposer's understanding of the stakeholders. The next step is updating the MRF model of the stakeholders and negotiation issues. The updated model is then solved using belief propagation, and a proposal is selected for each stakeholder based on its so-far learned preferences. Finally, a z-scoring technique is used to select a final proposal among the prepared proposals by considering the conflicts among the stakeholders. Figure 2.1 represents the steps of the proposed methodology. The following chapter discusses these steps in detail.

Using Belief Propagation-based Proposal Preparation for Automated Negotiation over Environmental Issues

by Faezeh Eshragh, Mozhdeh Shahbazi, Behrouz Far

Published in:

Theory and Application of Reuse, Integration, and Data Science

2018, Pages 69-95

2.1 Abstract

Automated negotiation is used as a tool for modeling human interactions with the aim of making decisions when participants have conflicting preferences. Although automated negotiation is extensively applied in different fields (e.g., e-commerce), its application in environmental studies is still unexplored. This paper focuses on negotiation in environmental resource-management projects. The primary objective of this study is to reach agreement/conclusion as fast as possible. To achieve this objective, an agent-based model with two novel characteristics is proposed. The first is automating the process of proposal offering using Markov Random Fields and belief propagation (BP). The second is the ability to estimate stakeholders' preferences through an argument handling (AH) model. The experiments demonstrated that the AH model and the BP-based proposal preparation (BPPP) approach improve the performance of the negotiation process. A combination of these two modules outperforms the conventional utility-based approach by decreasing the number of negotiation rounds up to 50%.

2.2 Introduction

Negotiation is known as one of the prominent ways of reaching an agreement when the decision makers have opposing interests Van Kleef et al. (2006). In the past few decades, negotiation has been studied from different perspectives such as psychology Pruitt and Carnevale (1993), economy Kreps (1990), and computer science Jennings et al. (2001). These studies investigate the complicated nature of a negotiation process to make it more efficient and reliable. However, the nature of the negotiation process highly depends on the context in which the negotiation occurs. While in some contexts the negotiation takes place between two parties (e.g., one buyer and one seller) over one single issue (e.g., the price of a product), the negotiation process involves multiple issues and multiple participants in other contexts. Environmental resource management is one of the domains that require negotiation among multiple stakeholders with different viewpoints over a variety of environmental and non-environmental issues. Particularly, negotiation can play a crucial role in

the context of managing common-pool environmental resources. These resources are shared by a group of stakeholders and can be overused or congested due to poor management. An instance of such resource-management problem occurs in developing a new electricity transmission line between two locations, where each possible transmission line is a proposal over which negotiation should occur. Each proposal is itself a function of several attributes (criteria), e.g., environmental damage in different forms, development cost, influence on population, and impact on land. Past studies have shown that modeling negotiation over common-pool environmental resources using simulation tools can facilitate the decision making and management processes considerably Ostrom et al. (1994); Bousquet and Trébuil (2005). Automated negotiation, which is facilitated by an agent-based model (ABM), is one of these tools. In this model, a set of intelligent agents interact with each other and search the space of potential agreements to find a (set of) mutually acceptable agreement(s). Automated negotiation has received considerable attention in several domains such as supply chain management Fink (2006), political studies Aragonès and Dellunde (2008), and e-commerce Dworman et al. (1996); Zeng and Sycara (1998); Faratin et al. (2002); Ramchurn et al. (2006); Shojaiemehr and Rafsanjani (2018). However, only a few studies have considered modeling stakeholders' negotiations over environmental issues using ABM Akhbari and Grigg (2013); Okumura et al. (2013); Pooyandeh and Marceau (2013, 2014); Alfonso et al. (2014). Apart from the lack of mutual collaborations between scientists in the areas of artificial intelligence (AI) and environmental studies, the other reasons for scarcity of automated negotiations over environmental issues include the high level of uncertainty in these issues, the high-stakes decisions, and the involved stakeholders who are not usually willing to reveal the details of their preferences. These factors may cause the negotiations to take too long, sometimes even without reaching an agreement. In this paper, new approaches are proposed to facilitate and accelerate automated dynamic negotiation over environmental issues. The following section provides a summary of the studies related to automated negotiation with a focus on environmental issues.

2.2.1 Related work

An agent-based model (ABM) is a computational model for simulating collective behaviors in which autonomous elements, called agents, with a predetermined set of goals and actions, imitate intelligent units of the environment Wooldridge (2001). One of the main characteristics of ABM is its autonomy as there is no global control over the agents, and they do not operate based on a globally consistent knowledge base. These agents communicate and interact with each other to exchange knowledge and achieve their goals Oprea (2004). ABM has been employed in various domains such as transportation management Jennings et al. (1995), telecommunication Bäumer and Magedanz (1999), business Jennings et al. (1996) and electronic commerce Cao et al. (2015). Automated negotiation is another application of ABM in which negotiation among different parties is modeled using agents and their interactions.

Three main approaches have been applied in automated negotiation: the game theory, the heuristic approach, and the argumentation-based negotiation (ABN) Jennings et al. (2001). Game theory originates from research conducted by Neumann and Morgenstern Von Neumann and Morgenstern (1945) and has its roots in economics. It provides the mathematical foundation for the analysis of interactions between self-interested agents MacKenzie and DaSilva (2006). The heuristic approach attempts to overcome limitations of the game theory techniques by finding a satisfactory, sub-optimal solution instead of an optimal one Kraus and Lehmann (1995); Barbuceanu and Lo (2000); An and Lesser (2012); Aydogan et al. (2013). Both game theory and heuristic approaches are proposal-based. That is, exchanging proposals is the only way through which those agents can gain information about their opponents. However, a proposal is only a point in the agreement space and contains information about neither the proposing agent nor the other parties. In this paper, the argumentation-based negotiation (ABN) approach is employed as it better fits the context of the stated problem. In ABN approaches, in addition to the proposals and reject/accept decisions, the agents can exchange supplementary information, called arguments, in support of their decisions or against their opponents' decisions Jennings et al. (1998). That is, in addition to

rejecting an offer (namely a proposal), an agent can explain which parts of the proposal are against his preferences. Different forms of arguments, rewards, treats or appeals are discussed in the literature Ramchurn et al. (2006). One of the benefits of sharing this information is facilitating the process of estimating the space of potential agreements and, therefore, expediting the negotiation process. The agents in ABN approaches can interpret the arguments and take advantage of them through an estimation mechanism. They are, therefore, designed with more components compared with the non-ABN agents; these components are in charge of generating and interpreting the arguments Rahwan et al. (2003). The argumentation-based approach has received significant attention within the last 20 years due to its potential in modeling real-world negotiations Sycara (1990); Parsons and Jennings (1996); Parsons et al. (1998); Kakas and Moraitis (2006); Rahwan et al. (2004); Karunatilake et al. (2009); El-Sisi and Mousa (2012). Wang et al. applied ABN on supply chain formation where the agents use argumentation to understand the preferences of the other participants Wang et al. (2009). In another study, ABN has been used in the design of complex infrastructure systems with many components and layers of subsystems Marashi and Davis (2006). They proposed a systematic way to decompose the system into subsystems and sub-processes by identifying the objectives and criteria of each process and then resolving the conflicts among them using argumentation-based negotiation. ABN has also been widely employed in e-commerce El-Sisi and Mousa (2012); Zhang et al. (2012); Jain and Dahiya (2012); Zhang et al. (2010). For example, El-Sisi and Mousa have proposed a bargaining negotiation framework in which classic non-ABN agents are compared to ABN agents in terms of the quality of the reached agreement and number of unsuccessful negotiations El-Sisi and Mousa (2012).

Research related to argumentation-based negotiation have mainly followed three major directions Karunatilake et al. (2009). In the first one, called argumentation-based defeasible reasoning, arguments about different alternatives are gathered and compared with other available arguments to find possible conflicts. The argumentation system then tries to resolve the conflicts by analyzing the relationship among the arguments (e.g., support, attack, conflict, etc.) and update the agent's

beliefs accordingly Chesnevar et al. (2000); Prakken and Vreeswijk (2002); Tamani et al. (2015); Thomopoulos et al. (2015). The second approach is used in case studies where the system needs to generate and send rhetorical statements to the user. Here, arguments are characterized as structures (schema) for generating persuasive feedback Gilbert et al. (2004). The third group of studies in this area employ arguments as a model of interaction when there are conflicts between negotiating parties. The arguments in this approach are generated under the structure of dialogue games using a common communication language and based on a set of pre-defined rules governing the negotiation Karunatillake et al. (2009). The argument handling in the current study follows the third direction as the agents communicate and generate arguments using a dialogue protocol, a set of negotiation rules and a communication language. The formal analysis of the relationships among arguments is therefore out of the scope of this paper. That is, arguments are represented as phrases justifying the agents' decisions for rejecting/accepting an offer and all the agents' arguments will be considered in advancing the negotiation regardless of possible conflicts among them. It is also assumed that the negotiation parties only submit reasonable and sensible arguments.

The order of offered proposals is another critical matter in the negotiation process that needs to receive more attention. The order by which the proposals are selected and offered to negotiation parties may dramatically influence the pace of the negotiation process and the quality of its outcomes. This is specifically important when the space of potential agreements is large and high-dimensional. However, for the proposing agent, determining the right proposal based on his limited knowledge of the preferences of the other parties is a challenging task. In the current study, we applied a belief propagation-based approach to tackle this problem. Belief propagation-based techniques have rarely been used in automated negotiation. In Robu et al. (2005); Hadfi and Ito (2015, 2016), a belief propagation-based technique is employed to represent the utility space and facilitate its exploration. The focus of these studies is the optimal solution for one participant using his utility graph and constraints. However, such an approach is not applicable to environmental negotiations, in which many participants with multiple perspectives are involved, and all

their preferences should be considered simultaneously.

2.2.2 Objectives and contributions

In this paper, we present an automated negotiation solution for modeling negotiations over multi-criteria, multi-participant decision-making problems, such as the ones that happen in environmental resource management. The main objective of this solution is to reach agreement/conclusion in as few rounds of negotiation as possible. This work is an extension of our previous conference presentation Eshragh et al. (2017).

The proposed solution is implemented as an ABM with several processing modules and knowledge bases. The novel characteristics of this ABM are twofold: further automation and intelligence. The first novelty is based on automating the proposal-offering process. This problem is graphically modeled as hidden Markov Random Fields (MRF), where attributes of the proposals are considered as observable nodes, and the possible values (states) of the attributes are the hidden variables. A joint probabilistic model, which depends on preferences of negotiation parties, is built over the attributes and hidden variables. The direct statistical dependencies between hidden variables are modeled by connecting hidden variables via undirected edges in the graph. A pair of states are connected if they can simultaneously be used to form a proposal, and the length of this edge describes their consistency. Inference on this graphical model is performed using min-sum loopy belief propagation. Having all available information in one MRF model, the best offer is selected in each round of negotiation. As the negotiation proceeds, the model is dynamically rebuilt based on re-estimating the preferences of negotiation parties, which is performed by the second contribution of this paper. To this end, an argument handling (AH) module is developed, using which the agents exchange information with the proposer agent in each round of negotiation. The module applies the exchanged information to update the knowledge of the proposer agent about the preferences of the other parties. The proposer agent, then, uses these new preference estimations to rebuild the MRF model and prepare a more mutually acceptable proposal. The dynamic nature of the MRF in combination with the AH model makes the system a very suitable match for the

target problem. The proposed method is successfully employed in the case study of energy-system planning. We have also evaluated this method in a real-estate case study to verify the performance when dealing with larger problems regarding the search space and the dimensionality of the involved criteria.

The rest of the paper is organized as follows. In section 2.3, the components of automated negotiation are described. Two case studies, energy-system planning in Alberta and King County House Sales, as well as experiments, implementation details, and the results are discussed in section 2.4. The conclusions are presented in section 2.5.

2.3 Automated Negotiation

Automated negotiation consists of three main components: negotiation protocol, negotiation object, and negotiation strategy (Lomuscio et al., 2001). The negotiation protocol defines the set of rules governing the interactions between agents. It determines the possible types of participants, the negotiation phases, the negotiation rules, and the possible actions for each participant in each phase. In this study we have two types of agents. The first one is the proposer agent, representing the party in charge of preparing and offering proposals in each round of negotiation. The second type of negotiating agents represents other stakeholders who receive the proposals and decide whether to agree or disagree with them. In accordance with most real-world negotiations, it is assumed that the proposer agent offers only one proposal in each round of negotiation and conceals his knowledge of other possible offers from other agents. The agents follow a simple protocol, shown in Fig. 2.2. Rectangles in this figure represent the set of actions and the lines between them show the order of possible moves between these actions. The proposer initiates the negotiation process by sending an offer to the stakeholder agent. The proposed offer will be either accepted or rejected by the agent. In case of rejection, the proposer will ask for the reason (Challenge) and receives an argument (or a set of arguments) in response (Assert). The process continues till the stakeholder accepts the offered proposal or the proposer calls for a withdrawal of the process. The

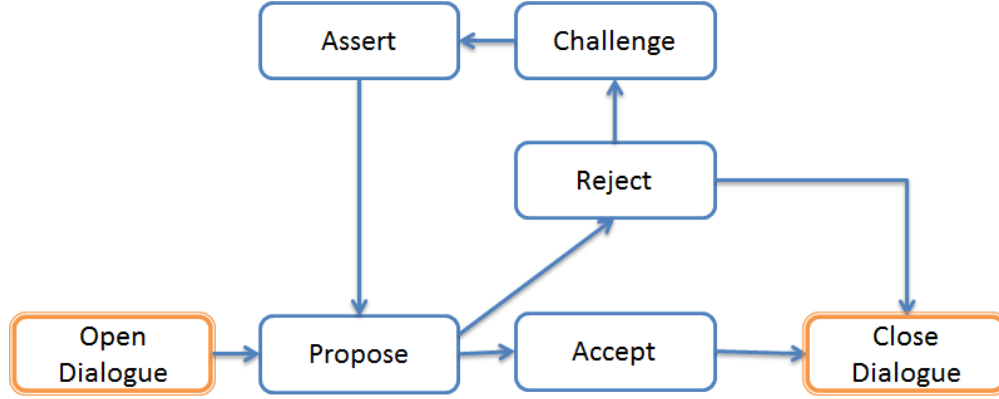


Figure 2.2: Negotiation Protocol

call for withdrawal happens when the proposer runs out of proposals (i.e. all possible proposals are already offered and rejected by the stakeholders). Here, we assume all agents are rational.

The second component of automated negotiation, called negotiation object, corresponds to the range of criteria over which the negotiation happens. The proposed methodology handles multi-criteria negotiations. In each round of negotiation, the proposer assigns values to this set of criteria which all together form a proposal. The stakeholders have to decide whether they accept this proposal or reject it. In this study, we assume that the criteria (attributes) can be numerically quantified. The third component is the agents' decision-making process, also called agents' strategy, used by the agents to act according to the negotiation protocol. It is basically the agent's plan for achieving its goals (Lomuscio et al., 2001). While the negotiation protocol is public and available to all participants, the agent's strategy is always private. In this study, the proposer agent's strategy is more about the way it prepares a proposal while the stakeholder agents' strategies explain how they evaluate a received proposal and generate arguments about their decisions. In the next sections, these strategies will be explained in more details.

In section 2.3.1, the ABM and its components are presented. In section 2.3.2, the problem of automated negotiation is broken down and formulated. Sections 2.3.3 and 2.3.4 describe the BPPP approach and the statistical solution to determine the final proposal in each round of the negotiation. Section 2.3.5 presents the details of argument handling model.

2.3.1 Agent-based model

In this study, an ABM has been employed to model a one-to-many type of negotiation among stakeholders with different perspectives. Each agent in the model has preferences over certain criteria that are designed based on the stakeholder's concerns and interests. These stakeholders represent the groups or individuals who are either involved in the decision-making process or affected by the final decision (Freeman, 2010).

Each agent in the ABM has access to a set of private and public knowledge bases containing information about the agent's goals, beliefs, and preferences. This information is gathered either from public documents and datasets or directly from the stakeholders. A part of the agents' knowledge that relates to the external environment and other participants may change during the negotiation as a result of re-estimation processes. In the proposed model, the re-estimation process happens in the argument handling model, explained in section 2.3.5. The proposer agent has a proposal database of all possible alternatives as well as a database for previously offered proposals. These records can be used for keeping track of negotiation rounds and future references (Rahwan et al., 2004). All agents have databases of their preferences and criteria. This database is utilized in assessing alternatives and generating arguments during the negotiation process.

A schematic representation of the negotiation system developed in this study is displayed in Fig. 2.3. The system includes an ABM as well as the databases and the external modules required for generating proposals, selecting among them, and handling the arguments.

2.3.2 Problem statement

Consider a set of $x + 1$ stakeholders, $S = \{s_p, s_1, s_2, \dots, s_x\}$, who are negotiating over a set of proposals, $P = \{p_1, p_2, \dots, p_y\}$. Here, s_p represents the proposer agent who initiates the negotiation and proposes a new offer in each round. Each proposal $p_j (j = 1, \dots, y)$ is described by a set of attributes $A = \{a_1, a_2, \dots, a_z\}$; that is, proposal p_j is a unique combination of the values of these attributes as $A_j = \{a_1^j, a_2^j, \dots, a_z^j\}$. The preference thresholds of stakeholder $s_i (i = 1, \dots, x)$ over

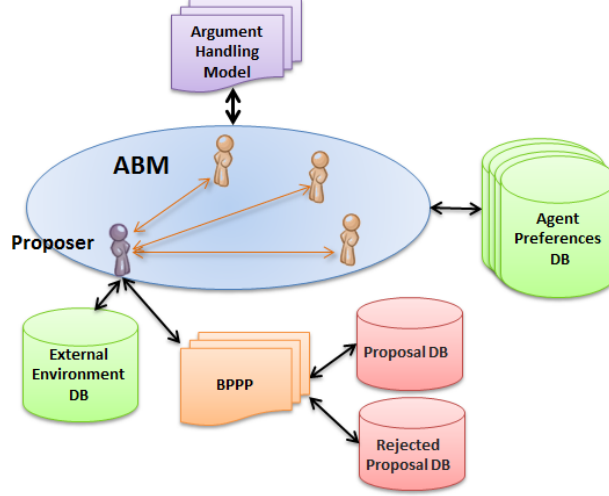


Figure 2.3: A schematic representation of the proposed negotiation system

attribute $a_k (k = 1, \dots, z)$ is represented as his minimum and maximum acceptable values for that attribute, denoted by Min_{a_k, s_i} and Max_{a_k, s_i} . The agents negotiate to find a proposal, if any, that meets all agents' preferences. Without loss of generality, we assumed that all agents (e.g., stakeholders) have the same importance in the negotiation process. The negotiation ends when either all agents agree on a proposal, or the proposer confirms that there is no more proposal to offer and terminates the negotiation without reaching an agreement.

2.3.3 Belief propagation-based proposal preparation (BPPP)

In each round of negotiation, a proposal is selected by the proposer agent to be offered to the stakeholders. The goal of the proposer is to find a proposal that is aligned with his preferences and has the highest probability of being accepted by other agents. In the beginning, the proposer agent knows nothing about the stakeholders' preferences; therefore, it acts based on its preferences as well as some assumptions about the ranges of acceptable values by other agents. However, as the negotiation proceeds and the feedback (e.g., agent's decisions and arguments) are received, he gradually learns about other participants' preferences. He will then utilize the recently learned knowledge about a particular stakeholder in finding a potential agreement with that stakeholder. From the probabilistic point of view, this proposal preparation process can be modeled as an infer-

ence problem; That is, the proposer infers the most probable proposal given the likelihood (based on the preferences of a specific stakeholder over each attribute of the proposal) and the prior (based on the interdependence of specific attributes). In the current study, the proposal preparation problem is modeled using Markov Random Fields (MRF) where belief propagation is applied to approximate the inference.

In the MRF, for each stakeholder s_i an undirected graph is defined with attributes $A = \{a_1, a_2, \dots, a_z\}$ as its nodes. Each attribute $a_k (k = 1, \dots, z)$ has a set of alternate states (values) $V_{a_k} = \{v_1^k, v_2^k, \dots, v_{n_k}^k\}$. The problem is finding the optimal state for each attribute and then identifying the most optimal proposal among all possible alternatives $P = \{p_1, p_2, \dots, p_y\}$ based on the selected states. To figure this problem out, the proposer needs to know the preferences of the stakeholder s_i , namely $\{Min_{a_k, s_i}, Max_{a_k, s_i}\}$. However, in real-world situations, stakeholders usually do not share this sort of information, and it is often considered private. Therefore, the proposer agent starts with an estimation of these thresholds and tries to learn them through the feedback he receives from others.

To solve this optimization problem, an energy minimization framework can be employed where a global energy function penalizes each alternative for being either dissatisfactory to the stakeholder (unary cost) or not consistent with other attributes' values (binary cost). The dissatisfaction level of a state of an attribute for stakeholder s_i is defined as the unary cost of that state. In mathematical terms, the cost of assigning state $v_l^k (l = 1, \dots, n_k)$ of node $a_k (k = 1, \dots, z)$ in negotiation with stakeholder s_i , the unary cost $U_{l,i}^k$, is defined as follows.

$$UC_{s_i}(v_l^k) = U_{l,i}^k = (v_l^k - Min_{a_k, s_i}) / (Max_{a_k, s_i} - Min_{a_k, s_i}) \quad (2.1)$$

This reverse-scoring function is selected to reflect the fact that the further the cost of an alternative is from the stakeholder's preferences, the lower the chance of its selection.

The binary cost is defined to enforce the compatibility between two states of two different attributes. It enforces the fact that two states of any two attributes should belong to the same

proposal for them to appear together. Quantitatively, the binary cost of the combination of the state $v_{l_1}^{k_1}$ of node a_{k_1} and state $v_{l_2}^{k_2}$ of node a_{k_2} ($k_1 = 1, \dots, z$ and $k_2 = 1, \dots, z$), denoted as $B_{l_1, l_2}^{k_1, k_2}$, is calculated as below.

$$B_{l_1, l_2}^{k_1, k_2} = \begin{cases} \infty, & \text{if } P_1 \cap P_2 = \emptyset \\ \frac{\sum_{j|p_j \in P_1 \cap P_2} c_j}{|P_1 \cap P_2|}, & \text{if } P_1 \cap P_2 \neq \emptyset \end{cases} \quad (2.2)$$

In Equation 2.2, $P_1 \subset P$ is a sub-set of all the proposals that are compliant with state $v_{l_1}^{k_1}$; i.e., $P_1 = \{p_j | a_{k_1}^j = v_{l_1}^{k_1}\}$. Similarly, $P_2 \subset P$ is a sub-set of all the proposals that are compliant with state $v_{l_2}^{k_2}$; i.e., $P_2 = \{p_j | a_{k_2}^j = v_{l_2}^{k_2}\}$. Also, the variable c_j refers to the cost of the proposal p_j according to the proposer himself.

Having the unary and binary costs, a belief propagation-based approach is used to find the solution with the minimum energy. Here, the belief of node a_{k_1} ($k_1 = 1, \dots, z$) about its state $v_{l_1}^{k_1}$ ($l_1 = 1, \dots, n_{k_1}$) during the negotiation with stakeholder s_i ($i = 1, \dots, x$) is defined as,

$$bel_{k_1}(v_{l_1}^{k_1}) = U_{l_1, i}^{k_1} + \sum_{a_{k_t} \in N_{k_1}} msg_{k_t \rightarrow k_1}(v_{l_1}^{k_1}) \quad (2.3)$$

where $U_{l_1, i}^{k_1}$ is the unary cost of state $v_{l_1}^{k_1}$ for node a_{k_1} ; the set N_{k_1} includes all the nodes connected to a_{k_1} ; and, $msg_{k_t \rightarrow k_1}(v_{l_1}^{k_1})$ is the message that node a_{k_1} receives from the node a_{k_t} about the state $v_{l_1}^{k_1}$, which is defined as follows.

$$msg_{k_t \rightarrow k_1}(v_{l_1}^{k_1}) = \min_{l_t=1, \dots, n_{k_t}} (U_{l_t, i}^{k_t} + B_{l_1, l_t}^{k_1, k_t}) + \sum_{a_{k_u} \in \{N_{k_t} \setminus a_{k_1}\}} msg_{k_u \rightarrow k_t}(v_{l_t}^{k_t}) \quad (2.4)$$

Here, $U_{l_t, i}^{k_t}$ is the unary cost of state $v_{l_t}^{k_t}$ for node a_{k_t} ; $B_{l_1, l_t}^{k_1, k_t}$ is the binary cost of the state $v_{l_1}^{k_1}$ of node a_{k_1} and state $v_{l_t}^{k_t}$ of node a_{k_t} ; and, the set $\{N_{k_t} \setminus a_{k_1}\}$ includes all the nodes connected to a_{k_t} except a_{k_1} .

The nodes pass messages to each other for several iterations till we ensure all information has been transmitted through messages, and the beliefs of all the nodes about their states reach a steady state. In this stage, the solution is the set of the attribute states with the lowest belief values.

2.3.4 Z-scoring approach for selecting among the prepared proposals

The proposer agent builds one MRF model, explained in section 2.3.3, for each stakeholder involved in the negotiation process to find a proposal with the highest acceptance likelihood for that particular stakeholder. By the end of the BPPP processes, the proposer has x-number of proposals (each prepared for a stakeholder). Therefore, he needs to choose one of these proposals as the final one to offer. In this study, we applied z-scoring to facilitate the selection among these proposals. A z-score indicates how many standard deviations a score is from the mean. For the score γ , the z-score, shown as ζ , is calculated as below,

$$\zeta = \frac{\gamma - \mu}{\sigma} \quad (2.5)$$

where γ and σ are the mean and standard deviation of the population, respectively. This is the way that a z-score is calculated for the values of one attribute. Here, however, we are dealing with proposals, which usually are consisted of more than one attribute. To calculate the z-score for each alternate proposal in our database, the process shown in Fig. 2.4 is used.

In the first step, the z-score for each attribute is calculated. The first requirement for analyzing the z-score is calculating the mean and the standard deviation of each attribute over all the proposals. Considering μ_{a_k} and σ_{a_k} as the mean and standard deviation of attribute a_k , the z-score of a_k^j (the value of attribute a_k in proposal p_j) is calculated as follows.

$$\zeta_k^j = \frac{a_k^j - \mu_{a_k}}{\sigma_{a_k}} \quad (2.6)$$

The next step is calculating the mean and standard deviation of the z-scores of all attributes for

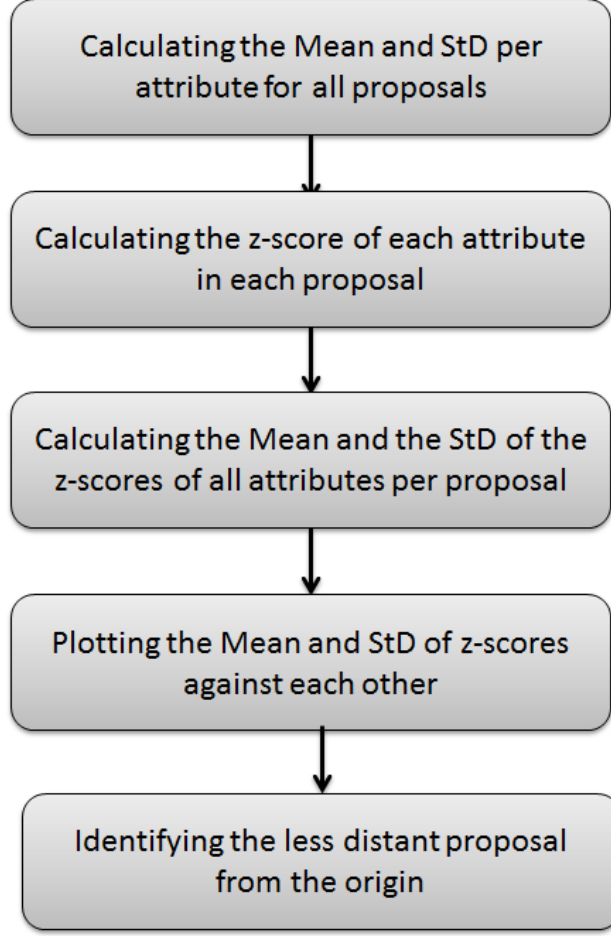


Figure 2.4: The method of calculating the z-score for each alternative proposal

each proposal. For proposal p_j , these values are calculated as in Equations 2.7 and 2.8,

$$\mu_j = \frac{\sum_{k=1}^z \zeta_k^j}{z} \quad (2.7)$$

$$\sigma_j = \sqrt{\frac{\sum_{k=1}^z (\zeta_k^j - \mu_j)^2}{z}} \quad (2.8)$$

where z is the total number of attributes.

In the final step, each proposal p_j is represented by a point in the 2D Euclidean space, with Cartesian coordinates of (μ_j, σ_j) . The proposal with the least distance to the origin will be then

selected and offered by the proposer agent to the other agents in the current round of negotiation.

2.3.5 Argument handling model

Once a proposal is offered, stakeholder agents start to evaluate it based on their preferences. Based on the results of these evaluations, the agents send feedback to the proposer. If the offered proposal meets all the preferences of a stakeholder, it will be accepted by that particular stakeholder agent with no further comments. However, if the agent evaluates the proposal as unacceptable, the proposer asks for the reason behind this decision and the stakeholder will support his decision by a (set of) relative argument(s) sent to the proposer. These steps are shown in Fig.2.2 with ‘Accept’, ‘Reject’, ‘Challenge’ and ‘Assert’ boxes. Definition 2.3.1 represents a more formal representation of these steps. This definition is a modified version of the communication language proposed in (Karunatilake et al., 2009).

Definition 2.3.1. Communication Language

ACCEPT

Usage: $\text{Accept}(s_p, s_i, p)$.

Informal Meaning: Accept the proposal p , proposed by proposer agent s_p to stakeholder agent s_i

REJECT

Usage: $\text{Reject}(s_i, s_p, p)$.

Informal Meaning: Reject the proposal p proposed by proposer agent s_p to stakeholder agent s_i

CHALLENGE

Usage: $\text{Challenge}(s_p, s_i, \text{Reject}(s_i, s_p, p))$

Informal Meaning: Proposer agent s_p Challenges the justification of stakeholder agent s_i for rejecting proposal p .

ASSERT

Usage: $\text{Assert}(s_i, s_p, \text{Challenge}(s_p, s_i, \text{Reject}(s_i, s_p, p))), \lambda$

Informal Meaning: s_i Asserts a particular set of arguments λ to s_p in response to challenging its

rejection decision.

The agents employ this notation to exchange proposals and feedback. Our focus, however, is on the way the argument received from the stakeholder agents affect the beliefs of the proposer agent.

2.3.6 Simple frequency-based argument processing¹

Based on the argument received from stakeholder s_i , the proposer improves his estimations of the preference thresholds of s_i about each attribute, $\{Min_{a_k, s_i}, Max_{a_k, s_i}\}$. If at the t^{th} round of negotiation, the argument received from stakeholder s_i states that the value of attribute a_k is too high, then the upper-bound of the preference limit should be revised. The current upper-bound, Max_{a_k, s_i}^t , is then re-estimated to Max_{a_k, s_i}^{t+1} as below,

$$Max_{a_k, s_i}^{t+1} = \begin{cases} a_k^j, & \text{if } Max_{a_k, s_i}^t > a_k^j \\ Max_{a_k, s_i}^t - (\eta \times Max_{a_k, s_i}^t), & \text{if } Max_{a_k, s_i}^t \leq a_k^j \end{cases} \quad (2.9)$$

where a_k^j is the value of attribute a_k in the current proposal p_j and η represents the ratio by which the upper-bound of the preference limit is updated.

On the other hand, if the feedback received from stakeholder s_i argues that the value of attribute a_k is too low, then the lower-bound of the preference limit should be readjusted. The current lower-bound, Min_{a_k, s_i}^t , is then updated to Min_{a_k, s_i}^{t+1} using the following equation.

$$Min_{a_k, s_i}^{t+1} = \begin{cases} a_k^j, & \text{if } Min_{a_k, s_i}^t \leq a_k^j \\ Min_{a_k, s_i}^t + (\eta \times Min_{a_k, s_i}^t), & \text{if } Min_{a_k, s_i}^t > a_k^j \end{cases} \quad (2.10)$$

Here, η represents the ratio by which the lower-bound of the preference limit is re-estimated.

¹This method is selected due to its popularity in the literature and will be replaced with a novel approach in Chapter

The new preference limits for attribute a_k according to stakeholder s_i will be then used to update the unary cost of its l^{th} state by substituting the preference thresholds in Equation 2.1 with the new preference limits from Equations 2.9 and 2.10. After updating the unary costs based on the received arguments, the BPPP approach will be used again to find the most appropriate proposal for the next round of negotiation. The process will continue till either all agents agree on a proposal or they terminate the negotiation with no agreement.

2.4 Experiments

We applied the proposed methodology in two different case studies; the first case, which discusses the energy- system planning in Alberta, deals with a set of GIS data, gathered from Alberta Biodiversity Monitoring Institute (ABMI) and Alberta Environment and Parks (AEP) public resources. The second case contains real estate data from King County, US, between May 2014 and May 2015. Although this case study is not related to environmental issues, it involves multi-criteria, multi-participant negotiations and gives us the opportunity to assess the performance and scalability of our proposed method on a problem with a larger solution space. The dataset includes 21614 records (comparing to 100 alternatives in the case study related to energy-system planning) and higher data dimensionality (eight attributes for each proposal comparing to five attributes in the case study related to energy-system planning).

In these experiments, the performance of a negotiation method is measured as the number of rounds within which the negotiation process terminates. The termination happens either when participants reach a mutually acceptable agreement or when searching for the possible agreement stops (e.g., due to time restriction or finding no agreement among the participants) (Jennings et al., 2001).

To examine the performance of the proposed techniques, a set of experiments has been designed. These experiments evaluate the effect of the AH model and BPPP approach on the performance of the negotiation process. We have also analyzed different settings in the BPPP approach

to find the optimal intervals for attribute-value discretization in the first case study. To compare the performance of the proposed negotiation strategy with the utility-based approach, we have conducted another set of experiments for both case studies. The utility-based approach is one of the most popular methods of evaluating proposals and is usually simplified to the weighted average of attribute values. That is, the utility of proposal p_j for stakeholder s_i is calculated as:

$$Utility_i(p_j) = \sum_{k=1}^z (w_k^i \times a_k^j) \quad (2.11)$$

where w_k^i is the weight of attribute a_k according to stakeholder s_i and a_k^j is the value of attribute a_k in proposal p_j .

The rest of this section explains the case studies and the conducted experiments for each case study.

2.4.1 Case study A: energy-system planning in Alberta

During the past decade, the electricity demand in Alberta has arisen, and more reliable electricity grids need to be developed to transfer the generated power to the consumers. Besides selecting among available technologies (e.g., transmission lines and substations), finding the most reasonable routing option (to link the supply source and the demand center) is a key problem in these sorts of projects where both environmental and non-environmental factors are involved. In the first case study of this research, we investigated an electricity transmission project, in which the supply source is a hydropower plant near Slave River and Forth Smith city at the border of Alberta and Northwest Territories. It is shown as a red star on the map of Fig. 2.5. “Thickwood Hills 951s” and “Ells River 2079s” are the substations nominated to receive the transferred electricity. Orange triangles illustrate these two substations in Fig. 2.5. There are many alternative routes to connect the hydropower plant and the proposed substations, and the goal is to find the one that satisfies every stakeholder involved in the project. To achieve this goal, a set of criteria has to be considered, e.g., the area, type, and the coverage of the land that will be affected, the development and construction costs, the environmental impacts (e.g., wildlife), and the population that will be

Table 2.1: Significant stakeholders in the project on energy-system planning

| Stakeholder category | Group name | Agent name | Primary concerns |
|-------------------------------------|---|----------------|--|
| First Nations | Community, Aboriginal and Native American Relations in TransCanada Treaty 8 First Nations of Alberta | FN | Damage to first nation reserves (FN value) |
| Environmentally focused groups | Alberta Environment and Sustainable Resource Development | AEP | Damage to forest areas (Forest value), wildlife (Wildlife value), and wetlands (Wetland value) |
| Industries on the transmission side | ATCO Electric AltaLink | TFO (Proposer) | Construction costs |

influenced by the route.

In our study, three categories of stakeholders are considered, including: first nations, industrial parties, and environmentally-focused groups (Table 2.1). Each category has preferences over a set of issues (attributes) based on their primary interests and concerns. For example, stakeholders in the environmentally-focused category are mostly concerned about specific issues including the disturbed area of forests, wetlands and wildlife, and, therefore, their preferences on these issues will influence their decisions during the negotiation process.

2.4.2 Data preparation and implementation

The developed ABM is implemented using thread processing in Java. The agents have some private and public knowledge bases. Some of these databases are developed using Microsoft SQL 2010, and the rest are shared files accessible to all the stakeholders.

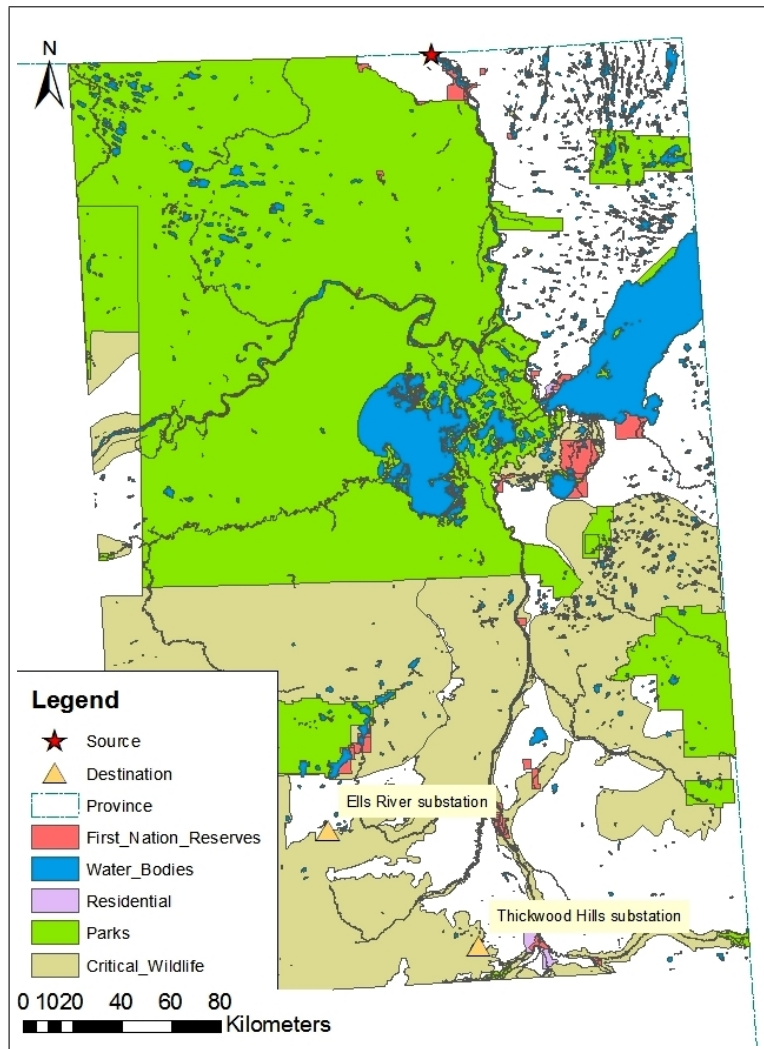


Figure 2.5: The study area located near the Slave River and Fort Smith city at the border of Alberta and the Northwest Territories

Specifying the search space

Several GIS data layers such as the maps of forests, lakes, rivers, caribou zones, roads, and first nations' reserves have been acquired through Alberta Biodiversity Monitoring Institute (ABMI) and Alberta Environment and Parks (AEP) public resources. A set of alternative routes from the supply source to the destination substations have been determined using a python script. This script uses Arcpy library Least Cost Path (LCP) analysis to find the alternative solutions. Each alternative path is characterized by several attributes including forest value, wildlife value, wetland value, FN value, and construction cost. These attributes are quantified using various environmental, ecological, cultural and economic measures. For example, the wildlife value of each path is calculated based on the intersection of the path with wildlife-sensitive areas. The construction cost of each path is determined based on the length of the route and the type of the land/water bodies it passes through. Fig. 2.6 shows the alternative routes selected for this case study, which approximately cover the whole area between the source and the destination points.

Discretizing attribute values

Since the attributes are quantified with continuous values, the number of possible states for each attribute would be infinite. Therefore, to make the problem solvable, the continuous values of the attributes should be reduced to finite discrete states.

The optimal number of discrete states should minimize the number of negotiation rounds. To determine this optimal value, we have empirically evaluated the influence of five different numbers of states (5, 10, 15, 20 and 25 states for each attribute) on the number of negotiation rounds with and without the AH model. As it is shown in Fig. 2.7, the case with 15 states for each attribute results in the minimum number of rounds for both settings (i.e., with/without argument handling module). Increasing the number of states beyond 15 results in degraded performance; using very small intervals to discretize attribute values leads to an increased number of negotiation rounds due to increasing the number of inter-state switches that the proposer should perform to find a suitable state. As Fig. 2.7 shows, the 5 and 10 states settings are not as efficient as 15 states setting either.

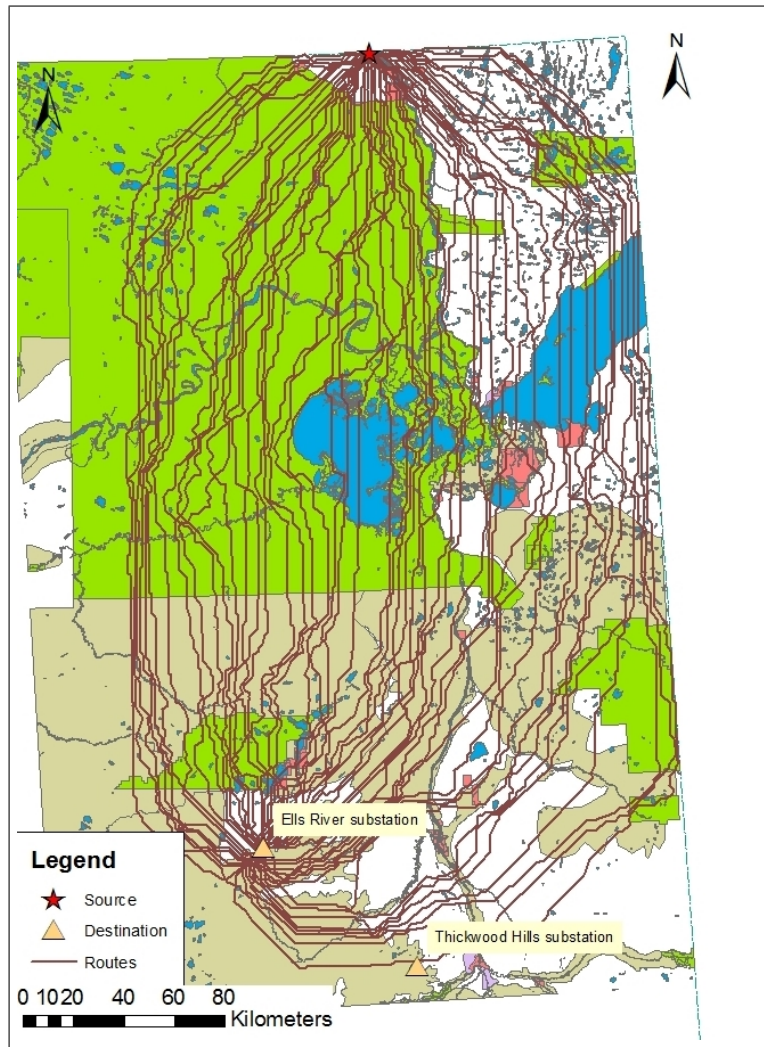


Figure 2.6: Selected routes in the data preparation phase

These experiments show that there is a tradeoff between the size of the search space (increasing by using small discretization intervals) and the chance of reaching states that have overlaps with stakeholders' preferences (increasing by using large discretization intervals). Therefore, setting the number of the states to 15 is a compromise between the extremes (5 states and 25 states).

As a result, the values of each attribute are clustered to 15 discrete states. Then, the unary cost of each state, as well as the binary cost of the combinations of every two states, is calculated using Equations 1 and 2. The results of these calculations are then employed by the ABM to find the most appropriate proposal in each round of negotiation using the BPPP approach.

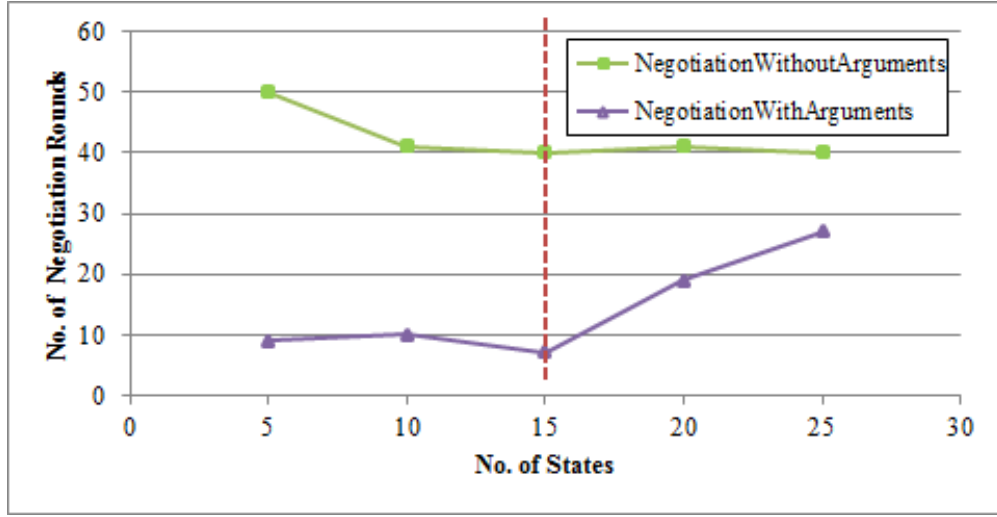


Figure 2.7: Effect of number of Ranges on rounds of negotiation

Determining the value of the parameter η

For this case-study the parameter η in equations 2.9 and 2.10 is empirically determined through experiments on the effects of this ratio on the number of negotiation rounds. Fig.2.8 indicates that 0.05, 0.1 values for this ratio result in less number of rounds. In the following experiments, the value of this ratio is set to 0.05 as it causes the smallest change to the preference limit while resulting in the least number of rounds of negotiation. Higher values of this ratio results in passing the threshold limits in the first few rounds of negotiation and increasing the number of negotiation rounds due to accuracy of estimations.

2.4.3 Results and discussions

The first experiment examines the effect of argument handling model on the performance of the negotiation process. Two different settings are tested in this experiment. In the first one, no argument is passed between each stakeholder and the proposer agent. The only response that the proposer receives from other parties is whether they accept or reject the proposal. With this setting, it takes 40 rounds of negotiation before the agents reach an agreement. In the second setting, the stakeholders provide arguments to justify their responses. In this case, the number of the negotiation

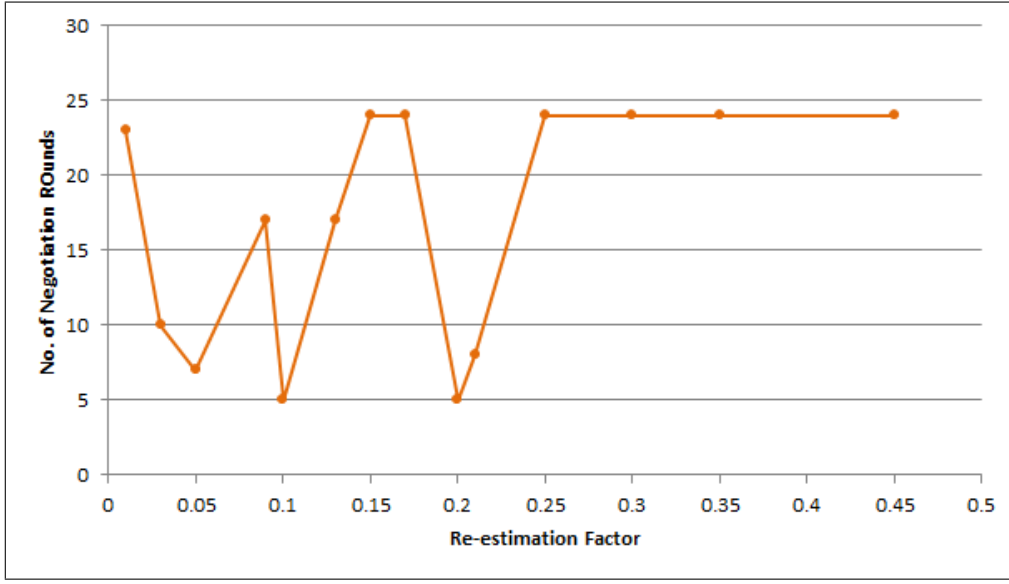
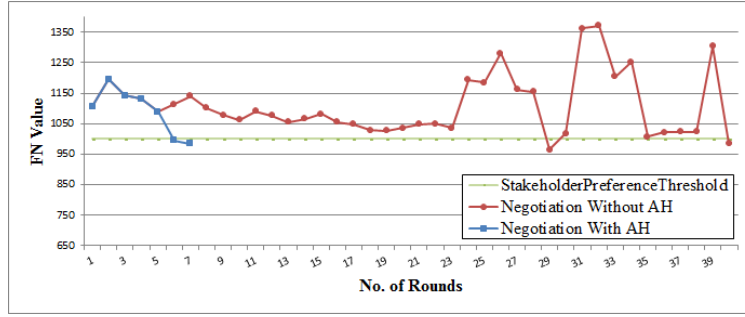


Figure 2.8: Effect of parameter η on rounds of negotiation

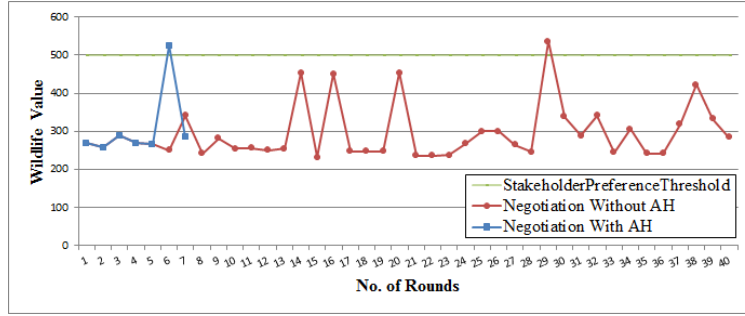
rounds decreases to 7 rounds. To have a better idea of the simulation time, the first experiment (i.e. without AH module) with 40 negotiation rounds takes about 0.425 seconds and the second one (i.e. with AH module) with 7 rounds takes around 0.15 seconds. About 40% of the running time is spent in BPPP module, which is the most time consuming element of the system. Note that the experiments were performed on a quad-core machine (Intel Core i7-860, 2.80 GHz, 8 GB RAM, Windows 7).

Fig. 2.9 shows the attribute values of the offered proposal in each round of negotiation before reaching an agreement with all the stakeholders. The attribute values of the offered proposals in the first and the second settings are illustrated by blue and red lines, respectively. In each sub-plot of Fig. 2.9, the green line represents the maximum value of the attribute which can be accepted by the relevant agent. For instance, the sign of reaching an agreement with the FN agent is that the “FN value” of the offered proposal becomes less than 1000 units. It should be noted that this value and other thresholds are completely hidden from the proposer agent.

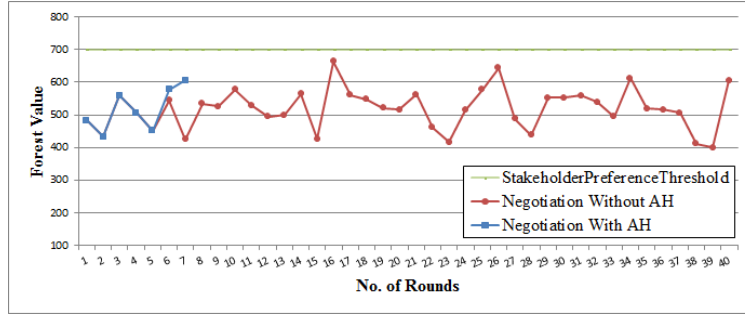
We have also conducted a set of experiments to examine how the BPPP approach affects the negotiation process. To this end, we introduced a disagreement distance measure for each proposal



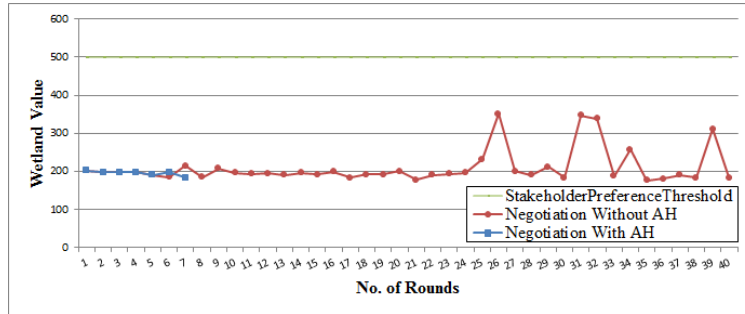
(a)



(b)



(c)



(d)

Figure 2.9: The attribute values of the offered proposals in two different settings- the case of 15 states; (a) First-nation values; (b) wildlife values; (c) forest values; (d) wetland values

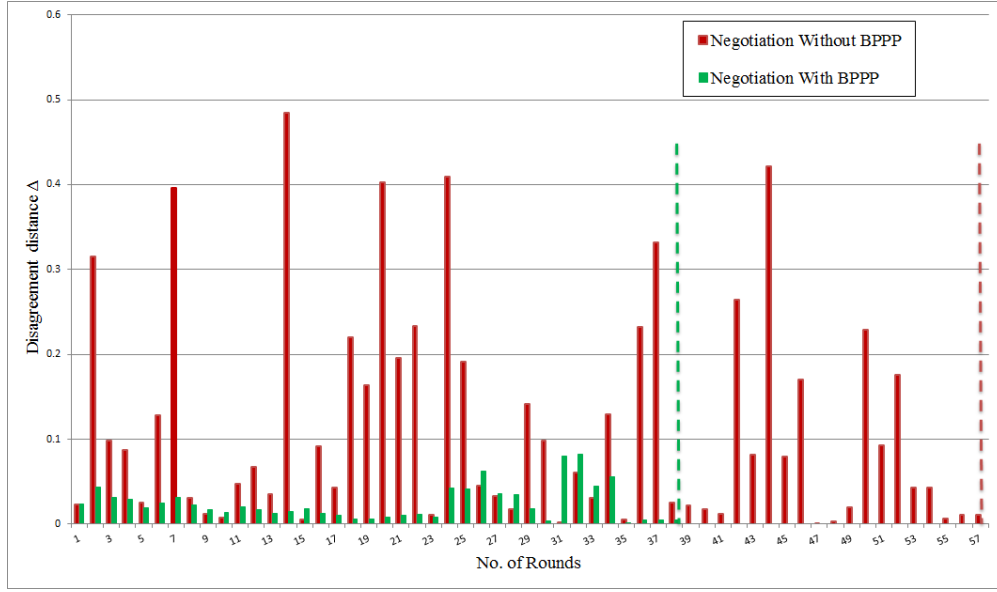


Figure 2.10: Disagreement distances in negotiation over the first case study with and without using the BPPP approach. The dashed lines represent the end of negotiation.

offered to the stakeholders. This measure is defined as the maximum of the average distances of normalized attribute values in the proposal to the normalized preference thresholds of the stakeholder. For proposal p_j , the disagreement distance measure is calculated as below.

$$\Delta(p_j) = \text{Max}_{i=1}^z \left(\frac{\sum_{k=1}^z (|a_k^j - T_{a_k, s_i}|)}{z} \right) \quad (2.12)$$

In Equation (12), a_k^j is the normalized value of attribute a_k in proposal p_j and T_{a_k, s_i} is the normalized threshold on attribute a_k according to the preference of stakeholder s_i . The distance $\Delta(p_j)$ measures the level of disagreement between the proposer and the stakeholders upon offering the proposal p_j ; the less the value of $\Delta(p_j)$, the higher the level of agreement. Accordingly, the negotiation terminates when $\Delta(p_j)$ reaches zero. Fig. 2.10 represents the disagreement distances when we run the negotiations with and without using the BPPP approach. This set of experiments has been conducted without using the AH model.

As illustrated in Fig. 2.10, without utilizing the BPPP approach, the disagreement distances fluctuate considerably. However, the BPPP approach reduces the number of fluctuations to a great

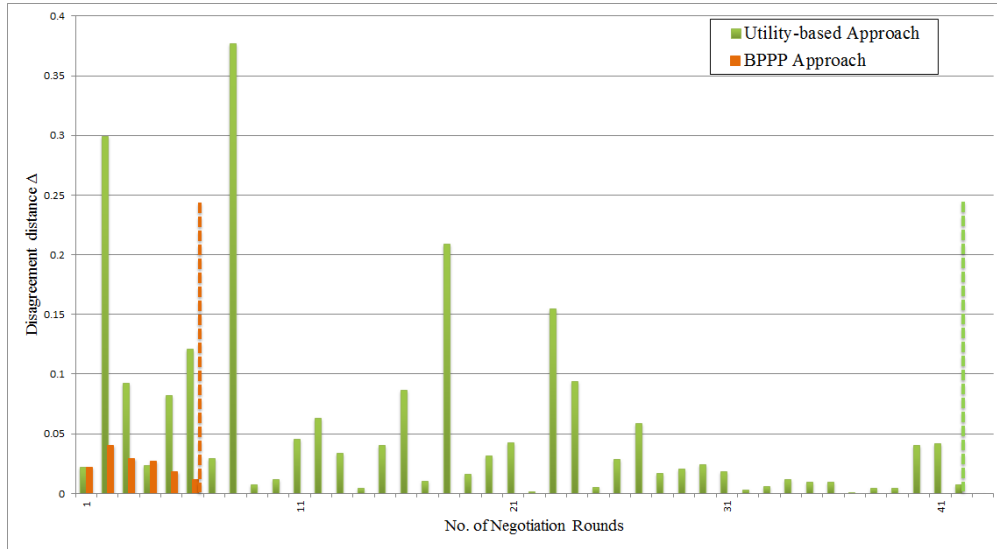
extent. This is because using the BPPP approach helps the proposer agent to take other stakeholders' preferences into account and therefore, reduces the disagreements to a great extent. It is shown in Fig. 2.10 that even without arguments, the BPPP approach can speed up the negotiation process up to approximately 1.5 times. The green and red dashed lines show where the negotiation ends with and without BPPP, respectively.

In another set of experiments the BPPP approach has been compared with the utility-based approach. Fig. 2.11 represents the results of these experiments with and without arguments. As illustrated in Fig. 2.11, the BPPP approach accelerates the negotiation process regardless of using the AH model. It is also concluded from both figures that the fluctuations in disagreement distances are smaller when we apply the BPPP technique. The dashed lines in the charts show where the negotiation ends for each setting ².

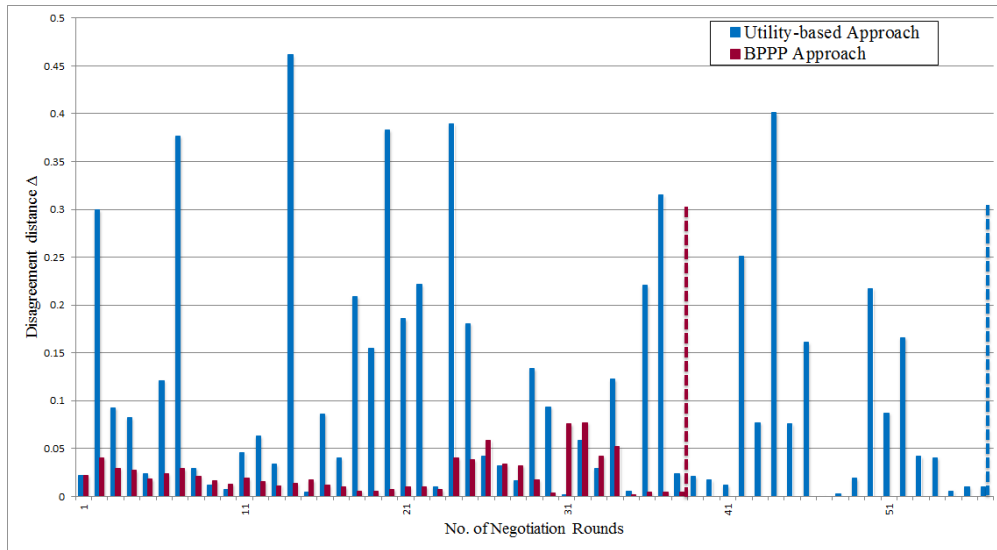
2.4.4 Case study B: King County House Sales

Although the first case study shows the efficiency of the proposed methodology, we introduced our second case study to ensure the performance, scalability, and applicability of the proposed methodology in other negotiation contexts with larger search spaces and higher dimensionalities. This dataset represents house sales in King County, US (Kaggle, 2017). The data provide different attributes of the houses sold between May 2014 and May 2015. The database contains 21614 records, which let us test the performance and efficiency of our proposed techniques on a broad set of data. We also have a larger set of attributes including price, number of bedrooms, number of bathrooms, the house area (in sq ft), number of floors, house condition, year of built and landscape view. Here, the negotiation happens between three agents: One representing the seller agent and two others representing the buyers who are, for example, a couple with different perspectives and preferences. For instance, while one of them is more concerned about the price (e.g PO agent), the other one cares more about the construction quality and the building area (e.g., QO agent).

²Figure 2.11b seems very similar to Figure 2.10, but the values in the charts are slightly different as they refer to two different experiments



(a)



(b)

Figure 2.11: Disagreement distances for a sequence of offered proposals in the Utility-based approach vs. the BPPP approach (a) With arguments (b) Without arguments. The dashed lines represent the end of negotiation.

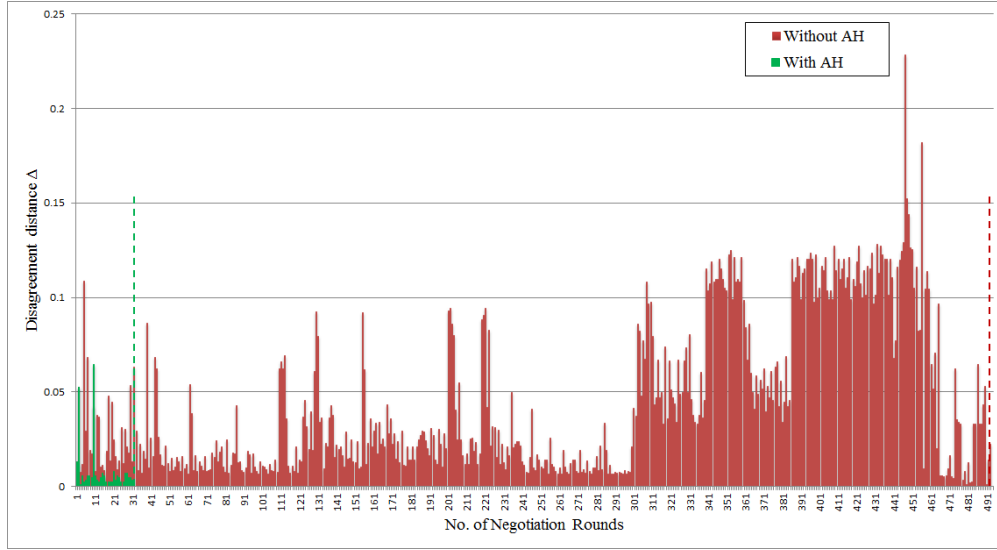


Figure 2.12: Disagreement distances in negotiation over the second case study with and without AH model. The dashed lines represent the end of negotiation.

Similar experiments have been conducted on this case study, and the results confirm what we have learned from our previous case study. Fig. 2.12 represents the level of disagreement between all agents through the negotiation. Here, the proposer agent employs the BPPP approach to prepare a proposal for each round. As depicted, without using the argument-handling module, it takes around 500 rounds for the negotiating agents to reach an agreement. However, when we applied the AH model, the agreement was reached in only 34 rounds.

We have also run a set of experiments to compare the BP-based proposal preparation approach with the utility-based approach. The results of the experiments, summarized in Fig. 2.13, confirms that the BPPP approach outperforms the utility-based approach when both of them benefit from the argument handling module. It takes 388 rounds for the negotiation with the utility-based approach to reach the agreement. However, the data shown in the chart is truncated to 300 rounds for improved legibility as the rest follows the same trend. The orange dashed line shows where the negotiation ends with the BPPP approach.

As Fig. 2.13 shows, there is a decreasing trend in the amount of normalized Δ when we use the utility-based approach. However, the result obtained from the BPPP approach does not follow

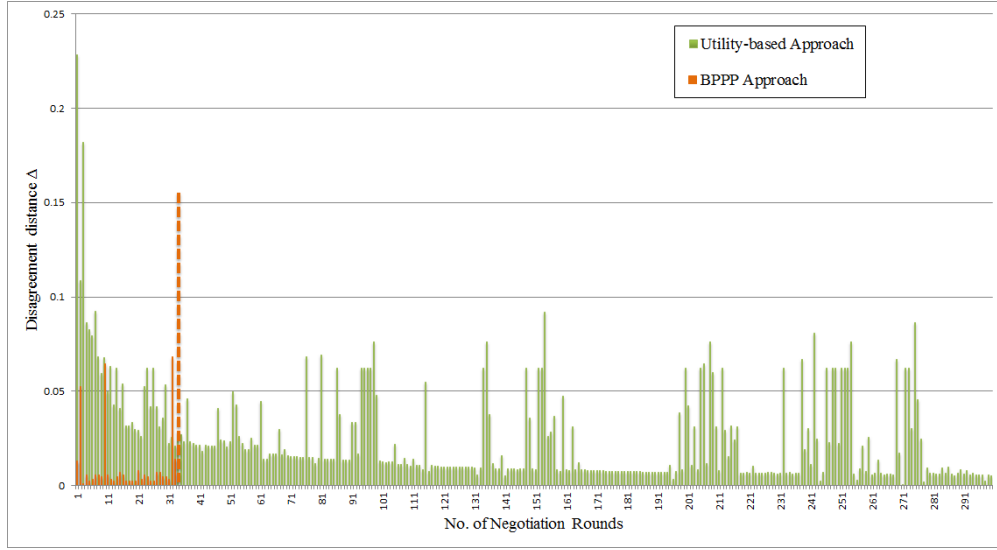


Figure 2.13: Disagreement distances for a sequence of offered proposals in the Utility-based approach vs. the BPPP approach. The dashed lines represent the end of negotiation.

any specific trend. In the utility-based approach, the proposer starts with his preferred attribute, which is the price attribute in this case study and ignores other attributes as well as other participant's preferences. Therefore, in the beginning of the negotiation, the level of disagreement is high. As the negotiation proceeds, the proposer improves its estimations about others' preferences through the arguments, and, therefore, the disagreement distance decreases. However, in the BPPP approach, the proposer starts as a neutral, giving all attributes equal importance. The only way he applies his own preferences is that he builds the unary cost tables based on his assumptions about the preferences of the other stakeholders. Then, the BPPP approach combined with the argument handling model allows him to adjust his assumptions in few rounds of negotiation.

2.5 Conclusion and Future Work

This paper represents a novel negotiation strategy in which a belief-propagation based technique, as well as argument handling, are employed to improve the efficiency of the negotiation process. While the belief-propagation-based approach improves the proposal selection in each round of

negotiation, the argument handling model provides the required information to feed this module. This information is gathered through agents communications about the past proposals and improve the proposer's learning rate to a great extent. To prove the ability of the developed methodology, a set of experiments has been conducted on two datasets, different in size and characteristics. The experiments on both environmental and non-environmental datasets confirm that argumentation-handling and the BP-based proposal preparation components facilitate and accelerate the negotiation process to a great extent. We have shown through experiments that the BP-based approach outperforms the utility-based approach regarding both the number of negotiation rounds and the fluctuations in disagreement distance measure.

In future, we will focus on more advanced models of stakeholders, where their decision-making process will be modeled using non-linear fuzzified functions. Also, the estimation process will be improved to adapt to the dynamic nature of the negotiations more appropriately. To build more complicated models for the stakeholders, more data ought to be gathered from real stakeholders; this will be done using a survey-based statistical technique such as the ones presented in (Truong et al., 2015) and (Cantillo et al., 2006).

Acknowledgements

This research project was partially supported by the funds available to the Schulich Chair in GeoSpatial Information Systems and Environmental Modeling held by Dr. Danielle Marceau at University of Calgary.

Chapter 3

Comprehensive Review of Machine Learning Approaches in Automated Negotiation: Examples in Environmental Resource Management

Article Presentation

Background

To achieve the second specific objective of the thesis, which is about learning the opponents' preferences in the context of automated negotiation, a comprehensive review is required to investigate and compare the learning approaches that has been used in the context of automation negotiation. This chapter is part of the thesis related to the this objective. It is published as "Automated Negotiation in Environmental Resource Management: Review and Assessment" in the Journal of Environmental Management in 2015.

This chapter investigates the learning techniques that have been applied in the context of automated negotiation with the aim of acquiring more knowledge about the participants. Four main machine learning approaches are discussed in this chapter and their potential in automated multi-issue multi-participant negotiations is evaluated. The negotiations over environmental issues are considered in this chapter as they usually occur among many stakeholders who, most of the time, are concerned about several issues and have very different viewpoints and preferences. This review allowed determining the shortcomings of existing automated opponent-learning approaches and applying such knowledge to develop a more functional solution to achieve the first and second objectives of this thesis in Chapter 4.

Automated Negotiation in Environmental Resource Management: Review and Assessment

by Faezeh Eshragh, Majeed Pooyandeh, Danielle J. Marceau

Published in:

Journal of Environmental Management

Volume 162, 2015, Pages 148-157

3.1 Abstract

Negotiation is an integral part of our daily life and plays an important role in resolving conflicts and facilitating human interactions. Automated negotiation, which aims at capturing the human negotiation process using artificial intelligence and machine learning techniques, is well-established in e-commerce, but its application in environmental resource management remains limited. This is due to the inherent uncertainties and complexity of environmental issues, along with the diversity of stakeholders' perspectives when dealing with these issues. The objective of this paper is to describe the main components of automated negotiation, review and compare machine learning techniques in automated negotiation, and provide a guideline for the selection of suitable methods in the particular context of stakeholders' negotiation over environmental resource issues. We advocate that automated negotiation can facilitate the involvement of stakeholders in the exploration of a plurality of solutions in order to reach a mutually satisfying agreement and contribute to informed decisions in environmental management along with the need for further studies to consolidate the potential of this modeling approach.

3.2 Introduction

Negotiation is one of the most common means for resolving conflicts in social interactions (Van Kleef et al., 2006). It can be defined as a discussion between two or more parties with conflicting interests aiming to reach an agreement (Pruitt and Carnevale, 1993). The involved participants may be individuals or groups of people who negotiate over single or multiple issues simultaneously. The agreement, which might be a mutually acceptable deal, new allocation of resources or new rules of behavior, has to satisfy all participants to some extent. A negotiation may also fail in case participants have nothing in common to agree on. In the past few decades, negotiation has been studied from different perspectives such as psychology (Pruitt and Carnevale, 1993), economics (Kreps, 1990) and computer science (Jennings et al., 2001). The aim is to understand the complicated nature of a negotiation process and make it more efficient and reliable in terms of exploring the

space of possible agreements, keeping track of negotiation rounds, and discovering negotiators' behavioral patterns. Environmental resource management is another domain that requires negotiation among stakeholders when one wishes to consider multiple viewpoints. This field of study deals with managing the effect of human activities on nature while guaranteeing the services provided by the natural resources (Pahl-Wostl, 2007). It is recognized that modeling tools designed to simulate negotiation of common-pool environmental resources, which are shared by a group of stakeholders and subject to overuse or congestion, can assist informed decision making (Ostrom et al., 1994; Bousquet et al., 1998). Involving stakeholders with different viewpoints helps reducing the complexity and uncertainties involved by providing an insight about the stakeholders' goals and preferences, and allows the capture of a diversity of interests to satisfy diverse expectations (Reed, 2008; Kenny et al., 2012). However, capturing the complexity of negotiation in such contexts is challenging. In addition to the conflicting preferences of the stakeholders, other factors such as power imbalance, time limitations, and the participants' attitude may also affect the results of the negotiation (Pruitt and Carnevale, 1993). Due to the large number of influential factors in the negotiation process, the space of all possible agreements can be hard to recognize and thus, difficult to be explored by human negotiators. Under such circumstances, some agreements, which could have been accepted by all participants, might have never been investigated. Additionally, stakeholders can act irrationally or have trouble keeping track of other parties' interests (Jonker et al., 2012). Considering all these concerns, a computational model can help to minimize the effect of biasing factors on the negotiation results and reach an agreement in a faster and more efficient manner. It can also be beneficial for understanding the complicated negotiation process and for engaging stakeholders into the decision-making process that could lead to better informed decisions. One of the modeling approaches in which stakeholders are involved from the early phases of the model development is participatory modeling. It includes companion (Barreteau et al., 2014) and mediated modeling (Van den Belt, 2004) that have been widely applied in environmental studies. In these modeling approaches, stakeholders get involved in the model construction (e.g., in mediated

modeling) as well as scenario simulations and result interpretation (e.g., in companion modeling). These approaches require a strong stakeholder's involvement over the whole modeling process, which might be difficult to obtain. Another important scientific approach in this domain, which has its roots in Artificial Intelligence (AI), is automated negotiation. It is a distributed search in the space of potential agreements, facilitated by an agent-based model (ABM), which consists of a set of intelligent elements, called agents, designed to mimic human behavior. Each agent represents a purposeful component of the system that acts autonomously in its environment to meet its predefined goals (Wooldridge, 1999). To better capture the complicated nature of human negotiation, machine learning (ML) techniques have been proposed to help the agents learn other participants' perspectives and utilize this information to enhance the negotiation. Automated negotiation was first employed in AI in the 1980s (Davis and Smith, 1983; Malone et al., 1983) where agents interact and negotiate to solve problems in a distributed way. With the widespread use of the internet and the World Wide Web, it has received a lot of attention in domains such as supply chain management (Fink, 2006), political studies (Aragonès and Dellunde, 2008), and especially e-commerce (Ramchurn et al., 2007; Jazayeriy et al., 2011). However, its application in modeling stakeholders' negotiation over environmental resources is still in its infancy (Akhbari and Grigg, 2013; Okumura et al., 2013; Pooyandeh and Marceau, 2013, 2014). This is largely due to the characteristics of environmental issues, such as the amount of uncertainty involved, the high-stakes decisions, the diversity of perspectives, and the inexistence of optimum solutions. This study was undertaken to better understand the challenges related to automated negotiation in order to exploit its full potential in environmental contexts. The objective of this paper is to review ML techniques currently employed in automated negotiation and evaluate their potential in terms of their compatibility with the nature of stakeholders' negotiation in the particular context of environmental resource management. It attempts at bridging the gap between the contributions made in Artificial Intelligence, Machine Learning, and Agent-Based Modeling in the field of stakeholders' negotiation. We advocate that due to the diversity of viewpoints when dealing with environmental issues,

automated negotiation can aid decision makers to explore uninvestigated solutions and therefore make more informed decisions. The remaining of the paper is organized as follow. In Section 3.3, the major concepts in automated negotiation are reviewed followed by a description of three major approaches, namely game theory, the heuristic approach, and the argumentation-based approach. In Section 3.4, four well-established ML techniques are described and compared based on a set of criteria that have been selected according to the most distinctive properties of negotiation contexts. This comparison is then used to evaluate their suitability for specific negotiation domains. Finally, in Section 3.5, guidelines are provided for the selection of appropriate learning techniques when modeling stakeholders' negotiation in the context of environmental resource management.

3.3 Automated negotiation

Automated negotiation consists of three main components: negotiation protocol, negotiation object, and negotiation strategy (Lomuscio et al., 2001). The negotiation protocol defines the set of rules governing the interactions between agents. It determines the possible types of participants, the negotiation states, the state transition rules, and the possible actions for each participant in each state. The negotiation object corresponds to the range of issues over which the negotiation happens. It may contain a single (single-issue negotiation) or multiple issues (multi-issue negotiation). When an agent makes an offer, in the simplest case, the set of issues and the range of values for each issue are fixed in the offer and the opponent agents can only reject or accept it. In a more complex form, in response to a proposal, negotiating agents are able to make a counter-offer by changing the issue values based on their own objectives. In more complex negotiations, agents are able to dynamically add or remove negotiation issues and make a proposal based on a new set of negotiation objects (Jennings et al., 2001). The third component is the agents' strategy, used by agents to act according to the negotiation protocol to reach a satisfactory agreement; it is basically the agent's plan for achieving its goals (Lomuscio et al., 2001). While the negotiation protocol is public and available to all participants, the agent's strategy is always private. Revealing the agent's

strategy can lead other participants to decipher its goals; in real-world negotiations stakeholders do not usually reveal their goals to negotiators to gain more benefits. Given the set of negotiation objects, the negotiation issues form the dimensions of the space of possible agreements. Automated negotiation can therefore be defined as a distributed search by negotiating agents in the space of potential agreements (Jennings et al., 2001). Each agent has its own mechanism for rating the points in the space and finds portions of the space that contain its acceptable agreements. Having an idea about other parties' agreement space helps the negotiating agents reach an agreement in a more efficient way. Three main approaches have been employed in automated negotiation: game theory, the heuristic approach, and the argumentation-based approach (Jennings et al., 2001). Game theory originates from a research conducted by Neumann and Morgenstern (von Neumann et al., 1944) and has its roots in economics. Games are well-defined mathematical objects with three main elements: the players of the game, the set of actions available to each player at each state of the negotiation, and the utilities assigned to possible outcomes. Game theory techniques use a set of rules, called solution concept, to find a strategy for each player to take the most rational action at each negotiation state (MacKenzie and DaSilva, 2006). To find the best choice of action, the agents assume that their opponents are rational (i.e. they try to optimize their outcome). These techniques have been used in the design of negotiation protocol and strategy. The designed protocols should be simple, Pareto efficient scalable, convergent to an agreement, and rational (Jennings et al., 2001; Lopes et al., 2008). A solution is called Pareto efficient when there is no other outcome that improves all participants' payoff (Kanbur, 2005). Game theoretic techniques have several advantages. They can be employed as a set of tools for the systematic analysis of negotiation contexts. They provide a clear view of different negotiation situations using mathematical analysis to determine the strategy that agents should follow to achieve the best possible outcome (Kraus, 2001a). When the negotiation context is stationary (i.e. occurs in a non-dynamic environment) and fully specified, it has been proven that game theoretic techniques can guarantee a solution that has all the desirable attributes such as Pareto efficiency as well as scalability and

rationality (Kraus, 1997). However, this approach assumes that all participants act rationally and that the set of agents' alternatives and the agreement space of each agent are known to other parties. These assumptions are not applicable in many negotiation contexts and thus, are restrictive. Moreover, this approach works very well for modeling interactions in special cases, for example in two-person games such as chess or poker (von Neumann et al., 1944), but is not applicable in more general situations (Zeng and Sycara, 1996). That is, if the details of the interactions (e.g., number of players or available actions in each state) change, the mathematical analyses are not applicable and the derived conclusions are not valid anymore (Binmore, 1992). The heuristic approach attempts to overcome the constraints of the game theory techniques. Rather than searching for the optimal solution, agents try to find a satisfactory, sub-optimal solution by reducing the search space to decrease the high computational complexity. Negotiating agents rate the points in the agreement space based on their utility function then exchange offers with other parties to find a mutually acceptable agreement. The process terminates either when the agreement is reached or the time limit for the negotiation is exceeded. Heuristic models are widely used (Barbuceanu and Lo, 2000; Aydogan et al., 2013; Costantini et al., 2013). For instance, Faratin et al. (1998) developed a heuristic model of multi-issue multilateral negotiation in which each agent employs a number of predefined tactics to gain more utility. An and Lesser (An and Lesser, 2012) applied a heuristic approach in the design of negotiating agents called Yushu in which agents employ simple heuristics for measuring the competitiveness of the negotiation and the time pressure, and use this information for making conservative concessions. Heuristic techniques also have their shortcomings. Since they do not search the whole agreement space, the outcome is not always the best. That is, better solutions may exist that have never been explored due to the approximation-based nature of this approach. Moreover, the results of these techniques are not reproducible, i.e. it is not guaranteed that the system produces the same results under the same conditions. Therefore, heuristic techniques should be evaluated through simulations and experimental analysis to prove that they act reasonably in different situations (Jennings et al., 1998; Baarslag et al., 2012). In game theory and

heuristic approaches, exchange of proposals is the only source of information that agents can use to know their opponents. However, a proposal is only a point in the agreement space that contains limited information about the proposing agent. Therefore, learning about the participants' agreement space and reaching an agreement is time consuming. In argumentation-based negotiation, in addition to the proposals, agents can also exchange supplementary information, called argument, which can help in reaching an agreement. An argument is defined as a piece of information that helps an agent to justify its stance or influence its opponents' stance by persuading them (Jennings et al., 1998; Rahwan et al., 2003). That is, instead of rejecting a proposal, an agent can explain which parts of the proposal are not acceptable and why. The agent can also provide arguments in form of rewards, treats or appeals to persuade other parties to accept an offer. As an example, an agent can threaten its opponent to terminating the negotiation. The reward is usually a promise about future interactions and the appeal is an explanation to clarify a situation and persuade other participants to accept the deal based on those rational explanations. Sierra and Jennings (Sierra et al., 1997) developed a model of negotiation between three types of agents involved in managing a business process in British Telecom. The model covers all forms of arguments (i.e. threat, reward and appeal). Ramchurn et al. (Ramchurn et al., 2007) developed a negotiation model in which an agent can persuade its opponent to accept an offer by giving it a reward. The agent can also ask for a reward in addition to the offered agreement to secure its benefits. Rahwan et al. (Rahwan et al., 2004) proved that arguments can change the portions of the agreement space that contain the agent's acceptable agreements and can improve the quality of the reached agreement. Since argumentation can help the agents to refine and reduce the search space by providing more meaningful information about other participants' viewpoint, the argumentation-based negotiations are more efficient compared to proposal-based approaches. Despite the considerable amount of research in the area of automated negotiation, several issues remain to be addressed. This is due to the complexity of automated negotiation in numerous aspects including the number of negotiation issues and the dependency between them, the shape and parameters of utility functions, the nego-

tiation protocol, strategy, and form (bilateral or multi-party), and the time limitations. To enhance automated negotiation, advanced AI techniques such as optimization and learning methods have been proposed (Marsa-Maestre et al., 2014). In the next section, the learning methods that have been widely applied in different contexts of automated negotiation are described and compared.

3.4 Machine learning techniques in automated negotiation

To ensure a successful negotiation, participants need to know their opponents and their preferences to propose more acceptable offers in limited time. However such knowledge is usually not available. One way to overcome this limitation is to investigate the exchanged proposals among participants to learn from them. A considerable amount of research has been devoted to the study of learning from opponent's moves. The proposed learning techniques vary in terms of learning objectives, the way knowledge is acquired as well as the required amount of historical information and the number of negotiation rounds. The ML techniques presented in this section are categorized based on the approach used to obtain the required information. They are then compared based on a set of criteria that includes: the number of negotiation issues, scalability of the learning technique, number of negotiation participants, required information at the beginning of the negotiation, learning objective, dynamics of negotiation (dynamic vs. stationary) and the ability to handle time constraints. The number of negotiation issues can be either single or multiple. In the reviewed literature, however, the number of issues does not exceed five. Because this number may rise in real-world negotiations, scalability is considered as an important factor in comparing learning approaches. It refers to the ability of a technique to manage a higher number of issues compared to the original number of issues that the technique is designed to deal with. The number of participants and the negotiation objective are other critical properties of negotiation that must be considered in the learning method. The required information for initiating a learning process is also important. In some negotiation applications, such as e-commerce, a huge amount of data about previous negotiations (within the same negotiation context) is available and can be used in

the training phase of the learning technique. However, in other contexts, such as environmental resources, it is difficult and sometimes impossible to include historical data as they simply do not exist. The ability of a learning technique to take into account the dynamics of the negotiation context is an important aspect to consider. Although many negotiation contexts are stationary and do not change over time, some others are dynamic in terms of negotiators' perspectives and other influential factors (e.g., market conditions). While some of the learning approaches can successfully deal with dynamic negotiation environments, others are not robust in these contexts. A final criterion is time limitation. In real-world negotiations, participants usually need to meet a deadline for achieving an agreement. Therefore, it is important that time is considered in the learning process so that the agents can achieve an agreement within a certain time framework. In the remaining of this section, the above criteria are used to evaluate the advantages and disadvantages of four well-established learning approaches that have been applied in automated negotiation: Bayesian learning (BL), reinforcement learning (RL), evolutionary learning (EL), and artificial neural network (ANN).

3.4.1 Bayesian learning

Bayesian learning employs Bayes rule to update the approximate probability of a hypothesis using the upcoming evidences (Langley, 1996). In the context of automated negotiation, the hypotheses set shows the information that an agent needs to learn about other parties and the initial probability of each hypothesis indicates the agent's belief about other negotiating agents. The learning process consists in using the recent evidences and information (e.g. a counteroffer from the other party) in the Bayes rule to update the probability of hypotheses. The hypothesis with the highest probability provides the learning agent with a model of its opponent, which helps it to make better offers in the next phases of the negotiation. Bayesian learning has been widely used in the context of automated negotiation because it provides the ability of learning new information from incomplete data. For example, Zeng and Sycara (Zeng and Sycara, 1998) presented Bazaar, which is a sequential decision-making negotiation model specifically designed for e-commerce. In this model,

agents use a Bayesian approach to learn the opponents' reservation price as a parameter of their strategy. Reservation price is a threshold for offer acceptance, i.e. the highest price that a buyer is prepared to pay for a product or the lowest price that a seller can accept. Buffett and Spencer (Buffett and Spencer, 2007) used a Bayesian classifier to learn the total order of preferences that can be employed for sorting the available proposals. The negotiation context is a special case of multi-issue negotiation, called multi-object, where participants negotiate over a set of objects and try to agree on a subset of the objects to be traded. The proposed approach is applicable in domains in which attributes of an offer can be stated as binary variables (e.g. having or not having a property) and a preference indicates the importance of an attribute over another one. More general cases of multi-issue negotiation have been modeled by Bui et al. (Bui et al., 1999) and Hindriks and Tykhonov (Hindriks and Tykhonov, 2008). Bui et al. (Bui et al., 1999) developed a model in which agents learn the evaluations of other participants of different proposals and then calculate a group evaluation function. This function provides the learning agent with an estimation of other agents' evaluation of different offers and is used to make the next offer. While this model provides a rough estimation of the offer evaluation value, the opponent model in Hindriks and Tykhonov (Hindriks and Tykhonov, 2008) contains the type of the evaluation function as well as the possible ranking of preferences. In this study, the opponents' most recent proposals are used to find the most probable model of its behavior using the Bayes' rule. This model is then employed to calculate the weights of preferences in the agent's utility function and the opponents' evaluation of an offer. Although Bayesian techniques have been used for different objectives such as learning the evaluation function (Hindriks and Tykhonov, 2008), the preferences (Buffett and Spencer, 2007), and reservation points (threshold) for a negotiation issue (Zeng and Sycara, 1998), they are mostly applied in stationary (non-dynamic) negotiation contexts. This means that in dynamic contexts in which the participants' perspective and attitude or other factors such as market conditions can change over time, the Bayesian approach has not been employed yet. A limitation of Bayesian learning is that it does not have an explicit tool for modeling time limitations as part of the learn-

ing process. Moreover, it requires information such as prior probabilities of the sets of possible outcomes, which has to be provided by a domain expert and might be difficult or impossible to acquire in some negotiation applications.

3.4.2 Reinforcement learning

Reinforcement learning (RL) is one of the most commonly-used learning techniques in automated negotiation. In this approach, the agents learn from the consequences of their actions rather than obtaining an explicit training. Based on the rewards received by the agent as a result of its previous actions, it selects the next action(s) in a way that maximizes its accumulated reward (Barto et al., 1990; Kaelbling et al., 1996). Different RL techniques have been employed to improve automated negotiation. One is the WoLF PHC algorithm, proposed by Bowling and Veloso (2001; 2002); it is a multi-agent reinforcement learning algorithm with variable learning rates, meaning that agents learn fast when they are losing, or cautiously when winning. The improved version of this algorithm was used by Shen et al. (2012) in an e-market place negotiation to make the negotiating agents capable of learning the current market factors and finding the optimal policy that maximizes their total rewards. In this study, an agent generates decisions after receiving offers. This decision, which can be accepting or rejecting the offer or proposing a counteroffer results in receiving a reward as a reinforcement to improve the agent's ability in making the next decisions. The acquired knowledge in the learning module is then used in the action selection module for future decision making. Another reinforcement learning approach that has been employed in automated negotiation is Q-learning. It is used for learning the evaluation function $Q(s,a)$, which helps to predict the expected utility of taking action a in a given state s . Having this function, an agent can simply select the action with the highest evaluation in each state to gain the maximum utility. Huang and Lin (2007) employed Q-learning in an argumentation-based negotiation between a seller agent and a human buyer. The Q-learning mechanism helps the seller agent learn the evaluation functions that can be used at different states to select the action. The learner agents, however, need to be trained using predesigned simulation games. In these games, a buyer agent, that mimics human buyer's be-

havior, is used to train the seller agent. The simulation has different rounds in which different types of buyer agents are used for training. This way, the seller agent is trained for dealing with different behaviors of human buyers. The model has to be also trained by humans. Monteserin and Amandi (2013) applied the Qlearning approach in an argumentation-based negotiation context, which is effective in both stationary and dynamic environments (e.g. when the participants characteristics change over time). In this research, the learning mechanism improves the argumentation selection policy using the agent's perceptions of its environment. That is, the agent is capable of handling dynamic environments by learning new selection rules or updating the existing ones. These rules shape a hierarchy, which can be used in selecting successful arguments. An advantage of reinforcement techniques is that they are able to handle dynamic contexts and time constraints (SHEN et al., 2012; Chen et al., 2013). However, they require training using historical data or simulations, while in many contexts (e.g., environmental resources) such information is not available. Another problem is the high number of required rounds for learning the utility of performing an action in a given state. Negotiators are supposed to learn the utility of an action in a given state, which is only possible when they can repeat that action several times in that particular state. That is rarely the case in real-world negotiations.

3.4.3 Evolutionary learning

Evolutionary learning is inspired by the evolution of species in nature. In evolutionary algorithms, a population of individuals evolves towards higher quality populations. The quality of a population, represented as a fitness function, is defined based on its value in the environment. In each round of the evolution, a set of individuals with higher fitness is selected and used for creating the next generation. This way the higher quality individuals have higher chance of reproduction compared to the lower quality ones. This approach has been implemented in three main groups of algorithms: Evolution Strategies (ESs) (Beyer and Schwefel, 2002), Genetic Algorithms (GAs) (Back, 1996), and Evolutionary Programming (EP) (Eiben et al., 2003). Among these, GA has been applied successfully in the field of automated negotiation. GA is a search between a set of

candidate solutions. It starts with a random population of solutions and evolves them to better solutions. Each solution has a set of chromosomes that represent its properties and can be mutated or altered in the next generation of solutions. Solutions are evaluated using a predefined fitness function. The evolution continues by producing new generations of solutions in each iteration and terminates when a satisfactory fitness level is reached (Back, 1996). Several researchers have applied GA to automated negotiation. Oliver (1996) employed GA to mimic human negotiators in finding Pareto-optimal agreements. In this study, strategies are modeled as simple sequential rules made up of offers and thresholds. The offer is presented by a set of properties (e.g. price, delivery date, etc.) and threshold represents the lowest acceptable payoff value of an offer. For example, a strategy with threshold T1 indicates that if the value of the upcoming offer is greater than T1, the offer should be accepted. Otherwise, the agent should reject the offer, make a counter-offer, and change the threshold to T2 to evaluate the next offer, if any. In each round of the GA process, a set of strategies with best payoffs is selected to generate a new population of strategies for the negotiating agents. The problem with this model is that the sequential rule model is not flexible enough to support different types of strategies. Tu et al. (2000) extended Oliver's research using a different representation of negotiation strategies. They employed Finite State Machine (FSM), which is a mathematical model for representing the behavior of a system. It consists of a set of states that represent the status of the system. The states of the system receive inputs; based on a set of predefined rules, the current state might change to another. Tu et al. (2000) used FSM as a representation of strategies in the GA algorithm to allow more expressiveness in terms of branching ability and memory. Lau et al. (2006) proposed a GA-based negotiation mechanism to find the fittest offer instead of the optimal strategy. The fittest offer is the offer that has the minimum distance to the agents' best offer (regarding the agent's own utility function) and the opponent's last counteroffer (which shows the opponent's point of view). In this study, chromosomes represent offers instead of strategies. The fact that preferences might change during the negotiation process is considered. Moreover, determining the fittest offer based on the opponent's last offer yields a gradual learning

from the other party's moves. The authors also provided models of different concession behaviors as well as negotiation under time pressure. The incorporation of learning, time management, and behavioral features make this model suitable for many real-world situations. Evolutionary learning techniques overcome the limitations of both Bayesian and reinforcement approaches. They can deal with dynamic negotiation contexts and time constraints and no historical data are required for their training (Lau et al., 2006). However, they are computationally expensive in terms of CPU and memory usage. Another problem with this learning approach is that the result might be only locally optimal. For example in Lau et al. (2006), while the negotiator looks for the best option for the next offer, the result of the GA search might not be a global optimum.

3.4.4 Artificial neural network

Artificial neural network (ANN) is a computational model that consists of several processing elements that receive inputs and deliver outputs based on their predefined activation functions. The network should be first trained to learn the hidden patterns in the input data represented in the model as the weights of the connections. It can then be used to predict the behavior of similar datasets (Gurney, 1997). Carbonneau et al. (2008) employed an ANN to find the opponent's counter-offer based on the previous offers. The proposed technique helps to predict the other party's next move before suggesting an offer and therefore makes the agents capable of performing what-if analysis. Roussaki et al. (2007) and Papaioannou et al. (2010) employed ANN in automated negotiation to discover the situations in which agreement is not achievable, in the early phases of the negotiation process. Lee and Ou-Yang (2009) developed an ANN to find the relationship between the previous bids and the supplier's next bid price. Using this relationship, the supplier's next bid price can be predicted, which helps the negotiation parties to reduce the negotiation time by making more informed decisions about their next offer. Although ANN is considered as a powerful tool for pattern recognition and prediction, it has its drawbacks. This learning technique requires huge amount of historical data for training that are not always available in real-world negotiation. Another problem is over-fitting. It occurs when a model is very

fit to a set of data because of excessive training; instead of finding underlying relationships in the data, it captures random noise and error in that specific dataset (Tetko et al., 1995). Therefore, the model is not general and has poor predictability for other datasets. In addition, because it is built in the training phase and cannot be changed afterward, it is not scalable in terms of the number of negotiation issues and cannot be employed in dynamic negotiation contexts. ANN is also unable to deal with time limitation, which is a common issue in real-world negotiations. A comparison of these learning techniques applied in automated negotiation is provided in Table 3.1. Studies listed in the table are considered among the most important ones based on their uniqueness and the number of citations they have received. Other learning techniques exist, such as volume-base measurement learning (Vetschera, 2009), ontology-based learning (Aydogan and Yolum, 2009), and fuzzy constraint-directed approach (Lai et al., 2010). They were not considered in this study because they have been used in only one context so far and therefore their general applicability and ability in dealing with different negotiation circumstances have not been demonstrated yet. The application of state-of-the-art AI and ML techniques remains limited in environmental resource management, in which negotiation plays an essential role. Studies that have been conducted in this field are presented in the next section, followed by guidelines for the selection of suitable techniques in this specific context.

3.5 Automated negotiation in environmental resource management

Bars et al. (2002) developed an ABM to simulate the negotiation between farmers and water suppliers. The negotiation process is a single-issue bilateral negotiation in which stakeholders negotiate over one issue (i.e., the amount of water) and the negotiation happens between two agents at a time. Unlike realistic situations, the negotiating agents do not know about the preferences of other participants and no learning procedure is employed. Another limitation of this study is that no convergence in negotiation is guaranteed. That is, negotiators exchange suggestions without taking into account the viewpoint of other stakeholders and in some cases negotiation may not

Table 3.1: Learning techniques applied in automated negotiation

| Learning Approach | Reference | Negotiation parameters | | | | | | |
|-------------------|-------------------------------|------------------------|-------------|----------------------------|--|---|-----------------|-----------------|
| | | No. of issues | Scalability | No. of parties | Required information | Learning objective | Dynamic context | Time constraint |
| Bayesian | (Hindriks and Tykhonov, 2008) | 4-5 | Yes | Bilateral | Initial probabilities of the sets of possible outcome + Certainty of agent about its opponent's idea | Weight ranking + Evaluation functions | No | No |
| | (Buffett and Spencer, 2007) | 5 | No | Bilateral | Initial probabilities of the sets of possible outcome | Participants' preferences | No | No |
| | (Bui et al., 1999) | 3 | Yes | Bilateral and multilateral | Not mentioned | Other participants' evaluation of a proposal | No | No |
| | (Zeng and Sycara, 1998) | 1 | Yes | Bilateral | A priori estimation of the opponent's reservation price | Opponents' reservation price | No | No |
| Reinforcement | (SHEN et al., 2012) | 3 | Yes | Multilateral | Historical negotiation data | Strategy (next action) | Yes | Yes |
| | (Chen et al., 2013) | 2 | No | Bilateral | Historical negotiation data + reward function | Strategy (next action) | Yes | No |
| | (Huang and Lin, 2007) | 1 | No | Bilateral | Historical negotiation data + simulation data | Evaluation functions for persuasion and negotiation + optimal functions | Yes | No |
| | (Radu et al., 2013) | 1 | No | Bilateral | Historical negotiation data + reward function | Preference coefficient | Yes | No |
| Evolutionary | (Lau et al., 2006) | 5 | Yes | Bilateral and Multilateral | N/A | Best offer for the next round | Yes | Yes |
| | (Oliver, 1996) | 4 | Yes | Bilateral | N/A | Strategy | Yes | No |
| | (Tu et al., 2000) | 4 | Yes | Bilateral | N/A | Strategy | Not mentioned | No |
| ANN | (Carbonneau et al., 2008) | 4 | No | Bilateral | Historical negotiation data | Other party's next offer | No | No |
| | (Roussaki et al., 2007) | 1 | No | Bilateral | Historical negotiation data | Other party's next offer | No | No |
| | (Lee and Ouyang, 2009) | 1 | No | Bilateral | Historical negotiation data | Supplier's bid prices | No | No |
| | (Park and Yang, 2006) | 4 | Yes | Multilateral | Historical negotiation data | Next offer | No | No |

converge to an agreement or even a disagreement. Belaqziz et al. (2011) tried to overcome this problem by considering the preferences of stakeholders in the negotiation process. They developed an ABM to simulate communications and social interactions between stakeholders in an irrigation system. In this model, the supervising agent who is in charge of managing the irrigation and the user of agricultural water agent negotiate over the amount of water. The supervising agent calculates the amount of required water for each farmer based on some parameters such as climate, annual amount of water, and soil occupation defined by the farmers. The supervising agent is aware of the other agent's preferences and tries to provide a proposal based on these preferences and his own interests. This model, however, does not represent the fact that in real-world negotiations, stakeholders have partial information about other negotiating parties and try to increase such information during the negotiation process. The convergence problem is not addressed either. That is, if the user of agricultural water agent keeps disagreeing with proposals, the negotiation will continue infinitely. To avoid the divergence problem and guarantee reaching a final decision, Akhbari and Grigg (Akhbari and Grigg, 2013) proposed a model in which only one agent (the state agent) is in charge of making the final decision considering the other agents' preferences and satisfaction. In this research, the state agent makes the final decision about a water allocation problem based on the agents' water demands, total available water, land area, and other environmental factors. The decision influences other agents' behavior in terms of their level of cooperation, which is how other agents take part in the decision making process. Here, the cooperation level is formalized as a function of the agent satisfaction and other parameters, such as social pressure and education in addition to incentives and penalties provided by the state agent. Another research that paid attention to the satisfaction of other agents as an influential factor in the decision making process was conducted by Kieser and Marceau (2011). In this study, each agent has its own preferences when evaluating proposed land development plans that have different weights corresponding to their importance in the decision making. A decision is made based on these evaluations and the level of satisfaction of the agents from previously made decisions. Although the negotiation pro-

cess is not mentioned explicitly in these two studies, the way in which these authors modeled the preferences and the satisfaction level of the agents provides useful insights when simulating real-world negotiations. In the aforementioned studies, the influence of stakeholders on each other's point of view is not considered. Regan et al. (2006) addressed this problem using a mathematical model proposed for a multi-issue negotiation process. In this model, each agent has an opinion about the expertise and rationality of other agents, which is represented as a weight for the respective agent, called respect weight. This weight indicates the ability of other agents to affect the agent's opinion about its preferences. At each iteration, the agents adjust the weight of their preferences based on the influence of the other agents and their own previous opinion. The process is repeated until all agents reach a same set of weights for the predefined preferences, which in turn leads to an agreement. The consensus is guaranteed in the ideal situation where the stakeholders are highly flexible and where reaching an agreement is mandatory, which is not often the case in real-world situations. A different issue that has to be considered in modeling the negotiation process is that in reality, stakeholders rarely have a crisp idea about their preferences. Pooyandeh and Marceau (2013) addressed this problem by integrating fuzziness in the modeling of stakeholders' preferences. In their model, the importance of each preference is defined as a fuzzy weight and a fuzzy analytic hierarchy process was employed to prioritize the stakeholders' preferences. During the negotiation process, agents can readjust the weights of their preferences in order to reach an agreement with other agents. Another aspect of real-world negotiation is dealing with the problem of incomplete information. Typically, in negotiations about environmental resources, information such as negotiation issues and the priority of an issue over another one are considered as public and can easily be shared by stakeholders. However, the negotiators are often not willing to share more private information, such as the way they evaluate an offer (i.e. their utility function) and the thresholds they use to accept or reject a proposal. In these negotiation contexts, so-called contexts with incomplete information (Bui et al., 1999), participants need to learn about other negotiating parties and the possibly changing negotiation environment (i.e. the conditions under which the

negotiation happens) in order to reduce the negotiation time and find the optimal or near optimal mutually accepted agreement. A few abovementioned studies consider the opponents' preferences as known information, but they do not take into account that negotiation issues might have different weights in the decision making process of a specific negotiator and that these weights are not usually known to the other parties (Belaqziz et al., 2011; Akhbari and Grigg, 2013). To address this problem, some researchers, such as Regan et al. (2006) use a mediator who has the required information about the negotiators including their preferences and the weight of each preference. Similarly, in Kieser and Marceau (2011), all the relevant information about the participant agents is known by a central party who is able to decide on behalf of negotiating agents. Although such a situation can happen in reality, in many real-world negotiations there is no mediator who can control the whole process. More realistic situations were modeled in which stakeholders' preferences and their weights are only known to themselves and not others. In the model developed by Pooyandeh and Marceau (2013), agents successfully reach an agreement by exchanging proposals. However, since the agents do not learn about each other during the negotiation process, the exchanged proposals may not improve over time and therefore, negotiation might take too long before reaching an agreement. Focusing on this issue, Okumura et al. (2013) built a model in which each agent estimates its utility function that reflects its preferences and their importance. Their proposed ABM uses a few sample points to generate park design proposals and then, using the agents' evaluations of these proposals, elicit the utility function of each agent. This function can be used later in the automated negotiation by improving the ongoing proposals and making them closer to the users' viewpoints (Ito et al., 2007; Fatima et al., 2009). Pooyandeh and Marceau (2014) extended their previous model (2013) and implemented a Bayesian learning mechanism so that the agents obtain information about other participants' perspectives throughout the negotiation and can learn from other parties' previous offers. Bayesian learning allows an agent to learn the evaluation function of other participants and then use the learned information to provide more acceptable proposals to them while maintaining its own interests. The results reveal that incorporating the learning

capability reduces the number of negotiation rounds. Moreover, it provides a model of human negotiation that adequately mimics real-world situations by explicitly considering the learning ability of humans. Table 3.2 provides a list of applications of automated negotiation in environmental resource management. It can be seen that very few of these studies incorporate learning mechanisms and that important criteria such as scalability, time constraint, and the dynamic nature of the negotiation context are also rarely considered. While there is an increasing need for improving automated stakeholders' negotiation in the particular context of environmental resource management, choosing an appropriate AI technique is not straightforward. Guidelines for the selection of appropriate learning techniques are provided in Table 3.3. It can be seen that all techniques are suitable when historical data are available. However, for applications in which acquiring such data is not feasible, ANN and reinforcement learning are not the best options while evolutionary learning and Bayesian techniques are more appropriate. Similarly, in applications where the context is not dynamic, all learning techniques are suitable. When the context is dynamic, ANN is particularly not recommended due to the fact that it is not flexible enough (i.e. by changing a parameter in the application the whole network should be rebuilt) while evolutionary learning has been proved to act successfully in changing environments. When time is a critical factor in the negotiation, selecting a learning technique able to handle time constraints is of high importance. Although negotiations related to environmental resource management issues are not usually under high time pressure compared to the ones in e-commerce, time constraints must still be satisfied. In that case, evolutionary and reinforcement learning techniques are recommended. Scalability is important in applications such as land development where the number of involved stakeholders and negotiation issues might greatly vary from one case to another. Evolutionary and Bayesian learning are more appropriate for such applications. When considering the largest combination of constraints, namely the lack of available historical data, ability to handle dynamic context, time constraint, and scalability, evolutionary learning is the most suitable technique. Bayesian learning can be used in the absence of historical data and can handle scalability while reinforcement learn-

Table 3.2: Applications of automated negotiation in environmental resource management studies

| Reference | Application | No. of issues | Scalability | | Learning | Dynamic context | | Time constraint | |
|-------------------------------|--------------------------|---------------|-------------|------------------------|----------|-----------------|-----------------|-----------------|-----------|
| | | | Need for | Is the model scalable? | | Need for | Addressed | Need for | Addressed |
| (Le Bars et al., 2002) | Water management | Single | Yes | No | No | Yes | No | Yes | No |
| (Belaqziz et al., 2011) | Water management | Single | Yes | No | No | Yes | No | Yes | No |
| (Akhbari and Grigg, 2013) | Water management | Single | Yes | No | No | Yes | No | Yes | No |
| (Kieser and Marceau, 2011) | Land development | Multiple | Yes | No | No | Yes | No | Yes | No |
| (Regan et al., 2006) | Environmental management | Multiple | Yes | Yes | No | Yes | No | Yes | No |
| (Pooyandeh and Marceau, 2013) | Land development | Multiple | Yes | Yes | No | Yes | Yes (fuzziness) | Yes | No |
| (Okumura et al., 2013) | Park design | Multiple | Yes | Yes | Yes | Yes | No | Yes | No |
| (Pooyandeh and Marceau, 2014) | Land development | Multiple | Yes | Yes | Yes | Yes | Yes (fuzziness) | Yes | No |

Table 3.3: Suitability of learning approaches in automated negotiation for environmental resource management

| Feature | Application characteristic | Suitable learning techniques |
|---------------------------------|----------------------------|------------------------------|
| Availability of historical data | Yes | ANN, RL, EL, BL |
| | No | EL, BL |
| Dynamic context | Yes | EL, RL |
| | No | ANN, RL, EL, BL |
| Time constraint | Yes | EL, RL |
| | No | ANN, RL, EL, BL |
| Scalability | Yes | EL, BL |
| | No | ANN, RL, EL, BL |

ing can deal with dynamic contexts and time constraints. ANN is the least suitable technique in dealing with the four features often encountered in environmental resource management.

3.6 Conclusion

While environmental resource management is a domain in which negotiation plays a critical role, the application of automated negotiation techniques to model stakeholders' negotiation remains limited due to the unique characteristics of the decision making process in this field. Beside the large number and sometimes conflicting perspectives that must be considered, the high risk often associated with the decisions and their potentially long lasting influence on the environment makes

the decision making process far more complicated than it is in other domains, such as e-commerce. The dynamic nature of the environment and its heterogeneity over time and space accentuate the complexity of the management issues, especially when decisions have to remain in force for long periods of time. Considering the complex and dynamic nature of environmental resource management, the large number of perspectives as well as the increasing need for transparency in the decision making process, involving stakeholders and taking their viewpoints into account is increasingly recognized as crucial for the success of environmental projects. However, in most of these projects, stakeholders are still either ignored or invited to play minor roles in the very last phases of the decision making process. Participatory techniques have been proposed to increase the engagement of stakeholders from the early stages of environmental decisions. These approaches, however, require a strong stakeholders' commitment over the whole modeling process, which is not always feasible. In comparison, automated negotiation based on AI and learning methods benefits from advanced search techniques and computational heuristics to handle the complexity of the negotiation space. While in participatory modeling, stakeholders are in charge of determining scenarios to be tested and investigated by the model, in automated negotiation it is the model that proposes different scenarios based on stakeholders' preferences through the exploration and filtering of the large agreement space. Besides, automated negotiation is capable of capturing aspects of human intelligence, such as emotions, ability to design a strategy, learn and adapt, that are crucial for the success of negotiation. The overall goal of this paper was to bridge the gap between the research contributions made in automated negotiation from the disciplines of AI, machine learning, and agent-based modeling to take advantage of the potential offered by automated negotiation in environmental resource management and lighten the burden of using these technologies for interested researchers. In such contexts, automated negotiation methods need to be selected cautiously to ensure their compatibility with the inherent complexity of environmental issues. Moreover, they should be considered as supplementary tools to human interactions and decision-making process, not a substitution. Gaining the stakeholders' trust is also essential, which requires involving them

in the definition of the problem and obtaining their feedback on the model design and outcomes. The modeling team must also guarantee the confidentiality of the information gathered from these stakeholders. We advocate the need of additional studies to further investigate the applicability of automated negotiation in environmental contexts. For instance, the use of threat arguments in argumentation-based negotiation is a representation of power imbalance, which might not be desirable in some negotiations. Also, a large amount of literature in the field of automated negotiation focuses on reaching an optimal or near-optimal agreement, which is not usually the case in environmental studies. It is mainly due to the fact that there is no best (optimal) solution when dealing with environmental resource management problems because of the high level of uncertainties involved and the dynamic nature of the context. Employing appropriate tools and processes such as adaptive management (Holling, 1978), an iterative decision making process that involves learning about the problem context and adapting over time, can improve automated negotiation models to tackle uncertainties and facilitate long-run management decisions.

Acknowledgments

This study was funded by a Discovery grant awarded by the Natural Sciences and Engineering Research Council of Canada (NSERC) (138414-2013) to D. Marceau and a post-doctoral fellowship awarded to M. Pooyandeh by MITACS.

Chapter 4

Real-time Opponent Learning in Automated Negotiation using Recursive Bayesian Filtering

Article Presentation

Background

This chapter is part of the thesis related to the first and second objectives, namely modeling the evaluation outcomes received from stakeholders and developing an approach that helps the proposer agent to learn about its opponents by estimating such models. This chapter is published in the Journal of Expert Systems with Applications in 2019.

Considering the fact that the proposer agent has no prior knowledge about its opponents and no data is available from previous similar negotiations, learning about the negotiation participants is really challenging. This problem is even more complicated when the participants have various preferences about multiple issues, and there is no way of knowing these preferences other than the feedback the proposer agent receives about the offered proposals. In this chapter, a recursive Bayesian filtering approach is proposed as a solution to this problem.

General Methodology

Given the stakeholders' decisions about an offered proposal, the proposer agent applies a learning mechanism based on unscented particle filtering to build a more accurate model of its opponents. Through several phases of prediction and correction, illustrated in figure 4.1 (opponent learning component), a new estimation of the opponents' criteria is acquired which is used later in the proposal-preparation component to update the MRF model. The article presented in this chapter explains the steps of modeling and learning the stakeholders' preferences in more details.

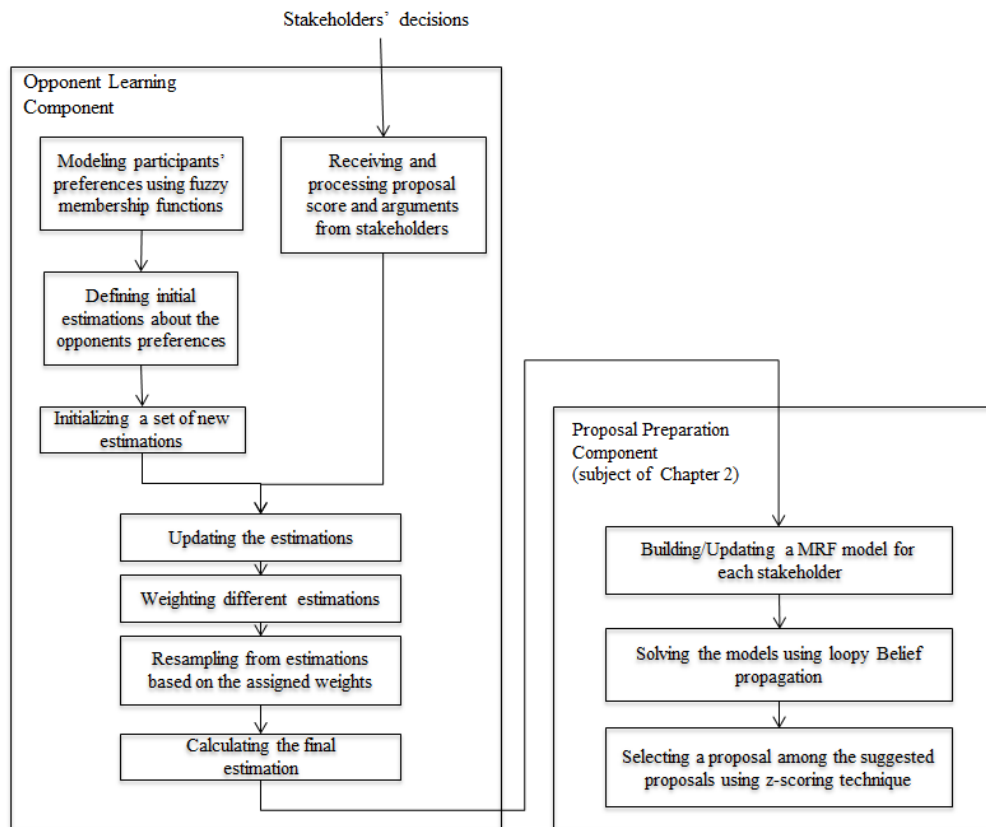


Figure 4.1: Opponent modeling and learning flowchart

Real-time Opponent Learning in Automated Negotiation using Recursive Bayesian Filtering

by Faezeh Eshragh, Mozhddeh Shahbazi, Behrouz Far

Expert Systems with Applications, Elsevier Journal

Volume 128, 2019, Pages 28-53

4.1 Abstract

Automated negotiation is a toolset to model human interactions during a negotiation process with the aim of improving the efficiency and quality of decision-making using advanced information analytics. During the negotiation, the participants share their viewpoints and concerns about the negotiation issues. However, in reality, they usually do not reveal the details of their preferences to one another. Therefore, modeling and learning opponents' behavior is a crucial component of automated negotiation. In this paper, we propose an estimation technique based on recursive Bayesian filtering to facilitate opponent-modeling and -learning in the context of multi-participant, multi-issue negotiations. In the proposed technique, opponents' preference profiles are modeled using fuzzy functions, which are very close to the way humans evaluate alternatives. As the negotiation progresses, the agents can recursively learn the parameters of these models in real time. The only required information for this learning process includes the feedback and the arguments the participants may provide in support of their decisions. At each round, a probabilistic graphical model is also implemented that utilizes the learned preference limits of the participants to offer a new proposal with a high probability of satisfying the participants and reaching an agreement. The proposed methodology is examined in two different negotiation contexts: energy-system development and real estate service. The experiments show that the proposed opponent modeling/learning approach increases the efficiency of the negotiation up to 85% and facilitates reaching an agreement in fewer rounds of negotiation without requiring any prior understanding of the negotiation participants.

4.2 Introduction

Negotiation, as a mean for conflict and disagreement resolution among multiple stakeholders, has been studied from different perspectives in the past few decades. Automated negotiation, implemented as a multi-agent system, is a tool for modeling human negotiations and facilitating decision-making. Each agent in an agent-based negotiation model represents a stakeholder that

interacts with others to reach an agreement. Automated negotiation has three main components: negotiation object, negotiation protocol, and negotiation strategy (Lomuscio et al., 2003). The negotiation object determines the range of issues discussed by the negotiators. These issues are the attributes of the exchanged proposals among the negotiators. The terms “negotiation issues” and “proposal attributes” are used in the text interchangeably. Depending on the context, the negotiation might be over a single issue (single-issue negotiation) or multiple issues (multi-issue negotiation). The second component of automated negotiation, called negotiation protocol, determines the number of participants, their roles in the negotiation process, different states of the negotiation, the state transition rules, and the set of possible actions for each agent in each state of the negotiation. The negotiation protocol is public and available to all the participants. The third component is the agent’s strategy that defines the agent’s behavior towards a satisfactory agreement based on the agent’s goals (Lomuscio et al., 2003).

Automated negotiation has been applied in different domains such as e-commerce (Ramchurn et al., 2007; Jazayeriy et al., 2011), supply chain management (Fink, 2006; Lee, 2014), e-marketplace development (Renna, 2010), natural resource management (Eshragh et al., 2018; Rodriguez-Fernandez et al., 2019), task and resource allocation (Bigham and Du, 2003; Lin et al., 2006a), event scheduling (Hossain, 2012), and politics (Aragonès and Dellunde, 2008). For example, Lee (2014) designed an interactive bidding strategy for both suppliers and customers in a supply chain. Another application of automated negotiation is studied by Renna (2010), where negotiation policies were proposed with the aim of positioning Small and Medium Enterprises (SMEs) in e-marketplaces. Different combinations of negotiation policies and customers’ tactics were then investigated to find out which combination leads to better negotiation outcomes.

Depending on the context in which the negotiation happens, the automated negotiation model and its components may vary. In this paper, our focus is on multi-issue, multi-participant negotiations where we have a one-to-many type of negotiation. That is, at each round, one of the participants is in charge of offering a proposal, and the rest of the participants can independently

agree or disagree with it. In the upcoming rounds, the proposal is revised and resent to the participants. This process ends when all parties agree over the offered proposal. It is important to mention that the negotiation environment studied in this paper is not an open environment as introduced in (Li et al., 2013) and (Resinas et al., 2012). The negotiating agents in such an environment can revise their evaluation criteria in each round of negotiation based on their new observations. For example, an agent can take part in two or more parallel negotiations and, depending on the deals it receives in other negotiation(s), it can change its current criteria. Although this framework is closer to real-world negotiation scenarios, it is not the approach followed by this paper.

Commonly, there are three approaches to implement automated negotiation: the game theory, the heuristic approach, and the argumentation-based negotiation (ABN). In this paper, we follow the third approach since it is a better fit for the context of our target case studies. In the first two approaches, interactions between negotiating agents are mainly through exchanging proposals. However, in the argumentation-based approach, in addition to proposals, agents exchange supplementary information, called arguments, which give them a chance to interact more effectively (Jennings et al., 2001). Argumentation-based negotiation has been studied in the literature using three major approaches (Karunatilake et al., 2009): The first one, called argumentation-based defeasible reasoning, is mostly about analyzing the relationships (e.g., support, attack, conflict) among the received arguments. The argumentation system gathers the received arguments and compares them with the previous arguments to find the possible conflicts. It, then, uses the relationship among the arguments to resolve those conflicts and update the agent's knowledge only based on reasonable arguments. (Chesnevar et al., 2000; Prakken and Vreeswijk, 2002; Tamani et al., 2015; Thomopoulos et al., 2015). The second approach works based on rhetorical statements that should be passed during the negotiation to persuade a negotiation party to accept a proposal. In such case studies, the arguments are mainly defined as structures (schema) for generating persuasive feedback (Gilbert et al., 2004). In the third group of studies, arguments are a way of interaction between negotiating parties. In this approach, the parties use a communication

language to generate the arguments under the structure of dialogue games. The arguments are usually defined based on a set of predefined rules governing the negotiation (Karunatilake et al., 2009). Argument handling in this article follows the third approach. Arguments are represented as phrases justifying the agents' decisions for rejecting an offer. These arguments are used in the negotiation process by the agents to acquire knowledge about their opponents.

One of the critical aspects of negotiation is acquiring knowledge about the participants' preferences. During the negotiation, the participants share their viewpoints and concerns about the negotiation issues so that they can settle the differences and reach an agreement. However, in many contexts, negotiators usually reveal only part of their preferences, and the rest remains hidden from the other parties (Niemann and Lang, 2009). For example, when negotiating over the price of a product, the buyer may give the seller some clues about his/her available budget but, he/she usually does not let the seller know his/her exact price preferences (Zeng and Sycara, 1998). In such cases, the negotiation parties need to learn about their opponents using the exchanged proposals, opponents' decisions and, in some cases, supplementary information such as the arguments. The learning/estimation process is called opponent modeling in the literature of automated negotiation (Baarslag et al., 2016).

Depending on the negotiation context, opponent modeling can be defined based on three different concepts: Opponent type (what type of characteristics the opponent has); Opponent's strategy (what the opponent will do); and Opponent's preference profile (what the opponent seeks) (Baarslag et al., 2016). The focus of this study is modeling the opponents' preference profiles. Each negotiating agent has a preference profile by which it ranks or scores the proposals. Scoring is stronger than ranking as ranking can be only used for comparing two proposals while scoring reflects the intensity of the difference between their importance as well. Preference profile is usually part of the information that the agents do not reveal to their opponents and can be very complicated, especially as the problem space enlarges. In such cases, the agents need to use more compact ways such as conditional preference networks (CP-Nets) or utility functions to represent

their preference profiles. CP-Nets (Boutilier et al., 2004) are graphical models in which negotiation issues are represented as nodes, the correlation between the nodes is modeled with edges, and the importance of one node over another one is denoted by the weight of the directed edge between them. Utility functions, on the other hand, assign a score to each proposal based on the values of its attributes. The most common type of utility functions is the linear additive one, where the utility of a proposal is the linear combination of the utilities of its attributes (Raiffa et al., 2002; Chen and Weiss, 2015). More complex non-linear versions of utility functions are also introduced in the literature (Ito et al., 2011; Marsa-Maestre et al., 2014).

In this study, a multi-agent system (MAS) has been employed to model a one-to-many type of negotiation among several stakeholders. Each agent in the model represents a (group of) stakeholder(s) with a set of criteria and preferences that are designed based on the stakeholders' concerns and interests. The stakeholders comprise the groups or individuals who are either involved in the decision-making process or affected by the final decision (Freeman, 2010). In this paper, we focus on estimating the stakeholders' preference profiles and particularly the utility functions required to evaluate the proposals (also known as proposal evaluation models). The utility function of each agent is built based on its criteria and preferences. Here, we consider a fuzzified version of participants' preferences based on soft weighted preference limits rather than binary rejection/acceptance models based on hard thresholds. That is, a stakeholder can assign a fuzzy score to a proposal to evaluate it.

Consider the following example in the case of selling/buying real estates. In this example, three stakeholders can be considered, one representing the seller and the other two representing the buyers (e.g., a couple). For the seller, it is important to persuade the buyers to buy an expensive place so that he can gain the highest possible profit. For one of the buyers, the price and the year built of the place are priorities while for the other one, the number of bedrooms, the square footage of the place, and the distance of the place to the nearest school are significant factors.

Two types of agents are defined in this type of multi-agent system. The first one is the proposer

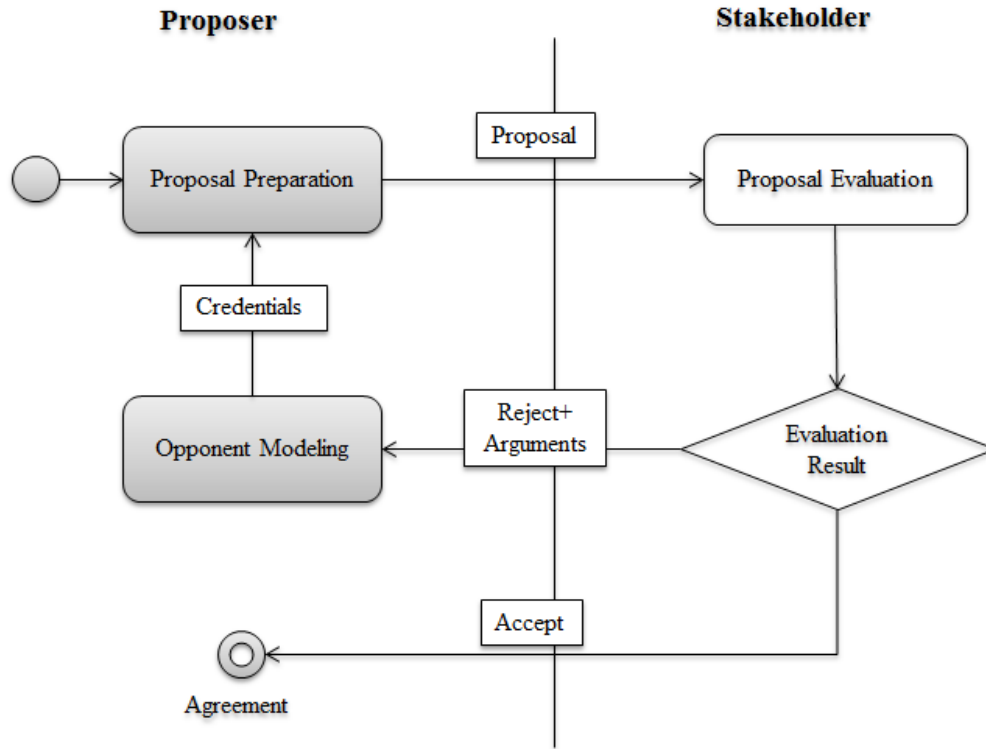


Figure 4.2: Negotiation Process

agent, who leads the negotiation and prepares and offers proposals in each round of negotiation. The second one includes other agents who receive the proposals and decide whether to agree or disagree with them; they may also provide arguments to the proposer agent. In the real estate example, the seller agent is the one who offers properties (proposals), and the buyer agents are the ones who decide about the received offers. Figure 4.2 shows the negotiation process considered in this paper.

Each agent in the MAS has access to a set of private and public knowledge bases containing information about itself and sometimes about the other agents' goals and preferences. A part of the agents' knowledge that relates to the other participants is gathered either from public documents and datasets or directly from the stakeholders. It may also change during the negotiation as a result of the learning process. The proposer agent has a database of all possible alternatives. For example, in the real estate case, the proposer agent has access to a dataset of all available real estates in the market and uses this dataset during the negotiation to find suitable proposals considering his

priorities and the preferences of the buyer agents. The proposer agent also has access to a database of previously offered proposals. These records can be used for future references (Rahwan et al., 2005).

The proposer agent needs to learn about other agents' preferences and criteria so that it can estimate their utility functions and reach an agreement with them in fewer rounds of negotiation. We, therefore, propose an estimation technique based on recursive Bayesian filtering (RBF) to approximate the parameters of the agents' utility functions. The only requirement of the proposed learning framework is the participants' feedback in each round of negotiation. This feedback includes the score assigned to the offered proposal as well as the stakeholders' arguments about it. Since the stakeholders' arguments are in the form of natural language sentences, they are first analyzed using a text processing module and then passed to the estimation module as inputs.

4.2.1 Objective and Contributions

The general goal of this study is making the negotiation process faster and more efficient by learning the participants' preferences. To this end, the opponents' preferences are modeled using fuzzy utility functions. Thus, the specific objective of this study is to estimate the parameters of these utility functions without requiring any before-hand knowledge about the negotiators or the negotiation context. The contributions of this paper are threefold.

1. Providing the proposer agent with the ability to recursively learn the parameters of the stakeholders' fuzzy evaluation models in short rounds of negotiations. The only required information is the scores the stakeholders assign to the offered proposals as well as any argument they may provide in support of their decision. The efficiency of this learning process is high even in multi-issue, multi-participant negotiations.
2. Processing the learned preference limits of the stakeholders and rebuilding the probabilistic graphical model of the stakeholders and negotiation issues. This graphical model is used by the proposal preparation module to find a new proposal that has

a high probability of satisfying the stakeholders. Although this inference approach is proposed in our previous work (Eshragh et al., 2018), its integration with the opponent-learning approach (contribution 1) is original to this paper.

3. Examining the proposed learning approach in two case studies of energy-system development and real estate service. These are two challenging negotiation cases in terms of the involvement of conflicting stakeholders and the large size/dimension of the solution space.

4.2.2 Assumptions

We assume that:

- There is one agent responsible for offering proposals, collecting answers and developing opponents' models. The other agents do not communicate with each other.
- Opponents' are logical and do not act randomly; i.e., their strategies do not change during the negotiation.
- The set of negotiation issues (attributes) is defined before the negotiation starts and cannot be modified during the negotiation.

4.3 Related Work

Opponent modeling has received much attention in the field of automated negotiation. Depending on what the negotiating agents need to know about each other, the objective of the modeling process can be categorized into three different groups: opponent type (the type of player the opponent is), opponent strategy (the way the opponent acts) and opponent preferences (what opponent is interested in). These categories may overlap as they are highly related aspects of the negotiating agents (Baarslag et al., 2016).

Learning the opponent type helps the proposer agent to predict the way it negotiates and how it should be replied. For example, in (Lin et al., 2006b) a finite set of agent types with different characteristics is considered, and an exact pre-defined utility function is assigned to each type. For this approach to be efficient, one needs to know a limited number of opponent types in advance, which is a limitation in many contexts.

Learning the opponent strategy is another approach in opponent modeling. It is about predicting the opponent's actions and their order of occurrence. Acceptance and bidding are the most common strategies that are modeled by this approach. Different learning techniques, including but not limited to, Bayesian learning (Ji et al., 2014), non-linear regression (Haberland et al., 2012; Brzostowski and Kowalczyk, 2006), artificial neural networks (Fang et al., 2008; Carbonneau et al., 2011) and kernel density estimation (Oshrat et al., 2009), are used to estimate such models. In the literature, modeling the acceptance strategy is usually simplified to estimating the probability of proposal acceptance (Lau et al., 2008; Oshrat et al., 2009; Chen et al., 2016). For example, in (Chen et al., 2016), the authors use the advice of the crowd (i.e., acceptance or rejection labels on proposals provided by a large group of related agents) to predict the chance of acceptance or rejection of a proposal. When we have more than one proposer agent, the bidding strategy is a popular objective in opponent modeling. The bidding strategy determines the proposal the agent offers in the next rounds of negotiation. Learning the opponent's bidding strategy helps the agent to decipher its goals and find the best possible deal based on them (Masvroula, 2013; Rajavel and Thangarathanam, 2016). For example, in (Rajavel and Thangarathanam, 2016), the authors propose an adaptive probabilistic behavioral learning system that works based on analyzing the sequence of proposals received during the negotiation process.

The third class of opponent modeling, which is also the focus of this paper, considers the opponent's preference profile (i.e., what the opponent cares most about). Preference profile describes the way an agent evaluates the received proposals. Learning the opponents' preference profile helps the proposer agent to offer proposals with a higher chance of acceptance and, therefore, im-

proves the efficiency of the negotiation process. The preference profile can be defined as simple as a reservation point for an issue, where everything beyond that point is rejected (Zeng and Sycara, 1998; Rodriguez-Fernandez et al., 2019). In a more complicated setting, the agents use a linear additive utility function, where the utility of a proposal is the weighted sum of the utility of its attributes (Raiffa et al., 2002; Chen and Weiss, 2015). There are more complicated ways of proposal evaluation modeling where the utility function becomes non-linear (Ito et al., 2011; Marsa-Maestre et al., 2014).

Researchers apply four different approaches to estimate the opponents' preference models (Baarslag et al., 2016):

1. Learning the importance of each issue to the opponent
2. Classifying the opponent's behavior where each class is associated with a set of known preferences
3. Using data from previous negotiations and pattern mining to find the agent's preferences
4. Using logical reasoning and heuristic search to learn the opponents' preferences

The first approach refers to the studies in which the proposer agent tries to estimate the rank or the weight of the issues according to each opponent. Estimating the ranks/weights of the issues is mainly based on the opponent's concession-making strategy. That is, the more an opponent concedes on an issue, the less important the issue is. The ranks of the issues are initialized with equal values at the beginning of the negotiation and are incrementally updated in each round of negotiation. For example, in (Niemann and Lang, 2009) a Bayesian learning technique is proposed to estimate the weight of negotiation issues. They use ten weight hypotheses, ranging from 0.5 to 0.95, for each issue. The initial probability weight for each issue is defined based on a uniform distribution. As the negotiation proceeds, a concession ratio is calculated for each issue. This ratio

has an inverse relation with the weight of the issue. Using the calculated ratios, the probability of the hypotheses and, therefore, the weight of each issue is updated.

Another approach in preference profile modeling is classifying the opponent's behavior to assign a predefined set of preferences to it (Hindriks and Tykhonov, 2008; Buffett and Spencer, 2005, 2007). In this approach, at first, several opponent classes are identified, and preferences are assigned to these classes. Then, a classification algorithm is applied to find out to which class the opponent belongs. For example, Hindriks and Tykhonov (Hindriks and Tykhonov, 2008) consider a set of possible hypotheses about the opponent's preference profile. These hypotheses are the Cartesian product of two sets of hypotheses: hypotheses about the ranks of the issues and hypotheses about the evaluation functions of the issues. The probability of each hypothesis is updated as the negotiation proceeds, and more pieces of evidence become available. This approach faces scalability challenges due to the large hypotheses space. To address this challenge, another version of this method has been proposed, in which the authors presume the ranks of the issues and, thus, the evaluation functions do not need to be learned simultaneously (Hindriks and Tykhonov, 2008).

The third approach in learning preference profiles uses available historical data from previous negotiations in a similar context with similar participants. Using such historical data, various data-mining techniques can be applied to predict the opponent's preference profile. Robu and La Poutre (Robu and La Poutre, 2006) use previous negotiation data to build the buyers' utility graphs. The utility graph is a structural model of a buyer that shows how the buyer evaluates the dependency of two different items. Using this graph, the seller can find a bundle of items that are considered of high value to the buyer and narrow them down to one item with the highest benefit for himself. This approach is not applicable when no previous negotiations are conducted in a specific context with a similar set of participants.

For the cases where no previous data is available, and it is not possible to limit the space of preference profiles to a finite set of known classes, the researchers apply heuristic search to learn the opponents' preferences. For example, Aydoğın and Yolum (Aydoğın and Yolum, 2012a,b)

use the proposals offered by the opponent as positive training instances and the counter-proposals rejected by the opponent as negative training instances. The agent then tries to learn the importance of the issues to the opponent and the type of the proposals that he may accept. Another example of heuristics applied in opponent modeling is frequency analysis where the agent estimate the importance of an issue based on the number of times the value of an issue has changed during the negotiation (van Galen Last, 2012; van Krimpen et al., 2013).

In this study, we assume that the agents use a utility function to evaluate the proposals. Each participant evaluates the attributes of a proposal (i.e., the negotiation issues) by assigning a fuzzy score to each attribute. Then, the proposal score is calculated by defuzzifying the union of the attributes' fuzzy scores using a weighted average approach. Our problem is more than finding the ranks and weights of the negotiation issues, as our utility models are functions of a more significant number of parameters than importance weights. Therefore, the proposed methods in the first category of preference profile modeling are not applicable. The classification approach is not followed in our research either as it limits the exploration of all the possibilities and makes the solution dependent on the negotiation context. For instance, in the real-estate case introduced in section 4.2, the seller agent needs to learn the preference models of the buyers. Based on the classification approach, a discrete number of classes (assumptions) should be defined, and the probability of each class should be updated based on the received pieces of evidence in each negotiation round. A buyer agent in this example is usually concerned about a large set of different attributes (e.g. price, number of bedrooms, number of bathrooms, square footage, and distance to nearest school) and each attribute has an importance weight and a utility function assigned to it. Setting the assumption pool for this example is challenging due to the number of attributes and the range of their possible values. For example, if the price can vary between \$150,000 to \$3,000,000, the buyers' preferences can be anywhere along this range. In such a situation, even a high number of assumptions (e.g. ten classes) will not cover the whole range of possible parameters for the buyers' utility functions. That is, even with ten classes, the gap between the defined classes is

as large as \$300,000. Missing the values that fall inside this gap can mislead the seller to wrong estimations of the buyers' preferences. Increasing the number of classes to reduce this gap will increase the computation complexity of the problem and will cause the scalability issue mentioned earlier in this chapter.

As acquiring historical data from previous negotiations is not feasible in many contexts including our case studies, the learning approaches relying on this sort of information are not applicable either. Therefore, in this paper, we propose a recursive Bayesian filtering technique that can estimate the parameters of these fuzzy models only using the feedback it receives from the participants in each negotiation round. Recursive Bayesian filtering (RBF) is a probabilistic approach for estimating the unknown state of the system using the received measurements over time (Särkkä, 2013). Different implementations of RBF are proposed in the literature such as Kalman filtering (Zarchan and Musoff, 2013) and particle filtering (Doucet and Johansen, 2011). In the specific problem of this paper, the state-space models comprise of elements of non-linearity, non-differentiability, and non-Gaussianity. Thus, techniques of Kalman filtering or their first-order approximations like extended Kalman filtering cannot be used. Therefore, we adopt a technique of unscented particle filtering (UPF), in which the posterior distribution of the state is recursively estimated by a set of samples (particles) that evolve by time conforming to the transition and observation models of the system. In automated negotiation, this method receives participants' feedback as observations in each negotiation round and estimates their preferences recursively. Besides, the state-space in automated negotiation might be tightly constrained. For instance, the number of bedrooms in a set of available houses might not be more than six. To consider such constraints, the proposed UPF adopts a gain projection approach as well. The proposed approach is described in more details in section 4.4.

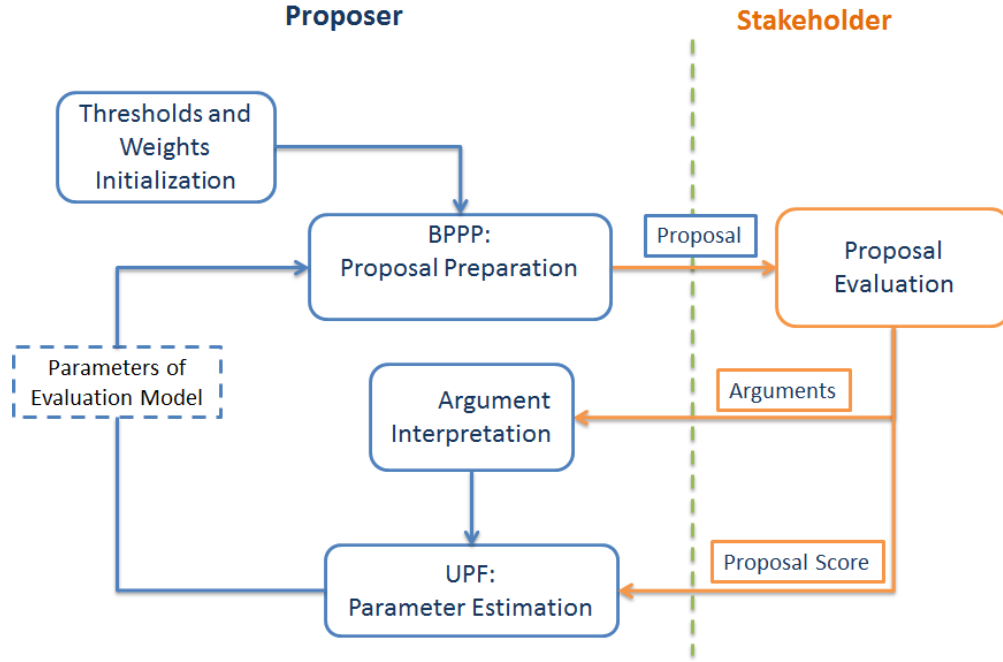


Figure 4.3: The system flowchart

4.4 Methodology

4.4.1 Multi-Agent System

The proposed model has two specific components in addition to the MAS. The first one, belief propagation proposal preparation (BPPP) module, is in charge of preparing proposals in each negotiation round. This component automates the process of proposal offering using Markov random fields (MRF) and belief propagation (BP) (Eshragh et al., 2018). The other component, called UPF module, helps the proposer to recursively learn about the stakeholders' preferences as the negotiation proceeds (Section 4.4). A schematic representation of the proposed negotiation system is displayed in Figure 4.3. The following sections describe the methodology specifically with a focus on the UPF component of the system. For more details about the BPPP component, the readers are referred to (Eshragh et al., 2018).

4.4.2 Problem Statement

Assume that a set of $x + 1$ stakeholders, $S = \{s_p, s_1, s_2, \dots, s_x\}$, are going to negotiate over a set of y proposals, $P = \{p_1, p_2, \dots, p_y\}$. Variable s_p represents the proposer agent who initiates the negotiation and continues it by offering a different proposal in each round. A set of z types of attributes (i.e. negotiation issues), $A = \{a_1, a_2, \dots, a_z\}$, are used to identify the proposals; i.e. each proposal $p_j (j = 1, \dots, y)$ is characterized by a unique combination of the values of these attributes as $A_j = \{v_1^j, v_2^j, \dots, v_z^j\}$ where v_k^j is the value of the k^{th} attribute in proposal p_j . The preference of a stakeholder, $s_i (i = 1, \dots, x)$, with respect to an attribute, $a_k (k = 1, \dots, z)$, is modeled by a weight factor and the minimum and maximum acceptable values for that attribute, respectively denoted by w_{k,s_i} , L_{k,s_i} and U_{k,s_i} . These parameters shape the agent's preference profile, which is in the form of a linear additive utility function described briefly in section 4.4.3.

The goal of the negotiation is to find a proposal, if any, that meets all agents' preferences. That is, the negotiation ends when either the agents reach a mutually acceptable agreement or the proposer confirms that there is no agreement and terminates the negotiation.

4.4.3 Proposal Evaluation Model

Each stakeholder $s_i (i = 1, \dots, x)$ has a set of preference limits, L_{k,s_i} and U_{k,s_i} , and a weight factor w_{k,s_i} over any attribute $a_k (k = 1, \dots, z)$. These parameters form the agent's preference profile that is used by the stakeholder to evaluate the offered proposal in each round of negotiation. That is, the stakeholder $s_i (i = 1, \dots, x)$ assigns a score, $\phi_{k,i}^j$, to the value of attribute a_k in a proposal p_j and then combines all the scores of different attributes to one score, Φ_i^j , to evaluate the whole proposal. Here, we have two different types of attributes: "less-is-better" attributes and "more-is-better" attributes. For example, when buying a house, the price is a less-is-better attribute (a preference reversal case) while square footage is usually a more-is-better attribute.

The way a stakeholder evaluates a less-is-better attribute is modeled with a monotonically

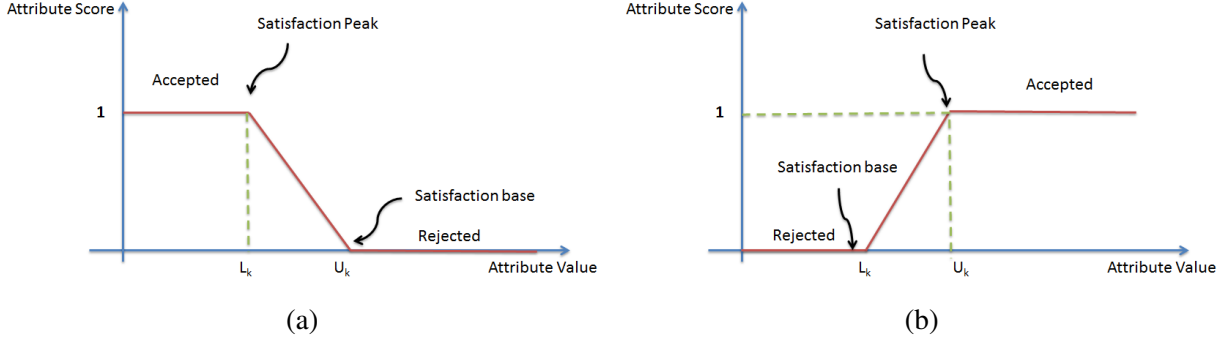


Figure 4.4: Objective function for a) less-is-better attributes b) more-is-better attributes

decreasing fuzzy membership function as in Equation 4.1 (Figure 4.4a).

$$\varphi_{k,i}^j(v_k) = \begin{cases} 0, & v_k > U_{k,s_i} \\ \frac{U_{k,s_i} - v_k}{U_{k,s_i} - L_{k,s_i}}, & L_{k,s_i} \leq v_k \leq U_{k,s_i} \\ 1, & v_k < L_{k,s_i} \end{cases} \quad (4.1)$$

In Equation 4.1, v_k is the value of attribute a_k that stakeholder s_i is evaluating; L_{k,s_i} is the lower bound of his preference and U_{k,s_i} is the upper bound of his preference with respect to attribute a_k .

The more-is-better attributes are evaluated by the stakeholder using a monotonically increasing fuzzy membership function as in Equation 4.2 (Figure 4.4b).

$$\varphi_{k,i}^j(v_k) = \begin{cases} 0, & v_k < L_{k,s_i} \\ \frac{v_k - L_{k,s_i}}{U_{k,s_i} - L_{k,s_i}}, & L_{k,s_i} \leq v_k \leq U_{k,s_i} \\ 1, & v_k > U_{k,s_i} \end{cases} \quad (4.2)$$

where v_k is the value of attribute a_k that stakeholder s_i is evaluating; L_{k,s_i} is the lower bound of his preference, and U_{k,s_i} is the upper bound of his preference with respect to attribute a_k .

Since a proposal p_j is made of a unique combination of the values of various attributes, the score a stakeholder s_i assigns to the proposal p_j must be defined as a combination of the fuzzy membership degrees, $\varphi_{k,i}^j (k = 1, \dots, z)$, assigned to the attributes of that proposal. That is, the proposal score is the result of defuzzifying the attributes' membership functions. In this study, the

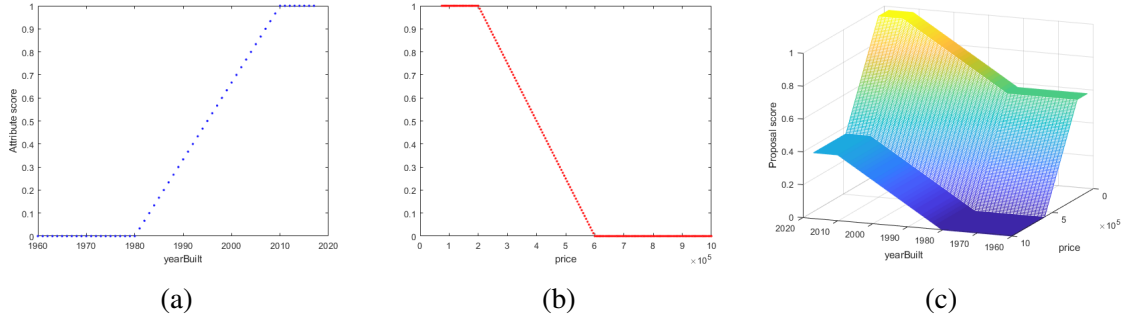


Figure 4.5: Attribute and proposal score functions a) Score function of the Year-Built attribute b) Score function of the Price attribute c) Score function of the proposal

weighted average method (Ross, 2005) is selected for this purpose as in Equation 4.3:

$$\Phi_i^j = \sum_{k=1}^z w_{k,s_i} * \phi_{k,i}^j(v_k^j) \quad (4.3)$$

where w_{k,s_i} is the weight of attribute a_k for stakeholder s_i and $\phi_{k,i}^j$ is the score of attribute a_k in proposal p_j for stakeholder s_i . v_k^j is the value of the k^{th} attribute in this proposal, and z is the number of the attributes involved in proposal p_j .

To understand the evaluation model more clearly, consider the price and year built attributes in the real estate example. For the buyer agent, year built is a “more-is-better” attribute, and therefore, its evaluation model is defined as a monotonically increasing fuzzy membership function. Assuming 1980 as the lower bound of the stakeholder’s preference and 2010 as his upper bound, the score of “year built” attribute can be modeled as Figure 4.5a. Price, on the other hand, is a “less-is-better” attribute for the buyer agent and is evaluated based on the model illustrated in Figure 4.5b. Assuming each proposal has only these two attributes (i.e. $z = 2$), the score of each proposal (i.e., a house in this example) is calculated as a result of defuzzifying the attributes membership functions, as shown in Figure 4.5c.

During the negotiation process, having the stakeholders’ proposal evaluation models can help the proposer agent to find a proposal that satisfies others and yet, meets his criteria to the highest possible extent. However, the stakeholders do not usually reveal this sort of information to one another. It is only through the stakeholders’ feedback that the proposer agent can devise the

participants' evaluation models. In this study, feedback consists of three components: 1) whether the proposal is rejected or accepted, 2) the score that the stakeholder assigns to the proposal if the proposal is rejected 3) the argument(s) about the rejected proposal.

In this research, recursive Bayesian filtering (RBF) (Särkkä, 2013) is applied to estimate stakeholders' preferences (thus their evaluation models) using stakeholders' feedback. Generally, filtering is a probabilistic estimation of the state of a system dynamically using the observations from the environment. In our study, the state consists of the parameters of stakeholders' evaluation models. The interaction with the environment is performed via the control data, i.e., the arguments made by the stakeholders, and the measurement data, i.e., the scores the stakeholders assign to the offered proposals. All RBF techniques are based on probabilistic generative laws. The state of a system can be characterized by a transition probability distribution that shows how the state evolves as a function of control data. The process by which the measurements are generated can be modeled using a generative probability distribution. Therefore, in RBF, two main steps repeat every time a new set of data is received from the environment. The first step, called "prediction", is about the stochastic update of the previous/initial estimations using the control data. In the second step, called "correction", the belief about the states from the previous step is updated using the measurement data. Different implementations of RBF exist in the literature, e.g., Kalman filtering, information filtering, histogram filtering, particle filtering, and their unscented and extended variants. In our case study, the belief of the state has an unknown distribution; i.e., the normal distribution might not be representative of the belief. Therefore, "particle" sampling can be used to approximate the belief. Also, the measurement-generation model (i.e. the proposal evaluation function in Equation 4.3) cannot be linearly approximated via Taylor expansions since it is not differentiable. As such, unscented particle filtering is the most appropriate RBF approach to apply. In this study, a modified UPF algorithm is proposed to address the learning process in automated negotiations.

4.4.4 Unscented Particle Filtering

Negotiation variables and their initialization

In the negotiation process, random variables are categorized to state, measurement data, and control data. Concerning the state, each attribute is associated with three elements of the state vector, i.e., the preference lower bound and upper bound as well as the weight factor. The state vector is, therefore, defined as:

$$X = \{U_k, L_k, w_k | k = 1, 2, \dots, z\} \quad (4.4)$$

where z is the total number of attributes.

We assume the knowledge (belief) of the proposer agent about the state vector (the preferences of the other agents) is dynamically changing as a function of two sorts of observations: first, the score each stakeholder assigns to the offered proposal at round t of the negotiation process, called measurement data (z_t) and second, a set of arguments about rejecting the proposal, called control data (\mathcal{U}_t). Measurement data includes some information related to the state at a distinct point t in time while control data includes information about the change of one/several elements of the state within the time interval $(t - 1, t]$.

Particle sampling is used to approximate the belief of the state using the particle sets. Each particle consists of a sample point (hypothesis) from the belief distribution over the state and an importance weight associated with each sample.

To initialize particle samples, a multivariate Gaussian probability distribution function (PDF), $N(\bar{X}, \Sigma_X)$, is considered using the initial knowledge of the proposer about the state, \bar{X} , and the uncertainty involved with this a priori knowledge represented by a covariance matrix, Σ_X . For instance, in the real estate case, from the initial interactions between the seller and the buyers, the seller understands that the buyers are looking for an “affordable” place. An affordable price can be translated to a range from \$100,000 to \$500,000. Therefore, for the upper bound of the price, the \bar{X} is set to \$300,000, but as we are uncertain about this initial value, an uncertainty of \$200,000

is assigned to this variable. The covariance matrix, Σ_X is built based on uncertainties assigned to each variable. The impact of the initial values of the state variables on the negotiation results is further discussed in section 4.5.2.

A total of N hypotheses are sampled from the PDF, $N(\bar{X}, \Sigma_X)$, and uniform weights are assigned to them. Equation 4.5 represents the particle sample set at time $t = 0$ (at the beginning of the negotiation process):

$$\bar{X}_0 = \left\{ \bar{X}_0^{(j)} | j = 1, 2, \dots, N \right\} \quad (4.5)$$

where $\bar{X}_0^{(j)}$ represents the j^{th} particle sample set at time $t = 0$.

$$\bar{X}_0^{(j)} = \left\{ U_k^j, L_k^j, w_k^j | k = 1, 2, \dots, z \right\} \quad (4.6)$$

As explained in previous sections, in each negotiation round, particles are updated and corrected using the feedback from the stakeholders. Then, an importance weight is given to each particle, and the particles are re-sampled based on these weights. The higher the weight, the more the chances of the particle to be resampled. The last step is calculating the weighted average of the resampled particles as the output of the UPF process in this round of negotiation. These outputs will be then used by the BPPP module to find a proposal with a higher probability of being accepted by all the stakeholders.

The next sections explain the rest of the UPF process in more details. Here, only the essential equations that are specifically modified to fit the negotiation problem are discussed in details. To read more about the details of a generic unscented particle filtering approach, the readers are referred to (Van der Merwe et al., 2000).

The first step of UPF is updating the particles. Each update has two stages: Prediction (using arguments) and correction (using proposal scores).

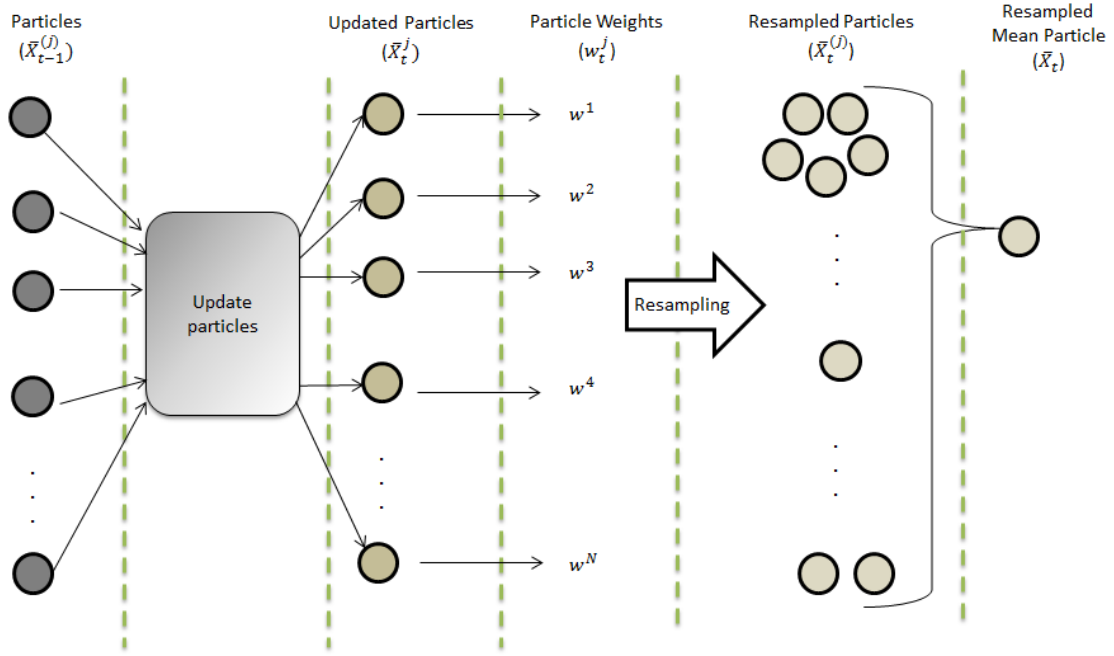


Figure 4.6: UPF process in each negotiation round

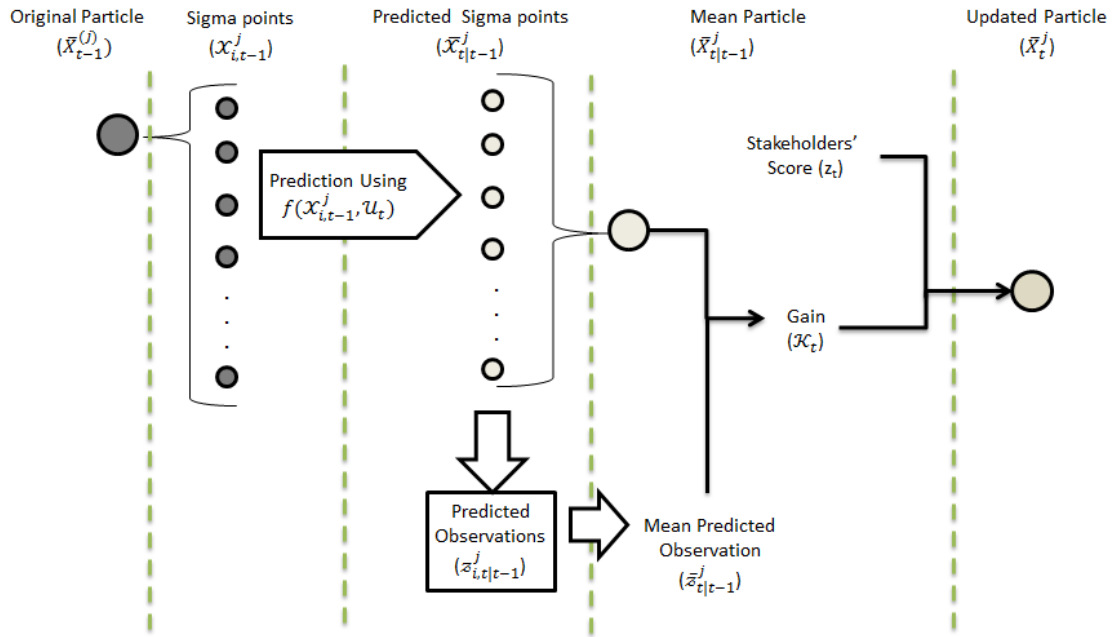


Figure 4.7: Update and correct a particle using stakeholders' feedback

Particle prediction

In the prediction step, each particle is transformed into a new particle (i.e., a set of weight factors and preference limits) using the arguments received from the stakeholders in the previous negotiation round. Here, we deal with a non-linear measurement model (Equation 4.3); therefore, an unscented transform is used to estimate the impact of applying this nonlinear measurement model on the state belief that, itself, is characterized only in terms of a finite set of particles.

In unscented transform, a number of weighted samples, called sigma points, are selected around each particle to capture its statistics, mean and covariance. For each particle, $\bar{X}_t^{(j)} (j = 1, \dots, N)$, there need to be $2n_x + 1$ sigma points, where $n_x = 3 * z$ is the length of the state vector and z is the number of negotiation issues. These sigma points are defined as:

$$\begin{aligned} \mathcal{X}_{i,t-1}^j &= \bar{X}_t^{(j)}, & i &= 0 \\ \mathcal{X}_{i,t-1}^j &= \bar{X}_t^{(j)} + (\sqrt{(n_x + \lambda)\Sigma_{t-1}^j})_i, & i &= 1, \dots, n_x \\ \mathcal{X}_{i,t-1}^j &= \bar{X}_t^{(j)} - (\sqrt{(n_x + \lambda)\Sigma_{t-1}^j})_i, & i &= n_x + 1, \dots, 2n_x \end{aligned} \quad (4.7)$$

where λ is a scaling parameter and $(\sqrt{(n_x + \lambda)\Sigma_{t-1}^j})_i$ is the i^{th} row or column of the matrix square root of $(n_x + \lambda)\Sigma_{t-1}^j$. Here, Σ_{t-1}^j is the covariance matrix of particle j estimated at round t .

The scaling parameter λ is calculated as:

$$\lambda = \alpha^2(n_x + \kappa) - n_x \quad (4.8)$$

where α and κ are positive scaling parameters. $\kappa \geq 0$ is recommended in (Van der Merwe et al., 2000) as it guarantees the positive semidefiniteness of the covariance matrix. α defines the size of the sigma point distribution. The smaller the α , the less non-local effects are sampled (Van der Merwe et al., 2000).

For each sigma point, a mean weight ($\mathcal{W}_i^{j(m)}$) and a covariance weight ($\mathcal{W}_i^{j(c)}$) are defined. These weights indicate the importance of the sigma points and are used to determine the mean of the sigma points and the covariance matrix associated with them. The mean and covariance

weights of sigma points are defined as:

$$\begin{aligned}\mathcal{W}_i^{j(m)} &= \frac{\lambda}{n_x + \lambda}, \mathcal{W}_i^{j(c)} = \frac{1}{2(n_x + \lambda)} + 1 - \alpha^2 + \beta, & i = 0 \\ \mathcal{W}_i^{j(m)} &= \frac{1}{2(n_x + \lambda)}, \mathcal{W}_i^{j(c)} = \frac{1}{2(n_x + \lambda)}, & i = 1, \dots, 2n_x\end{aligned}\tag{4.9}$$

Parameter β in this equation is used to control the weight of the zeroth sigma point for calculating the covariance (Van der Merwe et al., 2000).

Based on the arguments received from the stakeholders, the sigma points will be changed to generate new samples. The arguments, called control data (\mathcal{U}_t), give the model an idea of what needs to be changed to reduce the gap between the estimation and the truth. The arguments are provided about the attributes with values lower or higher than the preferred thresholds of the stakeholders. If the argument received from a stakeholder states that the value of attribute a_k is too low, then the upper-bound and the lower-bound of the preference limit, as well as the weight factor of the attribute a_k , should be increased. On the other hand, if the feedback received from the stakeholder argues that the value of attribute a_k is too high, then the upper-bound and the lower-bound of the preference limit should be decreased, but the weight factor of a_k should still be increased. For instance, in the real-estate case, when an argument is received about decreasing the price of a property, it means our previous estimations about the preference limits/weight factor of the price should be revised to better match the stakeholder's criteria. To achieve this goal, the sigma points should be changed to reflect our new understanding of the stakeholder's preferences. Considering \mathcal{X}_{t-1}^j as the set of sigma points of the j^{th} particle, the new set, called predicted sigma points $\tilde{\mathcal{X}}_{t|t-1}^j$, is calculated using a transition function that demonstrates how the control data can stochastically change the state from time $t - 1$ to time t .

$$\tilde{\mathcal{X}}_{t|t-1}^j = f(\mathcal{X}_{t-1}^j, \mathcal{U}_t)\tag{4.10}$$

Following Eshragh et al. (2018), the transition function f is described in Figure 4.8. This function takes a set of sigma points \mathcal{X}_{t-1}^j and arguments \mathcal{U}_t as the input and changes the sigma

Algorithm 1: Transition function f

```
1  $\underline{f}(\mathcal{X}_{t-1}^j, \mathcal{U}_t)$ 
   Input : Set of Sigma Points  $\mathcal{X}_{t-1}^j$  and Arguments  $\mathcal{U}_t$ 
   Output:  $\tilde{\mathcal{X}}_{t|t-1}^j$ 
2  $\tilde{\mathcal{X}}_{t|t-1}^j = \mathcal{X}_{t-1}^j$ ;
3 if (There is an argument  $u_t \in \mathcal{U}_t$  about attribute  $a_k, k = (1, \dots, z)$ ) then
4   if  $u_t = \text{increase } a_k$  then
5     for  $i \leftarrow 0$  to  $2n_x$  do
6       Increase weight factor of  $a_k$  ( $w_{k,s_i}$ ) in sigma point  $i$  of  $\tilde{\mathcal{X}}_{t|t-1}^j$ ;
7       Increase lower bound and upper bound of  $a_k$  ( $L_{k,s_i}$  and  $U_{k,s_i}$ )
         in sigma point  $i$  of  $\tilde{\mathcal{X}}_{t|t-1}^j$ ;
8     end
9   else if  $u_t = \text{decrease } a_k$  then
10    for  $i \leftarrow 0$  to  $2n_x$  do
11      Increase weight factor of  $a_k$  ( $w_{k,s_i}$ ) in sigma point  $i$  of  $\tilde{\mathcal{X}}_{t|t-1}^j$ ;
12      Decrease lower bound and upper bound of  $a_k$  ( $L_{k,s_i}$  and  $U_{k,s_i}$ )
        in sigma point  $j$  of  $\tilde{\mathcal{X}}_{t|t-1}^j$ ;
13    end
14  return  $\tilde{\mathcal{X}}_{t|t-1}^j$ ;
```

Figure 4.8: Transition function f pseudo-code

points based on the arguments. For each argument $u_t \in \mathcal{U}_t$, if the argument is about attribute $a_k, k = (1, \dots, z)$, then the weight and preference limits related to a_k should be revised in all sigma points. If the argument u_t is about the value of a_k being too high, it means that the preference limits of this attribute for the stakeholder are lower than the current estimation. That is, these preference limits should be increased. However, if the argument u_t is about the value of a_k being too low, it means that the preference limits of this attribute for the stakeholder are higher than the current estimation and should, therefore, be decreased. In both cases, the weight factor of the attribute a_k should be increased since receiving an argument about an attribute can be translated to its level of importance to the stakeholder. The amount of decreasing or increasing the preference limits is proportional to the standard deviation (uncertainty) of the current estimation of the respective state (e.g. 10% of the standard deviation). The amount of by which the weight factor states are increased is defined based on a percentage of the current weight value plus a percentage of the current standard deviation to its related state.

The next step is calculating the mean of the predicted sigma points which is called mean predicted particle, $\bar{X}_{t|t-1}^j$. Mean predicted particle is actually a refined version of the original particle (hypothesis). It is calculated as below:

$$\bar{X}_{t|t-1}^j = \sum_{i=0}^{2n_x} \mathcal{W}_i^{j(m)} * \tilde{X}_{i,t|t-1}^j \quad (4.11)$$

where $\mathcal{W}_i^{j(m)}$ is the mean weight of sigma point i defined as in Equation 4.9.

The uncertainty of the predicted particle is then calculated as:

$$\Sigma_{t|t-1}^j = \sum_{i=0}^{2n_x} \mathcal{W}_i^{j(c)} * (\tilde{X}_{i,t|t-1}^j - \bar{X}_{t|t-1}^j) * (\tilde{X}_{i,t|t-1}^j - \bar{X}_{t|t-1}^j)^T \quad (4.12)$$

where $\mathcal{W}_i^{j(c)}$, defined in Equation 4.9, is the covariance weight of sigma point i .

To measure how far the estimations are from the stakeholder's criteria, the score of the current proposal (P_{t-1}) should be estimated based on the predicted preference limits and weight factors (the predicted particle) and then be compared to the actual score received from the stakeholders. To estimate the score of the current proposal, first, the predicted measurement (score) of each predicted sigma point $\tilde{X}_{i,t|t-1}^j$ is calculated using the measurement model:

$$z_{i,t|t-1}^j = h(\tilde{X}_{i,t|t-1}^j) \quad (4.13)$$

where h is the evaluation function defined in Equation 4.3.

The average predicted score is then calculated as:

$$\bar{z}_{t|t-1}^j = \sum_{i=0}^{2n_x} \mathcal{W}_i^{j(m)} * z_{i,t|t-1}^j \quad (4.14)$$

where $\mathcal{W}_i^{j(m)}$ is the mean weight of sigma point i .

The differences of this score from the real scores of the proposal (received from the stakeholders) are used in the next step to correct the predicted state.

Particle correction

For correcting the so-far estimated state, a gain factor must be used to indicate to what extent the measurement data can affect the prediction results.

Using covariance weights defined in Equation 4.9, the variance of the predicted score is calculated as:

$$P_{\bar{z}_t, \bar{z}_t} = \sum_{i=0}^{2n_x} \mathcal{W}_i^{j(c)} * (z_{i,t|t-1}^j - \bar{z}_{t|t-1}^j) * (z_{i,t|t-1}^j - \bar{z}_{t|t-1}^j)^T + R_t \quad (4.15)$$

where R_t is the variance of the measurement noise. In our case, it is the variance of the proposal scores received from different stakeholders.

$$P_{\bar{x}_t, \bar{z}_t} = \sum_{i=0}^{2n_x} \mathcal{W}_i^{j(c)} * (\mathcal{X}_{i,t|t-1}^j - \bar{X}_{t|t-1}^j) * (z_{i,t|t-1}^j - \bar{z}_{t|t-1}^j)^T \quad (4.16)$$

$P_{\bar{x}_t, \bar{z}_t}$ is the cross-covariance matrix of the predicted sigma points and the scores associated with them.

The gain matrix is then calculated as:

$$\mathcal{K}_t = P_{\bar{x}_t, \bar{z}_t} * (P_{\bar{z}_t, \bar{z}_t})^{-1} \quad (4.17)$$

Having the gain matrix and the real scores received from the stakeholders, it is time to proceed to the correction step. Each particle is corrected as follows:

$$\bar{X}_t^j = \bar{X}_{t|t-1}^j + \mathcal{K}_t * (z_t - \bar{z}_{t|t-1}^j) \quad (4.18)$$

where z_t is the score (measurement data) received from the stakeholder for the last offered proposal P_{t-1} and $\bar{z}_{t|t-1}^j$ is the score predicted for this proposal using Equation 4.3. In the case of multiple stakeholders with multiple scores, z_t is the mean of their scores. In the case the proposer agent has more interest in satisfying one or a group of stakeholders compared to the others, then a weighted average of these scores can be used, where higher weights are assigned to more “important” stakeholders.

The covariance matrix of the updated particle is updated as follows:

$$\bar{\Sigma}_t^j = \Sigma_{t|t-1}^j - (\mathcal{K}_t * P_{\bar{z}_t, \bar{z}_t} * \mathcal{K}_t^T) \quad (4.19)$$

In our problem domain, the state values are limited; For example, the lower-bound preference limit for the price cannot be smaller than the price of the least expensive property in the database of the proposer agent. As such, we propose a constrained UPF technique. That is, a set of inequality constraints are enforced as follows:

$$D_k \bar{X}_t^j \leq d_k \quad (4.20)$$

where D_k is the design matrix of the constraints and $d_k \in R^{n_d}$ is a known vector representing the bounds of the inequality constraints. Using the defined constraints in the gain projection approach (Simon, 2010), the constrained updated particle is defined as:

$$\bar{X}_t^j = \bar{X}_t^j + D^T (DD^T)^{-1} (D_k \bar{X}_t^j - d_k) \quad (4.21)$$

Through the rules of error propagation, the covariance matrix is updated as follows:

$$\bar{\Sigma}_t^j = \bar{\Sigma}_t^j + D^T (DD^T)^{-1} D \bar{\Sigma}_t^j (D^T (DD^T)^{-1} D)^T \quad (4.22)$$

Having the updated particle, the predicted measurement (score) is updated as:

$$\bar{z}_t^j = h(\bar{X}_t^j) \quad (4.23)$$

Figure 4.7 briefly explains the prediction and correction stages to update a particle.

The importance weight of the particle indicates its importance compared to the other particles; the higher the importance weight, the closer the particle to the expected state. We use an exponential probability distribution to re-weight each particle as follows:

$$w_t^j = w_{t-1}^j * \lambda_p * \exp(-\lambda_p * |z_t - \bar{z}_t^j|) \quad (4.24)$$

where w_{t-1}^j is the weight of particle j from the previous negotiation round and λ_p is the rate parameter. The larger the λ_p , the higher the effect of $|z_t - \bar{z}_t^j|$ on the weight of the j^{th} particle;

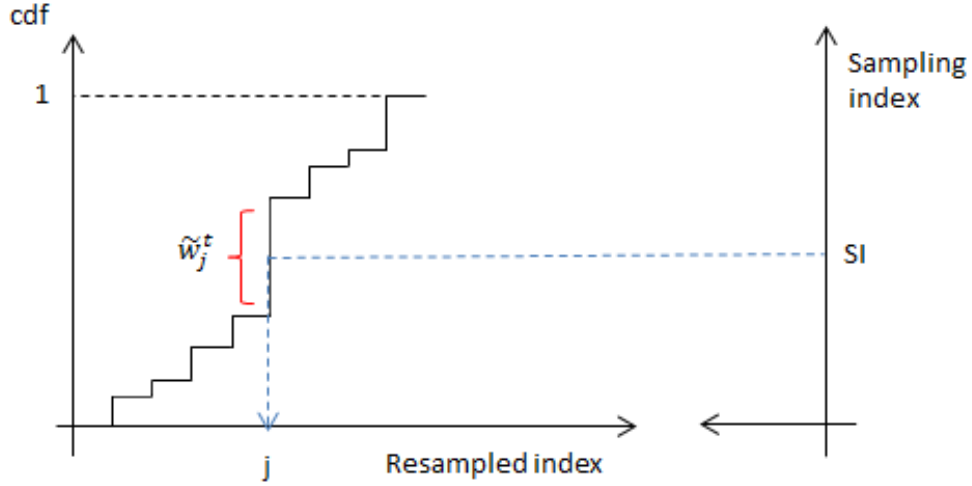


Figure 4.9: Mapping uniformly selected index SI to the cumulative distribution domain

i.e. the particles further from the state belief will have less weight. At $t = 0$, particle weights are uniformly distributed.

Once all the N particles are re-weighted, the weights of the particles need to be normalized as follows.

$$\tilde{w}_t^j = \frac{w_t^j}{\sum_{j=1}^N w_t^j} \quad (4.25)$$

State update by importance resampling

The next step is resampling in order to eliminate the particles with considerably low weights. There are various methods of resampling in the literature such as Bayesian importance sampling (Geweke, 1989), sequential importance sampling (SIS)(Liu and Chen, 1998) and sampling-importance resampling (SIR) (Rubin, 1988). In this study, SIR is employed as it does not have the degeneracy problem of the other algorithms. The degeneracy problem usually happens in sequential importance sampling when the variance of the importance weights increases stochastically over time.

In SIR, a Dirac random measure $\{\bar{X}_t^j, \tilde{w}_t^j\}$ is mapped into an equally weighted random measure $(X_t^j, 1/N)$. That is, N samples are uniformly drawn from the discrete set $(\bar{X}_t^j; j = 1, \dots, N)$ with probabilities $(\tilde{w}_t^j; j = 1, \dots, N)$.

```

1 Function FindresamplingIndices(weights: double[]) : int[]
2   Construct the cumulative weight distribution;
3   for  $i \leftarrow 0$  to  $N$  do
4     Set SI as a random number between 0 and 1;
5     Find SI in the the cumulative weight distribution;
6     Add SI's index to indices[] array;
7   end
8   return indices[];
9 end
10 Function resample(indices: int[]) : particle[]
11   Indices=FindresamplingIndices();
12   for  $i \leftarrow 0$  to  $N$  do
13     rasampledParticles(i)=Particles(indices(i));
14   end
15   return rasampledParticles;
16 end

```

Figure 4.10: Resampling algorithm pseudo-code

To do so, the cumulative distribution of the discrete set of the particles weights is constructed. Next, a sampling index SI is uniformly selected, and then, projected to the cumulative distribution function (Figure 4.9). The intersection with the function denotes the new sample at index j . That is, the vector \bar{X}_t^j is accepted as a new sample. Through this process, more copies of particles with higher weights are generated. Drawing N samples from the cumulative distribution $\sigma_{i=1}^j \tilde{w}_t^j(\delta \bar{X}_t)$ is the same as sampling $(N_j; j = 1, \dots, N)$ from a multinomial distribution with N number of trials and $\tilde{w}_t^j(j = 1, \dots, N)$ event probabilities. Figure 4.10 is the pseudo-code of the resampling algorithm.

As a result of resampling, some particles with higher weights might be duplicated too many times. To make sure there is still enough variation in the sample set, Markov Chain Monte Carlo (MCMC) solution (Andrieu et al., 1999; Doucet and Gordon, 1999) is applied. The idea is that applying a specific Markov chain transition kernel will result in the same distribution of the particles except that the new particles may move to more interesting areas of the state-space. That is, the variation of the current distribution can only increase using the Markov chain transition kernel. It is due to the fact that applying a Markov chain transition kernel $K(\bar{X}_{t-1}^j | \bar{X}_t^j)$ on particles distributed based on the posterior distribution $p(\bar{X}_t | z_t)$, only moves the particles to a new space and

does not change the posterior distribution. However, the new space to which the particles move might be more interesting in terms of variation and, therefore, may lead to more interesting results. If a move does not result in a more interesting outcome, then the move will not be accepted and the previous state of the particle will be conserved. The MCMC strategy avoids duplicating the particles whose probability of improving the state belief is less than their previous versions. The readers are referred to (Van der Merwe et al., 2000) for implementation details of this strategy.

The final step of the algorithm in the current negotiation round is finding the mean of all particles:

$$\bar{X}_t = \frac{1}{N} \sum_{j=0}^N \bar{X}_t^j \quad (4.26)$$

Finally, the average of the resampled particles is used as the estimation of the proposer about the stakeholders' preference limits and weights (e.g. $\{U_k, L_k, w_k | k = 1, 2, \dots, z\}$). Using these estimated values, the BPPP module readjusts its understanding of the stakeholders and prepares a new proposal to offer at the next round of negotiation.

The UPF process is repeated in every negotiation round after receiving the stakeholders' feedback until the negotiation terminates.

4.5 Experiments

The proposed methodology is applied to two different case studies. The first case study discusses negotiations in the context of an energy-system planning project in Alberta, Canada. The negotiations in this project involve multiple stakeholders and multiple issues. Various sets of geospatial data, gathered from Alberta Biodiversity Monitoring Institute (ABMI) and Alberta Environment and Parks (AEP) public resources are available to be used in these negotiations. The second case study models the negotiations in a real estate context. In this case study, three types of data have been used: real estate sale data from King County, US, between May 2014 and May 2015, the GIS data related to these properties from King County GIS Open data website (King County GIS

Center, 2018), and preference data from actual users gathered through online interviews. The sales dataset includes 21614 records and large problem domain caused by ten different negotiation attributes compared to the energy-system planning case study (with 100 alternatives and five attributes). Therefore, it is used to evaluate the performance and scalability of the proposed UPF approach in a large-scale problem. Besides, with the diversity of the data from actual users, more complicated negotiation scenarios could be tested on the real estate case.

4.5.1 Case study A: Energy-System Planning in Alberta

One of the most critical economic and environmental challenges in the world is developing energy resources sustainably and efficiently. The success of these decisions depends on two critical factors: first, the ability to process possible solutions effectively based on the economic, social and environmental aspects and second, including a large community of stakeholders so that the final decisions are acceptable to everyone. An example of these energy-system planning decisions is developing reliable electricity grids in Alberta due to the growing demand. Exploring the routing options (to find an alternative to link the supply source to the customers) is a key problem in these sorts of projects where both environmental and non-environmental factors are involved. Our first case study evaluates routing alternatives in an electricity transmission project. The supply source in this case study is a hydropower plant near Slave River and Fort Smith city at the border of Alberta and Northwest Territories, denoted by a red star on the map of Figure 4.11. The demand center is the Ells River 2079 substation, illustrated by an orange triangle in Figure 4.11. Among the possible electricity transmission routes between the hydropower plant and the substation, one should be selected through negotiations with various stakeholders. A set of criteria, including the area, type, and the coverage of the land that will be affected, the environmental impacts (e.g., wildlife and wetlands), the development costs, and the population that will be affected by the route, should be considered in this negotiation process. In the current study, three groups of stakeholders are identified and modeled: first nations, industrial parties, and environmentally-focused groups. While the industrial parties are more concerned about the economic cost of the project, the other

Table 4.1: Significant stakeholders in the project on energy-system planning; Source: (Eshragh et al., 2018)

| Stakeholder category | Group name | Agent name | Primary concerns |
|-------------------------------------|---|----------------|--|
| First Nations | Community, Aboriginal and Native American Relations in TransCanada Treaty 8 First Nations of Alberta | FN | Damage to first nation reserves (FN value) |
| Environmentally focused groups | Alberta Environment and Sustainable Resource Development | AEP | Damage to forest areas (Forest value), wildlife (Wildlife value), and wetlands (Wetland value) |
| Industries on the transmission side | ATCO Electric AltaLink | TFO (Proposer) | Construction costs |

stakeholders' criteria are mostly about the impact of the transmission line on humans and their natural and built environment. Table 4.1 summarizes the stakeholders and their primary concerns in more details.

Implementation

To implement the MAS and its additional components (i.e. BPPP and UPF modules), Java 8 thread processing is used. Proposal databases are implemented using Java Arrays as they are pretty small databases. Agents' interactions are implemented using shared files and synchronized Java variables. We have also used Python scripts to process the GIS data using GDAL (GDAL/OGR contributors, 2017) and ArcPy (Esri, 2017) libraries. These scripts generate alternative routes and calculate their costs using their geographic features. The next section explains the data preparation in more details.

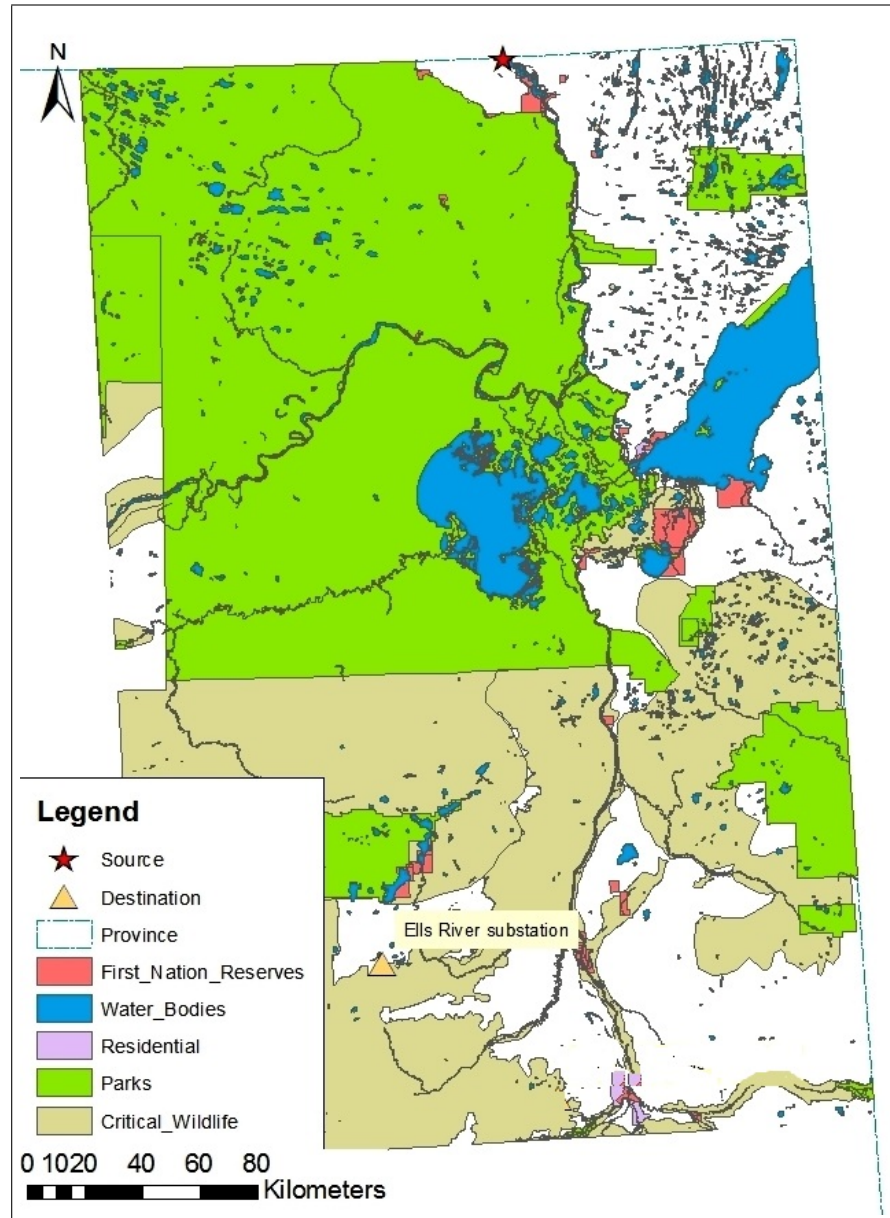


Figure 4.11: The study area located near the Slave River and Fort Smith city at the border of Alberta and the Northwest Territories; Source: (Eshragh et al., 2018)

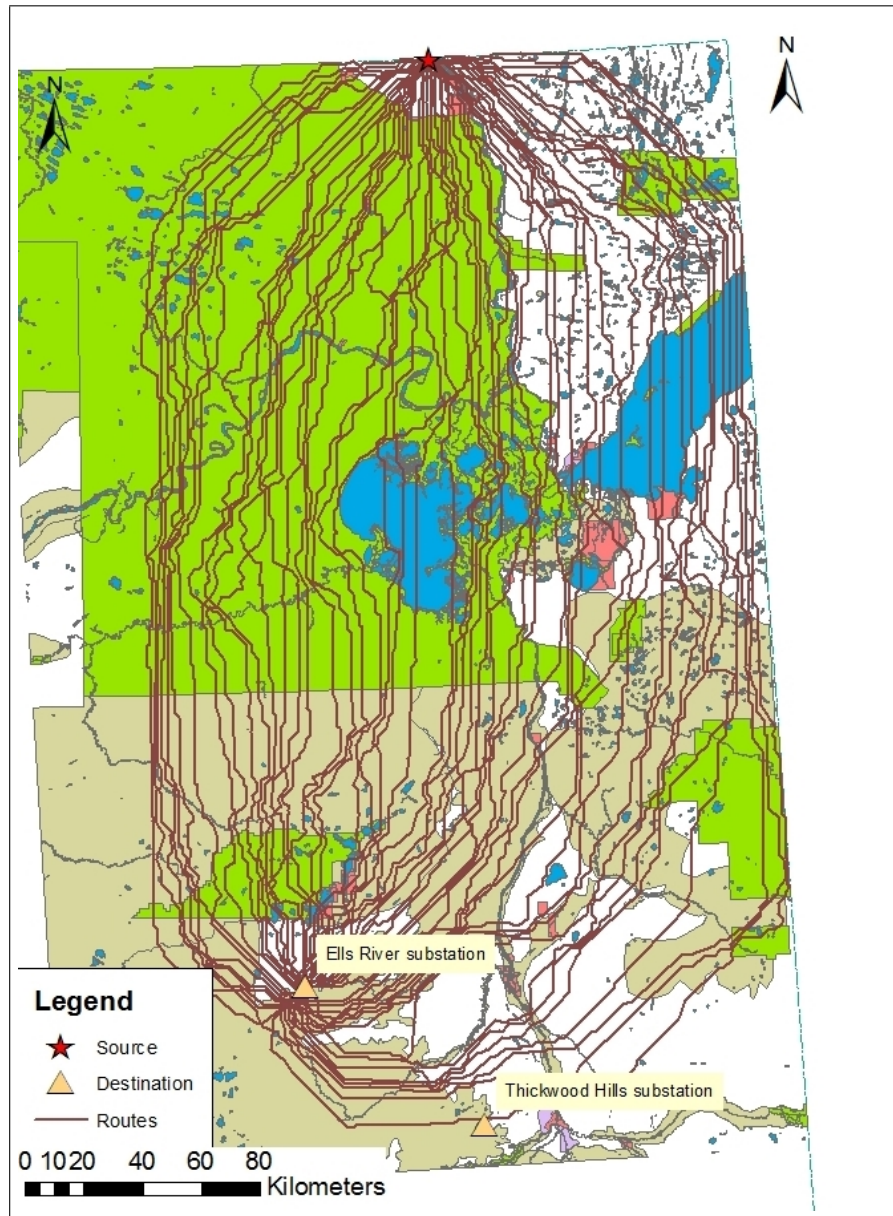


Figure 4.12: Selected routes in the data preparation phase; Source: (Eshragh et al., 2018)

Data preparation

In this case study, the required GIS data is acquired through the Alberta Biodiversity Monitoring Institute (ABMI) and Alberta Environment and Parks (AEP) public resources. The GIS data layers include the maps of different types of forests (e.g., Broadleaf and Coniferous), wetlands, wildlife areas (e.g., caribou zones, different fish areas), roads, and first nation reserves.

To retrieve a set of possible alternatives and build the proposal database, the GIS data layers are processed using a Python script. The Python script employs GDAL and ArcPy libraries to perform spatial data analysis to find a set of paths and assign attribute values to each path. These alternate paths are determined using least cost path analysis, where the cost is the construction cost of the path calculated based on the length of the path, its distance to the road network and its intersection with different types of land. Finally, each path, used as a proposal during the negotiation process, is described by a set of attributes including forest-destruction cost, wildlife-disturbance cost, wetland-damage cost, impact on First-nations' (FN) lands, and construction cost. These attributes are quantified by a combination of a variety of data sources and a set of environmental, ecological, cultural and economic measures. For instance, the forest-destruction cost is determined based on the intersection of the path with different types of forest area. For further details about this dataset, the readers are referred to (Eshragh et al., 2018). Figure 4.12 represents a hundred alternative routes that are selected based on the described approach. These routes approximately cover the whole area between the source power plant and the substations.

Results and discussions

The first experiment in this section is conducted to help determine the right value for the rate parameter λ_p in Equation 4.24. As figure 4.13 shows, this parameter does not have any major influence on the performance of the algorithm. Among the possible values between 0 and 1, $\lambda_p = 0.8$ and $\lambda_p = 1$ lead to less number of negotiation rounds. We, therefore, selected $\lambda_p = 1$ for the rest of the experiments.

In the rest of the experiments explained in this section, preference profile estimation using

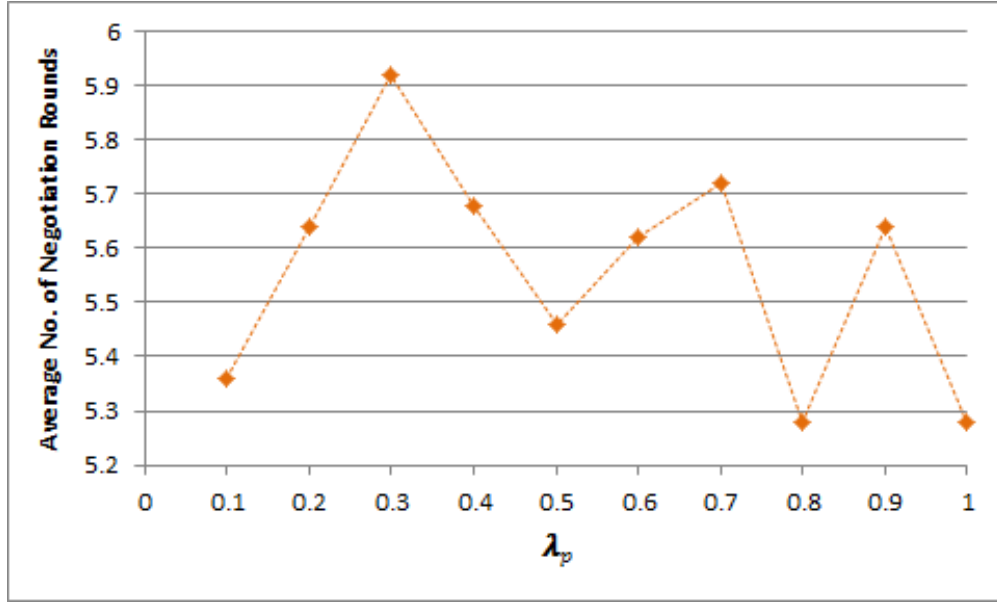


Figure 4.13: Effect of λ_p on the negotiation efficiency

the proposed UPF approach is compared to a frequency-based approach. All components of the negotiation system other than the opponent modeling component are identical in the comparisons to make the comparative analysis fair. That is, the opponents' utility functions are fuzzified in the same way, the initial knowledge-bases of the agents are identical, and the proposal preparation is performed using BPPP in both cases. The experiments, for both approaches, are conducted in two different settings: in the presence of the arguments and without arguments.

The reason for selecting the frequency-based approach is that it has been popular in related studies in automated negotiation (van Krimpen et al., 2013; Hao and Leung, 2012; van Galen Last, 2012). As explained in the related work section, the other available approaches in opponent modeling are not applicable in this case study. For example, applying the classification approach, introduced by Hindriks and Tykhonov (2008) is not feasible here as the size of the hypotheses space can grow really large considering the space of possible values for each attribute. Besides, these studies are designed for proposal-based negotiation and are not suitable for argumentation-based contexts. The main problem with heuristic opponent-learning approaches, such as the ones studied in (van Krimpen et al., 2013; Aydoğın and Yolum, 2012a; Restificar and Haddawy, 2004), is that they are mainly designed for specific negotiation contexts (e.g. bidding) and specific forms

of preferences (Restificar and Haddawy, 2004; Aydoğın and Yolum, 2012a). Thus, they cannot be generalized to other contexts. For example, in (Aydoğın and Yolum, 2012a), the preferences are defined either as a set of constraints in the form of conjunctives and disjunctives or via Conditional Preference Networks. The authors suggest a heuristic learning strategy that learns the opponents' preference models approximately. The other problem with heuristic approaches is that they are usually designed for and applied in proposal-based negotiations and cannot be easily adopted in argumentation-based negotiations. The reason is that the main part of the learning process in such studies is based on the way the opponent chooses the values of the attributes of the counter-proposal. With no access to such information in argumentation-based negotiation, heuristic approaches cannot proceed. For example, in (van Krimpen et al., 2013), the algorithm works based on the attribute values that the opponent has kept unchanged over all its offered bids. In our case studies, we receive no counteroffer from the stakeholders, and, therefore, such an approach is not applicable.

In the frequency-based approach (Eshragh et al., 2018), the parameters of an opponent's preference profile (e.g., the lower and upper preference limits as well as the weight factors) are readjusted every time an argument is received about a negotiation issue from that opponent. That is, if the argument is about increasing the value of a proposal attribute, the lower and upper preference limits and the weight factor associated with that attribute are increased to a pre-specified extent and vice versa if the argument recommends decreasing the value of an attribute.

In the experiments explained in this section, the performance of the proposed approach is measured in terms of the negotiation rounds required to reach a mutually acceptable agreement; i.e., the negotiation efficiency. Due to the stochastic nature of UPF, experiments using UPF are repeated $N_e = 50$ times. In cases where, due to graphics limitations, we cannot show the results from all the N_e trials, the demonstrated result corresponds to the mean outcome of N_e trials. To

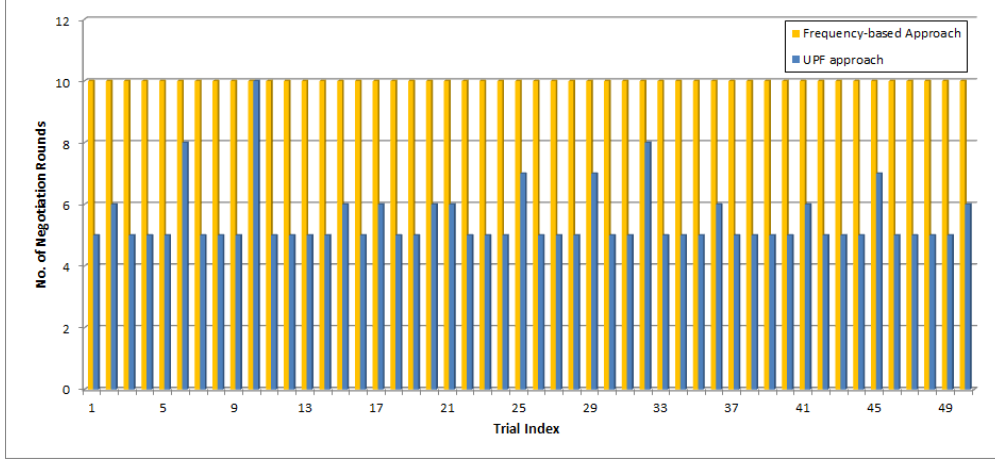


Figure 4.14: Number of negotiation rounds using UPF and frequency-based approaches

compare the results, a measure, called the improvement factor, is defined as:

$$i_n = \frac{(N_{UPF} - N_{Freq})}{N_{Freq}} * 100, n = 1, \dots, N_e \quad (4.27)$$

where i_n is the percentage of improvement for the n^{th} trial, N_{UPF} is the number of negotiation rounds required to reach an agreement using the UPF approach, and N_{Freq} is the number of negotiation rounds using the frequency-based approach. The average improvement is then calculated as:

$$I = \frac{(\sum_{n=1}^{N_e} i_n)}{N_e} \quad (4.28)$$

In a normal distribution, 99.73% of the values lie within three standard deviations of the mean. Therefore, in the experiments explained in this section, we set $\alpha = 0.727$, $\beta = 2.0$ and $\kappa = 2.0$ so that, $\lambda = -6$ based on Equation 4.8. Here, with 5 proposal attributes, $n_x = 5 * 3 = 15$ and therefore, $\sqrt{(n_x + \lambda)}$ is equal to 3 which leads us to sigma points located within three standard deviations of the original particle.

Figure 4.14 compares the number of negotiation rounds using the UPF and the frequency-based approaches.

Using the estimated preference profiles provided by the frequency-based approach, the agents reach an agreement in 10 rounds of negotiation. However, as illustrated in Figure 4.14, the UPF ap-

proach improves the performance of the negotiations process to a great extent. That is, the number of negotiation rounds using the UPF approach decreases to a minimum of five rounds. The reason for this improvement is the ability of the UPF approach to recursively estimate the preference profiles of the stakeholders based on their feedback. For example, when the UPF module receives an argument about the impact on first-nations' lands being too high, it quickly readjusts the predicted values for the lower and upper limits and the weight factor of the stakeholders' preference about this attribute and corrects the predictions using the scores assigned to the proposal. As described thoroughly in section 4.4.4, this correction process is more sophisticated than simply presuming the rate at which the correction should happen as in a frequency-based approach.

To determine the quality of the reached agreements using UPF and frequency based approaches, the agreement scores (Equation 4.3) are calculated. The agreement score according to the ESRD agent is 0.69 using both approaches. The reason the scores are the same for both approaches is that the number of alternatives in this case study is limited (100 possible solutions) and therefore, there is one alternative that is finally agreed upon no matter what approach is used by the negotiation system.

As shown in Figure 4.15, the UPF approach finds a close estimation of the upper preference limit for the FN attribute in the second round (the blue arrow) while it takes the frequency-based approach nearly eight rounds (the red arrow).

The other important difference between these two approaches is that after finding the right estimation, the UPF approach stays around the stakeholder's limit as the negotiation proceeds. However, in the frequency-based approach, the limit blindly decreases as many times as it receives arguments about the FN attribute even after reaching to the right estimation.

In another experiment, the sensitivity of the UPF approach to the arguments provided by the stakeholders is analyzed. As shown in Figure 4.16, the number of negotiation rounds using the UPF approach increases dramatically when the stakeholders send no arguments to the proposer. That is, the arguments help UPF to improve the estimations impressively, and therefore, reduce

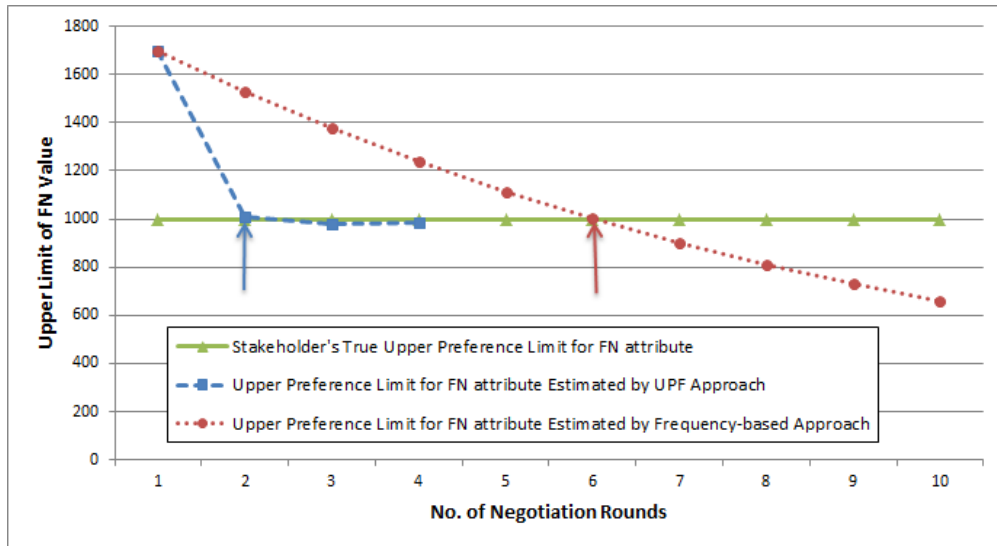


Figure 4.15: Estimated upper bound preference limit for FN attribute using UPF and frequency-based approaches

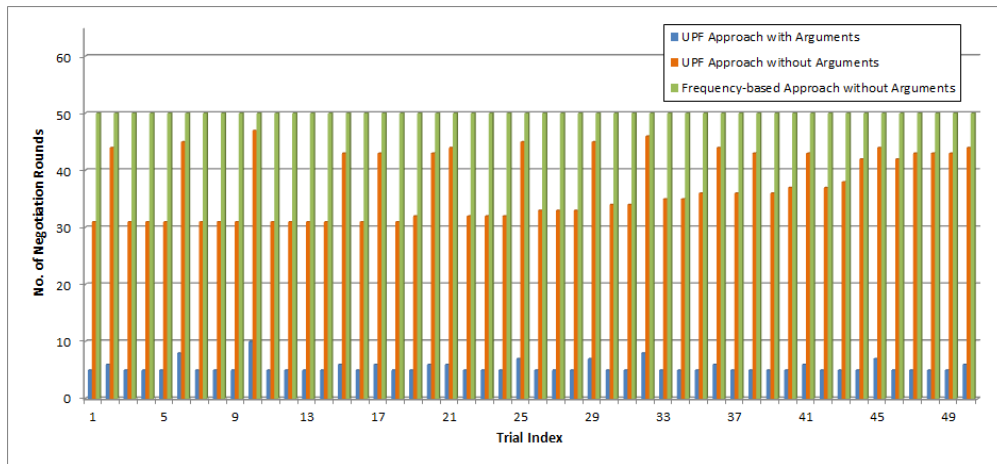


Figure 4.16: Number of negotiation rounds using UPF and frequency-based approaches with and without arguments

the number of negotiation rounds. Also, it is noticed that even without the arguments, the UPF approach still excels the frequency-based approach and reduces the required rounds of negotiation up to 38% (average improvement of 25.44%).

4.5.2 Case study B: King County House Sales

This case study is about the negotiations that occur in the context of purchasing a real estate. The dataset used in this case study contains 21614 records of house sale in King County, US, from May 2014 to May 2015 (Kaggle, 2017). Each record of the dataset represents a property and its attributes including the price, location, number of bedrooms, number of bathrooms, surface area (in sqft), the house condition (quantified by a number between 1 to 5 with five being in the best condition), and the year built. Figures 4.17 and 4.18 represent the study area of this case study and the houses available for negotiation. Through a new set of experiments, the performance, scalability, and applicability of the proposed methodology are tested.

Implementation

For this case study, a similar MAS as the first case study is used. The only difference in implementing this MAS is the proposals database that is developed using Microsoft SQL Server 2014 and MySQL 5.7 database management systems.

Data preparation

Using the location of a house and available GIS data about different facilities in King County (King County GIS Center, 2018), the distance of the house to the nearest school, hospital, park and subway station in the city is calculated and considered as additional attributes of the house. These geospatial calculations are performed using ArcGIS network analysis toolbox (Esri, 2018).

The negotiations in the second case study happen between two agents: One representing the realtor (i.e., seller) agent and other representing the buyer(s). For the realtor agent the price is the priority: the higher the price of the sold house, the more profit the realtor can make. However, the buyer tries to find a place that meets his criteria with the lowest possible price. People usually

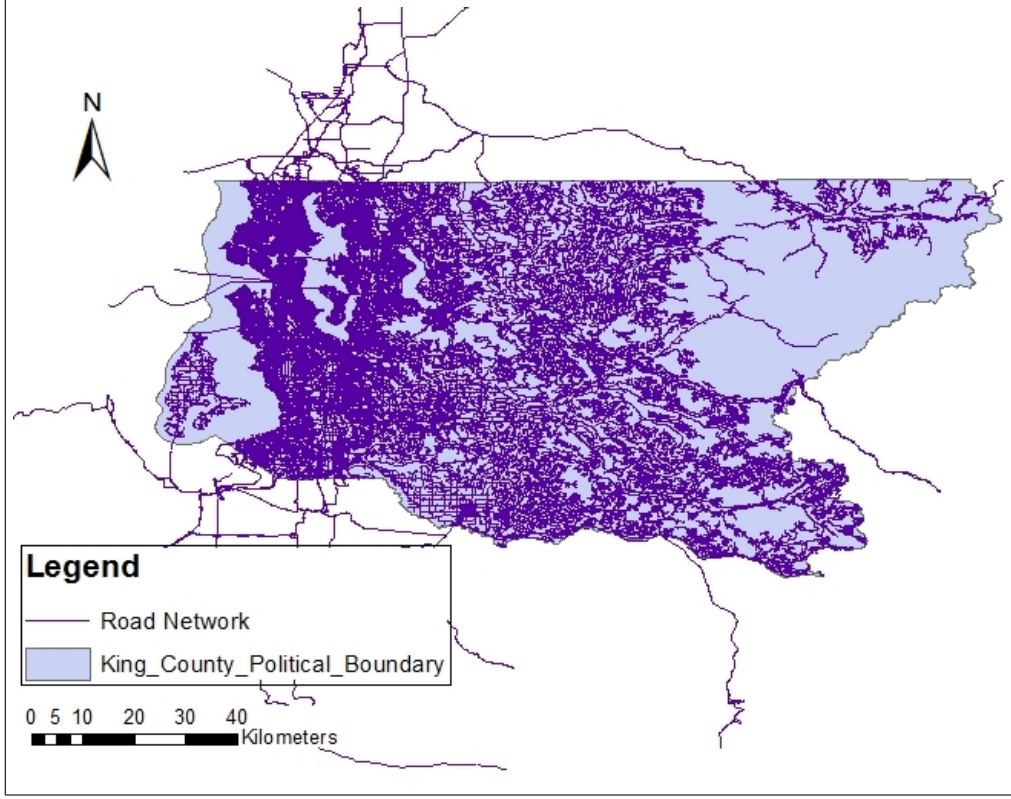


Figure 4.17: The study area-King County, Washington, US

have different preferences when buying a house. To model a wider range of buyers' behaviors, a web-based application is developed that interacts directly with different groups of people and acquires their concerns and preferences when buying a house (Figure 4.19). Through this website, we gathered data from 16 different users.

Results and Discussions

In a normal distribution, 99.73% of the values lie within three standard deviations of the mean. Therefore, in the experiments explained in this section, we set $\alpha = 0.53$, $\beta = 2.0$ and $\kappa = 2.0$ so that, $\lambda = -21$ based on Equation 4.8. Here, with 10 proposal attributes $n_x = 10 * 3 = 30$, and, therefore, $\sqrt{(n_x + \lambda)}$ is equal to 3 which leads us to sigma points within three standard deviations of the original particle.

In the first experiment, we used the users' data to build the model of the stakeholder agent and then run the MAS using UPF and frequency-based approaches. Figure 4.20 illustrates the

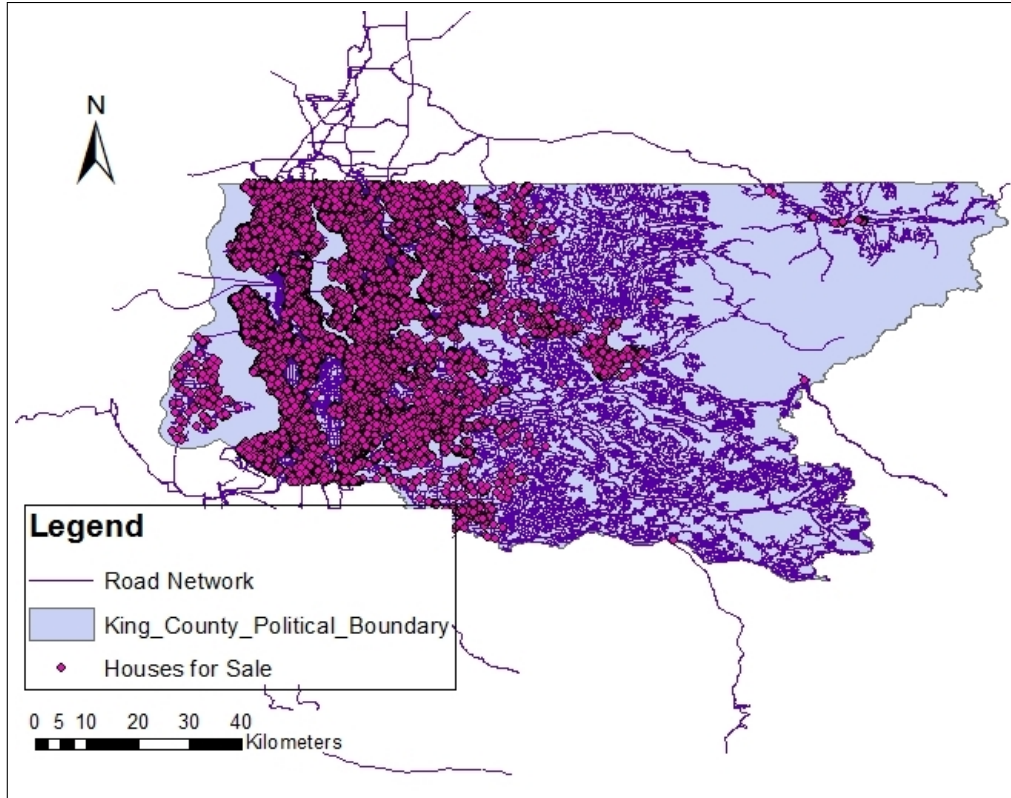
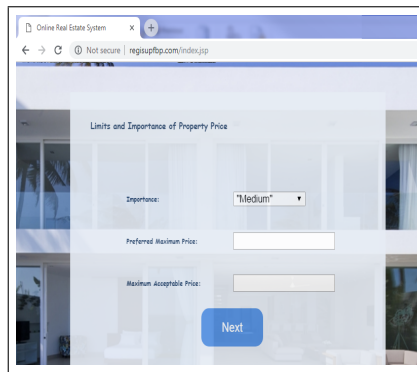


Figure 4.18: Houses for sale in King County

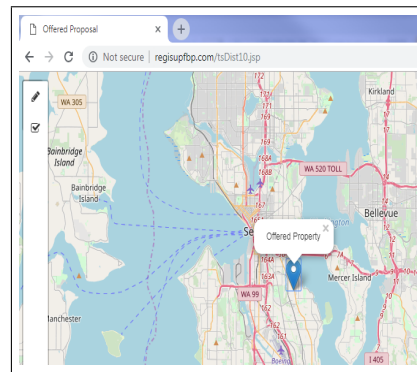
average improvement achieved by the UPF approach. As shown, the UPF approach can improve the efficiency of negotiation up to 85%. Two users are absent in this chart. That is because the criteria of these users were too easy to estimate, and the number of negotiation rounds required to reach an agreement with these users was less than three even without using the UPF approach.

The same set of experiments is repeated without passing arguments between the agents. As shown in Figure 4.21, even without passing arguments, the UPF approach can improve the negotiation efficiency up to 87.7%. For some users (e.g., U10, U11, and U12), however, the number of negotiation rounds can only be improved if the agents pass arguments. This is mainly because these users have very narrowed preference limits, too hard to estimate only based on the scores they assign to the offered properties.

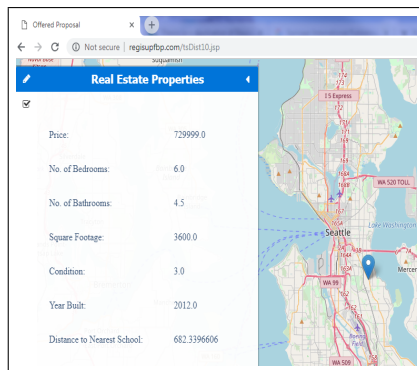
The quality of the reached agreements (measured by the agreement score) using the UPF and frequency-based approaches are compared in Figure 4.22. Although the goal of using the UPF



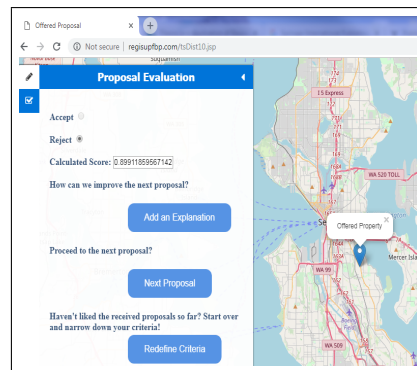
(a)



(b)



(c)



(d)

Figure 4.19: Developed website for collecting users' preference profiles; (a) Acquiring user's criteria; (b) Offering a property; (c) Representing attributes of the proposed property; (d) Acquiring user's feedback about the offered property

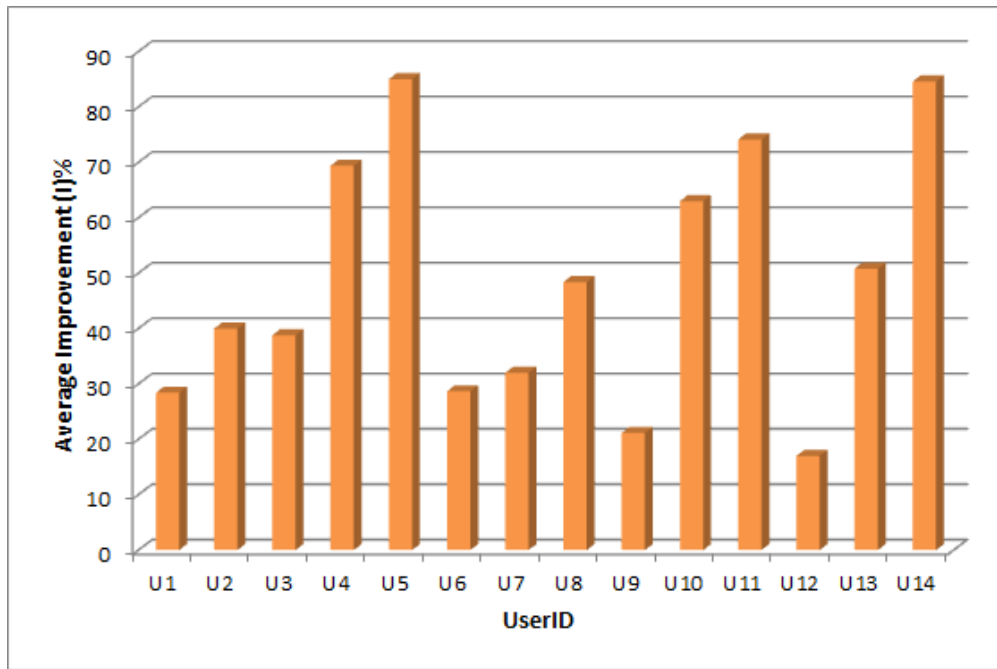


Figure 4.20: Average improvement using UPF approach with arguments over Frequency-based approach with arguments

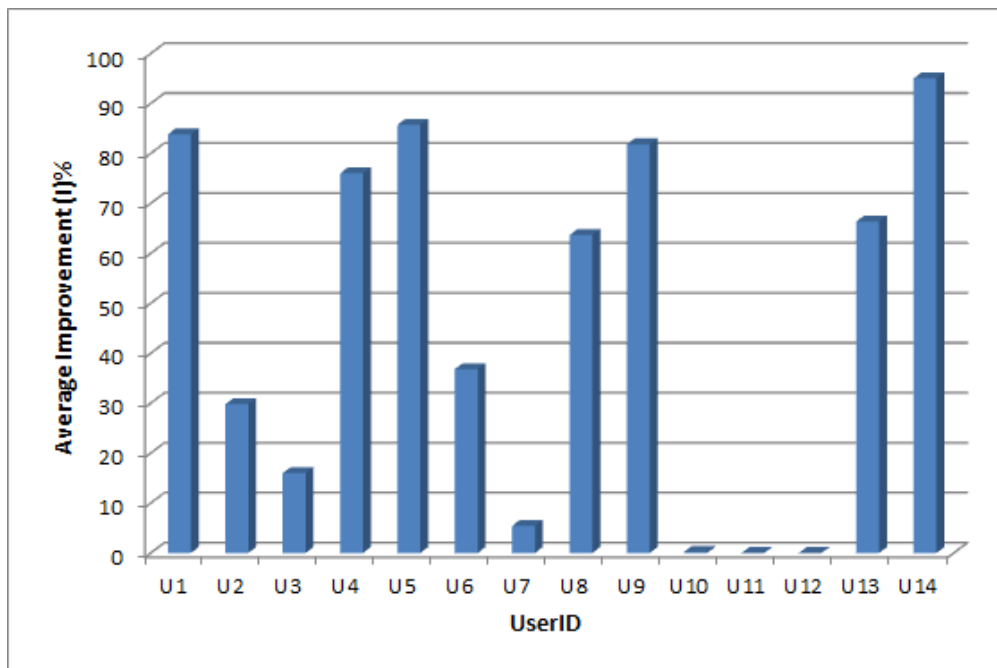


Figure 4.21: Average improvement using UPF approach without arguments over Frequency-based approach without arguments

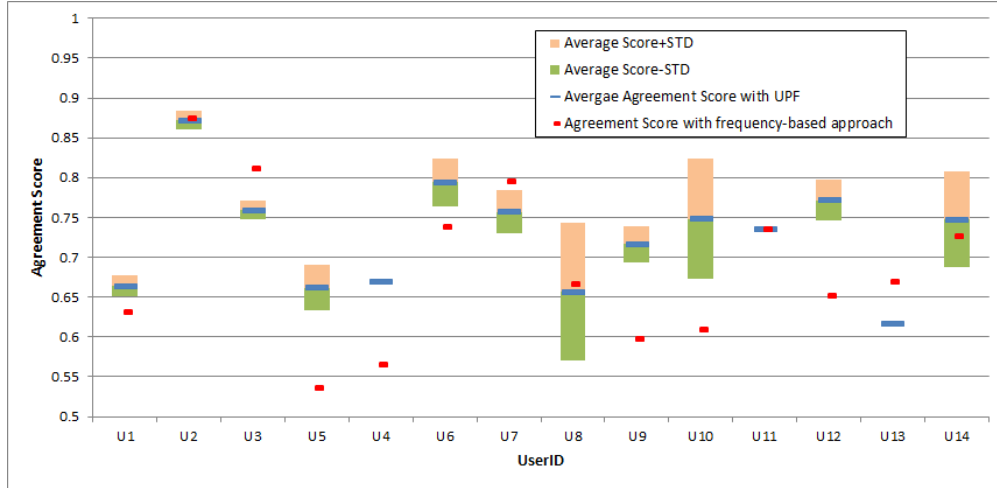


Figure 4.22: Quality of the final agreement using UPF approach and Frequency-based approach

approach is improving the efficiency of the negotiation process, the figure shows that the quality of the agreements achieved using this approach is also quite appealing comparing to the agreements reached by the frequency-based approach. As illustrated, for most of the users, the average agreement score achieved by the UPF approach is greater or equal to the agreement score reached using the frequency-based approach (shown by the red dash sign). The green and orange bars represent the range of the score for the agreements reached by the UPF approach in 50 trials. For those users with lower-quality agreements with the UPF approach, the difference is negligible.

In the negotiation context, the proposer agent must initiate the negotiation based on its limited knowledge of the other stakeholders. To investigate the effect of different initial estimations on the amount of improvement, two settings are tested. In the first setting, the buyer starts the negotiation with some realistic assumptions about the buyer's preference limits and weight factors. It means that the initial estimations are closer to the user's criteria. In the second setting, the seller does not know the buyer at all and, thus, makes random initial assumptions about the buyer's preference limits and weight factors. It means the initial estimations are very far from the user's criteria. This experiment is performed for four different users, Users 1, 4, 6, and 14. These users are particularly selected to have an example from each category of high, low and moderate improvements. For user U14, the UPF performance is considerably better than the frequency-based approach. For user

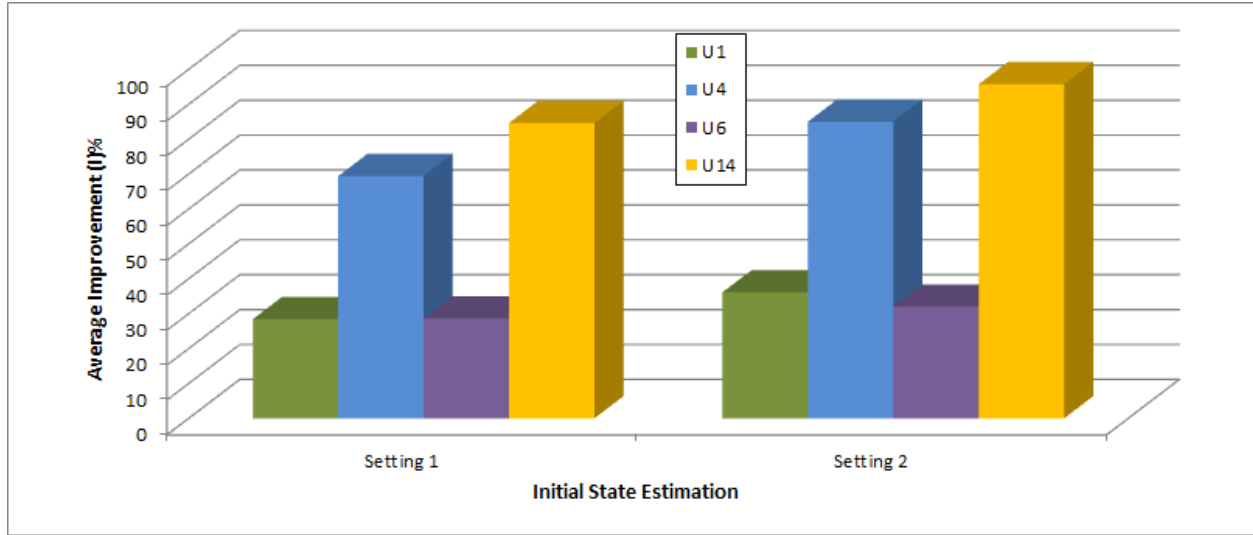


Figure 4.23: Average improvement with different initial state estimations for users U1, U4, U6, and U14

U4, the UPF improves the results to a moderate degree. For users U1 and U6, the improvement achieved by UPF is lower in comparison to the other users. We want to ensure that the high performance of UPF is not impacted by the initial approximations of the variables. Figure 4.23 illustrates the results.

In the experiment with user U1, with the first setting, 28.4% improvement is witnessed while with the second setting, 36.11% improvement is resulted. For user U4, the improvement in the case of the first setting is 69.4% which is then increased to 85% in the second setting. In the experiment with user U6, the first setting resulted in 28.66% improvement while the second setting resulted in 32% improvement. For user U14, the improvement with the first setting is 84.61% while the second setting resulted in 95.71%. That is, despite unreasonable, random initial approximations for the state variables, the UPF approach learns these parameters quickly and does not allow any decrease in performance compared to the frequency-based approach. On the other hand, the frequency-based approach depends on the frequency of the observations to rectify the initial state towards the true state; which means it is highly dependent on the initial approximations.

Figure 4.24 shows a specific example of the behavior of the two approaches for estimating the upper bound for the price attribute in the second setting for user U14. In this setting, the user's

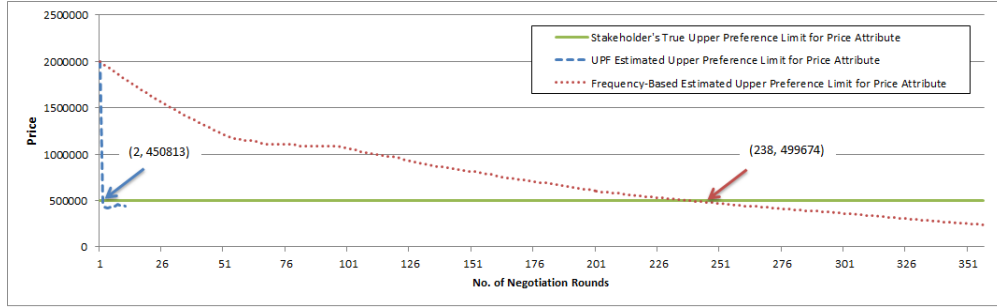
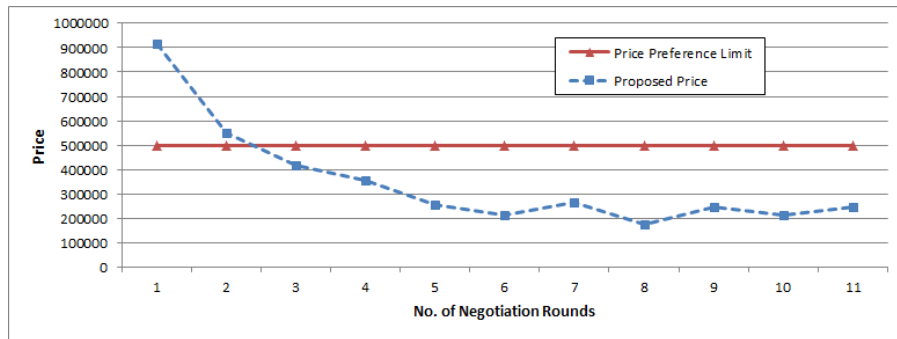


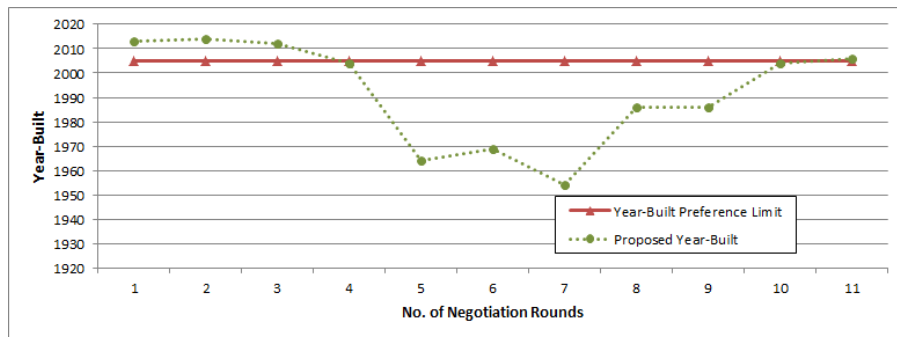
Figure 4.24: Estimated upper preference limit for price attribute using UPF approach and Frequency-based Approach

real upper bound for the price is \$500,000, but the proposer initiates the negotiation with setting this limit randomly to \$2,000,000. The UPF estimates the price upper bound right after receiving the first feedback; however, it takes the frequency-based approach 238 rounds to estimate the right limit. Also, it is noticed that the frequency-based approach keeps decreasing the limit even after finding the right value. However, the UPF approach stays around the estimated limit until the end of the negotiation. As such, the negotiation takes 375 rounds by the frequency-based approach, while the UPF approach learns the user's criteria in a few rounds and finds the appropriate proposal in only 11 rounds.

Another observation about the UPF approach is that the UPF applies the concepts of uncertainty (covariance matrix) in learning the preferences of the user. As such, in combination with BPPP module, it follows a logical trend in offering the proposals to the user. That is, when the user shows strong concern about an attribute, the uncertainty of the estimated limits/weights for that attribute increases and, therefore, the new offered proposal will be very different from the previous proposal (larger variance). However, when the user does not show any strong concern about many attributes, the uncertainty of the estimation decreases and, thus, there is no reason for the new proposal to be largely different from the last one. This type of covariance analysis is not possible in the frequency-based approach. To better understand this point, Figures 4.25 and 4.26 represents the price and the year built attribute values in the offered proposals by both the UPF and frequency-based approaches in the second setting of user U14. Although the proposed price trend seems to

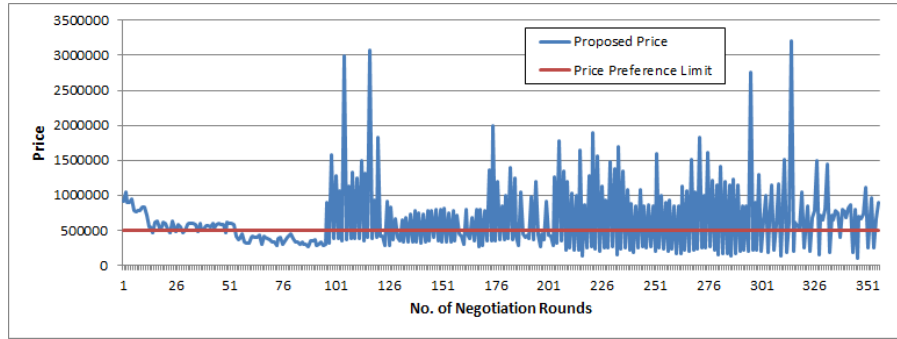


(a)

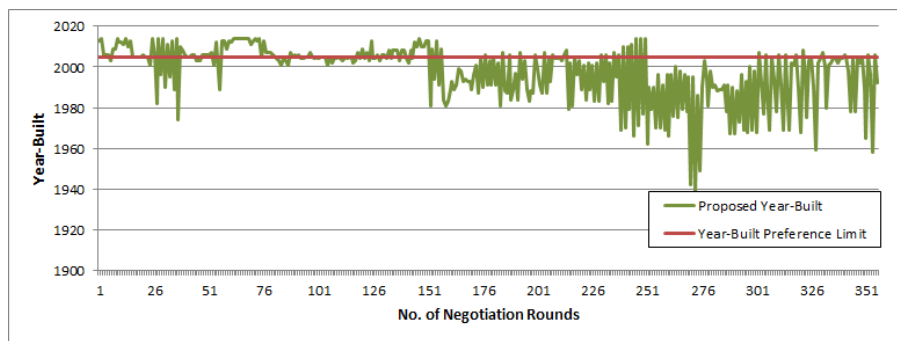


(b)

Figure 4.25: Proposed values for a) Price b)Year-Built attributes using the UPF estimations of stakeholder's criteria



(a)



(b)

Figure 4.26: Proposed values for a) Price b)Year-Built attributes using the Frequency-based estimations of stakeholder's criteria

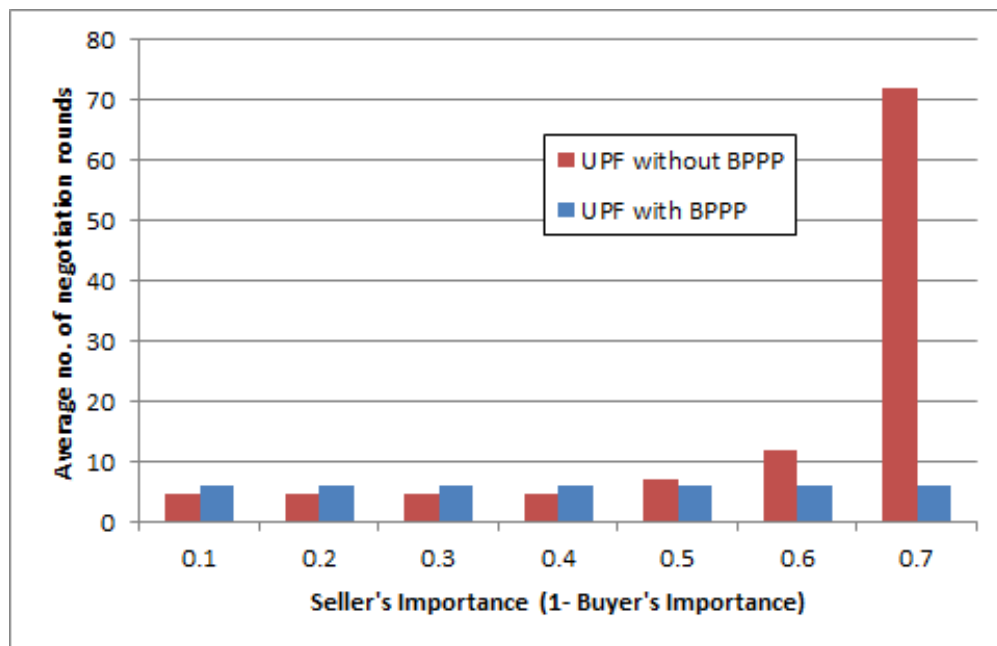


Figure 4.27: Negotiation results with UPF approach with and without BP proposal preparation

be similar for both UPF and frequency-based approach (i.e. they both approach the right solution), the pace of the trends seems different. This pace is particularly affected by the speed of learning the preferences. The other difference in the trends comes from the fact that, unlike the frequency-based approach, the UPF approach is capable of maintaining the trend because of its ability to measure and consider uncertainties.

Finally, an experiment is conducted to investigate the combination of the UPF module with the BP proposal preparation module explained in (Eshragh et al., 2018). As shown in Figure 4.27, the UPF approach works well either with or without the BPPP component. However, without using the BPPP component, the importance of different stakeholders' in the negotiation process can significantly affect the efficiency of the process. That is, using the BPPP approach removes the need for manual estimation of proper importance weights as opposed to what the utility-based approach imposes.

4.6 Conclusion and Future Work

In this paper, we presented a novel estimation technique for modeling and learning opponents' preferences in automated negotiation. Opponents' preference profiles are modeled using fuzzy functions, the parameters of which are estimated in real time using an approach based on unscented particle filtering. In each round of negotiation, a proposal is offered to stakeholders by the proposer agent. The stakeholders provide feedback to the proposer; the feedback includes a fuzzy score assigned to the proposal and some arguments about the offered proposal. Using the feedback received from the opponents and the proposal offered to them, the UPF approach recursively estimates the opponents' preference model. The estimated model is then used by the proposer agent to enable a graphical-modeling solution (BPPP module) to find the next proposal that is closer to opponents' preferences while meeting his criteria as well. That is, at each round of negotiation, the offered proposal has a higher probability of being accepted by all the stakeholders. The proposed estimation technique, therefore, facilitates the negotiation process and accelerates reaching

an agreement.

The proposed methodology was compared to a frequency-based approach in two different case studies: energy-system planning in Alberta and King County house sales. In both case studies, the UPF approach excels the frequency-based approach even when the arguments were removed from the agents' interactions. For the second case study, the experiments are repeated for 16 different real participants. The UPF approach accelerates the negotiation process up to 85% when realistic knowledge of the users is available and up to 95% where the negotiation starts with no prior knowledge about the users.

The proposed approach also considers the uncertainty of the estimations, which makes it capable of having a logical pattern in offering proposals to the opponents. For instance, when the stakeholders show strong concern about an attribute, the uncertainty of the preference profile estimated for that attribute increases; this causes the BPPP module to select a new proposal largely different from the previous one.

The performance of the proposed estimation technique depends on the initial approximations of the opponents' preferences. That is, the better the initial estimation, the quicker the UPF approach in learning the real preference profile of the user. In the future, we will focus on providing the estimation module with more accurate initial approximations using natural language processing techniques. Another future goal of this research is using the auxiliary data, such as geospatial neighborhood information, to generate pseudo-observations for the UPF module and improve the estimations even more.

Acknowledgements

This research project was partially supported by the funds available to the Schulich Chair in GeoSpatial Information Systems and Environmental Modeling held by Dr. Danielle Marceau at University of Calgary. This research was also funded by Natural Sciences and Engineering Research Council of Canada(NSERC) Discovery grant number RGPIN-2017-03881.

Chapter 5

General Conclusions and Discussions

5.1 Summary and Conclusions

Along with the widespread applications of artificial intelligence in every aspect of human life, automated negotiation has received a great deal of attention in the past few decades. There are several studies on modeling negotiation scenarios using computational models and advanced AI techniques in specific contexts such as e-commerce and supply chain management. Despite the advances made in this field, there are still challenges that need to be addressed due to the complicated nature of human interactions. Based on the context, there are two main types of negotiations: proposal-based negotiations and argumentation-based negotiations. Modeling argumentation-based negotiations that take place among multiple parties over a variety of criteria has been addressed less thoroughly compared to the other types of negotiation. Providing a way to model the participants' preferences, understanding the opponents' strategies and offering proposals with a high probability of being accepted are the most important challenges in modeling this type of negotiations. This thesis has addressed some of the problems in this domain with a focus on one-to-many argumentation-based negotiations between multiple parties over a variety of criteria.

In multi-issue multi-participant negotiations, it is essential to be able to offer an appropriate proposal in each round of negotiation based on all the available information about the negotiation participants. Considering a large set of solutions, while preparing the proposal, all negotiators' criteria need to be taken into account. The question is how to find the appropriate proposal with respect to all the negotiation issues especially when we have minimal knowledge about the opponents. This question can be broken into three detailed questions: 1- How to model stakeholders' preferences? 2- How to improve the agents' knowledge about the preferences of their opponents? 3- How to use this knowledge to find a proposal with a high likelihood of getting accepted. These

questions form the objectives of this study as: 1- Applying a close-to-reality mathematical model of stakeholders' preferences, 2- Developing a learning/estimation approach to learn this model and 3- designing a mechanism for finding the appropriate proposal based on the learned/estimated model of the stakeholders.

Starting by the third objective, initially, a basic model of stakeholders' preferences is considered and applied the frequency-based approach (already used in the literature) to estimate the stakeholders' preferences by processing their feedback. Based on this basic model, a mechanism for proposal preparation based on the estimated opponents' criteria is developed (presented in chapter 2) to address the third objective. The proposal-preparation approach represents the negotiation issues, their possible values and estimated opponents' criteria for each stakeholder as a graphical model and then approximates the inference of this model using min-sum loopy belief propagation. Having a proposal for every stakeholder using its specific model, a z-scoring approach is then used to select the final proposal. The proposed proposal preparation approach was applied to two different case studies. As the experiments show, the BP-based proposal preparation approach facilitate and accelerate the negotiation process to a great extent. It is also shown through experiments that the BP-based approach outperforms the utility-based approach regarding both the number of negotiation rounds and the fluctuations in disagreement distance measure. The experimental results indicate up to 50% improvement in the efficiency of the negotiation process using the BP-based approach.

To achieve the second specific objective of the thesis, a comprehensive review has been done to investigate the learning approaches applied in the context of automation negotiation. This review helped to determine the drawbacks of the existing automated opponent-learning approaches. We then applied this knowledge to develop a more functional solution for the opponent-learning problem.

Although the frequency-based approach for estimating opponent criteria has been widely used in the literature, it still has some problems such as defining the right rate parameter and having no

measure to check the correctness of the results. Therefore, to enhance the developed automated negotiation model, a more advanced model of stakeholders is employed to achieve the first objective of this thesis. However, the employed fuzzy model is not linear, not differentiable all time, and as the dimension of attribute space increases, the number of its parameters grows fast. Therefore, an estimation mechanism based on a recursive Bayesian filtering approach is proposed to address the second objective. In the proposed estimation approach, the agents can improve their estimations about their opponents in each round of negotiation solely based on the opponents' feedback to the offered proposal. This mechanism has a rigorous way of measuring the uncertainty of the calculated estimations, and therefore, improves the estimations if required and if not, maintains the right estimations till the proposal preparation module finds the right proposal. The proposed approach was applied to the same negotiation case studies to investigate its impact on the negotiation process. The results of the conducted experiments indicate up to 85% improvement in the efficiency and performance of the negotiation in regular negotiation scenarios while maintaining the quality of the agreement. It has been also shown through experiments that the estimation approach combines very well with the proposal preparation module.

The main contributions of the current thesis can be summarized as follows:

1. An estimation mechanism is proposed for learning the opponents' preference models considering its restrictions such as non-linearity and non-Gaussianity. This approach works based on unscented particle filtering technique and improves the estimation recursively based on the scores assigned to the offered proposals as well as provided arguments. The proposed approach reduces the number of negotiation rounds by providing the agents with extra knowledge about their opponents.
2. To prepare an appropriate proposal in every round of negotiation, a novel approach is proposed that represents the negotiation issues, proposals and current knowledge of the stakeholders' preferences via a Markov Random Field. It then uses belief propagation for probabilistic inference in this graphical model, which results in a

proposal with a high probability of being agreed upon.

3. The proposed approaches are examined using two case studies of energy system planning and real estate house buying. These are two challenging negotiation cases, which due to their complexity have not been frequently addressed in the literature.

Although the proposed techniques improve the negotiation process to a considerable extent, they still have some problems that need to be addressed in the future. For example, the belief-propagation-based proposal preparation approach provides and solves a separate graphical model for each stakeholder. That is, for each stakeholder we will have one proposal that best fits his so-far estimated needs. Because in each round only one proposal can be offered, one proposal needs to be selected among stakeholders' proposals. Currently, we apply the z-scoring method to solve this problem. However, it seems that building an aggregate probabilistic model for all stakeholders is a proper option that can improve the speed and results of the algorithm. The proposed estimation approach has some challenges as well. For example, the way the initial estimations are set affects the performance of the proposed methodology. Therefore, it is essential to design a mechanism for setting these values so that we can reach the maximum performance of the estimation approach.

5.2 Research Perspectives and Future Work

An important part of a negotiation is the initial phase where participants start to know each other and understand the atmosphere of the negotiation. Agents in automated negotiation are incapable of such understandings so far. In the future, a natural language processing component which is in charge of comprehending real stakeholders' concerns and viewpoints will be added to the automated model. This component can be helpful in building the agents more accurately and then equipping them with some knowledge about the others that can be extracted from initial talks between stakeholders. The better this component implemented, the more successful the model will be in capturing the characters and their behavior in real-world negotiation.


Emotions of the negotiation parties usually influence real-world negotiations. Despite the inevitable role of emotions in human negotiation, only a few studies have been conducted on the incorporation of emotions in automated negotiation. Jiang et al. (2006) proposed an automated negotiation model in which emotions have been incorporated. However, to get the real essence of actual negotiations, it is required to model and handle negotiators' emotions during the negotiation process. In the future, a model of emotion for agents will be developed their behavior during the negotiation will be investigated under the influence of their mood and feelings.

In the future, the estimation mechanism will also be improved to be applicable to other types of negotiation with different negotiators' behaviors and feedback. Moreover, some auxiliary data will be utilized to help the agents extend their knowledge about their opponents through external resources during the negotiation process.

Another possible future path for this research is investigating the correlations and interconnections between negotiation issues. Considering these sort of correlations, techniques such as principal component analysis (PCA) along with conditional random field (CRF), instead of MRF, can be used to improve the proposal preparation process. PCA can also be used to reduce the number of negotiation attributes and improve the efficiency of the UPF approach.

Chapter 6

Appendix

 ELSEVIER

About ElsevierProducts & SolutionsServicesShop & Discover

Search

Shopping cart icon

User icon

Home > About > Policies > Copyright > Permissions

Permissions

As a general rule, permission should be sought from the rights holder to reproduce any substantial part of a copyrighted work. This includes any text, illustrations, charts, tables, photographs, or other material from previously published sources. Obtaining permission to re-use content published by Elsevier is simple. Follow the guide below for a quick and easy route to permission.

[Permission guidelines](#)[ScienceDirect content](#)[Non-ScienceDirect content](#)[Tutorial videos](#)[Help and support](#)

Permission guidelines

For further guidelines about obtaining permission, please review our Frequently Asked Questions below:

[When is permission required?](#) +

[When is permission not required?](#) +

[From whom do I need permission?](#) +

[How do I obtain permission to use photographs or illustrations?](#) +

[Do I need to obtain permission to use material posted on a website?](#) +

[What rights does Elsevier require when requesting permission?](#) +

[How do I obtain permission from another publisher?](#) +

[What is Rightslink?](#) +

[What should I do if I am not able to locate the copyright owner?](#) +

[What is Elsevier's policy on using patient photographs?](#) +

[Can I obtain permission from a Reproduction Rights Organization \(RRO\)?](#) +

[Is Elsevier an STM signatory publisher?](#) +

[Do I need to request permission to re-use work from another STM publisher?](#) +

[Do I need to request permission to text mine Elsevier content?](#) +

[Can I post my article on ResearchGate without violating copyright?](#) +

[Can I post on ArXiv?](#) +

[Can I include/use my article in my thesis/dissertation?](#) –

Yes. Authors can include their articles in full or in part in a thesis or dissertation for non-commercial purposes.

[Which uses of a work does Elsevier view as a form of 'prior publication'?](#) +

SPRINGER NATURE LICENSE TERMS AND CONDITIONS

May 13, 2019

This Agreement between Faezeh Eshragh ("You") and Springer Nature ("Springer Nature") consists of your license details and the terms and conditions provided by Springer Nature and Copyright Clearance Center.

| | |
|--|---|
| License Number | 4587301150183 |
| License date | May 13, 2019 |
| Licensed Content Publisher | Springer Nature |
| Licensed Content Publication | Springer eBook |
| Licensed Content Title | Using Belief Propagation-Based Proposal Preparation for Automated Negotiation over Environmental Issues |
| Licensed Content Author | Faezeh Eshragh, Mozhddeh Shahbazi, Behrouz Far |
| Licensed Content Date | Jan 1, 2019 |
| Type of Use | Thesis/Dissertation |
| Requestor type | non-commercial (non-profit) |
| Format | electronic |
| Portion | full article/chapter |
| Will you be translating? | no |
| Circulation/distribution | <501 |
| Author of this Springer Nature content | yes |
| Title | Multi-Criteria Multi-Participant Automated Negotiation: Belief Propagation-based Proposal Preparation and Real Time Opponent Learning |
| Institution name | University of Calgary |
| Expected presentation date | May 2019 |
| Requestor Location | Faezeh Eshragh 500 Rocky Vista Gardens Nw 429 Calgary, AB T3G 0C3 Canada Attn: Faezeh Eshragh |
| Total | 0.00 CAD |
| Terms and Conditions | |

Springer Nature Terms and Conditions for RightsLink Permissions

Springer Nature Customer Service Centre GmbH (the Licensor) hereby grants you a non-exclusive, world-wide licence to reproduce the material and for the purpose and requirements specified in the attached copy of your order form, and for no other use, subject to the conditions below:

1. The Licensor warrants that it has, to the best of its knowledge, the rights to license reuse of this material. However, you should ensure that the material you are requesting is original to the Licensor and does not carry the copyright of another entity (as credited in the published version).

If the credit line on any part of the material you have requested indicates that it was reprinted or adapted with permission from another source, then you should also seek permission from that source to reuse the material.

2. Where **print only** permission has been granted for a fee, separate permission must be obtained for any additional electronic re-use.
3. Permission granted **free of charge** for material in print is also usually granted for any electronic version of that work, provided that the material is incidental to your work as a whole and that the electronic version is essentially equivalent to, or substitutes for, the print version.
4. A licence for 'post on a website' is valid for 12 months from the licence date. This licence does not cover use of full text articles on websites.
5. Where '**reuse in a dissertation/thesis**' has been selected the following terms apply: Print rights of the final author's accepted manuscript (for clarity, NOT the published version) for up to 100 copies, electronic rights for use only on a personal website or institutional repository as defined by the Sherpa guideline (www.sherpa.ac.uk/romeo/).
6. Permission granted for books and journals is granted for the lifetime of the first edition and does not apply to second and subsequent editions (except where the first edition permission was granted free of charge or for signatories to the STM Permissions Guidelines <http://www.stm-assoc.org/copyright-legal-affairs/permissions/permissions-guidelines/>), and does not apply for editions in other languages unless additional translation rights have been granted separately in the licence.
7. Rights for additional components such as custom editions and derivatives require additional permission and may be subject to an additional fee. Please apply to Journalpermissions@springernature.com/bookpermissions@springernature.com for these rights.
8. The Licensor's permission must be acknowledged next to the licensed material in print. In electronic form, this acknowledgement must be visible at the same time as the figures/tables/illustrations or abstract, and must be hyperlinked to the journal/book's homepage. Our required acknowledgement format is in the Appendix below.
9. Use of the material for incidental promotional use, minor editing privileges (this does not include cropping, adapting, omitting material or any other changes that affect the meaning, intention or moral rights of the author) and copies for the disabled are permitted under this licence.
10. Minor adaptations of single figures (changes of format, colour and style) do not require the Licensor's approval. However, the adaptation should be credited as shown in Appendix below.

Appendix — Acknowledgements:

For Journal Content:

Reprinted by permission from [the Licensor]: [Journal Publisher (e.g. Nature/Springer/Palgrave)] [JOURNAL NAME] [REFERENCE CITATION (Article name, Author(s) Name), [COPYRIGHT] (year of publication)]

For Advance Online Publication papers:

Reprinted by permission from [the Licensor]: [Journal Publisher (e.g. Nature/Springer/Palgrave)] [JOURNAL NAME] [REFERENCE CITATION (Article name, Author(s) Name), [COPYRIGHT] (year of publication), advance online publication, day month year (doi: 10.1038/sj.[JOURNAL ACRONYM].)]

For Adaptations/Translations:

Adapted/Translated by permission from [the Licensor]: [Journal Publisher (e.g. Nature/Springer/Palgrave)] [JOURNAL NAME] [REFERENCE CITATION (Article name, Author(s) Name), [COPYRIGHT] (year of publication)]

Note: For any republication from the British Journal of Cancer, the following credit line style applies:

Reprinted/adapted/translated by permission from [the Licensor]: on behalf of Cancer Research UK: : [Journal Publisher (e.g. Nature/Springer/Palgrave)] [JOURNAL NAME] [REFERENCE CITATION (Article name, Author(s) Name), [COPYRIGHT] (year of publication)]

For Advance Online Publication papers:

Reprinted by permission from The [the Licensor]: on behalf of Cancer Research UK: [Journal Publisher (e.g. Nature/Springer/Palgrave)] [JOURNAL NAME] [REFERENCE CITATION (Article name, Author(s) Name), [COPYRIGHT] (year of publication), advance online publication, day month year (doi: 10.1038/sj. [JOURNAL ACRONYM])]

For Book content:

Reprinted/adapted by permission from [the Licensor]: [Book Publisher (e.g. Palgrave Macmillan, Springer etc) [Book Title] by [Book author(s)] [COPYRIGHT] (year of publication)]

Other Conditions:

Version 1.1

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.

Bibliography

- Akhbari, M. and Grigg, N. S. (2013). A framework for an agent-based model to manage water resources conflicts. *Water Resources Management*, 27(11):4039–4052.
- Alam, M., Rogers, A., and Ramchurn, S. (2013). Interdependent multi-issue negotiation for energy exchange in remote communities. In *Proceedings of the Twenty-Seventh Conference on Artificial Intelligence (AAAI-2013)*, Bellevue, WA, page 25–31.
- Alfonso, B., Botti, V., Garrido, A., and Giret, A. (2014). A mas-based infrastructure for negotiation and its application to a water-right market. *Information Systems Frontiers*, 16(2):183–199.
- An, B. and Lesser, V. (2012). *Yushu: A Heuristic-Based Agent for Automated Negotiating Competition*, pages 145–149. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Andrieu, C., De Freitas, N., and Doucet, A. (1999). Sequential mcmc for bayesian model selection. In *Proceedings of the IEEE Signal Processing Workshop on Higher-Order Statistics, Caesarea, Israel*, pages 130–134.
- Aragonès, E. and Dellunde, P. (2008). An automated model of government formation. In *Proceeding of Workshop on the Political Economy of Democracy, Barcelona, Spain*, pages 279–303.
- Aydogan, R., Baarslag, T., Hindriks, K. V., Jonker, C. M., and Yolum, P. (2013). Heuristic-based approaches for cp-nets in negotiation. In Ito, T., Zhang, M., Robu, V., and Matsuo, T., editors, *Complex Automated Negotiations: Theories, Models, and Software Competitions*, pages 113–123. Springer, Berlin, Heidelberg.
- Aydogan, R. and Yolum, P. (2009). Ontology-based learning for negotiation. In *Proceeding of IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technologies, 2009*, volume 2, pages 177–184.

- Aydoğar, R. and Yolum, P. (2012a). The effect of preference representation on learning preferences in negotiation. In Ito, T., Zhang, M., Robu, V., Fatima, S., and Matsuo, T., editors, *New Trends in Agent-Based Complex Automated Negotiations*, pages 3–20. Springer, Berlin, Heidelberg.
- Aydoğar, R. and Yolum, P. (2012b). Learning opponent’s preferences for effective negotiation: an approach based on concept learning. *Autonomous Agents and Multi-Agent Systems*, 24(1):104–140.
- Azaria, A., Aumann, Y., and Kraus, S. (2012). Automated strategies for determining rewards for human work. In *Proceedings of the Twenty-Sixth Conference on Artificial Intelligence (AAAI-2012)*, Toronto, Canada.
- Baarslag, T., Hendriks, M. J. C., Hindriks, K. V., and Jonker, C. M. (2016). Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Autonomous Agents and Multi-Agent Systems*, 30(5):849–898.
- Baarslag, T., Hindriks, K., Jonker, C., Kraus, S., and Lin, R. (2012). The first automated negotiating agents competition (anac 2010). In Ito, T., Zhang, M., Robu, V., Fatima, S., and Matsuo, T., editors, *New Trends in agent-based complex automated negotiations*, pages 113–135. Springer.
- Back, T. (1996). *Evolutionary algorithms in theory and practice: evolution strategies, evolutionary programming, genetic algorithms*. Oxford university press.
- Barbuceanu, M. and Lo, W.-K. (2000). A multi-attribute utility theoretic negotiation architecture for electronic commerce. In *Proceedings of the fourth international conference on Autonomous agents, Barcelona, Spain*, pages 239–246. ACM.
- Barreteau, O., Bousquet, F., Étienne, M., Souchère, V., and d’Aquino, P. (2014). Companion modelling: a method of adaptive and participatory research. In Étienne, M., editor, *Companion modelling*, pages 13–40. Springer.

- Barto, A. G., Sutton, R. S., and Watkins, C. J. (1990). Learning and sequential decision making. In Gabriel, M. and Moore, J., editors, *Learning and computational neuroscience: Foundations of adaptive networks*. MIT Press, Cambridge.
- Bäumer, C. and Magedanz, T. (1999). Grasshopper: a mobile agent platform for active telecommunication networks. *Intelligent Agents for Telecommunication Applications(IATA 1999). Lecture Notes in Computer Science (Lecture Notes in Artificial Intelligence)*, 1699:690–690.
- Belaqziz, S. et al. (2011). An agent-based modeling approach for decision-making in gravity irrigation systems. In *Proceeding of International Conference for Internet Technology and Secured Transactions (ICITST)*, pages 673–680. IEEE.
- Beyer, H.-G. and Schwefel, H.-P. (2002). Evolution strategies—a comprehensive introduction. *Natural computing*, 1(1):3–52.
- Bigham, J. and Du, L. (2003). Cooperative negotiation in a multi-agent system for real-time load balancing of a mobile cellular network. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 568–575.
- Binmore, K. (1992). *Fun and games:A text on game theory*. DC Heath and Company, Lexington, MA.
- Bousquet, F., Bakam, I., Proton, H., and Le Page, C. (1998). Cormas: common-pool resources and multi-agent systems. In *Proceeding of International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems*, pages 826–837.
- Bousquet, F. and Trébuil, G. (2005). Introduction to companion modeling and multi-agent systems for integrated natural resource management in Asia. *Companion Modeling and Multi-Agent Systems for Integrated Natural Resource Management in Asia*, pages 1–20.
- Boutilier, C., Brafman, R. I., Domshlak, C., Hoos, H. H., and Poole, D. (2004). Cp-nets: A tool

- for representing and reasoning with conditional ceteris paribus preference statements. *Journal of artificial intelligence research*, 21:135–191.
- Bowling, M. and Veloso, M. (2001). Rational and convergent learning in stochastic games. In *Proceeding of International joint conference on artificial intelligence*, volume 17, pages 1021–1026.
- Bowling, M. and Veloso, M. (2002). Multiagent learning using a variable learning rate. *Artificial Intelligence*, 136(2):215–250.
- Brzostowski, J. and Kowalczyk, R. (2006). Adaptive negotiation with on-line prediction of opponent behaviour in agent-based negotiations. In *Proceedings of the IEEE/WIC/ACM international conference on Intelligent Agent Technology*, pages 263–269.
- Buffett, S. and Spencer, B. (2005). Learning opponents’ preferences in multi-object automated negotiation. In *Proceedings of the 7th international conference on Electronic commerce*, pages 300–305.
- Buffett, S. and Spencer, B. (2007). A bayesian classifier for learning opponents’ preferences in multi-object automated negotiation. *Electronic Commerce Research and Applications*, 6(3):274–284.
- Bui, H. H., Venkatesh, S., and Kieronska, D. (1999). Learning other agents’ preferences in multi-agent negotiation using the bayesian classifier. *International Journal of Cooperative Information Systems*, 8(04):275–293.
- Cantillo, V., Heydecker, B., and de Dios Ortúzar, J. (2006). A discrete choice model incorporating thresholds for perception in attribute values. *Transportation Research Part B: Methodological*, 40(9):807–825.
- Cao, M., Luo, X., Luo, X. R., and Dai, X. (2015). Automated negotiation for e-commerce decision

- making: A goal deliberated agent architecture for multi-strategy selection. *Decision Support Systems*, 73:1–14.
- Carbonneau, R., Kersten, G. E., and Vahidov, R. (2008). Predicting opponent’s moves in electronic negotiations using neural networks. *Expert Systems with Applications*, 34(2):1266–1273.
- Carbonneau, R. A., Kersten, G. E., and Vahidov, R. M. (2011). Pairwise issue modeling for negotiation counteroffer prediction using neural networks. *Decision Support Systems*, 50(2):449–459.
- Carraro, C., Marchiori, C., and Sgobbi, A. (2007). Negotiating on water: insights from non-cooperative bargaining theory. *Environment and Development Economics*, 12(2):329–349.
- Chen, L., Dong, H., Han, Q., and Cui, G. (2013). Bilateral multi-issue parallel negotiation model based on reinforcement learning. In *Proceeding of International Conference on Intelligent Data Engineering and Automated Learning*, pages 40–48.
- Chen, S. and Weiss, G. (2015). An approach to complex agent-based negotiations via effectively modeling unknown opponents. *Expert Systems with Applications*, 42(5):2287–2304.
- Chen, S., Weiss, G., and Zhou, S. (2016). Solving negotiation problems against unknown opponents with wisdom of crowds. In *Proceedings of Joint German/Austrian Conference on Artificial Intelligence (Künstliche Intelligenz)*, pages 126–133. Springer.
- Cheng, C.-B., Chan, C.-C. H., and Lin, K.-C. (2006). Intelligent agents for e-marketplace: Negotiation with issue trade-offs by fuzzy inference systems. *Decision Support Systems*, 42(2):626–638.
- Chesnevar, C. I., Maguitman, A. G., and Loui, R. P. (2000). Logical models of argument. *ACM Computing Surveys*, 32(4):337–383.
- Costantini, S., De Gasperis, G., Proveti, A., and Tsintza, P. (2013). A heuristic approach to proposal-based negotiation: with applications in fashion supply chain management. *Mathematical Problems in Engineering*, 2013:1–15.

- Davis, R. and Smith, R. G. (1983). Negotiation as a metaphor for distributed problem solving. *Artificial intelligence*, 20(1):63–109.
- Doucet, A. and Gordon, N. J. (1999). Simulation-based optimal filter for maneuvering target tracking. In *Proceedings of Signal and Data Processing of Small Targets*, volume 3809, pages 241–256.
- Doucet, A. and Johansen, A. M. (2011). A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of nonlinear filtering*, 12:656–704.
- Dworman, G., Kimbrough, S. O., and Laing, J. D. (1996). Bargaining by artificial agents in two coalition games: A study in genetic programming for electronic commerce. In *Proceedings of the first annual conference on genetic programming*, pages 54–62.
- Eiben, A. E., Smith, J. E., et al. (2003). *Introduction to evolutionary computing*. Springer Verlag, Berlin.
- El-Sisi, A. B. and Mousa, H. M. (2012). Argumentation based negotiation in multi-agent system. In *Proceedings of Seventh International Conference on Computer Engineering & Systems (ICCES)*, pages 261–266.
- Eshragh, F., Shahbazi, M., and Far, B. (2017). Automated dynamic negotiation over environmental issues. In *Proceedings of International Conference on Information Reuse and Integration, (IRI2017), San Diego, US*, pages 92–99.
- Eshragh, F., Shahbazi, M., and Far, B. (2018). *Using Belief Propagation-based Proposal Preparation for Automated Negotiation over Environmental Issues*, pages 69–95. Springer Switzerland.
- Esri (Accessed: July 2017). Arcpy library. <http://pro.arcgis.com/en/pro-app/arcpy/get-started/what-is-arcpy-.htm>.

- Esri (Accessed: July 2018). Network analysis toolbox. <http://desktop.arcgis.com/en/arcmap/10.3/tools/network-analyst-toolbox/an-overview-of-the-network-analyst-toolbox.htm>.
- Fang, F., Xin, Y., Yun, X., and Haitao, X. (2008). An opponent's negotiation behavior model to facilitate buyer-seller negotiations in supply chain management. In *Proceeding of International Symposium on Electronic Commerce and Security*, pages 582–587.
- Faratin, P., Sierra, C., and Jennings, N. R. (1998). Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 24(3-4):159–182.
- Faratin, P., Sierra, C., and Jennings, N. R. (2002). Using similarity criteria to make issue trade-offs in automated negotiations. *artificial Intelligence*, 142(2):205–237.
- Fatima, S., Wooldridge, M., and Jennings, N. R. (2009). An analysis of feasible solutions for multi-issue negotiation involving nonlinear utility functions. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems*, volume 2, pages 1041–1048.
- Fatima, S. S. and Wooldridge, M. (2001). Adaptive task resources allocation in multi-agent systems. In *Proceedings of the Fifth International Conference on Autonomous Agents, New York, USA*, pages 537–544.
- Fink, A. (2006). Supply chain coordination by means of automated negotiations between autonomous agents. *Multiagent-based supply chain management*, 28:351–372.
- Freeman, R. E. (2010). *Strategic management: A stakeholder approach*. Cambridge university press, Cambridge, UK.
- GDAL/OGR contributors (Accessed: July 2017). Geospatial Data Abstraction software Library. <https://www.gdal.org>.
- Geweke, J. (1989). Bayesian inference in econometric models using monte carlo integration. *Econometrica: Journal of the Econometric Society*, pages 1317–1339.

- Gilbert, M. A., Grasso, F., Groarke, L., Gurr, C., and Gerlofs, J. M. (2004). The persuasion machine. In Reed, C. and Norman, T. J., editors, *Argumentation Machines: New Frontiers in Argument and Computation*, pages 121–174.
- Gurney, K. (1997). *An introduction to neural networks*. CRC press, Florida, US.
- Haberland, V., Miles, S., and Luck, M. (2012). Adaptive negotiation for resource intensive tasks in grids. In Kersting, K. and Toussaint, M., editors, *Proceedings of the 6th starting AI researchers' symposium of frontiers in artificial intelligence and applications*, volume 241, pages 125–136.
- Hadfi, R. and Ito, T. (2015). Low-complexity exploration in utility hypergraphs. *Journal of information processing*, 23(2):176–184.
- Hadfi, R. and Ito, T. (2016). On the complexity of utility hypergraphs. In Fukuta, N., Ito, T., Zhang, M., Fujita, K., and Robu, V., editors, *Recent Advances in Agent-based Complex Automated Negotiation*, pages 89–105.
- Hao, J. and Leung, H.-F. (2012). Abines: An adaptive bilateral negotiating strategy over multiple items. In *Proceedings of the The 2012 IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technology-Volume 02*, pages 95–102. IEEE Computer Society.
- Hindriks, K. and Tykhonov, D. (2008). Opponent modelling in automated multi-issue negotiation using bayesian learning. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, volume 1, pages 331–338.
- Holling, C. S. (1978). *Adaptive environmental assessment and management*. John Wiley & Sons, Chichester, UK.
- Hossain, S. M. (2012). Selecting negotiation strategies for meeting scheduling using a model based approach. *Procedia Computer Science*, 10:1217–1220.

- Huang, C. C., Liang, W.-Y., Lai, Y. H., and Lin, Y. (2010). The agent-based negotiation process for b2c e-commerce. *Expert Systems with Applications*, 37(1):348–359. cited By 78.
- Huang, S. and Lin, F. (2007). The design and evaluation of an intelligent sales agent for online persuasion and negotiation. *Electronic Commerce Research and Applications*, 6(3):285–296.
- Ito, T., Hattori, H., and Klein, M. (2007). Multi-issue negotiation protocol for agents: Exploring nonlinear utility spaces. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence(IJCAI07), Hyderabad, India*, volume 7, pages 1347–1352.
- Ito, T., Zhang, M., Robu, V., Fatima, S., and Matsuo, T. (2011). *New trends in agent-based complex automated negotiations*, volume 383. Springer.
- Jain, P. and Dahiya, D. (2012). An intelligent multi agent framework for e-commerce using case based reasoning and argumentation for negotiation. *Information Systems, Technology and Management*, pages 164–175.
- Jazayeriy, H., Azmi-Murad, M., Sulaiman, N., and Izura Udizir, N. (2011). The learning of an opponent’s approximate preferences in bilateral automated negotiation. *Journal of theoretical and applied electronic commerce research*, 6(3):65–84.
- Jennings, N. R., Corera, J. M., and Laresgoiti, I. (1995). Developing industrial multi-agent systems. In *Proceedings of the first International Conference on Multiagent Systems, San Francisco, USA*, pages 423–430.
- Jennings, N. R., Faratin, P., Johnson, M., Norman, T. J., O’Brien, P., and Wiegand, M. (1996). Agent-based business process management. *International Journal of Cooperative Information Systems*, 5(02n03):105–130.
- Jennings, N. R., Faratin, P., Lomuscio, A. R., Parsons, S., Wooldridge, M. J., and Sierra, C. (2001). Automated negotiation: prospects, methods and challenges. *Group Decision and Negotiation*, 10(2):199–215.

- Jennings, N. R., Faratin, P., Norman, T. J., O'Brien, P., Odgers, B., and Alty, J. L. (2000). Implementing a business process management system using adept: A real-world case study. *Applied Artificial Intelligence*, 14(5):421–463.
- Jennings, N. R., Parsons, S., Norriega, P., and Sierra, C. (1998). On augmentation-based negotiation. In *Proceedings of the international workshop on multi-agent systems (IWMAS-98), Boston, US*, pages 1–7.
- Ji, S.-j., Zhang, C.-j., Sim, K.-M., and Leung, H.-f. (2014). A one-shot bargaining strategy for dealing with multifarious opponents. *Applied Intelligence*, 40(4):557–574.
- Jiang, H., Vidal, J. M., and Huhns, M. N. (2006). Incorporating emotions into automated negotiation. In *Proceedings of the Agent Construction and Emotions Workshop*.
- Jonker, C. M., Hindriks, K. V., Wiggers, P., and Broekens, J. (2012). Negotiating agents. *Artificial Intelligence Magazine*, 33(3):79.
- Kaelbling, L. P., Littman, M. L., and Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4:237–285.
- Kaggle (Accessed:August 2017). House Sales in King County, USA. <https://www.kaggle.com/harlfoxem/housesalesprediction>.
- Kakas, A. and Moraitis, P. (2006). Adaptive agent negotiation via argumentation. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 384–391.
- Kanbur, R. (2005). Pareto’s revenge. *Journal of Social and Economic Development*, 7:1–11.
- Karunatilake, N. C., Jennings, N. R., Rahwan, I., and McBurney, P. (2009). Dialogue games that agents play within a society. *Artificial intelligence*, 173(9-10):935–981.

- Kenny, B., Vredenburg, H., and Lucas, A. (2012). The new role of law in stimulating industrial innovation and regional development: The canadian experience with reflexive law in reconciling economic development, environmental protection and entrepreneurship in the energy industry. *International Journal of Innovation and Regional Development*, 4(1):8–27.
- Kieser, M. and Marceau, D. J. (2011). Simulating a land development planning process through agent-based modeling. In Alkhateeb, F., Al Maghayreh, E., and Abu Doush, I., editors, *Multi-Agent Systems-Modeling, Control, Programming, Simulations and Applications*, pages 416–450.
- King County GIS Center (Accessed: May 2018). King County GIS Open Data. <https://gis-kingcounty.opendata.arcgis.com>.
- Kraus, S. (1997). Negotiation and cooperation in multi-agent environments. *Artificial intelligence*, 94(1-2):79–97.
- Kraus, S. (2001a). Automated negotiation and decision making in multiagent environments. In *ECCAI Advanced Course on Artificial Intelligence*, pages 150–172.
- Kraus, S. (2001b). *Strategic negotiation in multiagent environments*. MIT press, Cambridge, Massachusetts, US.
- Kraus, S. and Lehmann, D. (1995). Designing and building a negotiating automated agent. *Computational Intelligence*, 11(1):132–171.
- Kreps, D. M. (1990). *Game theory and economic modelling*. Oxford University Press, Oxford, UK.
- Lai, K. R., Lin, M. W., and Yu, T. J. (2010). Learning opponent’s beliefs via fuzzy constraint-directed approach to make effective agent negotiation. *Applied Intelligence*, 33(2):232–246.
- Langley, P. (1996). *Elements of machine learning*. Morgan Kaufmann, San Francisco, US.

- Lau, R. Y., Li, Y., Song, D., and Kwok, R. C. W. (2008). Knowledge discovery for adaptive negotiation agents in e-marketplaces. *Decision Support Systems*, 45(2):310–323.
- Lau, R. Y., Tang, M., Wong, O., Milliner, S. W., and Chen, Y.-P. P. (2006). An evolutionary learning approach for adaptive negotiation agents. *International Journal of Intelligent Systems*, 21(1):41–72.
- Le Bars, M., Attonaty, J.-M., and Pinson, S. (2002). An agent-based simulation for water sharing between different users. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1*, pages 211–212.
- Lee, C.-C. (2014). Development and evaluation of the many-to-many supplier negotiation strategy. *Comput. Ind. Eng.*, 70:90–97.
- Lee, C. C. and Ou-Yang, C. (2009). A neural networks approach for forecasting the supplier’s bid prices in supplier selection negotiation process. *Expert Systems with Applications*, 36(2):2961–2970.
- Li, M., Vo, Q., Kowalczyk, R., Ossowski, S., and Kersten, G. (2013). Automated negotiation in open and distributed environments. *Expert Systems with Applications*, 40(15):6195–6212.
- Lin, R., Dor-Shifer, D., Kraus, S., and Sarne, D. (2006a). Local negotiation in cellular networks: From theory to practice. In *Proceedings of the National Conference on Artificial Intelligence*, volume 21, pages 1801–1807.
- Lin, R., Kraus, S., Wilkenfeld, J., and Barry, J. (2006b). An automated agent for bilateral negotiation with bounded rational agents with incomplete information. *Frontiers in Artificial Intelligence and Applications*, 141:270–274.
- Liu, J. S. and Chen, R. (1998). Sequential monte carlo methods for dynamic systems. *Journal of the American statistical association*, 93(443):1032–1044.

- Lomuscio, A. R., Wooldridge, M., and Jennings, N. R. (2001). A classification scheme for negotiation in electronic commerce. In Dignum, F. and Sierra, C., editors, *Agent Mediated Electronic Commerce: The European AgentLink Perspective*, pages 19–33. Springer, Berlin, Heidelberg.
- Lomuscio, A. R., Wooldridge, M., and Jennings, N. R. (2003). A classification scheme for negotiation in electronic commerce. *Group Decision and Negotiation*, 12(1):31–56.
- Lopes, F., Wooldridge, M., and Novais, A. Q. (2008). Negotiation among autonomous computational agents: principles, analysis and challenges. *Artificial Intelligence Review*, 29(1):1–44.
- MacKenzie, A. B. and DaSilva, L. A. (2006). Game theory for wireless engineers. *Synthesis Lectures on Communications*, 1(1):1–86.
- Malone, T. W., Fikes, R. E., and Howard, M. T. (1983). *Enterprise: A market-like task scheduler for distributed computing environments*. ChiZine Publications.
- Marashi, E. and Davis, J. P. (2006). An argumentation-based method for managing complex issues in design of infrastructural systems. *Reliability Engineering & System Safety*, 91(12):1535–1545.
- Marsa-Maestre, I., Lopez-Carmona, M. A., Ito, T., Zhang, M., Bai, Q., and Fujita, K. (2014). *Novel insights in agent-based complex automated negotiation*, volume 535. Springer.
- Masvroula, M. (2013). Forecasting negotiation counterpart’s offers: A focus on session-long learning agents. In *Proceedings of the fifth international conference on advanced cognitive technologies and applications (COGNITIVE 2013), Valencia, Spain*, pages 71–76.
- McBurney, P., Parsons, S., and Wooldridge, M. (2002). Desiderata for agent argumentation protocols. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1*, pages 402–409.

- Modi, P. J., Veloso, M., Smith, S. F., and Oh, J. (2005). Cmradar: A personal assistant agent for calendar management. In Bresciani, P., Giorgini, P., Henderson-Sellers, B., Low, G., and Winikoff, M., editors, *Agent-Oriented Information Systems II*, pages 169–181. Springer.
- Monteserin, A. and Amandi, A. (2013). A reinforcement learning approach to improve the argument selection effectiveness in argumentation-based negotiation. *Expert Systems with Applications*, 40(6):2182–2188.
- Niemann, C. and Lang, F. (2009). Assess your opponent: A bayesian process for preference observation in multi-attribute negotiations. In Ito, T., Zhang, M., Robu, V., Fatima, S., and Matsuo, T., editors, *Advances in Agent-Based Complex Automated Negotiations*, pages 119–137. Springer, Berlin, Heidelberg.
- Norman, T. J., Jennings, N. R., Faratin, P., and Mamdani, E. (1996). Designing and implementing a multi-agent architecture for business process management. In *Proceeding of International Workshop on Agent Theories, Architectures, and Languages*, pages 261–275.
- Okumura, M., Fujita, K., and Ito, T. (2013). An implementation of collective collaboration support system based on automated multi-agent negotiation. *Complex Automated Negotiations: Theories, Models, and Software Competitions*, pages 125–141.
- Oliver, J. R. (1996). On artificial agents for negotiation in electronic commerce. In *Proceedings of the Twenty-Ninth Hawaii International Conference on System Sciences*, volume 4, pages 337–346.
- Oprea, M. (2004). Applications of multi-agent systems. In Reis, R., editor, *Information Technology. IFIP International Federation for Information Processing*, volume 157, pages 239–270. Springer.
- Oshrat, Y., Lin, R., and Kraus, S. (2009). Facing the challenge of human-agent negotiations via

- effective general opponent modeling. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems*, volume 1, pages 377–384.
- Ostrom, E., Gardner, R., and Walker, J. (1994). *Rules, games, and common-pool resources*. University of Michigan Press, Michigan, US.
- Pahl-Wostl, C. (2007). The implications of complexity for integrated resources management. *Environmental Modelling & Software*, 22(5):561–569.
- Papaioannou, I., Roussaki, I., and Anagnostou, M. (2010). Using neural networks to minimize the duration of automated negotiation threads for hybrid opponents. *Journal of Circuits, Systems, and Computers*, 19(01):59–74.
- Park, S. and Yang, S.-B. (2006). An automated system based on incremental learning with applicability toward multilateral negotiations. In *Proceeding of SICE-ICASE International Joint Conference*, pages 6001–6006.
- Parsons, S. and Jennings, N. R. (1996). Negotiation through argumentation—a preliminary report. In *Proceedings of the 2nd international conference On multi agent systems*, pages 267–274.
- Parsons, S., Sierra, C., and Jennings, N. (1998). Agents that reason and negotiate by arguing. *Journal of Logic and computation*, 8(3):261–292.
- Paurobally, S., Turner, P. J., and Jennings, N. R. (2003). Automating negotiation for m-services. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 33(6):709–724.
- Pooyandeh, M. and Marceau, D. J. (2013). A spatial web/agent-based model to support stakeholders’ negotiation regarding land development. *Journal of environmental management*, 129:309–323.
- Pooyandeh, M. and Marceau, D. J. (2014). Incorporating bayesian learning in agent-based simulation of stakeholders’ negotiation. *Computers, Environment and Urban Systems*, 48:73–85.

- Prakken, H. and Vreeswijk, G. (2002). Logics for defeasible argumentation. In Gabbay, D. M. and Guenther, F., editors, *Handbook of Philosophical Logic*, pages 219–318. Springer, Dordrecht, Netherlands.
- Pruitt, D. G. and Carnevale, P. J. (1993). *Negotiation in social conflict*. Thomson Brooks/Cole Publishing Co., Pacific Grove, US.
- Radu, S., Kalisz, E., and Florea, A. M. (2013). Automatic negotiation with profiles and clustering of agents. *International Journal of Intelligence Science*, 3(02):69–76.
- Rahwan, I., Ramchurn, S. D., Jennings, N. R., McBurney, P., Parsons, S., and Sonenberg, L. (2003). Argumentation-based negotiation. *The Knowledge Engineering Review*, 18(4):343–375.
- Rahwan, I., Sonenberg, L., and McBurney, P. (2004). Bargaining and argument-based negotiation: Some preliminary comparisons. In *Proceeding of International Workshop on Argumentation in Multi-Agent Systems*, pages 176–191.
- Rahwan, I., Sonenberg, L., and McBurney, P. (2005). Bargaining and argument-based negotiation: Some preliminary comparisons. In *Proceedings of International Workshop on Argumentation in Multi-Agent Systems(ArgMAS)*, pages 176–191, Berlin, Heidelberg. Springer.
- Raiffa, H., Richardson, J., Metcalfe, D., et al. (2002). *Negotiation analysis: The science and art of collaborative decision making*. Harvard University Press.
- Rajavel, R. and Thangarathanam, M. (2016). Adaptive probabilistic behavioural learning system for the effective behavioural decision in cloud trading negotiation market. *Future Generation Computer Systems*, 58:29 – 41.
- Ramchurn, S. D., Sierra, C., Godo, L., and Jennings, N. R. (2006). Negotiating using rewards. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 400–407.

- Ramchurn, S. D., Sierra, C., Godo, L., and Jennings, N. R. (2007). Negotiating using rewards. *Artificial Intelligence*, 171(10-15):805–837.
- Reed, M. S. (2008). Stakeholder participation for environmental management: a literature review. *Biological conservation*, 141(10):2417–2431.
- Regan, H. M., Colyvan, M., and Markovchick-Nicholls, L. (2006). A formal model for consensus and negotiation in environmental management. *Journal of environmental management*, 80(2):167–176.
- Renna, P. (2010). Negotiation policies and coalition tools in e-marketplace environment. *Computers and Industrial Engineering*, 59(4):619–629.
- Resinas, M., Fernández, P., and Corchuelo, R. (2012). A bargaining-specific architecture for supporting automated service agreement negotiation systems. *Science of Computer Programming*, 77(1):4–28.
- Restificar, A. and Haddawy, P. (2004). Inferring implicit preferences from negotiation actions.
- Robu, V. and La Poutre, H. (2006). Retrieving the structure of utility graphs used in multi-item negotiations through collaborative filtering of aggregate buyer preferences. In *Proceedings of the 2nd international workshop on rational, robust and secure negotiations in MAS. Berlin*. Springer.
- Robu, V., Somefun, D., and La Poutre, J. A. (2005). Modeling complex multi-issue negotiations using utility graphs. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 280–287.
- Rodriguez-Fernandez, J., Pinto, T., Silva, F., Praca, I., Vale, Z., and Corchado, J. (2019). Context aware q-learning-based model for decision support in the negotiation of energy contracts. *International Journal of Electrical Power & Energy Systems*, 104:489 – 501.

- Ross, T. J. (2005). *Fuzzy logic with engineering applications*. John Wiley & Sons.
- Roussaki, I., Papaioannou, I., and Anagnostou, M. (2007). Building automated negotiation strategies enhanced by mlp and gr neural networks for opponent agent behaviour prognosis. In *International Work-Conference on Artificial Neural Networks*, pages 152–161. Springer.
- Rubin, D. B. (1988). Using the sir algorithm to simulate posterior distributions. *Bayesian statistics*, 3:395–402.
- Särkkä, S. (2013). *Bayesian Filtering and Smoothing*. Institute of Mathematical Statistics Textbooks. Cambridge University Press.
- SHEN, C., LIU, L., LUO, F., Lu, Y., and Lu, S. (2012). An adaptive market-driven agent based on multi-agent reinforcement learning for automated negotiation. *Inter-national Journal of Digital Content Technology and Its Applications*, 6(2):43–51.
- Shojaiemehr, B. and Rafsanjani, M. K. (2018). A supplier offer modification approach based on fuzzy systems for automated negotiation in e-commerce. *Information Systems Frontiers*, 20(1):143–160.
- Sierra, C., Jennings, N. R., Noriega, P., and Parsons, S. (1997). A framework for argumentation-based negotiation. In *International Workshop on Agent Theories, Architectures, and Languages*, pages 177–192. Springer.
- Simon, D. (2010). Kalman filtering with state constraints: a survey of linear and nonlinear algorithms. *IET Control Theory & Applications*, 4(8):1303–1318.
- Sycara, K. P. (1990). Negotiation planning: An AI approach. *European Journal of Operational Research*, 46(2):216–234.
- Tamani, N., Mosse, P., Croitoru, M., Buche, P., Guillard, V., Guillaume, C., and Gontard, N. (2015). An argumentation system for eco-efficient packaging material selection. *Computers and Electronics in Agriculture*, 113:174 – 192.

- Tetko, I. V., Livingstone, D. J., and Luik, A. I. (1995). Neural network studies. 1. comparison of overfitting and overtraining. *Journal of chemical information and computer sciences*, 35(5):826–833.
- Thomopoulos, R., Croitoru, M., and Tamani, N. (2015). Decision support for agri-food chains: A reverse engineering argumentation-based approach. *Ecological Informatics*, 26:182 – 191. Information and Decision Support Systems for Agriculture and Environment.
- Thoyer, S., Simon, L. K., Rausser, G. C., and Goodhue, R. E. (2008). Negotiating over the allocation of water resources: The strategic importance of bargaining structure. In *Game Theory and Policy Making in Natural Resources and the Environment*, pages 152–174. Routledge.
- Truong, T. D., Wiktor, L., and Boxall, P. C. (2015). Modeling non-compensatory preferences in environmental valuation. *Resource and Energy Economics*, 39:89–107.
- Tu, M. T., Wolff, E., and Lamersdorf, W. (2000). Genetic algorithms for automated negotiations: A fsm-based application approach. In *Database and Expert Systems Applications, 2000. Proceedings. 11th International Workshop on*, pages 1029–1033. IEEE.
- Van den Belt, M. (2004). *Mediated modeling: a system dynamics approach to environmental consensus building*. Island press.
- Van der Merwe, R., Doucet, A., de Freitas, N., and Wan, E. (2000). The unscented particle filter technical report cued. Technical report, F-INFENG/TR 380 (Cambridge: Cambridge University Engineering Department).
- Van Der Putten, S., Robu, V., La Poutré, H., Jorritsma, A., and Gal, M. (2006). Automating supply chain negotiations using autonomous agents: a case study in transportation logistics. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 1506–1513. ACM.

- van Galen Last, N. (2012). *Agent Smith: Opponent Model Estimation in Bilateral Multi-issue Negotiation*, pages 167–174. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Van Kleef, G. A., De Dreu, C. K., and Manstead, A. S. (2006). Supplication and appeasement in conflict and negotiation: The interpersonal effects of disappointment, worry, guilt, and regret. *Journal of personality and social psychology*, 91(1):124.
- van Krimpen, T., Looije, D., and Hajizadeh, S. (2013). *HardHeaded*, pages 223–227. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Vetschera, R. (2009). Learning about preferences in electronic negotiations—a volume-based measurement method. *European Journal of Operational Research*, 194(2):452–463.
- Von Neumann, J. and Morgenstern, O. (1945). *Theory of games and economic behavior*. Princeton University Press Princeton, US.
- von Neumann, J., Morgenstern, O., and Rubinstein, A. (1944). *Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition)*. Princeton University Press.
- Wang, M., Wang, H., Vogel, D., Kumar, K., and Chiu, D. K. (2009). Agent-based negotiation and decision making for dynamic supply chain formation. *Engineering Applications of Artificial Intelligence*, 22(7):1046–1055.
- Wolski, R., Brevik, J., Plank, J. S., and Bryan, T. (2003). Grid resource allocation and control using computational economies. *Grid computing: making the global infrastructure a reality*, 772.
- Wooldridge, M. (1999). Multiagent systems. chapter Intelligent Agents, pages 27–77. MIT Press, Cambridge, MA, USA.
- Wooldridge, M. (2001). Intelligent agents: The key concepts. *Multi-Agent Systems and Applications*, 2322:3–43.

- Zarchan, P. and Musoff, H. (2013). *Fundamentals of Kalman filtering: a practical approach*. American Institute of Aeronautics and Astronautics.
- Zeng, D. and Sycara, K. (1996). How can an agent learn to negotiate? In *International Workshop on Agent Theories, Architectures, and Languages*, pages 233–244. Springer.
- Zeng, D. and Sycara, K. (1998). Bayesian learning in negotiation. *International Journal of Human-Computer Studies*, 48(1):125–141.
- Zhang, G., Jiang, G.-R., and Huang, T.-Y. (2010). Design of argumentation-based multi-agent negotiation system oriented to e-commerce. In *Proceedings of International Conference on Internet Technology and Applications*, pages 1–6. IEEE.
- Zhang, G., Sun, H., and Jiang, G. (2012). Adding argument into multi-agent negotiation in the context of e-commerce. In *IEEE Symposium on Robotics and Applications (ISRA)*, pages 517–520. IEEE.