

ASSOCIATIVE AND FREQUENTIST OPPONENT MODELING APPROACHES IN AUTOMATED BILATERAL NEGOTIATIONS

A Thesis

by

Okan Tunalı

Submitted to the
Graduate School of Sciences and Engineering
In Partial Fulfillment of the Requirements for
the Degree of

Master of Computer Engineering

in the
Department of Computer Engineering

Özyeğin University
September 2021

Copyright © 2021 by Okan Tunalı

ASSOCIATIVE AND FREQUENTIST OPPONENT MODELING APPROACHES IN AUTOMATED BILATERAL NEGOTIATIONS

Approved by:

Assistant Professor Reyhan Aydoğan,
Advisor
Department of Computer Engineering
Özyeğin University

Assistant Professor İsmail Arı
Department of Computer Science
Özyeğin University

Associate Professor Ömer Korçak
Department of Computer Engineering
Marmara University

Date Approved: 24 August 2021



To my dearest family

ABSTRACT

This thesis mainly focuses on the problem of learning opponent's preferences during the negotiation in bilateral automated negotiation in which agents negotiate with each other to reach an agreement. Accordingly, it addresses the problems with the classical frequentist approach and advances the state-of-the-art in opponent modeling in automated negotiation by introducing a novel frequency opponent modeling mechanism, which updates some of the assumptions introduced by classical frequency approaches. Moreover, this thesis also proposes adopting association rule mining techniques to learn the opponent's preferences in bilateral negotiation. An extensive evaluation of those proposed approaches show that the proposed approaches outperform the classical frequency model.

In addition, this thesis argues that while optimizing one's utility function is essential, agents in a society should not ignore the opponent's utility in the final agreement to improve the agent's long-term interests in the system. It aims to show whether or not it is possible to design a social agent (i.e., one that aims to optimize both sides' utility functions) while performing efficiently in an agent society. Accordingly, we propose a social agent supported by a portfolio of strategies, a novel tit-for-tat concession mechanism, and a frequency-based opponent modeling mechanism capable of adapting its behavior according to the opponent's behavior and the state of the negotiation. The results show that the proposed social agent does not only maximize social metrics such as the distance to the Nash bargaining point or the Kalai point but also is shown to be a pure and mixed equilibrium strategy in some realistic agent societies.

ÖZETÇE

Bu tez esas olarak, etmenlerin bir anlaşmaya varmak için birbirleriyle pazarlık yaptığı otomatikleştirilmiş ikili pazarlık sırasında rakibin tercihlerini öğrenme sorununa odaklanmaktadır. Buna göre, klasik frekansa dayalı yaklaşımla ilgili sorunları ele alır ve klasik frekans yaklaşımlarının getirdiği bazı varsayımların güncellendiği yeni bir frekansa dayalı rakip modelleme mekanizması sunarak, otomatik pazarlıkta alanında rakip modellemede literatürdeki güncel çalışmaları ileri taşır. Ayrıca bu tez, ikili müzakerede rakibin tercihlerini öğrenmek için birliktelik kuralları madenciliği tekniklerinin uygulanmasını önermektedir. Önerilen bu yaklaşım ların kapsamlı bir değerlendirmesi, önerilen yaklaşımların klasik frekans modelinden daha iyi performans gösterdiğini göstermektedir.

Buna ek olarak, bu tez aynı zamanda kişinin fayda fonksiyonunu optimize etmenin gerekli olduğunu, ancak bir toplumdaki etmenlerin sistemdeki uzun vadeli çıkarlarını iyileştirmek için nihai anlaşmada rakibin faydasını göz ardı etmemesi gerektiğini savunur. Bir etmen toplumunda verimli bir şekilde performans gösteren (yani her iki tarafın fayda fonksiyonlarını optimize etmeyi amaçlayan) bir sosyal etmen tasarlamının mümkün olup olmadığını göstermeyi amaçlar. Buna göre, bir strateji portföyü, yeni bir kısasa kısas ödün mekanizması ve davranışını rakibin davranışına ve müzakerenin durumuna göre uyarlayabilen frekansa dayalı bir rakip modelleme mekanizması tarafından desteklenen bir sosyal etmen önerilmektedir. Sonuçlar, önerilen sosyal etmenin yalnızca Nash pazarlık noktasına veya Kalai noktasına olan uzaklık gibi sosyal ölçütleri maksimize etmediğini, aynı zamanda bazı gerçekçi sosyal etmen toplumlarında saf ve karma bir denge stratejisi olduğu da gösterilmiştir.

ACKNOWLEDGEMENTS

”A creative man is motivated by the desire to achieve, not by the desire to beat others.”

– Ayn Rand

First of all, I would like to state my gratitude to my advisor Assistant Professor Reyhan Aydoğan and Assistant Professor Victor Sanchez-Anguix for their limitless support, time, guidance, and motivation whenever I need. I would also like to thank all members of Interactive Intelligent Systems for sharing their valuable ideas and resources. I am deeply grateful to Assistant Professor İsmail Arı and Associate Professor Ömer Korçak for participating in my thesis jury and giving me feedback. Last but not least, I would love to thank to my dearest family. It would be impossible to complete without their unconditional support over those years.

TABLE OF CONTENTS

DEDICATION	iii
ABSTRACT	iv
ÖZETÇE	v
ACKNOWLEDGEMENTS	vi
LIST OF TABLES	x
LIST OF FIGURES	xii
LIST OF SYMBOLS	xiii
I INTRODUCTION	1
II BACKGROUND	8
2.1 Automated Negotiation	8
2.2 Association Rule Mining	11
III DISTRIBUTION-BASED OPPONENT MODELING	13
3.1 Value function estimation	13
3.2 Issue weight estimation	14
3.3 Evaluation of the Distribution-Based Frequency Approach	17
3.3.1 Experimental Design	18
3.3.2 Results	23
IV SOCIAL NEGOTIATION STRATEGY	26
4.1 Concession Mechanism	28
4.1.1 Bid Windows	29
4.1.2 Target utility update	30
4.2 Bidding Mechanism	32
4.2.1 Strictly Exploring & Exploratory Tit-For-Tat	33
4.2.2 Nash optimizer	34
4.2.3 Similarity-Based TFT	34
4.2.4 Eager	35
4.3 Acceptance Mechanism	35

4.3.1	Strictly Exploring	35
4.3.2	Exploratory TFT & Similarity-Based TFT	35
4.3.3	Nash Optimizer	35
4.3.4	Eager	36
V	EVALUATION OF NEGOTIATION STRATEGY	37
5.1	General experimental setting	38
5.1.1	Quality metrics	38
5.1.2	Negotiation domains	39
5.1.3	Agents	40
5.1.4	Social agent configuration	43
5.2	Analyzing the impact of the window size	44
5.3	Performance in a stationary agent society	45
5.4	Performance in a non-stationary and utilitarian agent society . . .	48
VI	ASSOCIATIVE OPPONENT MODELING	51
6.1	Adapting Association Rule Mining for Negotiation	51
6.2	Preference Modeling Based on Association Mining	54
VII	EVALUATION OF OPPONENT MODELING APPROACHES	57
7.1	General Experimental Settings	57
7.1.1	Quality Metrics	57
7.1.2	Negotiation Domains	57
7.1.3	Negotiation Agents	59
7.1.4	Negotiation Tournament Setup	64
7.2	Experimental Results	65
7.2.1	Final Round Based Results	66
7.2.2	Round Based Results	67
VIII	RELATED WORK	76
8.1	Negotiation Strategies	76
8.2	Opponent Modeling Approaches	81
IX	CONCLUSION	84
APPENDIX A	— ADDITIONAL RESULTS	85

REFERENCES	87
VITA	95



LIST OF TABLES

1	Issue values' smooth frequency calculation given an example window of bids.	16
2	The domains chosen for testing our opponent modeling mechanism .	19
3	Domain aggregated results for the given metrics.	23
4	Opponent aggregated results for the given metrics.	23
5	General configuration for our social agent during the experiments .	44
6	Effect of window size on the given metrics.	45
7	Agent aggregated results in the specified domains.	46
8	Payoff matrix for negotiation.	49
9	Mixed Nash equilibrium for the dynamic agent society.	50
10	Negotiation Domains	60
11	Negotiation Agents in Groups	63
12	Last round utility-space prediction based aggregated medians for Agent Group 1 Tournament with 100 rounds deadline.	66
13	Last round utility-space prediction based aggregated medians for Agent Group 2 Tournament with 100 rounds deadline.	67
14	Final round utility-space prediction based <i>domain</i> aggregated medians from Agent Group 1 Tournament with 100 rounds deadline. .	68
15	Final round utility-space prediction based <i>domain</i> aggregated medians from Agent Group 2 Tournament with 100 rounds deadline. .	68
16	Final round utility-space prediction based <i>domain opposition level</i> aggregated medians from Agent Group 1 Tournament with 100 rounds deadline.	68
17	Final round utility-space prediction based <i>domain opposition level</i> aggregated medians from Agent Group 2 Tournament with 100 rounds deadline.	69
18	Utility-space prediction based <i>number of received bids</i> aggregated medians from Agent Group 1 Tournament with 100 rounds deadline.	70
19	Utility-space prediction based <i>number of received bids</i> aggregated medians from Agent Group 2 Tournament with 100 rounds deadline.	71
20	Utility-space prediction based <i>number of received bids and domain opposition level</i> aggregated medians from Agent Group 1 Tournament with 100 rounds deadline.	85

21	Utility-space prediction based <i>number of received bids and domain opposition level</i> aggregated medians from Agent Group 2 Tournament with 100 rounds deadline.	86
----	--	----



LIST OF FIGURES

1	The issue weight update for various levels of β	18
2	Finite state machine representing the behavior of the proposed social agent	28
3	The Concept of Bid Window. In the figure, the reader can observe 4 consecutive windows of k offers	30
4	The bid windows that split the negotiation history into partitions of K bids	30
5	An example of the <i>top median</i> mechanism for $z = 80$	33
6	Experimental Setup Diagram	65
7	Agent Group 1, Pearson correlations over the <i>rounds</i> for the selected domains.	72
8	Agent Group 1, Pearson correlations over the <i>unique bids</i> for the selected domains.	73
9	Agent Group 2, Pearson correlations over the <i>rounds</i> for the selected domains.	74
10	Agent Group 2, Pearson correlations over the <i>unique bids</i> for the selected domains.	75

LIST OF SYMBOLS

$\Delta(t)$	Issue weight update parameter
γ	Exponential filter
$\mathcal{AT}$	Negotiation issues
$\mathcal{D}$	Domain values
$\mathcal{L}$	A set transactions including item L
$\mathcal{O}$	A list of offers
$\mathcal{O}_t$	A bid history until time t such that
$\mathcal{T}$	The set of all transactions
$\mathcal{U}(o)$	Agents' preferences are represented by means of linear additive utility functions
$\mathcal{W}$	Issue weights for opponent model
$\mathbf{L}$	The number of occurrences of item I
$\mathbf{N}$	The number of transactions
$\vec{\mathcal{F}}_i$	Frequency distribution of the issue values in the window
$E[\mathcal{U}_i(\mathcal{O})]$	Expected utility obtained in issue i for the current window
$Fr_i(j, \mathcal{O})$	Occurrence frequencies of each of the issue i 's negotiation values j
H^t	Bid history received from an opponent agent
o	A negotiation offer
o_i	represents the value for issue i in offer o
S	Partial bid which is a subset of value assignments

$V_i(.)$	The valuation function for issue i
w_i	Represents the importance of the negotiation issue i
W_r	Recent complete bid window



CHAPTER I

INTRODUCTION

Automated negotiation is an iterative and distributed search process between multiple intelligent agents exchanging offers, with the goal of finding a mutual agreement that allows for cooperation between the different parties [1, 2, 3, 4, 5]. In the last few years, we have seen an increasing interest on the study of agreement technologies [6]. This increasing interest goes hand in hand with an incipient acceptance of autonomy and delegation in technology, with some technologies such as self-driven cars [7] being the prime example of this trend. As delegation and autonomous systems become the norm, so will agreement technologies. The reason is simple: autonomous agents are driven by real users' preferences, and, as we all know, conflict is inherent in our world. As a consequence, we need technologies that allow autonomous agents to solve preferential conflicts and, hence, make delegation and autonomy as transparent for the user as possible. Automated negotiation [8, 9, 2] is considered as one of the core technologies in agreement technologies, as it provides autonomous entities with protocols and algorithms to reach agreements in a distributed way. The applications of automated negotiation range from electronic commerce [10, 11], coordination in robotics [12], energy markets [13], computer networks [14], video games [15], and even traffic control [16].

Despite this recent and increasing interest in automated negotiation, research has been carried out for decades. Researchers have proposed a number of negotiation protocols [17, 18] and negotiation strategies that guide autonomous agents on how to act in a distributed negotiation process [19, 20, 21, 22]. There are two main families of strategies in automated negotiation process: game theoretic and heuristic approaches. On the one hand, the former focuses on achieving optimal

negotiation results under the assumption of full rationality, unbounded computational resources, and, often, full disclosure of preferences. On the other hand, heuristic approaches assume that agents' resources are limited and partial or nil knowledge about the others' preferences, precluding agents from guaranteeing optimal results. This present work is categorized as a heuristic approach.

While optimal negotiation outcomes cannot be guaranteed, it is still crucial for agents to reach outcomes that are as close as possible to the optimal outcomes. There are several ways that agents can resort to optimizing the resulting negotiation outcomes, but perhaps opponent modeling is one of the most important mechanisms. Opponent modeling [23] allows us to build an approximate model of the opponents' preferences, which can be used to propose outcomes that result in win-win situations for involved parties. Hence, making outcomes more appealing and maximizing the odds of reaching an agreement.

One of the most popular opponent modeling mechanisms in automated negotiation is the frequency model [24]. The frequency model aims to build a model of the opponents' preferences assuming linear additive utility functions and steady concession towards lower utilities. For that, the frequency model uses the frequency of negotiation issue values as an indicator of both negotiation issue and value importance. Due to its simplicity and wide acceptance, the frequency model has been used in a myriad of scenarios [25, 24, 26, 27, 28]

Despite its popularity, there is no informed research on the robustness of the frequency model in a wide variety of scenarios. Most of the work in this field focuses only on the quality of the agreements and/or percentage of successful negotiations when the given negotiation strategy uses this model as opponent modeling. However, there are a number of factors having a significant impact on the negotiation outcome such as bidding strategy, acceptance strategy, how the estimated opponent's preference model is used in the underlying negotiation strategy, and so on. Therefore, gaining high utility agreements does not indicate by itself how good the opponent model is. Accordingly, this thesis analyzes how

well the frequency model predicts opponent's preferences elaborately by comparing the estimated opponent model with the real preferences.

Furthermore, we are inspired by the association rule mining and present an association based opponent modeling approach in the field of automated negotiation. Here, we take the bids as the sample transactions and use their co-occurrences to infer opponent agent's preferences. The application resembles to learning customer behavior by analyzing items in his/her purchase history. Using the opponent model as a recommendation engine, we aim to evaluate bids' relative importance for the opponent; to generate offers that are likely to be accepted in the automated negotiation. This thesis provides an elaborate analysis and evaluation of the proposed approaches by comparing them with the classical frequency-based opponent model.

Apart from the opponent modeling approaches, this thesis addresses the problem of bidding strategy as follows. Open multi-agent systems are systems where heterogeneous agents may enter and leave the system [29, 30]. Consequently, one may face a variety of agent behaviors when interacting with other agents in the system. This heterogeneity precludes agents from assuming a fully cooperative behavior by other agents, and it forces agents to do their best effort to try to achieve a good performance and guarantee a high utility. Therefore, in negotiation scenarios, agent designers have focused on developing the best performing agent among all competitors, that is, the agent that gets the best utility for itself [31], with less focus on the utility achieved by its opponent. While achieving a high utility for itself is an interesting property for an automated agent, it may not be the best strategy in an agent society [32], where agents may need to collaborate on multiple occasions. For instance, an agent may act selfishly and get the maximum utility for itself and exploit the opponent, ending up in an agreement that may portray a low utility for its opponent. In that case, what would it be the incentive for the opponent to collaborate with us? One should take into consideration that, in an agent society, agents may share opinions about their counterparts. This may

make it difficult for selfish agents to collaborate with other agents in the future. As a result, this may harm the agent’s long-term interests.

Although an agent may need to achieve the best utility for itself, it cannot ignore the utility achieved by its opponent. Therefore, one may be tempted to think that the solution to this problem is for all agents to take a fully cooperative approach to negotiation. In the past, there have been proposals of negotiation models and strategies that assume cooperative behaviors during the negotiation. In many of these proposals, the authors either assume that both parties play the same cooperative strategy, or they assume that the other party also employs a cooperative strategy [33, 34, 35, 36]. While this is a feasible assumption in some domains, it is not appropriate for open environments where agents may display a plethora of behaviors and strategies. Previous works have shown how social or cooperative agents may be exploited by competitive agents [37, 31, 38]. Ideally, a good balance should be struck between both, taking a social attitude and preserving one’s own utility. In open environments, one may negotiate with a wide range of strategies, and a social agent should be prepared to attain mutually attractive utilities independently of the other agents’ attitudes. The hypothesis that we aim to answer in this article is whether or not we can design a social agent that performs well against a wide range of opponents in an open and dynamic society. Several difficulties hinder the design of such an agent, as we discuss in the following paragraphs.

Firstly, agents should reach a consensus before their deadline, and that typically entails some type of reciprocal concessions. Constant and timely concessions regardless of the opponent’s moves mostly end up in situations where the agent can be exploited, such as the case of ANAC2011 [39, 40], where the Hardheaded agent [41] won the competition by just waiting for its opponents’ concessions. When the opponent concedes, a social agent should reciprocate to some extent to build a good long-term relationship and encourage further cooperation. In contrast, if the opponent’s strategy is designed in a way that it does not respond to a nice

move, it may not be beneficial to sustain the negotiation or make any further concessions. In those cases, it may be good to signal when we are not pleased with an opponent's attitude during the negotiation. For example, given a selfish opponent, a counter selfish move may help the opponent to notice that we are not inclined to concede against its selfish attitude [42]. Therefore, concessions should be carried out attending to one's own and the opponent's moves.

Secondly, achieving a social agreement, one that is appealing for both parties, requires understanding the other side's preferences/needs; thus, the agent needs to follow up its opponent's moves and accordingly make mutually beneficial offers. It has been shown that learning an opponent's preferences in negotiation is not trivial by any means [43, 3, 44]. For instance, non-exploratory behaviors, such as an opponent constantly making the same bids, prevents the agent from learning its opponent's preferences (especially in large-sized domains). The repetition results in overestimated preferences for repeated bids while underestimating preferences for less repeated or non-explored bids. Even when the opponent concedes, it should be taken into consideration that a concession at the start of the negotiation may not imply the same about the opponent's preferences as a concession towards the end. That is, while one may not concede on important issues at the start of the negotiation, this may prove false as the negotiation advances. As the opponent explores the bid space, one may be able to utilize the learned information to make better estimations, but a social agent should also be prepared for situations where the exploration of the bid space carried out by the opponent is scarce.

In this thesis, we also aim to find whether or not it is possible to design a social agent capable of negotiating efficiently in a competitive agent society. That is, while the agent should attempt to ensure an appropriate utility for itself, it should also strive to find a good agreement for the opponent. This should improve the long-term perspectives of the agent in a dynamic agent society. More specifically, we aim to answer the following research questions:

- Q1. *Can a social agent obtain a high utility in heterogeneous environments?*

- Q2. *Can a social agent achieve socially efficient agreements in heterogeneous environments?*
- Q3. *Can social negotiation strategies be an efficient choice for self-interested agents in agent societies?*

With these questions in mind, we have designed an adaptive social strategy that considers both the opponent’s moves and the context of the negotiation to adapt its social behavior accordingly. More specifically, our agent employs a portfolio of strategies supported by a new variant of the Tit-For-Tat strategy [45] that is based on the robust analysis of windows of opponent’s bids and smooth transitions between target utility updates, a new frequency-based opponent modeling mechanism that aims to be robust against some of the aforementioned problems, and an adaptive behavior strategy that depends on the time pressure, the opponent’s moves, and the information available about the opponent during the negotiation. The assumptions made by the proposed negotiation strategy are the following: (i) it assumes that agents’ preferences are portrayed by linear additive utility functions; (ii) it assumes no previous knowledge or model about the opponents’ preferences during the negotiation, with all the learning taking place from the information collected during the negotiation; (iii) it assumes bounded rationality by the agents and, therefore, agents employ heuristics rather than game-theoretic mechanisms to make decisions during the negotiation; (iv) it assumes no particular behavior for the opponents; Thus, this negotiation strategy is designed to work on heterogeneous and imperfect knowledge negotiations where no prior information is available about the opponent, and where preferences are represented through linear additive utility functions

Considering our social agent proposal, the contributions of this thesis are twofold. Firstly, we pose the problems faced by the frequency model in realistic scenarios and propose a new opponent modeling approaches that deals with some of these problems and outperforms the classic frequency model mechanism. Secondly, we propose a social agent supported by a portfolio of strategies, a novel

tit-for-tat concession mechanism, and a frequency-based opponent modeling mechanism capable of adapting its behavior according to the opponent's behavior and the state of the negotiation.

Rest of this thesis is structured as follows. In Chapter 2, we provide necessary background information. Chapter 3 covers our methodology of Distribution-Based Frequency Model and shows experimental results. Chapter 4 we explain our proposed Social Negotiation Strategy and in Chapter 5 we evaluate the strategy and show experimental results. Chapter 6 covers Associative Opponent Modeling approach and Chapter 7 consists of comparative experimental results for the proposed models. In Chapter 8, we present related studies with our work. Finally in Chapter 9, we conclude this thesis.

CHAPTER II

BACKGROUND

In this chapter, we provide the necessary background on automated negotiation such formalization, protocol and opponent modeling (Section 2.1) and association rule mining (Section 2.2).

2.1 Automated Negotiation

Automated negotiation is an iterative and distributed search process between multiple intelligent agents exchanging offers, with the goal of finding a mutual agreement that allows for cooperation between the different parties [1, 2, 3, 4, 5]. The applications of automated negotiation range from electronic commerce [10, 11], coordination in robotics [12], energy markets [13], computer networks [14], video games [15], and even traffic control [16]. Despite the large and growing body of literature about automated negotiation, there are still challenges and issues to be solved.

For the sake of simplicity we assume that two agents negotiate following the alternating offers protocol. Nevertheless, the model can also be extrapolated to other protocols, including multilateral scenarios. The agents negotiate in a time-bounded scenario where T delimits the end of the negotiation. If the deadline is reached without any agreement, the agents get their reservation utility.

The negotiation scenario consists of $\mathcal{AT} = \{1, 2, \dots, n\}$ negotiation issues whose domain values are represented by $\mathcal{D} = \{D_1, \dots, D_n\}$. An offer is represented by o , while $\mathcal{O}^*$ represents the set of all the offers in the negotiation domain. The agents' preferences are represented by means of linear additive utility functions in the form of $\mathcal{U}(o) = \sum_{i \in \mathcal{AT}} w_i \times V_i(o_i)$, where w_i represents the importance of the negotiation issue i , o_i represents the value for issue i in offer o , and $V_i(\cdot)$ is the valuation function for issue i , which returns the desirability of the issue value.

Without losing generality, it is assumed that $\sum_{i \in \mathcal{AT}} w_i = 1$ and the domain of $V_i(.)$ is $(0,1)$ for any i .

As mentioned in the previous sections, frequency models have been widely used as opponent modeling mechanisms in automated negotiation [25, 24, 26, 27, 28]. Apart from being reported as one of the most effective families of opponent modeling techniques, frequency approaches have the advantage of being simple and offering a good balance of time/exploration [28].

There have been multiple implementations and variations of frequency models in the automated negotiation literature, but perhaps the most popular implementation is the frequency model in HardHeaded agent [24]. The model was proposed with the following assumptions in mind: (1) the opponent steadily restricts the offers proposed to a possibly moving and decreasing utility range; (2) the opponent prefers to explore the negotiation space rather than repeating the same offer(s) over and over; (3) opponents tend to concede less on the most preferred issues, keeping them unchanged.

Briefly, this model works as follows. To estimate the weight of an issue value (e.g. Dell, HP, MAC for “laptop brand”), the frequency model computes how often each issue value appears in the opponent’s bids. The weight of the issue value is then normalized by the most repeated issue value. For instance, consider that Dell, HP and MAC appear 20 times, 10 times and 15 times respectively. In that case, the model estimates the issue value weights as $V(\text{Dell}) = \frac{20}{20}$, $V(\text{HP}) = \frac{10}{20}$ and $V(\text{MAC}) = \frac{15}{20}$. The frequency model analyzes how often the value of an issue changes. At the beginning, the assumption is that each issue has the same importance. For example, if we have four issues, the weight of each issue is set as 0.25. For each successive pair of offers made by the opponent, if the value of an issue did not change, then the model increases the weight of that issue.

While these assumptions may sound appropriate for some scenarios, the truth is that many state of the art agents do not fully comply with those assumptions. Firstly, agents may prefer to follow a more flexible concession and bidding strategy

that allows them to stochastically explore a wide portion of the negotiation space. Despite the fact that the general trend for the opponent is conceding, consecutive offers may not reflect this general trend due to the stochastic and flexible nature of agents. As a consequence, opponents may make a range of negotiation steps (e.g., concession, trade-off, unfortunate move, etc.), misleading the learning mechanism in the classic frequency model.

Secondly, another common behavior of the state-of-the-art agents is repeating the same set of offers for a long period of time. This is true for those agents that try to avoid exploitation by not leaking significant and full information about their utility functions. This goes against the classic frequency's assumptions and the original model is not ready for dealing with these cases. In fact, we experimentally found that, when the same bid is repeated for a significant number of rounds, the update and normalization rules in the classic frequency implementation presents a convergence problem: all the issue weights converge towards $\frac{1}{n}$, where n is the number of issues in the negotiation. Being hard headed and repeating the same offer does not mean that all of the issues are equally valued. Hence, other mechanisms are necessary to tackle these situations.

Last, but not least important, it is true that in the first negotiation rounds most agents do not tend to vary the value for those issues that are the most important. The reason for this is because most agents start by demanding the best offers for themselves, and these offers entail very little changes in the most important issues. However, as the negotiation proceeds, opponents may concede. At some point, it is possible to reach one's own aspirations by varying the values for the most important issues and maximizing less important issues. In fact, this behavior can be observed in many state of the art agents who trade-off issues to achieve one's own aspirations. Therefore, the assumption that opponents tend to concede less on the most preferred issues may hold for hardheaded agents or agents that do not steadily concede, but it may result fruitless in other scenarios.

2.2 Association Rule Mining

Association rule mining has become one of the most prominent issues in sales and marketing operations. Especially in e-commerce, retail and recently the fast-food sector, extracting valuable information, customer preferences, from sales transactions is an inevitable way of having higher revenue by inferring what a customer may also like based on previous purchases. A transactional dataset is mainly composed of the item and quantity features given a period of time. So, in a transactional dataset, each transaction can be defined as an itemset, list of items or an item bundle and the association analysis methods are commonly based on items' coexistence. In the literature, one may find various studies based on frequent itemset analysis for learning dependent item preferences of a customer or an agent; basically any source of transactional dataset, given a slice of history [46, 47, 48]. We should also note that association learning is an essential task in recommendation systems [49, 50], where item association can be utilized as context based similarity between items.

In Chapter 6, we elaborately explain and exemplify the usage of association rule learning idea in automated negotiation field where the bid values correspond to items and an offered bid matches with a sample market basket or an itemset proposed by the negotiating agent. We will also show the structural resemblance between negotiation history of an agents and a transactional datasets.

We mentioned that item association is commonly evaluated as an item coexistence frequency, which is turned into a tree search problem. The Apriori algorithm is one of the well-known approaches which applies breadth-first-search for association rule mining approach [46]. Here, we must note that, as it is explained in the opponent modeling chapter; our problem definition requires a tree traverse instead of searching for a specific itemset or *negotiation bid*. So, we will show that we use Apriori algorithm to calculate all possible element-wise associations; leading us to opponent preference learning. A brief introduction and necessary definitions for the Apriori algorithm is demonstrated in the following lines.

$$support(L) = \frac{\mathbf{L}}{\mathbf{N}} \quad \text{or} \quad support(L) = \frac{|\mathcal{L}|}{|\mathcal{T}|} \quad (1)$$

Here, $\mathbf{L}$ denotes the number of occurrences of item L and $\mathbf{N}$ is the number of transactions. Also, $\mathcal{L}$ is a set transactions including item L , and similarly, $\mathcal{T}$ is the set of all transactions where $|\mathcal{T}|$ is equal to $\mathbf{N}$. In multi-item case, *support* is calculated as a number of coexistence over the number of transactions:

$$support(L, R) = \frac{|\mathcal{L} \cap \mathcal{R}|}{|\mathcal{T}|} \quad (2)$$

Another metric *confidence*($L \rightarrow R$) is based on related item *support* ratio, which gives the purchasing likelihood of item R (right hand side) given the item L (left hand side) which is calculated as in (3).

$$confidence(L \rightarrow R) = \frac{support(L, R)}{support(L)} \quad (3)$$

The last metric *lift* can be described as an extension of *confidence*, since it additionally takes *support* of the right hand side into account in the calculation as (4, 5).

$$lift(L, R) = \frac{support(L, R)}{support(L) \cdot support(R)} \quad (4)$$

$$= \frac{confidence(L \rightarrow R)}{support(R)} \quad (5)$$

CHAPTER III

DISTRIBUTION-BASED OPPONENT MODELING

For modeling the preferences of the opponent, we have proposed a frequency-based opponent modeling mechanism, whose earlier version was presented in [44]. The model is developed for a time-bounded negotiation setting where each of the agents has a linear additive utility function, as described in Chapter 2. The opponent modeling mechanism proposed for our social agent covers the modeling and update mechanisms for both relative issue preferences and the issue value evaluations.

3.1 Value function estimation

The value function estimation aims to estimate the $V_i(\cdot)$ function for each of the negotiation issues. It assumes that the most preferred issue values appear more frequently in offers than others. Hence, the frequent issue values are assumed to present more utility than those that are less frequent. The estimation of the issue value evaluation is made by a max-normalization function. The function estimates the relative importance of each of the issues' values compared to the most frequent value of that issue. Formally, the estimation is calculated by the following formula:

$$\hat{V}_i(j) = \frac{(1 + \sum_{o \in \mathcal{O}_{0 \rightarrow t}} \delta_i(j, o))^\gamma}{\max_{k \in \mathcal{AT}} (1 + \sum_{o \in \mathcal{O}_{0 \rightarrow t}} \delta_i(k, o))^\gamma} \quad (6)$$

where $\delta_i(j, o)$ is 1 if the value j is used for issue i in offer o and 0 otherwise.

We propose the use of Laplace smoothing to the value occurrences. The reason behind this approach is to initialize the issue values that have never been received. Besides, the ratio is filtered by an exponential filter with coefficient γ , having the range of $0 < \gamma \leq 1$. The proposed filter is one of the main differences of this value function estimation component with respect to the classic frequentist opponent modeling approaches. It controls the relative importance of repeatedly received issue values by suppressing their growth. When the γ constant is assigned to be

1, the estimation function turns into a classic frequency model with a Laplace smoothing.

3.2 Issue weight estimation

The proposed model seeks statistically significant changes in issue value distributions of the two consecutive, disjoint and time-sorted window of offers as shown in Figure 4. The underlying assumption is that when a concession is carried out, the distribution of issue values changes from window to window. Thus, if one detects changes between in the distribution of issue values sent by the opponent, one may detect concession moves.

Before going further, we explain the concept of issue value distribution in more detail. Equation 7 formulates the issue value distributions in a list of offers $\mathcal{O}$. Given a window of bids, the formula calculates the occurrence frequencies of each of the issue i 's negotiation values j , $Fr_i(j, \mathcal{O})$. The δ_i returns 1 if the value j of issue i exists in the offer o . The $|\mathcal{O}|$ denotes the length of the bid window. The formula applies Laplace smoothing to handle the erroneous outcomes of the statistical significance test in the Algorithm 1. Table 1 demonstrates how this calculation is done for a given example.

$$Fr_i(j, \mathcal{O}) = \frac{1 + \sum_{o \in \mathcal{O}} \delta_i(j, o)}{n + |\mathcal{O}|} \quad (7)$$

The proposed algorithm departs from classic frequentist opponent modeling approaches in two ways: it employs non-sliding windows rather than two consecutive offers to decide how to update weights, and it is based on statistical tests of significance to compare the issue value distribution in windows of bids. The proposed algorithm iterates over every single negotiation issue i and calculates the frequency distribution of the issue values in the previous window $\vec{\mathcal{F}}'_i$ and the frequency distribution of the issue values in the current window $\vec{\mathcal{F}}_i$ (lines 3-5). Then, a Chi-squared test is carried out with the null hypothesis being that both frequency distributions, $\vec{\mathcal{F}}_i$ and $\vec{\mathcal{F}}'_i$, are statistically equivalent (line 6). The main

Data: t : The current time in the negotiation, $\mathcal{O}'$: The previous partition of k offers, $\mathcal{O}$: The current partition of k offers, $\mathcal{O}_{0 \rightarrow t}$: All the offers received so far, $\mathcal{W}' = \{w'_1, \dots, w'_n\}$: The current weights for the opponent model

Result: $\mathcal{W} = \{w_1, \dots, w_n\}$: The new weights for the opponent model

```

1  $e \leftarrow \emptyset$ ;
2  $\text{concession} \leftarrow \text{False}$ ;
3 foreach  $i \in \mathcal{I}$  do
4    $\vec{F}'_i \leftarrow (Fr_i(1, \mathcal{O}'), \dots, Fr_i(n, \mathcal{O}'))$ ;
5    $\vec{F}_i \leftarrow (Fr_i(1, \mathcal{O}), \dots, Fr_i(n, \mathcal{O}))$ ;
6    $p_{val} \leftarrow \chi^2\text{-test}(\vec{F}_i = \vec{F}'_i)$ ;
7   if  $p_{val} > 0.05$  then
8      $e \leftarrow e \cup \{i\}$ ;
9   else
10     $\vec{\mathcal{V}}_i \leftarrow (\hat{V}_i(1), \dots, \hat{V}_i(n))$ ;
11     $E[\mathcal{U}_i(\mathcal{O}')] \leftarrow \vec{\mathcal{V}}_i \times \vec{F}'_i$ ;
12     $E[\mathcal{U}_i(\mathcal{O})] \leftarrow \vec{\mathcal{V}}_i \times \vec{F}_i$ ;
13    if  $E[\mathcal{U}_i(\mathcal{O})] < E[\mathcal{U}_i(\mathcal{O}')] \text{ then}$ 
14       $\text{concession} \leftarrow \text{True}$ ;
15    end
16  end
17 end
18 if  $|e| \neq n$  and  $\text{concession} = \text{True}$  then
19   foreach  $i \in e$  do
20      $w_i \leftarrow w'_i + \Delta(t)$ 
21   end
22 else
23    $\mathcal{W} \leftarrow \mathcal{W}'$ 
24 end

```

Algorithm 1: The issue weight update mechanism

Table 1: Issue values' smooth frequency calculation given an example window of bids.

Fr(j, $\mathcal{O}$)		Issue 1		Issue 2		
		Value1	Value2	Value1	Value2	Value3
W I N D O W	Bid 1	1	0	1	0	0
	Bid 2	1	0	0	1	0
	Bid 3	0	1	1	0	0
	Bid 4	1	0	1	0	0
	Bid 5	1	0	0	1	0
	Bid 6	1	0	0	1	0
	Bid 7	0	1	1	0	0
	Bid 8	1	0	1	0	0
	Bid 9	1	0	1	0	0
	Bid 10	0	1	0	1	0
Value Count		7	3	6	4	0
Padded Count		8	4	7	5	1
Smoothed Frequency		0.67	0.33	0.54	0.38	0.08

goal behind this test is checking whether or not the distribution of issue values for i has changed from the previous window of offers to the current one. This information will help us to determine if, overall, the opponent has changed the type of offers sent. In the case that the null hypothesis cannot be rejected, we add the issue i to the set of issues e whose distribution did not change from the previous to the current window (lines 7-8).

When the null hypothesis is rejected (line 9), it means that the frequency distribution for issue i has been different from the past to the current window. The question is in what direction the change points for that issue (e.g., concession, increase of utility). More specifically, inspired by classic frequency approaches, we are interested in checking if the opponent has conceded in the issue because then we can update the weights for those issues that remained the same. Again, the assumption is that opponents tend not to change the most important issues more often than less preferred issues. In order to estimate if the opponent has conceded in the issue, we employ the frequency distribution for issue i during the whole negotiation $\mathcal{V}_i$ as an approximation of the actual valuations, as specified in Equation 6. Then, the expected utility obtained in issue i for the previous window of opponent offers $E[\mathcal{U}_i(\mathcal{O}')] is calculated (line 10). The same procedure is applied to obtain the expected utility obtained in issue i for the current window $E[\mathcal{U}_i(\mathcal{O})]$$

(line 11). Then, both expected utilities are compared to assess if a concession has been carried out in the issue i (lines 12). If the value distribution of an issue does not change significantly over the consecutive windows and the expected utility of the previous window is higher than that of the current window, it is considered as a concession (lines 13).

We propose to take an aggressive strategy to detect overall concessions over two consecutive windows of opponent offers. We consider that if there is a concession as long as the opponent has conceded in any of the issues (line 14). In that case, we update the importance for those issues that stayed in the same frequency distribution by adding a parameter that is calculated by a time-dependent decay function denoted in Equation 8 where β tunes the speed of the decay (lines 15-16). We understand that there are other strategies to detect an overall concession, and we are currently exploring the performance of more conservative approaches and probabilistic approaches. If there is no concession in any of the issues, then issue weights stay the same (line 17).

$$\Delta(t) = \alpha \times (1 - t^\beta) \quad (8)$$

Figure 1 shows how the issue weight update parameter, Δ , is calculated for varying β values over negotiation time.

3.3 Evaluation of the Distribution-Based Frequency Approach

In this section we evaluate the performance of the proposed frequency mechanism, the distribution-based frequency model, and compare it with the classic frequency opponent modeling mechanism. The goal of this section is assessing whether or not the proposed strategy is capable of overcoming the shortcomings highlighted in the previous sections. First, we describe how the experiments were designed, and then we analyze the results gathered.




Figure 1: The issue weight update for various levels of β

3.3.1 Experimental Design

Given the fact that the goal of this paper is comparing the performance of two learning mechanisms (classic frequency model, and distribution-based frequency model), we decided to use the same bidding and concession strategies for both strategies. This setting allows us to study both learning mechanisms in fair and equal conditions. More specifically, we chose HardHeaded’s concession and bidding strategy. The rationale for selecting this strategy is twofold. First, as the agent employs a Boulware strategy [51], it guarantees that the agent will not rapidly end the negotiation. A quick and abrupt end of the negotiation would preclude learning mechanisms from being studied effectively, as they have not been exposed to sufficient bids. Second, the HardHeaded’s bidding strategy actively employs the opponent model to propose bids to the opponent. This is important, as many times opponent modeling will have an impact on the opponent’s actions and bidding steps. By taking this realistic setting, we are also able to observe whether or not one’s opponent model influences the negotiation towards actions that further improve one’s opponent model. For the sake of simplicity, we decided to employ the alternating offers protocol although the opponent modeling mechanism should

Table 2: The domains chosen for testing our opponent modeling mechanism

	Laptop	CypressI.	EngvsZim.	Grocery	Amsterdam	Camera	S.market	Travel
Size	Small	Small	Medium	Medium	Medium	Large	Large	Large
Conflict	Low	High	Medium	Low	Low	Low	High	Medium

be applicable to other settings as long as multiple offers are exchanged between parties over time. Accordingly, we employed bilateral negotiation domains to experimentally test the performance of our opponent modeling mechanism. More specifically, we decided to test our modeling mechanism under a wide range of domain characteristics. These characteristics include different domain sizes (i.e., number of possible outcomes) and different degrees of competition between agents, measured by the distance from the Kalai point to the complete satisfaction point (1,1) [23]. The list of domains can be found in Table 3.3.1. Discount factors were ignored (i.e., removed) in this experimental setting and they are regarded as a matter of future study.

Another important decision to take for the experimental setting was deciding on the opponent agents to negotiate with. There were some factors that influenced our decision in this matter. First of all, any agent with offline opponent modeling was discarded as these may introduce interdependences between the outcomes of different negotiations, leading to effects in our opponent model that may be the by-product of past negotiation interactions. Second, we wanted to expose our opponent modeling mechanism to a variety of concession behaviors and bidding strategies representing the state of the art in negotiation. For that, we employed the following opponent agents:

- **AgentK [52]:** This agent was the winning agent of the 2010 ANAC negotiation competition. It is a conceder agent whose concession speed is regulated by the average utility of all received bids and its standard deviation. In terms of bid proposal, it just selects any offer from above the current aspiration. Hence, pairs of consecutive offers may not present an obvious decreasing trend. This behavior is in conflict with assumption 1 in Section 2.1.

- **IAmHaggler2011** [53]: A negotiating agent that uses Gaussian processes to predict the future concession of its opponent, and then adjust its concession rate accordingly to get the most from the negotiation. The goal of this agent is that of optimizing one’s own utility while also trying to reduce the utility received by the opponent. Bids are only selected from a small range around the target utility. This aggressive stance usually results in the agent repeating the same offers. This is contradiction with assumption 2 in Section 2.1.
- **TheNegotiatorReloaded** [54]: This agent was the best performing agent in undiscounted domains for the ANAC 2012 agent competition. The agent divides the negotiation into non-sliding windows, similarly to our approach. For each window, the agent estimates the type of agent behavior that it is facing and adjusts its concession rate accordingly. The most similar bids to the current target utility are sent back to the opponent. This agent was selected for the experiment for similar reasons to IAmHaggler2011.
- **Boulware agent** [51]: A classic negotiation agent that adjusts its aspiration levels according to time, only conceding in the later stages of the negotiation. Bids close to the target utility are selected and sent to the opponent. This agent was included as an example of scenarios where none of the aforementioned assumptions are strictly violated.
- **Conceder agent** [51]: A classic negotiation agent that adjusts its aspiration levels according to time. However, concessions are carried out early on in the negotiation process. Bid are randomly selected from above the threshold defined by the concession strategy. Due to the rapid concessions at the start of the negotiation, assumption 3 in Section 2.1 may be invalid very quickly.

The platform that supported our experiments was Genius[55]. We compared our opponent modeling with the performance of the classic frequency model. For

that, both opponent models faced all of the opponents in every single domain, which included two preference profiles per domain. In order to capture stochastic variations in negotiations, each possible case was repeated a total of 20 times. This gives a total of 3200 negotiations¹.

In order to assess the quality of our opponent modeling mechanism, we employed the following quality metrics:

- **Pearson correlation of bids:** It aims to compare the estimated outcome space with the real outcome space. For that, the Pearson correlation of bids is calculated and averaged. This metric is employed due to the fact that it has a strong correlation with overall opponent modeling performance [28].
- **Difference in surface of Pareto frontiers:** Another metric that is employed to assess the overall performance of an opponent modeling mechanism is the absolute difference between the area under the real Pareto optimal curve and the area under the estimated Pareto optimal curve. The rationale behind this metric is that some claim that it is enough to accurately estimate the Pareto optimal frontier for successful negotiations with the opponent. Again, this metric was shown to have a strong correlation with overall opponent modeling performance [28].
- **Spearman rank correlation of the issue weights:** The previous two metrics offer an insight into how the opponent modeling performs overall. However, both our opponent modeling mechanism and the classic frequency model have two components: the weight and the issue value update mechanisms. Therefore we decided to include extra metrics to assess the performance of each of these individual components. This metric compares the rank correlation between the issue weights learned by an opponent modeling mechanism and the target issue weights. The value is between -1 and 1, with 1 being used for a perfect ranking of the issue weights and -1 for a

¹2 models $\times$ 5 opponents $\times$ 8 domains $\times$ 2 profiles $\times$ 20 repetitions

completely opposite ranking. The rationale for selecting a ranking metric is that a ranking of the issues is normally enough to trade-off [56, 57].

- **Weighted Root Mean Squared Error of the issue values:** Given an estimated model and the target model, this metric computes a weighted version of the root mean squared error (RMSE) per issue. Predicted issue values are compared against the target issue values and weighted according to the importance of the issue value. This can be observed in Equation 9

$$\text{WRMSE}(i) = \sqrt{\sum_{j \in D_i} w_j \times (\hat{V}_i(j) - V_i(j))^2} \quad (9)$$

In our case, the weights for issue values w_j were set to $\frac{V_i(j)}{\max_{k \in D_i} V_i(k)}$ so that more weight is given to those issue values that provide more utility to the opponent. Then, after each negotiation, the metric is averaged with all the issues. In this case we employed a metric that both captures ranking and value accuracy. Although a ranking of issues is enough for carrying out trade-offs, one needs to have an accurate estimation of values for successfully providing appealing offers to the opponent.

As for the parameters of our model, we set $\alpha = 10$ and $\beta = 5$ in Equation 8. This means that weight updates will have a greater magnitude at the start of the negotiation, and they will gradually be reduced as the negotiation finishes. This type of update is meant to avoid incorrect updates when the opponent starts changing the most important issues relatively soon in the negotiation. With regards to Equation 6, γ was set to 0.25 to slow the growth of value importance when the opponent tends to repeat the same offer repeatedly. These values were found as good in a previous experimental setup. However, no exhaustive search was carried out over them. Therefore, the performance depicted in these experiments should be considered as a lower bound for the best achievable performance with this opponent modeling mechanism.

Table 3: Domain aggregated results for the given metrics.

	Distribution-based Frequency Model				Frequency Model			
	Prs. B.	Par. Fr. D.	Spr. W.	WRMSE	Prs. B.	Par. Fr. D.	Spr. W.	WRMSE
EngZimb	0.91	0.003	0.35	0.053	0.80	0.007	0.32	0.080
Cypress	0.83	0.043	0.46	0.042	0.60	0.092	0.30	0.063
Travel	0.85	0.015	0.73	0.043	0.70	0.019	0.21	0.057
Amst.	0.91	0.004	0.61	0.044	0.86	0.015	0.31	0.072
Grocery	0.90	0.005	0.96	0.041	0.86	0.010	0.54	0.046
Laptop	0.87	0.006	0.84	0.094	0.89	0.006	0.59	0.10
Camera	0.89	0.006	0.77	0.035	0.86	0.005	0.57	0.041
S.Market	0.87	0.044	0.93	0.030	0.69	0.067	0.59	0.050

Table 4: Opponent aggregated results for the given metrics.

	Distribution-based Frequency Model				Frequency Model			
	Prs. B.	Par. Fr. Dis.	Spr. W.	WRMSE	Prs. B.	Par. Fr. Dis.	Spr. W.	WRMSE
AgentK	0.91	0.007	0.76	0.040	0.83	0.019	0.58	0.065
Haggler	0.93	0.005	0.78	0.035	0.88	0.014	0.32	0.036
TNR	0.79	0.024	0.63	0.079	0.74	0.018	0.50	0.092
Boulw.	0.91	0.007	0.70	0.033	0.78	0.017	0.35	0.072
Conc.	0.84	0.035	0.67	0.050	0.68	0.070	0.38	0.053

3.3.2 Results

In this section we analyze the performance of our opponent modeling mechanism with respect to the classic frequency model. As mentioned we employ four metrics: two that measure the overall quality of the model (i.e., Pearson correlation of bids - Prs. B., and the difference between the surfaces defined by the Pareto optimal frontiers - Par. Fr. D.) and two other metrics that assess the quality of the two individual components of the opponent modeling (i.e., Spearman rank correlation of the issue weights - Spr. W., and the weighted root mean squared error - WRMSE). Table 3 aggregates the results obtained by domain, while Table 4 aggregates the results obtained by opponent. Those results that are statistically better than their counterparts are highlighted with a bold font. For statistical significance, a one-tailed Mann-Whitney test was carried out with $\alpha = 0.05$.

First, we will analyze the results per domain. Both the Pearson correlation of bids, and the difference between Pareto optimal frontiers tend to indicate that our opponent modeling provides a more accurate and overall estimation of the opponent’s preferences. In the case of the Pearson correlation of bids, our approach is statistically better for all domains except for the Laptop domain, where

the classic frequency model obtains a better estimation. However, the difference between both metrics is small for that domain. Similarly, our model outperforms when it comes to estimating the Pareto optimal frontier in all domains except for the Laptop and Camera domain, where there is no difference between our and the classic frequency model. Both the Laptop and the Camera domain are two of the less competitive domains. This may suggest that our opponent modeling may not necessarily outperform the classic frequency model for domains with low conflict. However, further experiments will be needed to make that conclusion. In this very same table, we can also observe that both our weight and issue value estimation are statistically more accurate than the classic frequency model. This is supported by statistically better results in both the Spearman rank correlation of issue weights and the weighted root mean squared error for issue values. Only in the Laptop domain the issue value estimation is no better, but also no worse, than the classic frequency model.

We can observe very similar results if we focus on the results aggregated by opponent. Overall, our opponent model produces a statistically better and more accurate model of the opponent’s preferences (i.e., Pearson correlation of bids, and Pareto frontier estimation). Only in the case of The Negotiation Reloaded (TNR) the estimation of the Pareto optimal frontier is no better, but also no worse than the classic frequency model. Component by component, we can appreciate that our weight update mechanism is consistently more accurate at detecting the relative importance of issues (i.e., Spearman rank correlation of weights). We also tend to produce a better value estimation for issue values for most opponent agents (i.e., WRMSE). Only we produce a statistically equivalent estimation against IAmHaggler2011 and the Conceder agent.

Overall, it can be appreciated that our opponent modeling mechanism tends to produce more accurate models of the opponent’s preferences, regardless of opponent and domain. The differences tend to be more acute in the weight update

mechanisms (i.e., Spearman rank correlation of weights) than in the issue value update mechanism (i.e., WRMSE). This suggests that, most likely, an important part of our improvement is due to the weight update mechanism. Other issue value update mechanisms may be necessary to further improve the classic frequency model. Further exploring other issue value update mechanism is highlighted as future areas of improvement for our current opponent modeling.



CHAPTER IV

SOCIAL NEGOTIATION STRATEGY

Our social agent employs a range of negotiation behaviors. The agent shifts behaviors depending on the negotiation conditions and the information that is gradually available about the opponent's preferences. In this chapter, we first describe the general ideas behind the negotiation strategy of our social agent and then explain our concession strategy elaborately. After that, we describe how our agent switches bidding and acceptance behaviors during the negotiation.

The behavior of our social agent is determined by a finite state machine that specifies the negotiation behavior to be followed during the negotiation. Each state corresponds to a behavior adopted by the agent in particular circumstances, depending on some factors such as attitude of the opponent, remaining time, or agent's current state. To sum up, our agent's negotiation strategy is composed of five negotiation behaviors.

The representation of the finite state machine is shown in Figure 2, and the state transition rules can be explained as follows. In the beginning, our agent aims to get to know its opponent from the bid exchanges. The goal is obtaining knowledge to use in the rest of the negotiation. Therefore, the initial negotiation behavior is *Strictly Exploring*, which is employed until a strict exploration time threshold (t_{exp}) has been surpassed. After the strict exploration phase, the agent may have learned to some extent about its opponent's preferences, but it may be the case that the learned model is not reliable yet to accurately estimate the bids maximizing the social welfare. In that case, until receiving a certain proportion of unique bids (α_{exp}), the agent employs a *Exploratory Tit-For-Tat* behavior. Whenever the agent has received the specified proportion of unique bids, the agent's negotiation behavior turns to be *Nash Optimizer*. In that behavior, the estimated

Nash product bid is offered. If the negotiation time surpasses the time threshold (t_{sim}) and the proportion of unique bids has not been received, the *Similarity-Based Tit-For-Tat* behavior is followed. During this phase, the agent may transit to *Nash Optimizer* as it receives enough bids. The final behavior *Eager*, occurs only when the negotiation time approaches the deadline (t_{eager}).

The next list briefly describes each of the behaviors:

- *Strictly Exploring*: At the start of the negotiation, the agent aims to explore the negotiation space. Therefore, it is not inclined to accept the opponent's bids. The aspiration of the agent is controlled by a special window-based tit-for-tat mechanism that is described in Section 4.1.
- *Exploratory TFT*: After a certain time threshold t_{exp} in the negotiation, the agent transits to this state. Even though the agent still explores the negotiation space, the agent may accept incoming offers if they are above the current aspiration of the agent (see Section 4.1).
- *Nash optimizer*: As mentioned, our agent's orientation is social. Thus, the agent aims to find a fair deal for both parties, regardless of the opponent's behavior. The strategy of the proposed agent is finding a deal that is as close as possible to the Nash bargaining point. To achieve this, the agent needs to learn an accurate model of the opponent's preferences. Learning an accurate model about the opponent's preferences involves receiving several unique bids from the opponent. When the agent has received a proportion α_{exp} of the negotiation space, the agent considers the opponent model to be approximately accurate. Note that we adopt the distribution-based frequency model to estimate the opponent's preference profile (Chapter 3). Then, the agent is ready to change to the *Nash optimizer* behavior. In this state, the bidding and acceptance strategy try to be as close as possible to the estimated Nash point.
- *Similarity-based TFT*: Although our agent aims to be social, and it will try




Figure 2: Finite state machine representing the behavior of the proposed social agent

to get as close as possible to the Nash bargaining point, in some situations, the opponent may not reveal many unique bids. Then, in such a case, learning an accurate opponent model becomes hard. Consequently, the agent starts using a similarity-based heuristic when the negotiation is approaching a threshold (i.e., $t > t_{sim}$). The aspiration of the agent is again controlled by a special tit-for-tat concession mechanism (Section 4.1).

- *Eager*: In the last moments of the negotiation (i.e., $t > t_{eager}$), the agent tries to close an agreement by proposing back the best offer sent by the opponent or accepting any offer that improves or equalizes the best offer received so far.

The reader can find the specific details of our agent strategy in Sections 4.1, 4.2 and 4.3.

4.1 Concession Mechanism

Concession mechanisms generally determine a target utility for an agent in a given negotiation round. The concession mechanism is tightly linked to what offers are considered acceptable in a negotiation round, as well as the minimum utility demanded when bidding, both typically relying on such target utility. In this article, we propose a new and robust variant of the Tit-for-Tat tactic [37].

Tit-for-Tat tactics respond to the concession moves carried out by the opponent in a reciprocal way. In other words, when the opponent concedes, so does in a

certain magnitude the agent. If the opponent does not concede, neither does the agent. While basic and classic Tit-for-Tat considers the opponent’s current move (i.e., by comparing the last two offers made by the opponent), our new Tit-for-Tat strategy analyzes a non-sliding window of the opponent’s consecutive offers (e.g., windows formed by k offers). The primary motivation of considering ”windows of negotiation history” rather than the current move is to avoid mimicking the irregular fluctuations in utility changes that prevent us from understanding the opponent’s general concession trend. Next, we describe the non-sliding window mechanism.

4.1.1 Bid Windows

The proposed agent employs non-sliding windows of bids in many of its components. In general, the negotiation history is split into disjoint and successive windows of bids received from the opponent. These windows can contain at most k bids. The most recent window may contain less than k bids as it is filled with new offers received from the opponent. Both Figures 4 and 3 contain examples of this concept. On the one hand, Figure 3 is an example of a sequence of received bids by an opponent. The horizontal axis shows the negotiation rounds while the vertical axis denotes the received utility of the opponent’s bids at each round. Each bid window is shown by colored rectangles. One may think of windows as consecutive slices of received bids. In this figure, there are four completed windows of $k = 20$ offers and one that is currently being filled with opponent offers (the last one to the right). On the other hand, Figure 4 shows a conceptual example with $k = K$ offers and three complete bid windows, and one that is currently being filled at $t_{current}$.

In the new Tit-for-Tat concession mechanism, the agent considers the most recent complete window pairs: W_{r-1} and W_r (just filled), in the history of bid windows $\mathcal{W} = \{W_1, \dots, W_r\}$. If we consider an example with $k = 25$, the concession mechanism would work with two consecutive windows of bids of size $k = 25$. In this example one can formulate the current window of bids as $W_r = [o_{t-24}, \dots, o_t]$




Figure 3: The Concept of Bid Window. In the figure, the reader can observe 4 consecutive windows of k offers




Figure 4: The bid windows that split the negotiation history into partitions of K bids

where t is current time step and o_i denotes the offer sent by the opponent at round i . For instance, the agent would take into account Window 2 and Window 3 at $t_{current}$ for the given example in Figure 4. Only when the current window is filled with k bids, the concession mechanism is executed back again. Therefore, the concession and update of the target utility occur in window-sized mini-batches. After the update, the current window is added to the history of previous bid windows, and it is emptied to start a new bid window.

4.1.2 Target utility update

After explaining the concept of a bid window, let us formalize some parameters and data employed by the target utility update mechanism, which is carried out only when the concession mechanism is executed (i.e., after comparing two non-sliding windows as described in the previous subsection). Coupled with the non-sliding

window mechanism, the target utility update mechanism is what makes our proposed Tit-for-Tat strategy different from classic variants. Firstly, o_{best} , represents the best offer received from the opponent, considering our utility function. Secondly, c_{max} represents the maximum utility update at each concession step. This new parameter, which is not included in classic Tit-for-Tat variants avoids abrupt concessions and forces a more thorough exploration of the utility space. The details of our proposed target utility update mechanism can be found in Algorithm 2.

2. Next, let us explain how this mechanism works in detail.

Data: W_{recent} : The most recent bid window received from the opponent,
 $\mathcal{W}$: the previous bid windows received from the opponent, o_{best} :
the best offer received from the opponent

Result: $\mathcal{U}_{target}$: the current target utility for our agent.

```

1  $U_r = \{\mathcal{U}(o) \mid o \in \mathcal{W}_{recent}\};$ 
2  $P_{r,z} = percentile(U_r, z);$ 
3  $U_{best,r} = \{o \mid o \in U_r \wedge o \geq P_{r,z}\};$ 
4  $M_r = median(U_{best,r});$ 
5  $M_p = 0;$ 
6 foreach  $W_i \in \mathcal{W}$  do
7    $U_i = \{\mathcal{U}(o) \mid o \in \mathcal{W}_i\};$ 
8    $P_{i,z} = percentile(U_i, z);$ 
9    $U_{best,i} = \{o \mid o \in U_i \wedge o \geq P_{i,z}\};$ 
10   $M_i = median(U_{best,i});$ 
11  if  $M_i > M_p$  then
12     $M_p = M_i;$ 
13  end
14 end
15  $\Delta = min(c_{max}, M_r - M_p);$ 
16  $U_{target} = min(1, max(\mathcal{U}(o_{best}), U_{target} - \Delta));$ 

```

Algorithm 2: The target utility update mechanism

The first steps (lines 1-4) consist of estimating a central measure of utility received from the opponent in the current window of bids. For that, bids in a window first need to be transformed to a utility space. The agent uses a median utility calculation mechanism as a measure of centrality due to the abrupt and unintended changes in utilities of bids sent by the opponent. In fact, instead of considering the whole window of bids, we propose for the agent to just consider those offers whose utility are above a given percentile (i.e., $P_{r,z}$, z being the specific

percentile). We experimentally found that many windows of bids only contain a few reasonable offers for the agent even though the opponent may have conceded. In some scenarios, this is caused by inaccurate opponent models in the early and middle stages, and in others just due to the exploratory nature of the opponent. Therefore, we believe that concessions should only be detected with the top offers from the window of bids $U_{best,r}$, rather than just two consecutive offers or from the whole set of offers, as proposed by classic Tit-for-Tat variants. After filtering out those offers above the percentile in the windows of bids, its median M_r is calculated. From this point on, we refer to this median calculation on the top bids in a window as the *top median*. Figure 5 shows a bid window collected from a real negotiation and illustrates how calculating the median on the offers above a given percentile ($z = 80$, 20% top offers) may better represent concessions carried out by the opponent.

After that, the agent checks the previous window of bids to calculate the best *top median* received so far from previous windows (lines 5-12). Please, notice that this median calculation follows the same procedure as described in the paragraph above (i.e., median on the top offers from the window of bids). We consider that a concession has been carried out by the opponent if the *top median* in the current window is above the best *top median* from previous bid windows. The concession step is the difference between both *top medians* (line 13) and, if necessary, truncated to the maximum concession step c_{max} . Finally (line 14), the target utility is recalculated as the previous target utility minus the concession step calculated before. Then, the calculated target utility is compared with the utility of the best offer o_{best} received so far, and the higher value is set to be the current target utility.

4.2 Bidding Mechanism

While the concession and target utility update mechanism are the same regardless of the underlying behavior followed by our agent, the bidding behavior changes




Figure 5: An example of the *top median* mechanism for $z = 80$

accordingly. This section describes the bidding strategy followed by our agent depending on the underlying negotiation behavior. The general description of each behavior will not be repeated to avoid redundancy; the reader can find it the introductory part of Chapter 4.

4.2.1 Strictly Exploring & Exploratory Tit-For-Tat

Both the Strictly Exploring and the Exploratory Tit-For-Tat behaviors employ the same bidding strategy. This bidding strategy allows the agent to explore the negotiation space with the opponent. It should be highlighted, that, in this case, the target utility U_{target} is employed to filter which offers can be sent to the opponent. That is, in these strategies, offers sent to the opponent will be at least as good as the target utility. The agent sorts the outcome/offer space in descending order according to its utility function. This bidding mechanism iteratively sends offers in descending order from the best possible offer in the domain to the current target utility. Then, once all the offers in the interval have been sent at least once, the agent starts sending random offers from the utility

range. It should be highlighted that once the target utility is updated, the agent starts back sending new and unexplored offers in descending order until no more new offers are available, turning back again to the random bidding in the utility range.

4.2.2 Nash optimizer

In Nash Optimizer behavior, the agent believes that the opponent model gained significant information and it can be used to estimate what the opponent gets from any bid. In addition, this is the behavior which makes the agent social since it offers the bid that maximizes the product of the utility it gets and what the opponent is estimated to get. The selection of prospective bids is, again, regulated by the current target utility. More specifically, the agent always sends this offer given a current target utility:

$$o_{nash} = \underset{o \in \{o' \mid \mathcal{U}(o') \geq U_{target}\}}{\operatorname{argmax}} \quad \mathcal{U}(o) \times \bar{\mathcal{U}}_{op}(o) \quad (10)$$

where $\bar{\mathcal{U}}(.)$ is the estimated utility function of the opponent, estimated as described in Section 3, U_{target} and is the current target utility.

4.2.3 Similarity-Based TFT

The *Similarity-Based TFT* is employed when the agent abandons the exploratory behaviors and it considers that it does not have enough data to model the opponent correctly. In that case, a simple similarity mechanism is employed to propose offers to the opponent. Again, the set of available offers to be sent to the opponent is regulated by the current target utility U_{target} . From that set of offers, the agent selects the offer that maximizes the similarity to the best offer sent by the opponent so far o_{best} . Notably, the mechanism can be described as follows:

$$\underset{o \in \{o' \mid \mathcal{U}(o') \geq U_{target}\}}{\operatorname{argmax}} \quad \mathcal{S}(o, o_{best}) \quad (11)$$

where $\mathcal{S}(.)$ is a similarity function between 0 and 1 that indicates how similar two offers are. In the case of our experiments, as we worked with discrete issues,

this similarity was calculated as the percentage of common values in both offers. It should be highlighted that this strategy includes a small exploration mechanism that allows it to propose random offers in the current utility range. This exploration mechanism is activated with a probability p_{esc} that depends on the remaining negotiation time:

$$p_{esc} = 1 - \frac{t}{T} \quad (12)$$

where t is the current time, and T is the total negotiation time.

4.2.4 Eager

Since the negotiation is about to end, the agent offers the best offer received from the opponent o_{best} in an attempt to successfully close the negotiation.

4.3 Acceptance Mechanism

As it happened with the bidding mechanism, the behavior of the agent changes with regard to bidding acceptance. Next, we briefly describe each of these acceptance mechanisms.

4.3.1 Strictly Exploring

As the agent solely aims to explore the negotiation space and get to know the opponent, it rejects all the incoming offers from the opponent.

4.3.2 Exploratory TFT & Similarity-Based TFT

Unlike the *Strictly Exploring* behavior, in these behaviors, the agent accepts the offer when $\mathcal{U}(o) \geq U_{target}$. That is, when the utility received in the offer is at least as good as the current target utility.

4.3.3 Nash Optimizer

Considering that this behavior is active when the agent has explored the negotiation space and it can rely on its opponent model, the agent accepts an offer if it is at least as good as the next offer to be sent. That is, if the received offer is at least as good as the bid sent by the Nash bidding mechanism: $\mathcal{U}(o) \geq \mathcal{U}(o_{nash})$.

4.3.4 Eager

As it is mentioned in the bidding strategy, in this behavior, the negotiation time is about to end. So, the agent aims to reach an agreement on a bid that provides at least the highest utility it has received (i.e. $\mathcal{U}(o) \geq \mathcal{U}(o_{best})$).



CHAPTER V

EVALUATION OF NEGOTIATION STRATEGY

In this chapter, we assess the empirical performance of the proposed agent, and aim to answer the initial question raised at the Chapter 1: Can social agents efficiently perform in automated negotiation? To answer that general question, we propose three different research questions:

- Q1. *Can a social agent obtain a high utility in heterogeneous environments?*
- Q2. *Can a social agent achieve socially efficient agreements in heterogeneous environments?*
- Q3. *Can social negotiation strategies be an efficient choice for self-interested agents in agent societies?*

To answer these questions, we have designed three types of experiments. The first experiment serves to analyze the impact of the length of the window of bids on the performance of our social agent. The second experiment aims to analyze the behavior and performance of our social agent in a static society of agents. By a static society of agents, we mean a society where multiple types of agents exist, but they do not change their strategies. With this experiment we aim to find if our social agent can achieve socially efficient agreements in heterogeneous environments. Lastly, we carry out a theoretical analysis using game theory to find out about the prospective use of our social strategy in a society where agents can decide their negotiation strategies. This experiment aims to find if a social strategy such as the one proposed in this article may be efficient even for self-interested agents. The structure of this section is organized as follows. First, we describe the experimental settings common to all our experiments. Then, we describe and analyze the experiments carried out to identify an appropriate window size for the

proposed agents. After that, we provide the empirical results for the experiment carried out to determine the behavior and performance of our agent in a static agent society. Finally, we describe the insights arising from the theoretical analysis carried out in a dynamic setting.

5.1 *General experimental setting*

In this section, we describe the settings and design decisions common to all the experiments carried out to assess the performance of our social agent. To be more specific, we describe the metrics employed to evaluate the performance of the agents, as well as the agents that form the agent society in which our agent interacts.

5.1.1 **Quality metrics**

As mentioned throughout the article, the goal of the proposed agent is to be social and reaching an agreement that is as fair as feasible for both agents. For that reason, most of the metrics employed to assess the performance of the agent aim to measure the fairness/quality of the agreement from both agents' perspectives. Next, we describe each of the metrics that we use in our experimental setting:

- **Individual utility:** This is the utility received by the agent at the end of the negotiation. When the negotiation ends with no agreement, then the agent gets an individual utility of 0.0. Otherwise, the utility obtained by the agent is that of the individual utility reported by the agreement.
- **Nash bargaining point distance:** The Nash bargaining point is the Pareto optimal point that maximizes the product of both agents' utilities. Traditionally, as the product of the two utilities penalizes scenarios where one of the agents gets a very high utility while the other agent gets a low utility, the Nash bargaining point has been considered as an indicator of the fairness of the resulting agreement. Thus, in practice, the closer an agreement is to the Nash bargaining point, the fairer it should be for both

involved parties.

- **Kalai point distance:** The Kalai point also seeks to find an outcome that is fair for both parties. However, and differently from the Nash bargaining point, the Kalai point considers the initial reservation utilities of the agents (i.e., the status quo solution) when computing the fairness of the agreement. In essence, both parties should gain in the same proportion from their status-quo point. Thus, the Kalai point is the Pareto optimal solution that maximizes the ratio of the gains of both agents. In practice, the closer an agreement is to the Kalai point, the fairer it should be for both involved parties.
- **Negotiation time:** Although it is not related to the fairness of the agreement found by negotiating agents, the negotiation time helps to see the performance of the different strategies from a different perspective. The sooner agreements are closed, the more time that the spare time and computational power can be allocated to other relevant tasks. Therefore, other aspects being equal, one should prefer negotiation agents that close deals more rapidly.

The aforementioned metrics will be employed to analyze the empirical performance of our social agent and other agents employed in the experiments.

5.1.2 Negotiation domains

In these experiments, we have considered four different negotiation domains with a medium or large number of possible negotiation outcomes to reduce the odds of the Nash point or agreements close by mere chance. Next, we briefly describe the negotiation domains chosen for the experiments:

- **Itex vs Cypress (ANAC 2010, [58]):** This is a classic buyer-seller setting where a bike manufacturing company negotiates on the sale of a commodity with a bike retail shop. There are a total of 180 possible outcomes.

- **England vs Zimbabwe (ANAC 2010, [58]):** This is a negotiation domain consisting of the national negotiations between England and Zimbabwe on the application of the World’s Health Organization Framework on Tobacco Control. The negotiation domain consists of 576 possible negotiation outcomes and 5 negotiation issues.
- **Grocery domain (ANAC 2011, [59]):** This domain involves agents negotiating on the type of groceries to purchase from a supermarket. There are a total of 1,600 possible outcomes distributed across 5 negotiation issues.
- **Amsterdam Trip (ANAC 2011, [59]):** This is a negotiation domain where two parties of friends negotiate on the details of a trip to Amsterdam. The negotiation domain contains 3,024 possible outcomes distributed across 6 negotiation issues.
- **Party domain:** This is a negotiation party where two hosts negotiate on how to organize a party for their friends. The negotiation scenario contains 3,072 possible negotiation outcomes and 6 negotiation issues.
- **Camera domain (ANAC 2012, [60]):** The camera domain contains preference profiles that describe the purchase of a camera and its accessories for a specific budget. More specifically, the negotiation domain contains 3,600 prospective outcomes distributed across 6 negotiation issues.

5.1.3 Agents

Another critical decision to make for the experimental setting is deciding on the opponent agents to negotiate with. Some factors influenced our decision in this matter. First of all, our proposed agent is an agent that models the opponent from the data leaked in the current negotiation process, and we assume that no other data from the opponent is available. This is the situation faced in case of a cold start in an agent society, in case that the preferences of agents change from negotiation to negotiation, or in the case that negotiation partners rarely

interact more than once. Therefore, negotiating agents that learn from multiple negotiations are excluded from this experimental setting, as introducing these agents would result in an unfair comparison. Second, we want to evaluate the performance of our agent in an agent society that shows a range of behaviors and bidding strategies, as it would probably happen in a real agent society. For that reason, we employ the following state-of-the-art agents to represent a realistic agent society:

- **AgentK (K)** [52]: This is the winning agent of the 2010 ANAC negotiation competition. It is an agent whose concession speed is regulated by the average utility of all received bids and its standard deviation. In terms of offer proposal, it selects any offer from above the current aspiration. Therefore, it does not model the opponent’s preferences on offers.
- **IAmHaggler2011 (IH2011)** [53]: A negotiating agent that uses Gaussian processes to predict the future concession of its opponent, and then adjust its concession rate accordingly to get the most from the negotiation. The goal of this agent is that of optimizing one’s utility while also trying to get more utility than the opponent. Bids are only selected from a small range around the target utility. This aggressive stance usually results in the agent repeating the same offers. This agent placed 3rd in the ANAC 2011 competition.
- **Gahboninho (GAH)** [61]: This agent finished second in the ANAC 2011 competition. Gahboninho employs an adaptive concession strategy that attempts to detect if agents are susceptible to being pressured or if they are stubborn. For that, Gahboninho limits most of the outcomes sent to outcomes that report a low utility value for the opponent.
- **Hardheaded (HH) agent** [41]: The winning agent of the 2011 ANAC competition. This agent follows a time-based bidding strategy that consists

of not conceding until almost the end of the negotiation. While not conceding, the agent performs opponent modeling on the opponent’s preferences by using a frequency modeling mechanism. Then, in the last states of the negotiation, the agent attempts to find an agreement by using the trained model on the opponent’s preferences and finding outcomes that have the same target utility for one’s own but maximize the preferences of the opponent.

- **IAmHaggler2012 (IH2012) [60]:** The 2012 version of IAmHaggler ranked as the most social agent in the ANAC 2012 competition (i.e., maximization of the sum of the opponent’s and one’s utility). This agent was employed to compare the performance of our social agent against some state-of-the-art social agents.
- **TheNegotiatorReloaded (TNR) [54]:** This agent is the best performing agent in non-discounted domains for the ANAC 2012 agent competition. The agent divides the negotiation into non-sliding windows, similarly to our approach. For each window, the agent estimates the type of agent behavior that it is facing and adjusts its concession rate accordingly. The agent employs Bayesian learning for modeling the opponents’ preferences and having an updated estimation of the Kalai point.
- **CUHK Agent (CUHK) [62]:** This agent won the ANAC 2012 agent competition for discounted domains. The agent adopts an adaptive concession strategy that sets a minimum target utility depending on both the behavior of the opponent and the characteristics of the domain (i.e., discount factor). In general, the agent will attempt to exploit the opponent as much as possible by not conceding. Compromises are only taken at the end of the negotiation after a time threshold that is set dynamically according to the opponent’s behavior. In this case, the agent attempts to model the opponent preferences by attempting to estimate a rank for proposed outcomes.

From the agents selected, one can observe that we have both agents that carry out opponent modeling (i.e., IAmHaggler2011, Hardheaded, TheNegotiatorReloaded, CUHK Agent, IAmHaggler2012), agents that show an adaptive concession behavior (e.g., AgentK, IAmHaggler2011, TheNegotiatorReloaded, CUHK Agent, Gahboninho, IAmHaggler2012), and agents that consider the opponent’s preferences when constructing bids (e.g., Hardheaded, TheNegotiatorReloaded, IAmHaggler2012). They are agents that performed among the best in the International Automated Negotiating Agents Competition (ANAC) editions of 2010, 2011, and 2012. It should be noted that agents from the 2013 competition onward were discarded as they specialize in different problems to the ones covered in this article (e.g., learning over multiple negotiations, non-linear utility functions, multi-lateral negotiations, etc.). Thus, we argue that a realistic agent society would be formed by top-performing agents as most of the ones employed in this setting.

5.1.4 Social agent configuration

As the reader may have realized, our agent’s behavior is controlled by many hyperparameters. After an initial exploration of our agent, we found a set of parameter values that generally performed well in practice. To avoid overfitting of our agent to the agents selected for the experiments, we carried out manual tests using *GENIUS* with our agent and Agent K in one of the domains: the party domain. During the tests, we observed both the outcome obtained by our agent and the evolution during the negotiation. After a few manual tests, we found a combination of hyperparameters (see Table 5) that obtained good joint utility and distances to both Nash point and the Pareto frontier. Of course, we do not claim that the configuration results in an optimal configuration of the agent, but, as the reader will appreciate, these values are enough to prove the main contribution of this article: how a social agent can efficiently perform in an agent society. As only one agent and domain were employed for finding this combination of hyperparameters, we argue that there is no risk of over-fitting, as there are five negotiation domains

and six opponents that were not employed during this exploration. Therefore, we used the parameter values found in Table 5 for the subsequent experiments carried out to assess the performance of our social agent.

Table 5: General configuration for our social agent during the experiments

Parameter	Description	Value
t_{exp}	Ratio of time spent using the <i>Strict Exploration</i> strategy	0.33
α_{exp}	Ratio of offers explored that is necessary to change to <i>Nash optimizer</i>	0.05
t_{sim}	Ratio threshold to jump from the <i>Exploratory</i> to the <i>Similary TFT</i> strategy	0.9
t_{eager}	Ratio threshold to jump to the <i>Eager</i> strategy	0.999
z	Percentile employed to apply the <i>top median</i> mechanism	80
c_{max}	Maximum concession step at each update	0.03
α and β	Time-dependent decay for the issue weight update mechanism	5 and 3

5.2 Analyzing the impact of the window size

In this first batch of experiments, we study to what extent the length of the window, employed by our agent to carry out the different behavioral analyses, impacts the social performance of our agent. The main goal of this experiment is to identify an adequate window size for properly and socially reacting to the moves of the agents in the agent society. For this first experiment, we set the value of the window size to $= \{25, 50, 75\}$. The rest of the parameters of the agent were set according to the values represented in Table 5.

Each configuration of our agent faced a subset of four agents (i.e., AgentK, IAmHaggler2011, Hardheaded agent, and TheNegotiator Reloaded) in a subset of three negotiation domains (i.e., party, Amsterdam trip, and England vs Zimbabwe domains) and negotiation profiles. The reason to carry out these experiments with just a subset of the opponents and domains is avoiding overfitting. In these experiments, each instance of our agent faced each of the 4 opponent agents for all 3 negotiation domains, using both preference profiles, and repeating each negotiation setting 10 times¹. In addition, we measured the average distance to the Nash bargaining point (Nash Distance), the distance to the Kalai point (Kalai Distance), individual utility (Ind. Utility) and the negotiation time at the end

¹A total of $4 \times 3 \times 2 \times 10 = 240$ negotiations per our social agent configuration.

of each negotiation, i.e. ratio of negotiation time consumed (% time). The experiments were conducted using *Genius* simulator [55] using bilateral negotiations with a common deadline of 90 seconds.

Table 6: Effect of window size on the given metrics.

Window size	Nash Distance	Kalai Distance	Ind. Utility	% time
25	0.050	0.045	0.80	0.82
50	0.057	0.052	0.80	0.82
75	0.056	0.052	0.80	0.83

Table 6 contains the results of this first experiment. As the reader may observe, a window size of 25 is capable of obtaining a lower distance to the Nash bargaining point and the Kalai point, resulting in a more social outcome than the other two configurations. More interestingly, there were no observable differences in the individual utility obtained by the proposed agent. This means that smaller window sizes also impacted the learning carried out by the agent. In fact, the larger the window sizes, the fewer updates that are carried out on the opponent model. We could not observe a notable difference between the three studied configurations in terms of speed, with negotiations taking approximately 82% of the total time.

5.3 Performance in a stationary agent society

As mentioned at the start of this section, the second experiment studies the behavior and performance of the proposed agent in a static society. The society is formed by the top-performing agents described in Section 5.1.3 and our social agent. For this analysis, we have assumed that agents may face any other agent in society, including versions of themselves. Another assumption for this first experiment is that the agents face the same number of times each other in the experiment. The game-theoretic analysis carried out in Section 5.4 will help us to answer what would happen if the agent society started from other different agent distributions in the society.

To carry out an experiment with this type of setting, we run a *GENIUS* tournament between all of the agents described in 5.1.3 and our agent. Each agent

faces each other agent (including itself) 10 times in every single domain, and for every possible profile in that domain. This means that each pair of agents faces in $10 \times 6 \times 2 = 120$ negotiations. In order to analyze the behavior and performance of each agent, we record the utility achieved at the end of the negotiation, the distance to the Nash bargaining and Kalai point, and the time taken to finish the negotiation. The parameters of our social agent are set at the ones described in Table 5 and the appropriate negotiation window size highlighted in the previous experiment (i.e., 25 bids).

Firstly, we focus on analyzing the overall performance of our social agent in this tournament. The overall results of this experiment can be found in Table 7. The table shows the average individual utility, distance to the Nash bargaining point, distance to the Kalai point, and negotiation time in the negotiations carried out by each of the agents in the society. As it can be observed, our social agent obtains the best average utility, distance to the Nash bargaining point, and distance to the Kalai point. These overall results point to various insights. The best performing values are in boldface.

Table 7: Agent aggregated results in the specified domains.

Agent	Utility	Nash Dist.	Kalai Dist.	Time
Social	0.81	0.08	0.07	0.85
Agent K	0.76	0.13	0.12	0.84
GAH	0.78	0.21	0.20	0.97
HH	0.76	0.22	0.21	0.95
IH2011	0.70	0.17	0.16	0.78
TNR	0.75	0.27	0.26	0.97
CUHK	0.79	0.27	0.27	0.99
IH2012	0.73	0.12	0.12	0.72

First of all, the distance to the Nash and Kalai points corroborates that our agent is indeed a social agent. It is, in fact, the most social agent among the agents in this society, including IAmHaggler2012 which was awarded as the most social agent in the ANAC 2012 competition. The reader should remind that both the Nash bargaining and the Kalai point are considered as fair outcomes in

negotiation settings. Therefore, the closer that the agreements are to these points, the fairer outcome that is achieved. A social agent should aim to be fair, as it tends to maximize the utility of both parties at the same time. Hence, our agent outperforms the current state of the art in terms of social agents. These results corroborate our second research question Q2 (i.e., *Can a social agent achieve socially efficient agreements in heterogeneous environments?*).

Secondly, one can also observe that our agent is capable of obtaining competitive results concerning the individual utility of the agent. Our social agent obtains the best average utility for the agents included in the experiment. This includes winning agents from the ANAC competition like Agent K, CUHK Agent, and HardHeaded. This positive result for the individual utility of the agent is related to the fact that the agent achieves deals that are close to the Kalai and Nash bargaining point. Usually, these points contain agreements that have a high utility for both parties, and no party suffers from exploitation. Therefore, by securing the socially optimal outcome, one can also guarantee a high utility for oneself. In this case, it was the highest overall average utility. This helps us answer our initial research question regarding the prospective performance of social agents in automated negotiation. Particularly, these results support one of our initial research questions, Q1 (i.e., *Can a social agent obtain a high utility in heterogeneous environments?*).

These results regarding the individual utility of agents are surprising at first, as most of the rest of the agents have been designed to be competitive and maximize one's utility as much as possible. Therefore, we carried out a more in-depth analysis of the negotiations of the tournament. This analysis threw some light into why competitive agents were not able to guarantee the highest average utility. Many negotiations where these competitive agents participated finished without any agreement due to the competitive nature of agents participating in the tournament. For instance, 12.5% of the negotiations for The Negotiator Reloaded, 8.33% of the negotiations for HardHeaded, 5.52% of the negotiations for Gahboninho,

and 4.68% of the negotiations for CUHK Agent finished without an agreement. The numbers were particularly worrying when some of these competitive agents faced themselves (e.g., 26.67% of negotiations failed when TheNegotiator Reloaded faced itself). No negotiation in which our agent participated finished in failure.

We argue that in an agent society, achieving agreements is of utmost importance as, otherwise, business opportunities and partnerships may be missed, and the reputation of the agent in the society may also decrease among other negotiators. In the end, if other agents in the society have had bad experiences negotiating against another agent, they may avoid selecting the agent as a negotiating partner in the future, further decreasing the performance of that agent in the society.

From these initial experiments, one can conclude that (i) the proposed agent indeed shows a social behavior by minimizing the distance of agreements achieved to both the Nash and the Kalai point, both measures of fairness in negotiation; (ii) although the design of the agent focuses on social performance and behavior, the agent is capable of achieving efficient results in terms of the individual utility of the agent. This latter result supports our initial hypothesis that social agents can also perform efficiently in automated negotiation. These results have been obtained based on the assumption that there is a static agent society composed of the agents included in the experiment. In the next experiment, we show how our agent would perform in case that the agent society is dynamic and agents can decide on what strategy to play in negotiations.

5.4 Performance in a non-stationary and utilitarian agent society

In this case, we plan to study the appropriateness of the proposed agent in a dynamic agent society, where agents can change their negotiation strategy prior to engaging in a negotiation with another agent. More specifically, assuming that agents can choose based on a range of available negotiation strategies, is our proposed strategy attractive even for self-interested agents? For that, we are going to analyze the agent society from a game-theoretic perspective. First, we formulate

a game where two agents have to negotiate, and each of them can decide to play according to one of the strategies mentioned in Section 5.1.3 or the proposed social agent. Then, the agents negotiate, and by the end of the negotiation, they achieve a utility. In this game, we assume that agents are rational, and their choice is purely driven by the individual utility expected at the end of the game. Therefore, we formulate a game with two players and eight possible strategies for each agent. The payoff of each strategy combination is obtained from the average utility achieved by that pair of agents across all domains and profiles in the previous tournament. The 9×9 payoff matrix arising from the definition of this game can be observed in Table 8, where bold pair represents the pure equilibrium of the game.

Table 8: Payoff matrix for negotiation.

	Social	K	GAH	HH	IH2011	TNR	CUHK	IH2012
Social	0.79/0.79	0.81/0.81	0.79/0.81	0.80/0.82	0.80/0.82	0.77/0.82	0.88/0.67	0.81/0.80
K	0.81/0.81	0.81/0.81	0.67/0.89	0.64/0.91	0.80/0.80	0.62/0.91	0.81/0.77	0.86/0.75
GAH	0.81/0.79	0.89/0.67	0.80/0.80	0.73/0.61	0.94/0.60	0.59/0.58	0.55/0.94	0.90/0.66
HH	0.82/0.80	0.91/0.64	0.61/0.73	0.72/0.72	0.92/0.63	0.60/0.64	0.56/0.79	0.91/0.65
IH2011	0.82/0.80	0.80/0.80	0.60/0.94	0.63/0.92	0.81/0.81	0.55/0.95	0.50/0.97	0.84/0.78
TNR	0.82/0.77	0.91/0.62	0.58/0.59	0.64/0.60	0.95/0.55	0.59/0.59	0.59/0.64	0.91/0.64
CUHK	0.67/0.88	0.77/0.81	0.94/0.55	0.79/0.56	0.97/0.50	0.64/0.59	0.73/0.73	0.82/0.75
IH2012	0.80/0.81	0.75/0.86	0.66/0.90	0.65/0.91	0.78/0.84	0.64/0.91	0.75/0.82	0.81/0.81

First, we carried out a pure and a mixed equilibrium analysis on the resulting game using the algorithm described in [63]. In the game, there is a single pure equilibrium consisting of one of the players using our social strategy, and the other agent employing The Negotiator Reloaded strategy. A total of 4 mixed equilibria were found for the game, and they can be observed in Table 9 where each row represents a different mixed equilibrium; and in each cell, the first number represents the ratio of use of the strategy for the first player, while the second represents the ratio of use of the same strategy for the second player. As it is observable both in Tables 8 and 9, our social agent has a significant contribution (i.e., most used strategy) to the strategy profile of at least one of the two players in the negotiation. In fact, it is the most used negotiation strategy for the two players in 3 out of the 4 mixed strategies of the game. Overall, even in a society

where agents can rationally adapt their strategies for the negotiation, our social agent has a prevalence over the rest of most competitive strategies. Both game-theoretic results support the fact that our proposed agent can be an efficient choice even for self-interested agents, answering our initial research question Q3 (i.e., *Can social negotiation strategies be an efficient choice for self-interested agents in agent societies?*).

Table 9: Mixed Nash equilibrium for the dynamic agent society.

Social	K	GAH	HH	IH2011	TNR	CUHK	IH2012
0.72/0.72	0/0	0/0	0.28/0.28	0/0	0/0	0/0	0/0
0.77/0.77	0/0	0.05/0.05	0.19/0.19	0/0	0/0	0/0	0/0
0.50/0.50	0/0	0.47/0.47	0/0.0	0/0	0/0	0.04/0.04	0/0
0.92/0	0/0.48	0.06/0.32	0/0	0/0	0/0.21	0.02/0	0/0

CHAPTER VI

ASSOCIATIVE OPPONENT MODELING

In this section, we propose an opponent modeling based on learning associations between issue values appeared in opponent's bid history. Particularly, we aim to propose a model that is capable of modelling linear additive utility functions based solely on the opponent's bids during the negotiation, inspired by association rule mining approaches.

6.1 Adapting Association Rule Mining for Negotiation

In automated negotiation, agents exchange bids to reach agreements based on their strategy. During that time, agents may keep track of what the opponent offered previously; forming a collection of bids with a timestamp, i.e. a bid history until current time t such that $\mathcal{O}_{0 \rightarrow t} = \mathcal{O}_t = \{o_0, \dots, o_t\}$ where $t \in (0, 1)$. A complete bid is composed of combinations of issue values. For example, we have a seller and a buyer agent negotiating on Laptop domain having issues of hard disk and monitor. Let's assume that, hard disk issue has values/options of *(128GB, 256GB)* and monitor has list of values *(19", 23", 27")*. According to our assumption, to form a complete bid, agents must pick one value from each of the issues such as *(128GB hard disk, 19" monitor)* or *(256GB hard disk, 27" monitor)* monitor. Depending on the issue weights and value evaluations, agents calculate a utility of each of bids in the negotiation space.

Unlike the frequency-based models explained in the previous chapters, our association-based opponent modeling approach is not based on the heuristic of "least variation implies the most importance". Specifically, those models track the changes in issue values; penalizing the issues if their selected value changed from one bid to another. Here, we focus on the content of bid history rather than their received order. We investigate the co-occurrence patterns or associations of

the bid components (i.e. issue values). For this purpose, we adapt three indicators of association from association rule mining field, namely, *support*, *confidence* and *lift*.

As we briefly explained in Section 2.2 and the beginning of this chapter, a list of bids or a bid history received from an opponent agent, H^t , has the same structure with transactional datasets, which lets us to apply association rule mining approaches. In the rest of this section, three indicators of association are explained in detail.

The first indicator of association, $support^t(v_i)$ indicates the occurrence frequency of i^{th} issue value, v_i , e.g. *256GB hard disk*, in the bid history H^t . So, if that bid exists in eight bids in H^t out of 10 bids at time t ; then $support^t(256GB\ hard\ disk)$ would be 80%. We calculate co-existence of issue values such as $support^t(v_i, v_j) \mid i \neq j$, where the second element indicates another issue value. An example for v_j would be *27" monitor*, and if (*256GB hard disk*, *27" monitor*) pair occurred 4 times in the same bid history; $support^t(256GB\ hard\ disk, 27"\ monitor)$ would be 40%. In general case, it would be enough to indicate that the associating items should be different but since we assume additive utility function; bid history can never have association of values from the same issue. We may think that as a set of specific constraints imposed for negotiation domains. In order to generalize the notation and calculate the marginal frequency of a set of items - support, we introduce partial bid or offer which is a subset of value assignments $S \subset o$ where $|S| < |o|$. So, We formally define the support as:

$$support(S) = \frac{\#\{S \subset o \mid o \in \mathcal{H}^t\}}{|\mathcal{H}^t|} \quad (13)$$

where $\#\{S \subset o \mid o \in \mathcal{H}^t\}$ denotes the number of times that a partial bid S appears as part of the history of bids $\mathcal{H}^t$, and $|\mathcal{H}^t|$ as the number of bids in the history until time t . One may think of support as an estimation for the marginal probability of a partial bid at the population level.

Another relevant metric in association mining is confidence. Confidence provides us a measure of certainty for the appearance of an item given that another item is present. When applied to collections of bids, we need to introduce some additional definitions. We say that two partial bids S_i and S_j are compatible when they set values for different negotiation issues. For instance, $23''$ is compatible with $500GB$, but $23''$ is not compatible with $19''$ as they set the value for the same negotiation issue, the screen size. The confidence between partial bids S_j and S_i , both compatible, gives us the certainty of finding partial bid S_j in the same bid where S_i is also present. Formally, this can be defined as:

$$confidence(S_i \rightarrow S_j) = \frac{support(S_i \cup S_j)}{support(S_i)} \quad (14)$$

That is, the joint frequency of both partial bids, divided by the marginal frequency of S_i . Hence, one may think of confidence between S_j and S_i as an estimation for the conditional probability of S_j given S_i , $P(S_j|S_i)$.

The last important term in association mining is lift. The lift between two sets of items is a measure of the strength of the dependency between both item sets [64]. It is calculated as the probability of seeing both items divided by the multiplication of prior probabilities of those items. If it is less than 1, it indicates a negative dependency between those items while a positive dependency is indicated when it is greater than 1. When analyzing collection of bids, it is indicating the strength of dependency between two partial bids for appearing in the same bid. Let us formalize lift as:

$$lift(S_i, S_j) = \frac{confidence(S_i \rightarrow S_j)}{support(S_j)} \quad (15)$$

Since it gives some insights which partial bid contents appear with others, it could be used in the estimation of the opponent's preferences. Accordingly, we propose an opponent modeling based on lift estimation in the next section.

6.2 Preference Modeling Based on Association Mining

In this section, we describe how one can employ the information extracted from association bid mining in order to carry out opponent's preference modeling. It should be highlighted that the proposed model can be employed during the negotiation, based on the history of bids from the current negotiation, or after the negotiation to estimate the opponent's preferences for future interactions. The latter estimation can be carried out from the full interaction history with the opponent.

Intuitively, the lift values calculated by using the opponent's offer history, give some insights about which issue values are likely to appear in the opponent's offers so far. Our model also implicitly considers how often each value individually appears in the bid history while estimating the overall utility. Recall that the opponent's preferences are represented in terms of additive utility function. To find the weighted evaluation for *each issue value*, we calculate the lift values of all related values that can co-exist with that issue value in the bid history.

By definition, association is defined between at least two elements. In our setup, we calculate association between two values ($v_i \in D_i$, $v_j \in D_j$) of different issues where D_i and D_j denote sets of possible values for negotiation issues i and j respectively. Instead of estimating the evaluation values and issue weights separately, we aim to estimate weighted evaluation values. Recall that each issue value is associated with a constant weighted evaluation value according to the additive utility function. For this purpose, we consider the co-occurrences of a particular issue value with all other valid issue values and take the average of them. In other words, average lift values corresponds to the estimated weighted evaluation values in our case.

According to the frequentist approaches, the most preferred values would appear more often in the opponent's bid history. As mentioned in the previous chapters, this heuristics has some drawbacks especially when the opponent repeats the same bids. Also, a negotiation space may have bids with close utilities

and the opponent may adopt exploratory behavior in a utility range. In this case, opponent model decreases the weight of varying issues; even they are frequently asked. Finally, bids that have close utilities might have different content (i.e., combination of different issue values). Therefore, the variation on the content of the consequent offers may not be a sign of insignificance.

Our proposed association approach aims to amend the weaknesses of the frequentist heuristics. Lift calculation basically has three elements; first one is the joint frequency (co-occurrence) of the given items in numerator. In the denominator, we have the product of the items' marginal frequencies. Following that, lift value increases when the items are frequently seen together and especially when any of them is received rarely. Notice that, this formulation ignores the order of the observations making the algorithm insensitive against exploratory behavior of the opponent. In addition, if the opponent insists on the same bid for a period of time, joint frequency and the marginal frequencies of that bid's issue values increase simultaneously; over time, ending up with lift of 1 (greater than 0 - no association). For example, assume that we have three issues with some possible values as follows: "Harddisk: 128GB and 256GB", "Monitor: 19", 23" and 27"" and "Brand: Dell, HP and Lenovo". In order to estimate weighted evaluation of the issue value of 19", we calculate all possible pairwise lift combinations of 19" Monitor (i.e., $\text{lift}(19", \text{Dell})$, $\text{lift}(19", \text{HP})$, and $\text{lift}(19", \text{Lenovo})$, $\text{lift}(19", 128\text{GB})$, $\text{lift}(19", 256\text{GB})$) and takes their averages.

Algorithm 3 illustrates how we estimate the weighted evaluation values of each issue values. For each issue value , we calculate lift values with the other issue values (Lines 1-8) while counting the number of lift calculations (Line 7). Afterwards, we calculate the average of those values dividing by the total number of calculations (Line 9). In Lines 10-11, we update the maximum lift value for each issue. Finally, we scale the estimated values in the range of zero and one by dividing the maximum value (i.e., max normalization, see Lines 12-13).

Based on the estimated weighted evaluation values, the utility of a given offer

is estimated as the sum of the evaluation values for each issue. That is $\sum_{i=1}^n E_{v_i}/n$ when there are n issues. For instance, the estimated utility of the following offer based on the laptop scenario $\langle 19'', 128GB, Dell \rangle$ would be $(E_{19''} + E_{128GB} + E_{Dell})/3$.

Data: D : Values of domain issues, $Lift$: lift calculation for the value pair v_i and v_j

Result: E_D : Estimated evaluation for each issue value $v_i \in D$.

```

1 foreach  $i \in Issues$  do
2    $E_{v_i}^{max} \leftarrow 0$ ;
3   foreach  $v_i \in D_i$  do
4      $N_i \leftarrow 0$ ;
5     foreach  $j \neq i, j \in Issues$  do
6       foreach  $v_j \in D_j$  do
7          $N_i += 1$ ;
8          $E_{v_i} += lift(v_i, v_j)$ ;
9       end
10     end
11      $E_{v_i} /= N_i$ ;
12     if  $E_{v_i}^{max} < E_{v_i}$  then
13        $E_{v_i}^{max} \leftarrow E_{v_i}$ ;
14     end
15   end
16   foreach  $v_i \in D_i$  do
17      $E_{v_i} \leftarrow \frac{E_{v_i}}{E_{v_i}^{max}}$ ;
18   end
19 end

```

Algorithm 3: Weighted Value Evaluation Estimation Algorithm

CHAPTER VII

EVALUATION OF OPPONENT MODELING APPROACHES

In this section, we assess the empirical performance of the proposed opponent modeling methods in various domains given agents' real negotiation history.

7.1 General Experimental Settings

In this section, we provide the detailed explanation of our experimental setup.

7.1.1 Quality Metrics

In this study, our main focus is on the ranking of the offers in the outcome space, rather than their true utility values. Therefore, we consider the following metrics in our evaluation.

- **Pearson Correlation:** This metric shows the linear relationship of two utility-spaces. The correlation coefficient $\mathbf{r}$ ranges between -1 and 1; where sign of the coefficient shows the direction and the magnitude is the strength of the linear relationship. In our experiment, we aim to predict utility-spaces with strong positive $\mathbf{r}$ values.
- **Kendall Tau Rank Correlation:** Kendall Tau coefficient can get values such that $-1 \leq \tau \leq 1$ and gives strong positive value when the ranks of the given utility-spaces are similar. If the utility spaces have reversed ranking of each other, τ gets value of -1.

7.1.2 Negotiation Domains

In these experiments, we selected 13 negotiation domains with various outcome sizes and opposition classes to evaluate the models with significant sample size

from each criterion. In the next lines, we briefly explain the domains listed in Table 10.

- **Ultimatum (ANAC 2013, [65]):** A small but highly competitive domain that has only 2 issues having 3 values each; ending up with 9 bid options.
- **Laptop (ANAC 2011, [59]):** This is a relatively small domain with 27 possible outcomes which has three issues for buyer and seller parties: brand of laptop, size of hard disk, and the screen size of external monitor.
- **Rental House(ANAC 2012, [60]):** This domain is highly competitive and has 4 issues: price and style of the house, acceptable locations and finally water heater type. Domain is composed of 60 possible outcomes, so it is relatively larger than Ultimatum and Laptop domains but still a small domain.
- **Itex vs Cypress (ANAC 2010, [58]):** This is a classic buyer-seller setting where a bike manufacturing company negotiates on the sale of a commodity with a bike retail shop. There are a total of 180 possible outcomes.
- **England vs Zimbabwe (ANAC 2010, [58]):** This is a negotiation domain consisting of the national negotiations between England and Zimbabwe on the application of the World’s Health Organization Framework on Tobacco Control. The negotiation domain consists of 576 possible negotiation outcomes and 5 negotiation issues.
- **Barbecue (ANAC 2012, [60]):** This domain is about preferences on a barbeque gathering. Domain has 5 negotiation issues: meat, drinks, location, vegetables, and finally barbeque type. It has 1,440 possible outcomes and medium opposition.
- **Grocery domain (ANAC 2011, [59]):** This domain involves agents negotiating on the type of groceries to purchase from a supermarket. There are a total of 1,600 possible outcomes distributed across 5 negotiation issues.

- **Amsterdam Trip (ANAC 2011, [59]):** This is a negotiation domain where two parties of friends negotiate on the details of a trip to Amsterdam. The negotiation domain contains 3,024 possible outcomes distributed across 6 negotiation issues.
- **Music Collection (ANAC 2012, [60]):** This one is one of the large domains that has 4,320 possible outcomes and low opposition. It has 6 negotiation issues which are genres: classical, rock, jazz, pop, Brazil, and Latin where their values correspond to famous artists known for given genre.
- **Supermarket (ANAC 2012, [60]):** This domain is one of the largest ones among others where two parties decide on a market basket; having 98,784 possible outcomes and high level of opposition. The domain has 6 negotiation issues: bread type, fruit, snacks, spreads, vegetables, and drinks.
- **Travel (ANAC 2012, [60]):** The domain is a large and highly competitive domain that has 7 negotiation issues: atmosphere, amusement, culinary, shopping, culture, sport, environment. Travel is one of the large domains with medium competition that consists of 188,160 possible outcomes.
- **Energy (ANAC 2011, [59]):** This domain is a highly competitive and it is about demand and supply side of an energy grid which includes 390,625 possible agreements making it the largest among others. It has 8 negotiation issues that are basically time intervals that the power transmission will occur.

7.1.3 Negotiation Agents

Among ANAC competition agents, we consider the winning agents that are designed for bilateral or multi-lateral negotiation. We picked two set of winner agents based on years and bilateral-multiparty focus in those years' automated negotiating agents competition. Ideally, we picked top 3 agents from each year considering individual utility; however, agents in the second place (CaucusDC16

Table 10: Negotiation Domains

Domain	Year	# of Issues	# of Outcomes	Opposition Value	Opposition Class
Ultimatum	2013	2	9	0.545	High
Laptop	2011	3	27	0.160	Low
Rental House	2012	4	60	0.327	High
Barter	2012	3	80	0.492	High
Itex vs Cypress	2012	4	180	0.431	High
EnglandvsZimbabwe	2012	5	576	0.272	Medium
Barbecue	2012	5	1440	0.239	Medium
Grocery	2011	5	1600	0.191	Low
Amsterdam	2011	6	3024	0.223	Medium
Music Collection	2012	6	4320	0.150	Low
Supermarket	2012	6	112896	0.347	High
Travel	2012	7	188160	0.230	Medium
Energy	2011	8	390625	0.445	High

[66]) and third place (Rubick [66]) faced problems in the tournaments on GENIUS negotiation application. Therefore, we left them out.

- **PonPokoAgent (ANAC 2017, [66]):** This agent is a time-based conceder and it has various bidding strategies but since ANAC 2017 had data persistency option; it picks a random bid to evade its opponents. Agent is designed to reach the highest individual utility and proved its success by being the winner of that year’s competition in this category.
- **Caduceus (ANAC 2016, [67]):** This agent is the winner of ANAC 2016.
- **AgentYX (ANAC 2016, [67]):** This agent is got the second place in ANAC 2016 competition.
- **ParsCat (ANAC 2016, [67]):** Agent that received the same score with YXAgent and shared the second place in ANAC 2016.
- **Atlas3 (ANAC 2015, [68]):** This agent follows a concession strategy that is previously developed by authors which is based on expected utility of evolutionary stable strategy bilateral closed bargaining problem to set a proper minimum point to stop conceding.

- **ParsAgent (ANAC 2015, [69]):** This agent adopts mixture of bidding strategies that includes random, time-dependent and frequency-based behavior. To increase the chances to agree on a better bid, agent offers that are similar to opponent's offers.
- **RandomDance (ANAC 2015, [69]):** This agent took the third place in ANAC 2015 competition based on individual utility criteria.
- **TheFawkes (ANAC 2013, [70]):** This agent adopts the strategy used by OMACagent [71] and improves its bidding strategy by offering bids considering opponent's utility, instead of sending random bids around a target utility level.
- **MetaAgent (ANAC 2013, [72]):** This is a meta-learning agent that picks best strategy based on opponent's bidding history.
- **TMFAgent (ANAC 2013, [65]):** This agent took the third place in ANAC 2013 competition and it is one of the best performing agents in preliminary rounds.
- **CUHK Agent (ANAC 2012, [62]):** This agent won the ANAC 2012 agent competition for discounted domains. The agent adopts an adaptive concession strategy that sets a minimum target utility depending on both the behavior of the opponent and the characteristics of the domain (i.e., discount factor). In general, the agent attempts to exploit the opponent as much as possible by not conceding. Compromises are only taken at the end of the negotiation after a time threshold that is set dynamically according to the opponent's behavior. In this case, the agent attempts to model the opponent preferences by attempting to estimate a rank for proposed outcomes.
- **AgentLG (ANAC 2012, [60]):** This agent got the second place in discounted domains of the final round of the Automated Negotiating Agents Competition 2012.

- **OMACagent (ANAC 2012, [71]):** The agent is discounted domain compatible and it calculates an expected utility considering the discount factor not to fall below reservation value. Agent follows an adaptive concession pattern based on time and expected utility from opponent and applies standard wavelet decomposition and cubic smoothing spline for opponent modeling.
- **Hardheaded (ANAC 2011, [41]):** The winning agent of the 2011 ANAC competition. This agent follows a time-based bidding strategy that consists of not conceding until almost the end of the negotiation. While not conceding, the agent performs opponent modeling on the opponent's preferences by using a frequency modeling mechanism. Then, in the last states of the negotiation, the agent attempts to find an agreement by using the trained model on the opponent's preferences and finding outcomes that have the same target utility for one's own but maximize the preferences of the opponent.
- **Gahboninho (ANAC 2011, [61]):** This agent was the runner-up in the ANAC 2011 competition. Gahboninho employs an adaptive concession strategy that attempts to detect if agents are susceptible to being pressured or if they are stubborn. For that, Gahboninho limits most of the outcomes sent to outcomes that report a low utility value for the opponent.
- **IAmHaggler2011 (ANAC 2011, [53]):** A negotiating agent that uses Gaussian processes to predict the future concession of its opponent, and then adjust its concession rate accordingly to get the most from the negotiation. The goal of this agent is that of optimizing one's utility while also trying to get more utility than the opponent. Bids are only selected from a small range around the target utility. This aggressive stance usually results in the agent repeating the same offers. This agent took third place in the ANAC 2011 competition.
- **Nice TFT (ANAC 2011, [59]):** This agent adopts Bayesian learning for

opponent modeling and follows a concession mechanism such that it reacts in opponent moves in reciprocal direction and a proportionate amount. In decision making, agent also considers the information gain of its opponent model and adjusts its moves accordingly.

- **Conceder (CN): agent [51]:** A classic negotiation agent that adjusts its aspiration levels according to time, but large concessions are done at the start of the negotiation. This type of agent behavior is found when agents are eager to close an agreement no matter the utility of that agreement (e.g., urgency to get a service).
- **Boulware (BW) agent [51]:** A classic negotiation agent that adjusts its aspiration levels according to time, only conceding in the later stages of the negotiation. Bids close to the target utility are selected and sent to the opponent. This agent does not take into consideration the opponent's preferences when generating bids.

We grouped the these agents according to the competition settings, in terms of bilateral / multilateral negotiation challenges.

Table 11: Negotiation Agents in Groups

Year	Agent Group 1	Year	Agent Group 2
2013	TheFawkes	2017	PonPokoAgent
2013	MetaAgent	2016	Caduceus
2013	TMFAgent	2016	AgentYX
2012	CUHKAgent	2016	ParsCat
2012	AgentLG	2015	Atlas3
2012	OMACagent	2015	ParsAgent
2011	HardHeaded	2015	RandomDance
2011	Gahboninho		
2011	IAMhaggler2011		
2011	Nice TFT		
	Conceder		
	Boulware		

7.1.4 Negotiation Tournament Setup

In this section, we briefly explain the experimental setup. As it is visualized in Figure 7.1.4, we run tournaments including all agents shown in Table 11, on the negotiation scenarios listed in Table 10. In the yearly Automated Negotiating Agents Competitions (ANAC), traditional preference for the deadline selection is 180 seconds; however, it would take too long to run the entire tournament because most of the agents have difficulty in large domains [60]. Following that we set the deadline as 100 rounds.

The diagram shows also that the logs of each negotiation session generated during the tournaments are passed to *Negotiation History Decomposition* module which parses the GENIUS negotiation application's session logs and prepares them for offline learning. Please note that the opponent models can operate online but since bidding or acceptance strategies are out of scope of this work; we intended to provide a real negotiation history dataset to the literature. Finally, we have the *Opponent Model Training & Evaluation* module, where each negotiation history is provided by single rounds to the target learning model as if it is an online learning. Here, in each round, we record utility-space predictions of the models.

Since each agent plays two different profiles in a scenario, the number of sample sessions for each of the *profile* in a single scenario can be calculated as (Group Size -1) x R repetitions. So, in each domain, an agent has (Group Size -1) x R x 2 session records. Following that, any agent in Group 1 has been a part in each scenario $(12-1) \times 20 \times 2 = 440$ times; and for any agent in Group 2 it is $(7-1) \times 20 \times 2 = 240$. Finally, when we consider the number of domains, which is 13; any agent in Group 1 end up with 5720 sessions and for Group 2, it is 3120.




Figure 6: Experimental Setup Diagram

7.2 *Experimental Results*

In this part, we present our experimental results of three opponent models where Associative Opponent Model (ASSOC) and Distribution-Based Frequency Model (DBFQ) and Frequency Model (FQ) as mentioned in the Chapter 3. Note that because the deadline round limit, DBFQ model's window size parameter is set to five which is a small value and expected to harm the model's predictions.

In the following part, we analyze experiment evaluation into from perspectives: model success on final rounds of the negotiations and progress of the models over the rounds.

7.2.1 Final Round Based Results

In this section, we show comparative results of the models based on success measures of Kendall Tau and Pearson.

Here, we applied the non-parametric statistical test called Kruskal-Wallis to see whether there is any statistically significant difference between distributions of the negotiation outcomes given the criteria. Afterwards, we carried out the Posthoc-Dunn test to detect pairwise differences. In the tables, a boldfaced result indicate that it is the highest median value in the given metric and the distribution of median scores are also significantly different than others.

In Table 12, we see the aggregated results for Agent Group 1. Results for Kendall Tau metric show that proposed ASSOC model is commonly winner except for the TMFAgent, where our DBFQ model took the lead. We also see that results for Gahboninho agent are close and none of them is significantly different than other. In Pearson metric, ASSOC model shows dominant success.

Table 12: Last round utility-space prediction based aggregated medians for Agent Group 1 Tournament with 100 rounds deadline.

Metric	Kendall Tau			Pearson		
Agent / Model	ASSOC	DBFQ	FQ	ASSOC	DBFQ	FQ
AgentLG	0.700	0.593	0.617	0.867	0.769	0.810
Boulware	0.676	0.601	0.585	0.831	0.741	0.772
CUHKAgent	0.571	0.542	0.535	0.769	0.731	0.726
Conceder	0.686	0.467	0.414	0.875	0.653	0.602
Gahboninho	0.550	0.552	0.551	0.740	0.732	0.736
IAMhaggler2011	0.755	0.603	0.661	0.926	0.804	0.846
KLH	0.630	0.572	0.555	0.788	0.721	0.756
MetaAgent2013	0.688	0.615	0.602	0.858	0.785	0.800
NiceTitForTat	0.587	0.562	0.539	0.777	0.726	0.731
OMACagent	0.628	0.563	0.581	0.829	0.730	0.767
TMFAgent	0.549	0.552	0.497	0.728	0.716	0.692
TheFawkes	0.529	0.457	0.446	0.725	0.644	0.641

In Table 13, we see the aggregated results for Agent Group 2. Here, we have some variations where the results for Kendall Tau metric show that proposed ASSOC model has significance dominance for half of the cases and for agent RandomDance and YXAgent, it has the highest median score but not significant. In

Table 13: Last round utility-space prediction based aggregated medians for Agent Group 2 Tournament with 100 rounds deadline.

Metric	Kendall Tau			Pearson		
Agent / Model	ASSOC	DBFQ	FQ	ASSOC	DBFQ	FQ
Atlas3	0.640	0.586	0.532	0.838	0.786	0.715
Caduceus	0.542	0.544	0.546	0.728	0.719	0.736
ParsAgent	0.544	0.544	0.549	0.722	0.721	0.738
ParsCat	0.784	0.606	0.648	0.926	0.780	0.835
PonPokoAgent	0.684	0.567	0.605	0.870	0.734	0.792
RandomDance	0.513	0.511	0.475	0.714	0.702	0.673
YXAgent	0.648	0.608	0.638	0.835	0.812	0.830

Pearson metric, ASSOC model has the highest number of significant results but this time FQ model also gets two of them.

In Table 14 and Table 15 we demonstrate the domain aggregated results for Agent Group 1 and Agent Group 2 respectively. In both of the table, we see the ASSOC model comes first in most of the cases and for both of the metrics. We also see that in the Ultimatum domain, multiple models shares the highest median scores and still their distributions are significantly different. This may happen since the statistical tests consider the shape of the distribution, instead the median values. So, we conclude that the models commonly have significant and highest results have prediction distributions having same median but in different shape.

In Table 16 and Table 17 we demonstrate the domain opposition aggregated results for Agent Group 1 and Agent Group 2 respectively. In the first table, ASSOC model gets the highest score for every opposition level in both of the metrics and the results the Agent Group 2 show that even ASSOC model got the highest score in Low Opposition domain case Pearson metric, it is not statistically significant.

7.2.2 Round Based Results

In this section of the experiments, we analyze the results based on negotiation rounds or the success of models depending on number of received bids. Here, we did not explicitly apply statistical tests since the first section provides enough

Table 14: Final round utility-space prediction based *domain* aggregated medians from Agent Group 1 Tournament with 100 rounds deadline.

Metric	Kendall Tau			Pearson		
Domain / Model	ASSOC	DBFQ	FQ	ASSOC	DBFQ	FQ
Amsterdam	0.650	0.601	0.595	0.848	0.785	0.790
Barbecue	0.548	0.481	0.471	0.752	0.648	0.663
Barter	0.739	0.647	0.592	0.877	0.760	0.753
Energy	0.406	0.367	0.375	0.568	0.516	0.558
EnglandvsZimbabwe	0.591	0.556	0.529	0.742	0.739	0.730
Grocery	0.626	0.595	0.645	0.816	0.780	0.841
ItexvsCypress	0.740	0.648	0.618	0.899	0.808	0.793
Laptop	0.714	0.705	0.698	0.850	0.836	0.832
MusicCollection	0.543	0.537	0.525	0.733	0.707	0.723
RentalHouse	0.774	0.705	0.619	0.926	0.847	0.796
Supermarket	0.554	0.533	0.554	0.736	0.711	0.744
Travel	0.412	0.390	0.387	0.600	0.552	0.555
Ultimatum	0.816	0.816	0.725	0.964	0.964	0.964

Table 15: Final round utility-space prediction based *domain* aggregated medians from Agent Group 2 Tournament with 100 rounds deadline.

Metric	Kendall Tau			Pearson		
Domain / Model	ASSOC	DBFQ	FQ	ASSOC	DBFQ	FQ
Amsterdam	0.653	0.652	0.637	0.854	0.836	0.838
Barbecue	0.513	0.488	0.480	0.714	0.673	0.684
Barter	0.655	0.642	0.609	0.814	0.762	0.754
Energy	0.535	0.421	0.424	0.743	0.595	0.619
EnglandvsZimbabwe	0.654	0.589	0.581	0.854	0.757	0.782
Grocery	0.602	0.578	0.636	0.796	0.775	0.840
ItexvsCypress	0.677	0.621	0.619	0.853	0.791	0.803
Laptop	0.718	0.693	0.668	0.846	0.831	0.833
MusicCollection	0.546	0.526	0.519	0.740	0.711	0.722
RentalHouse	0.708	0.683	0.622	0.887	0.857	0.814
Supermarket	0.583	0.562	0.568	0.781	0.743	0.759
Travel	0.421	0.396	0.420	0.619	0.571	0.610
Ultimatum	0.816	0.816	0.725	0.963	0.964	0.964

Table 16: Final round utility-space prediction based *domain opposition level* aggregated medians from Agent Group 1 Tournament with 100 rounds deadline.

Metric	Kendall Tau			Pearson		
Opposition / Model	ASSOC	DBFQ	FQ	ASSOC	DBFQ	FQ
Low	0.652	0.575	0.620	0.830	0.765	0.826
Medium	0.562	0.489	0.494	0.742	0.660	0.704
High	0.688	0.598	0.572	0.862	0.760	0.755

Table 17: Final round utility-space prediction based *domain opposition level* aggregated medians from Agent Group 2 Tournament with 100 rounds deadline.

Metric	Kendall Tau			Pearson		
Opposition / Model	ASSOC	DBFQ	FQ	ASSOC	DBFQ	FQ
Low	0.627	0.575	0.612	0.82	0.770	0.809
Medium	0.568	0.522	0.535	0.776	0.715	0.737
High	0.651	0.595	0.591	0.838	0.761	0.767

information about the model significance and overall model success.

The Table 18 and Table 19 show the success of the models that we applied aggregation on rounds or received bids to the sessions of each agent and for each domain. Here, we see that ASSOC model is more successful than others and as we previously mentioned, DBFQ model comes after in some cases because of its small window size parameter. Notice that DBFQ model has no predictions until it receives double size of its windows size, which is 10 for this case. For that reason, round 2 and 5 cell of the tables are left blank for DBFQ model. The following result tables demonstrate the further details of the given grouping, aggregation and metrics. Correlation values at various rounds are given in Appendix A.

Figure 7, 8, 10 and 10 visualize the results for the selected domains and Pearson metric for each of the Agent Group. Here, we selected Amsterdam, EnglandvsZimbabwe and Travel domains since they share the opposition level but vary in size. In the Figure 7 and 8 are the results for the round or the number of received bids; on the other hand, Figure 9 and 10 show the effect of *received unique bid*. We wanted to include this distinction because models' learning success affected by the agents' bidding strategy. Some of the agents insist on the same bids and do not explore the negotiation space, e.g. Caduceus agent, so the models can't learn because no information is provided. As a final note, notice that the FQ model suffers from fast conceding behavior, especially in small domains as in Figure 7: sub-figure top right.

Table 18: Utility-space prediction based *number of received bids* aggregated medians from Agent Group 1 Tournament with 100 rounds deadline.

Metric	Kendall Tau			Pearson		
Round / Model	ASSOC	DBFQ	FQ	ASSOC	DBFQ	FQ
2	0.579		0.554	0.719		0.719
5	0.579		0.541	0.726		0.728
10	0.579	0.569	0.559	0.734	0.722	0.734
15	0.579	0.569	0.550	0.739	0.722	0.732
20	0.579	0.569	0.550	0.739	0.720	0.732
25	0.579	0.569	0.550	0.742	0.722	0.730
30	0.579	0.565	0.547	0.744	0.723	0.732
35	0.579	0.565	0.546	0.744	0.725	0.732
40	0.579	0.564	0.553	0.744	0.726	0.732
45	0.584	0.558	0.552	0.747	0.726	0.734
50	0.589	0.559	0.542	0.753	0.728	0.734
55	0.591	0.561	0.547	0.763	0.730	0.736
60	0.592	0.561	0.550	0.774	0.733	0.734
65	0.593	0.562	0.547	0.781	0.734	0.739
70	0.595	0.562	0.553	0.787	0.735	0.739
75	0.610	0.565	0.559	0.799	0.739	0.739
80	0.613	0.565	0.555	0.805	0.739	0.739
90	0.610	0.567	0.556	0.804	0.736	0.737
95	0.606	0.565	0.548	0.795	0.734	0.736
100	0.618	0.564	0.556	0.807	0.739	0.748

Table 19: Utility-space prediction based *number of received bids* aggregated medians from Agent Group 2 Tournament with 100 rounds deadline.

Metric	Kendall Tau			Pearson		
Round / Model	ASSOC	DBFQ	FQ	ASSOC	DBFQ	FQ
2	0.547		0.535	0.694		0.711
5	0.547		0.541	0.711		0.732
10	0.552	0.552	0.557	0.725	0.728	0.739
15	0.561	0.552	0.552	0.739	0.729	0.736
20	0.571	0.558	0.56	0.739	0.733	0.737
25	0.579	0.558	0.562	0.745	0.734	0.739
30	0.58	0.56	0.559	0.753	0.734	0.739
35	0.584	0.563	0.562	0.763	0.739	0.743
40	0.591	0.569	0.574	0.769	0.739	0.745
45	0.596	0.569	0.577	0.779	0.742	0.752
50	0.604	0.569	0.572	0.789	0.744	0.753
55	0.604	0.57	0.572	0.791	0.744	0.758
60	0.614	0.571	0.583	0.796	0.746	0.764
65	0.618	0.57	0.581	0.799	0.747	0.768
70	0.623	0.572	0.584	0.807	0.749	0.772
75	0.623	0.572	0.584	0.808	0.749	0.772
80	0.627	0.574	0.587	0.811	0.752	0.773
85	0.629	0.574	0.591	0.815	0.753	0.777
90	0.635	0.577	0.601	0.817	0.755	0.782
95	0.639	0.583	0.605	0.822	0.758	0.786
100	0.667	0.589	0.609	0.842	0.76	0.795




Figure 7: Agent Group 1, Pearson correlations over the *rounds* for the selected domains.




Figure 8: Agent Group 1, Pearson correlations over the *unique bids* for the selected domains.




Figure 9: Agent Group 2, Pearson correlations over the *rounds* for the selected domains.




Figure 10: Agent Group 2, Pearson correlations over the *unique bids* for the selected domains.

CHAPTER VIII

RELATED WORK

In this section, we overview the related work in the field of automated negotiation literature. We first mention the state-of-the-art social negotiation strategies (Section 8.1) and then discuss some opponent modelling approaches used in bilateral negotiations (Section 8.2).

8.1 Negotiation Strategies

The literature on cooperative agents is extensive, and it includes works as relevant as collaborative knowledge networks [73, 74, 75], multi-agent teams [11, 76, 77], cooperative agent planning [78, 79], coalition formation [80], and so forth. Despite social and cooperative negotiating agents have links with many other areas of cooperative agents, in this section we aim to review those works that are the most related to our social agent proposal. Mainly, we aim to discuss cooperative and adaptive negotiation agent proposals and highlight the differences between these agents and our proposal.

Lai *et al.* [34] propose a multi-issue bilateral negotiation agent that is based on a bidding mechanism that proposes offers in the iso-utility curve. The agent sends up to k offers per round to its counterpart, who may accept or propose new counteroffers. The offers sent to the opponent are the most similar to the previous offers sent by the opponent. The results show that when both agents employ this model, they are capable of reaching agreements that are very close to Pareto optimality. Our agent also employs similarity metrics when there is insufficient information to estimate the other agent's preferences correctly. However, there are several differences between our proposal and [34]. Our agent focuses on reaching social agreements, and it does not require the counterpart to be cooperative. In addition to it, our agent changes its behavior depending on available information

and the opponent's moves.

The Nice Tit-For-Tat [45] strategy was proposed in one of the Automated Negotiating Agents Competitions to reciprocate the opponents' behavior. Similar to our proposal, the agent is equipped with an opponent model, one based on Bayesian learning, that helps to estimate the opponent's preferences on bids. This opponent model guides the concession strategy carried out by the tit-for-tat agent. When the opponent carries out a concession, the agent reciprocates some proportion of the opponent's concession and selects the bid that is the most likely to be accepted by the opponent according to the opponent model. Our agent also employs an opponent model to drive bidding, but it is aware when the model may not be trusted due to the lack of information about the opponent. In addition to this, concessions are not detected solely based on a single pair of bids, but a window of bids is employed to ensure that the concession is a byproduct of a general and genuine change by the opponent, and it is not a byproduct of mere chance.

Ilany *et al.* [81] propose an interesting approach that combines machine learning and statistics for appropriately selecting a negotiation strategy. Two mechanisms are proposed: one based on supervised learning that selects a negotiation depending on the negotiation domain characteristics, and another based on multi-armed bandits that selects the negotiation strategy during the negotiation. The results showed that the agent was capable of over-performing agents for the previous negotiation competitions. Our agent also adapts its behavior during the negotiation. Even though the adaptive mechanism is not as sophisticated, it is engineered based on the prospective problems that may arise during the negotiation. On top of that, the main contribution of the article is not the adaptation mechanism itself, but the proposal of a social strategy that is capable of efficiently performing in an agent society. Moreover, the focus of the agent proposed in [81] is not necessarily social utility maximization, but individual utility maximization.

Another agent that is supported by a variety of negotiation strategies [82]. The authors present a multi-party negotiation agent whose decisions (e.g., bidding,

acceptance, and so on) are built by aggregating the decisions of multiple winning agents from previous ANAC competitions. While our agent is also supported by several strategies, the behavior of the agent proposed in [82] does not adapt according to the negotiation context, as the aggregation rules remain the same. In addition, the objective of our agent is that of being social, while the work presented in Güneş *et al.* [82] solely strives for utility maximization.

Katsuhide [83] proposes an adaptive agent that changes the concession speed according to previous negotiation sessions with the opponent and the current negotiation state. The strategy uses an opponent modeling mechanism that maps the opponent bids to its utility function by considering time, mean received utility, and the deviation of the received utility. Then, the agent classifies opponents according to the Thomas-Killman Conflict Mode Instrument and adapts its concession speed accordingly. If the opponent is classified as cooperative or passive, the proposed agent demonstrates a more cooperative behavior. In addition to this, the agent also employs a genetic algorithm to search for Pareto optimal bids by generating variations of the opponent's first bid. It assumes, with a high likelihood, that bids generated from the first opponent's bids are the best for the opponent. The work presented in [83] also adapts its behavior and may respond cooperatively to other agents' cooperation. However, the agent requires multiple negotiation sessions to learn about the opponent's style, while our agent is designed to work for single negotiations where we may not have information about the opponent.

Pan *et al.* [84] propose a distributed multi-objective genetic algorithm for solving bilateral negotiations. In their proposal, each agent carries out an NSGA-II genetic algorithm that employs one's utility function and the evaluation of the opponent for a subset of offers (i.e., partial disclosure on one's preferences). Once near Pareto optimal bids have been generated, the agents negotiate on which bid from the estimated Pareto frontier they should agree on. The agents proposed by Pan *et al.* are cooperative, as they disclose information about their preferences. However, the negotiation carried out on the estimated Pareto frontier does not

necessarily need to be cooperative and ends up in a social agreement. In addition to our agent aiming for social agreements, the other main difference between Pat *et al.* and in our proposal we do not assume any information disclosure in the negotiation apart from exchanged bids. Therefore, our negotiation strategy is also adequate for less cooperative scenarios.

Liu *et al.* [85] presents a negotiating agent that searches high joint utility agreements around the Nash bargaining solution. This agent also employs a frequentist opponent model, assuming that the most frequently received issue values are the most important for the opponent. In their work, these values are named as prior values. On the other hand, issues with high deviation are priors for the opponent. In their study, the agent calculates standard deviations for each of the issue values, and then it estimates the standard deviation of the aforementioned deviations, coined as the Issue Standard Deviation (ISD). ISD is used as a bridge to estimate the opponent's preferences on issues. In contrast to our work, they do not have a bid/time window approach in the opponent model, which may face the problems mentioned throughout our article. Moreover, our social agent strategy adapts its bidding and acceptance strategy based on the information available during the negotiation, while the proposal in Liu *et al.* does not.

Jonker and Aydoğan present a negotiating agent aiming not only to optimize its utility but also optimizing an outcome in combination with acceptability for human negotiators [42]. The agent determines its moves by following a set of predefined rules considering the opponent's last three moves. Similar to our work, that agent deploys an adaptive behavior-based mechanism. However, in contrast to our strategy, it has been mainly designed for human-agent negotiation settings where the negotiators have fewer rounds with limited bid exchanges.

Another related negotiating agent is proposed by M. Amini *et al.* [86]. Their agent follows an extension of the Boulware time-based concession tactic where the agent does not concede at the beginning of the negotiation while adjusting its target utility based on the highest estimated utility to be received. Unlike

our time-independent target utility update mechanism with bid windows analysis, they strictly follow a time-based concession strategy lead by the best-estimated bid for opponents given the list of acceptable bids for the agent itself. The opponent model is also implemented in a frequentist fashion, prioritizing received bid issues with the highest repetition having the least variation of values; but, as mentioned, opponents' bid histories are not evaluated in windows. On the other hand, the opponent model is designed as an ensemble of individual models that learn each opponent separately, unlike our work. They also introduce the "attention coefficient", which is used to aggregate opponent models with higher values for the more conceding opponents. Given the adaptive expected utility-based target utility setting and weighted opponent modeling components, they claim that the agent seeks to maximize social welfare.

Azaria *et al.* [87] propose a social agent aiming to provide strategic advice in repeated interactions between humans and self-interested software agents. In the proposed model, the agent has complete information about the state of the world, but the human participant may not. The challenge of advice provision is choosing proposals that have some utility for the software agent and have high probabilities of being accepted by the human counterpart. The authors propose an adaptive model capable of gradually changing the weights assigned to each participant when providing advice. The experiments carried out by the authors show how the proposed social agent outperforms behavioral economics and psychological models. Both our agent and the agent proposed in [87] are social, and they show some form of adaptability. The main difference between [87] and our proposed agent is that we focus on bilateral negotiations between software agents where multiple bids may be exchanged between involved parties and where both parties have imperfect knowledge about the world. In addition to this, software agents may exhibit different behavior to that of human beings.

8.2 Opponent Modeling Approaches

A variety of opponent modeling mechanisms have been proposed in the automated negotiation literature. Some opponent modeling mechanisms aim to provide an educated guess over the opponents' reservation value or the opponent's concession strategy, while other opponent modeling approaches take an educated guess on the opponents' preferences with respect to outcomes. This thesis is enclosed in the latter family. As far as learning techniques for opponent's preferences are concerned, two main approaches namely, probabilistic models (e.g. Bayesian) and frequency approaches, come to the forefront.

Bayesian approaches [88, 89, 57, 21] usually employ Bayes' update rule and a set of hypotheses to model the opponents' preferences. For instance, Bui *et al.* [88] propose a multi-party cooperative negotiation mechanism for the distributed meeting scheduling domain. Agents follow an iterative process that gradually partitions the negotiation space into acceptable areas by expressing their preferences on suggested partitions. In order to speed up the negotiation process, the agents employ Bayesian classifiers to learn other agents' preferences according to the information gathered from the current and past negotiations. Another example of the use of Bayesian learning in negotiation is presented by Buffett *et al.* [89]. In the proposed model, agents negotiate over a set of limited objects that can either be included or excluded from the final deal. Bayesian classifiers are employed to classify opponent's preferences into classes of preference relations over the objects in the negotiation domain. Bayesian learning was also used by Hindriks and Tykhonov [57] in order to predict the shape of the opponent's utility function (i.e, downhill, uphill and triangular), as well as the corresponding rank of issue values and issue weights. Sanchez-Anguix *et al.* [21] used Bayesian classifiers to learn the acceptability of partial offers for each team member in a negotiation team, and their opponent.

On the other hand, frequency approaches [25, 24, 26, 27, 28] usually model opponents' preferences by counting the frequency of issue values and the frequency of

changes in negotiation issues of the given bids, without considering an explicit set of initial hypotheses in mind. The most popular frequency model was introduced by the HardHeaded agent [24], whereby issue weights are updated when issue values do not change in consecutive pairs of opponent offers, and issue value weights are estimated by counting the occurrences of values in opponent's offers. A more detailed description of this model can be found in Section 2.1. In [25], the authors propose a frequency model similar to HardHeaded's frequency model. The main difference between those approaches is how they estimate the issue weights. The approach in [25] estimates the issue weights based on the relative frequency of the most offered values. Afiouni [26] adapted the classic frequency model to a real-time strategy for a video game (i.e., Civilization IV) where bilateral negotiation is used to exchange resources between parties. The frequency model showed to be applicable in real time, while also shortening the negotiation time/interactions between parties. HardHeaded agent, and thus the frequency opponent modeling, was also employed in [27] to study the efficiency of different agents in cloud computing negotiations. Furthermore, Ikarashi and Fujita proposed a weighted counting method, which aims to learn opponent's preferences by taking what time the bids made by the opponent into account [90]. That study focuses on learning from past negotiations so that the proposed approach was applied on the history of their opponent's bids in the previous negotiation.

There are also other remarkable approaches that are not classified under those two families. Aydoğan *et. al* proposed a concept based learning algorithm to figure out what offers are more likely acceptable for the opponent during the negotiation [91]. Kernel density function was used to predict issue weights of the opponent's preferences by Coehoorn and Jennings [92].

Recently, Baarslag [23, 28] showed that, despite their simplicity, frequency models tend to outperform in practice more complex approaches like Bayesian opponent models. Part of this success can be attributed to the fact that frequency approaches tend to make less assumptions about the opponent behavior, and the

fact that frequency models allow for more exploration of the negotiation space due to its quicker computation with respect to Bayesian approaches. In this thesis, we further study how to improve the efficiency of the classic frequency model by alleviating the effect of some of its assumptions.



CHAPTER IX

CONCLUSION

This thesis mainly focuses on opponent modeling, particularly opponent's preferences during the negotiation in automated bilateral negotiations. First, it poses the problems related to the classical frequentist opponent modeling approaches and introduces an extension of this model to deal with those issues. Then, inspired by recommendation systems and association rule mining practices, we advocate that negotiating agents' bid preferences can be predicted by revealing the associative patterns between issue values. Our elaborate evaluation shows that Associative Model significantly outperforms the classical frequentist model as well as our distribution based extension considering prediction performance of the opponent's preference ranking in a variety of negotiation scenarios with different domain sizes and oppositions.

Furthermore, this thesis investigates whether it is possible to design a social agent aiming to maximize the mutual utility of both sides in a bilateral negotiation. To this aim, the proposed agent follows a novel negotiation strategy by adopting a reciprocating bidding tactic with a set of complex behaviours. The evaluation results show that our agent makes agreements on the bids closest to the Nash Bargaining and Kalai point, both measures of social fairness; and still, it obtains the highest average individual utility.

APPENDIX A

ADDITIONAL RESULTS

Further round based comparative evaluations are presented as follows:

Table 20: Utility-space prediction based *number of received bids and domain opposition level* aggregated medians from Agent Group 1 Tournament with 100 rounds deadline.

Round	Metric	Kendall Tau			Pearson		
	Opposition Model	ASSOC	DBFQ	FQ	ASSOC	DBFQ	FQ
10	High	0.604	0.589	0.577	0.744	0.744	0.744
	Low	0.569	0.569	0.561	0.745	0.728	0.772
	Medium	0.478	0.478	0.473	0.653	0.643	0.671
20	High	0.604	0.592	0.579	0.753	0.744	0.744
	Low	0.569	0.569	0.565	0.759	0.737	0.764
	Medium	0.482	0.478	0.473	0.675	0.641	0.673
40	High	0.613	0.589	0.577	0.763	0.744	0.747
	Low	0.569	0.569	0.569	0.782	0.735	0.771
	Medium	0.489	0.478	0.497	0.685	0.654	0.700
60	High	0.634	0.590	0.579	0.809	0.753	0.748
	Low	0.589	0.561	0.567	0.781	0.726	0.764
	Medium	0.523	0.489	0.511	0.711	0.661	0.711
80	High	0.657	0.593	0.571	0.836	0.754	0.747
	Low	0.611	0.561	0.581	0.806	0.729	0.784
	Medium	0.581	0.527	0.523	0.739	0.688	0.721
100	High	0.668	0.582	0.577	0.852	0.757	0.767
	Low	0.594	0.560	0.585	0.799	0.725	0.777
	Medium	0.581	0.531	0.527	0.758	0.715	0.726

Table 21: Utility-space prediction based *number of received bids and domain opposition level* aggregated medians from Agent Group 2 Tournament with 100 rounds deadline.

Round	Metric	Kendall Tau			Pearson		
	Opposition Model	ASSOC	DBFQ	FQ	ASSOC	DBFQ	FQ
10	High	0.582	0.579	0.577	0.748	0.743	0.753
	Low	0.564	0.569	0.576	0.731	0.745	0.776
	Medium	0.478	0.478	0.495	0.662	0.658	0.695
20	High	0.598	0.579	0.583	0.768	0.744	0.744
	Low	0.569	0.569	0.577	0.758	0.758	0.775
	Medium	0.502	0.494	0.505	0.694	0.684	0.711
40	High	0.606	0.575	0.579	0.788	0.744	0.747
	Low	0.600	0.576	0.612	0.798	0.768	0.793
	Medium	0.558	0.541	0.538	0.739	0.717	0.732
60	High	0.628	0.585	0.599	0.811	0.751	0.770
	Low	0.611	0.569	0.611	0.799	0.767	0.796
	Medium	0.591	0.558	0.557	0.766	0.737	0.739
80	High	0.659	0.585	0.607	0.831	0.753	0.778
	Low	0.609	0.574	0.622	0.808	0.771	0.819
	Medium	0.591	0.561	0.558	0.781	0.739	0.743
100	High	0.697	0.607	0.621	0.863	0.764	0.807
	Low	0.632	0.572	0.635	0.832	0.772	0.842
	Medium	0.585	0.526	0.542	0.791	0.717	0.739

Bibliography

- [1] V. Sanchez-Anguix, V. Julian, V. Botti, and A. García-Fornes, “Tasks for agent-based negotiation teams: Analysis, review, and challenges,” *Engineering Applications of Artificial Intelligence*, vol. 26, no. 10, pp. 2480–2494, 2013.
- [2] S. Fatima, S. Kraus, and M. Wooldridge, *Principles of Automated Negotiation*. Cambridge University Press, 2014.
- [3] T. Baarslag, M. J. Hendriks, K. V. Hindriks, and C. M. Jonker, “Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques,” *Autonomous Agents and Multi-Agent Systems*, vol. 30, no. 5, pp. 849–898, 2016.
- [4] V. Sanchez-Anguix, R. Chalumuri, R. Aydoğan, and V. Julian, “A near pareto optimal approach to student–supervisor allocation with two sided preferences and workload balance,” *Applied Soft Computing*, vol. 76, pp. 1 – 15, 2019.
- [5] R. Aydoğan, P. Öztürk, and Y. Razeghi, “Negotiation for incentive driven privacy-preserving information sharing,” in *PRIMA 2017: Principles and Practice of Multi-Agent Systems* (B. An, A. Bazzan, J. Leite, S. Villata, and L. van der Torre, eds.), (Cham), pp. 486–494, Springer International Publishing, 2017.
- [6] S. Ossowski, C. Sierra, and V. Botti, “Agreement technologies: A computing perspective,” in *Agreement Technologies*, pp. 3–16, Springer, 2013.
- [7] M. Gerla, E.-K. Lee, G. Pau, and U. Lee, “Internet of vehicles: From intelligent grid to autonomous cars and vehicular clouds,” in *IEEE World Forum on Internet of Things*, pp. 241–246, 2014.
- [8] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, M. J. Wooldridge, and C. Sierra, “Automated negotiation: Prospects, methods and challenges,” *Group Decis. and Negot.*, vol. 10, pp. 199–215, 2001.
- [9] V. Sanchez-Anguix, V. Julian, V. Botti, and A. García-Fornes, “Tasks for agent-based negotiation teams: Analysis, review, and challenges,” *Eng. Appl. of Artif. Intel.*, vol. 26, no. 10, pp. 2480–2494, 2013.
- [10] M. Cao, X. Luo, X. R. Luo, and X. Dai, “Automated negotiation for e-commerce decision making: a goal deliberated agent architecture for multi-strategy selection,” *Decision Support Systems*, vol. 73, pp. 1–14, 2015.
- [11] V. Sanchez-Anguix, R. Aydoğan, V. Julian, and C. Jonker, “Unanimously acceptable agreements for negotiation teams in unpredictable domains,” *Electronic Commerce Research and Applications*, vol. 13, no. 4, pp. 243–265, 2014.
- [12] S. C. Botelho and R. Alami, “M+: a scheme for multi-robot cooperation through negotiated task allocation and achievement,” in *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No. 99CH36288C)*, vol. 2, pp. 1234–1239, IEEE, 1999.

- [13] F. Lopes, T. Rodrigues, and J. Sousa, “Negotiating bilateral contracts in a multi-agent electricity market: a case study,” in *2012 23rd International Workshop on Database and Expert Systems Applications*, pp. 326–330, IEEE, 2012.
- [14] E. De la Hoz, J. M. Gimenez-Guzman, I. Marsa-Maestre, and D. Orden, “Automated negotiation for resource assignment in wireless surveillance sensor networks,” *Sensors*, vol. 15, no. 11, pp. 29547–29568, 2015.
- [15] G. Iuhasz, V. I. Munteanu, and V. Negru, “Evolutionary approach to negotiation in game ai,” in *2013 15th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing*, pp. 296–302, IEEE, 2013.
- [16] J. Takahashi, R. Kanamori, and T. Ito, “Evaluation of automated negotiation system for changing route assignment to acquire efficient traffic flow,” in *2013 IEEE 6th International Conference on Service-Oriented Computing and Applications*, pp. 351–355, IEEE, 2013.
- [17] B. An, N. Gatti, and V. Lesser, “Extending alternating-offers bargaining in one-to-many and many-to-many settings,” in *Proc. of the 2009 IEEE/WIC/ACM Int. Conf. on Intelligent Agent Technology*, vol. 2, pp. 423–426, 2009.
- [18] R. Aydoğan, D. Festen, K. V. Hindriks, and C. M. Jonker, “Alternating offers protocols for multilateral negotiation,” in *Modern Approaches to Agent-based Complex Automated Negotiation* (K. Fujita, Q. Bai, T. Ito, M. Zhang, F. Ren, R. Aydoğan, and R. Hadfi, eds.), pp. 153–167, Springer, 2017.
- [19] X. Luo, N. R. Jennings, N. Shadbolt, H.-f. Leung, and J. H.-m. Lee, “A fuzzy constraint based model for bilateral, multi-issue negotiations in semi-competitive environments,” *Artificial Intelligence*, vol. 148, no. 1, pp. 53–102, 2003.
- [20] S. S. Fatima, M. Wooldridge, and N. R. Jennings, “An agenda-based framework for multi-issue negotiation,” *Artificial Intelligence*, vol. 152, no. 1, pp. 1–45, 2004.
- [21] V. Sanchez-Anguix, R. Aydoğan, V. Julian, and C. Jonker, “Unanimously acceptable agreements for negotiation teams in unpredictable domains,” *Electronic Commerce Research and Applications*, vol. 13, no. 4, pp. 243–265, 2014.
- [22] R. Aydoğan, T. Baarslag, K. V. Hindriks, C. M. Jonker, and P. Yolum, “Heuristics for using cp-nets in utility-based negotiation without knowing utilities,” *Knowledge and Information Systems*, vol. 45, no. 2, pp. 357–388, 2015.
- [23] T. Baarslag, M. J. Hendrikx, K. V. Hindriks, and C. M. Jonker, “Learning about the opponent in automated bilateral negotiation: A comprehensive survey of opponent modeling techniques,” *Auton. Agent Multi Agent Syst.*, vol. 30, pp. 849–898, 2016.

- [24] T. van Krimpen, D. Looije, and S. Hajizadeh, “Hardheaded,” in *Complex Automated Negotiations: Theories, Models, and Software Competitions* (T. Ito, M. Zhang, V. Robu, and T. Matsuo, eds.), pp. 223–227, Springer, 2013.
- [25] N. van Galen Last, “Agent smith: Opponent model estimation in bilateral multi-issue negotiation,” in *New Trends in Agent-Based Complex Automated Negotiations*, pp. 167–174, Springer, 2012.
- [26] E. N. Afiouni and L. J. Øvreliid, “Negotiation for strategic video games,” Master’s thesis, NTNU, 2013.
- [27] F. Alsrheed, A. El Rhalibi, M. Randles, and M. Merabti, “Intelligent agents for automated cloud computing negotiation,” in *IEEE Int. Conf. on Multimedia Computing and Systems*, pp. 1169–1174, IEEE, 2014.
- [28] T. Baarslag, “Measuring the performance of online opponent models,” in *Exploring the Strategy Space of Negotiating Agents: A Framework for Bidding, Learning and Accepting in Automated Negotiation*, pp. 111–127, Springer, 2016.
- [29] T. Payne, M. Paolucci, R. Singh, and K. Sycara, “Communicating agents in open multi agent systems,” 08 2002.
- [30] V. Sanchez-Anguix, A. Espinosa, L. Hernandez, and A. Garcia-Fornes, “Mamsy: A management tool for multi-agent systems,” in *7th International conference on practical applications of agents and multi-agent systems (PAAMS 2009)*, pp. 130–139, Springer, 2009.
- [31] T. Baarslag, R. Aydoğan, K. V. Hindriks, K. Fujita, T. Ito, and C. M. Jonker, “The automated negotiating agents competition, 2010–2015,” *AI Magazine*, vol. 36, no. 4, pp. 115–118, 2015.
- [32] H. Billhardt, V. J. Julian Inglada, J. Corchado, and A. Fernández, “Human-agent societies: challenges and issues,” *International Journal of Artificial Intelligence*, vol. 13, no. 1, pp. 28–44, 2015.
- [33] J.-H. Chen, K.-M. Chao, N. Godwin, and V.-W. Soo, “A multiple-stage cooperative negotiation,” in *IEEE International Conference on e-Technology, e-Commerce and e-Service, 2004. EEE’04. 2004*, pp. 131–138, IEEE, 2004.
- [34] G. Lai, K. Sycara, and C. Li, “A decentralized model for automated multi-attribute negotiations with incomplete information and general utility functions,” *Multiagent and Grid Systems*, vol. 4, no. 1, pp. 45–65, 2008.
- [35] V. Sánchez-Anguix, S. Valero, V. Julián, V. Botti, and A. García-Fornes, “Evolutionary-aided negotiation model for bilateral bargaining in ambient intelligence domains with complex utility functions,” *Information Sciences*, vol. 222, pp. 25–46, 2013.
- [36] S. Zheng, Y. Xiang, L. Shi, K. Liao, J. Yang, and Z. He, “A cooperative game-based negotiation strategy of electric vehicle discharging price,” in *2019 IEEE Sustainable Power and Energy Conference (iSPEC)*, pp. 2437–2442, IEEE, 2019.

- [37] P. Faratin, C. Sierra, and N. R. Jennings, “Negotiation decision functions for autonomous agents,” *Robotics and Autonomous Systems*, vol. 24, no. 3-4, pp. 159–182, 1998.
- [38] T. Baarslag, E. H. Gerding, R. Aydoğan, and m.c. schraefel, “Optimal negotiation decision functions in time-sensitive domains,” in *Proceedings - 2015 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology, WI-IAT 2015*, vol. 2, pp. 190–197, Institute of Electrical and Electronics Engineers Inc., February 2016.
- [39] C. M. Jonker, R. Aydoğan, T. Baarslag, K. Fujita, T. Ito, and K. Hindriks, “Automated negotiating agents competition (anac),” in *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [40] R. Aydoğan, O. Keskin, and U. Çakan, “Let’s negotiate with jennifer! towards a speech-based human-robot negotiation,” in *Advances in Automated Negotiations* (T. Ito, R. Aydoğan, and M. Zhang, eds.), pp. 3–16, Singapore: Springer, 2020.
- [41] T. Van Krimpen, D. Looije, and S. Hajizadeh, “Hardheaded,” in *Complex Automated Negotiations: Theories, Models, and Software Competitions*, pp. 223–227, Springer, 2013.
- [42] C. M. Jonker and R. Aydoğan, “Deniz: A robust bidding strategy for negotiation support systems,” in *Advances in Automated Negotiations* (T. Ito, M. Zhang, and R. Aydoğan, eds.), pp. 29–44, Springer, 2020.
- [43] T. Baarslag, M. Hendriks, K. Hindriks, and C. Jonker, “Predicting the performance of opponent models in automated negotiation,” in *2013 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*, vol. 2, pp. 59–66, IEEE, 2013.
- [44] O. Tunalı, R. Aydoğan, and V. Sanchez-Anguix, “Rethinking frequency opponent modeling in automated negotiation,” in *International Conference on Principles and Practice of Multi-Agent Systems*, pp. 263–279, Springer, 2017.
- [45] T. Baarslag, K. Hindriks, and C. Jonker, *A Tit for Tat Negotiation Strategy for Real-Time Bilateral Negotiations*, vol. 435, pp. 229–233. Springer, 01 2013.
- [46] R. Agrawal, T. Imieliński, and A. Swami, “Mining association rules between sets of items in large databases,” *SIGMOD Rec.*, vol. 22, no. 2, p. 207–216, 1993.
- [47] O. Tunalı and A. T. Bayrak, “Generic association rule engine optimization using item support projection,” in *2021 2nd International Conference on Computing and Data Science (CDS)*, pp. 524–529, 2021.
- [48] M. Zaki and K. Gouda, “Fast vertical mining using diffsets,” pp. 326–335, 2003.

- [49] O. Tunalı and A. T. Bayrak, “İşbirlikçi tercih Öğrenimi: Bir vaka analizi collaborative preference learning: A case study,” in *2020 4th International Symposium on Multidisciplinary Studies and Innovative Technologies (ISM-SIT)*, pp. 1–4, 2020.
- [50] H. Zhou and K. Hirasawa, “Evolving temporal association rules in recommender system,” *Neural Computing and Applications*, vol. 31, pp. 2605–2619, Jul 2019.
- [51] P. Faratin, C. Sierra, and N. R. Jennings, “Negotiation decision functions for autonomous agents,” *Robotics and Autonomous Systems*, vol. 24, no. 3-4, pp. 159–182, 1998.
- [52] S. Kawaguchi, K. Fujita, and T. Ito, “Agentk: Compromising strategy based on estimated maximum utility for automated negotiating agents,” in *New Trends in Agent-Based Complex Automated Negotiations*, pp. 137–144, Springer, 2012.
- [53] C. R. Williams, V. Robu, E. H. Gerding, and N. R. Jennings, “Iamhag-gler2011: A gaussian process regression based negotiation agent,” in *Complex Automated Negotiations: Theories, Models, and Software Competitions*, pp. 209–212, Springer, 2013.
- [54] A. Dirkzwager and M. Hendriks, “An adaptive negotiation strategy for real-time bilateral negotiations,” in *Novel Insights in Agent-based Complex Automated Negotiation*, pp. 163–170, Springer, 2014.
- [55] K. Hindriks, C. M. Jonker, S. Kraus, R. Lin, and D. Tykhonov, “Genius: negotiation environment for heterogeneous agents,” in *Proc. of the 8th Int. Conf. on Autonomous Agents and Multiagent Systems*, pp. 1397–1398, 2009.
- [56] P. Faratin, C. Sierra, and N. R. Jennings, “Using similarity criteria to make issue trade-offs in automated negotiations,” *Artificial Intelligence*, vol. 142, no. 2, pp. 205–237, 2002.
- [57] K. Hindriks and D. Tykhonov, “Opponent modelling in automated multi-issue negotiation using bayesian learning,” in *7th Int. Joint Conf. on Autonomous Agents and Multiagent Systems*, pp. 331–338, 2008.
- [58] T. Baarslag, K. Hindriks, C. Jonker, S. Kraus, and R. Lin, “The first automated negotiating agents competition (anac 2010),” in *New Trends in agent-based complex automated negotiations*, pp. 113–135, Springer, 2012.
- [59] K. Fujita, T. Ito, T. Baarslag, K. Hindriks, C. Jonker, S. Kraus, and R. Lin, “The second automated negotiating agents competition (anac2011),” in *Complex Automated Negotiations: Theories, Models, and Software Competitions*, pp. 183–197, Springer, 2013.
- [60] C. R. Williams, V. Robu, E. H. Gerding, and N. R. Jennings, “An overview of the results and insights from the third automated negotiating agents competition (anac2012),” in *Novel Insights in Agent-based Complex Automated Negotiation*, pp. 151–162, Springer, 2014.

- [61] M. B. Adar, N. Sofy, and A. Elimelech, “Gahboninho: Strategy for balancing pressure and compromise in automated negotiation,” in *Complex Automated Negotiations: Theories, Models, and Software Competitions*, pp. 205–208, Springer, 2013.
- [62] J. Hao and H.-f. Leung, “Cuhkagent: An adaptive negotiation strategy for bilateral negotiations over multiple items,” in *Novel Insights in Agent-based Complex Automated Negotiation*, pp. 171–179, Springer, 2014.
- [63] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [64] S. Bagui and P. C. Dhar, “Positive and negative association rule mining in hadoop’s mapreduce environment,” *Journal of Big Data*, vol. 6, no. 1, pp. 1–16, 2019.
- [65] Y. Gal and L. Ilany, “The fourth automated negotiation competition,” in *Next Frontier in Agent-Based Complex Automated Negotiation*, 2015.
- [66] R. Aydogan, K. Fujita, T. Baarslag, C. Jonker, and T. Ito, *ANAC 2017: Repeated Multilateral Negotiation League*, pp. 101–115. 01 2021.
- [67] R. Aydogan, K. Fujita, T. Baarslag, C. Jonker, T. Ito, and K. Hindriks, “Nac2016 - automated negotiating agents competition 2016,” 2016.
- [68] A. Mori and T. Ito, *Atlas3: A Negotiating Agent Based on Expecting Lower Limit of Concession Function*, vol. 674, pp. 169–173. 04 2017.
- [69] K. Fujita, R. Aydogan, T. Baarslag, K. Hindriks, T. Ito, and C. Jonker, *The Sixth Automated Negotiating Agents Competition (ANAC 2015)*, vol. 674, pp. 139–151. 04 2017.
- [70] V. Koeman, K. Boon, J. Oever, M. Dumitru-Guzu, and L. Stanculescu, *The Fawkes Agent—the ANAC 2013 Negotiation Contest Winner*, vol. 596, pp. 143–151. 03 2015.
- [71] S. Chen and G. Weiss, “An efficient and adaptive approach to negotiation in complex environments,” vol. 242, 01 2012.
- [72] L. Ilany and Y. Gal, “Algorithm selection in bilateral negotiation,” *Autonomous Agents and Multi-Agent Systems*, vol. 30, p. 697–723, July 2016.
- [73] E. Ioannidis, N. Varsakelis, and I. Antoniou, “Communication policies in knowledge networks,” *Physica A: Statistical Mechanics and its Applications*, vol. 492, pp. 360–374, 2018.
- [74] E. Ioannidis, N. Varsakelis, and I. Antoniou, “Experts in knowledge networks: Central positioning and intelligent selections,” *Physica A: Statistical Mechanics and its Applications*, vol. 509, pp. 890–905, 2018.
- [75] E. Ioannidis, N. Varsakelis, and I. Antoniou, “Intelligent agents in co-evolving knowledge networks,” *Mathematics*, vol. 9, no. 1, 2021.

- [76] F. M. Hafızoğlu and S. Sen, “The effects of past experience on trust in repeated human-agent teamwork,” in *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 514–522, 2018.
- [77] B. Dunin-Keplicz and R. Verbrugge, *Teamwork in multi-agent systems: A formal approach*, vol. 21. John Wiley & Sons, 2011.
- [78] A. Torreno, E. Onaindia, and O. Sapena, “Fmap: Distributed cooperative multi-agent planning,” *Applied Intelligence*, vol. 41, no. 2, pp. 606–626, 2014.
- [79] R. Nissim and R. Brafman, “Distributed heuristic forward search for multi-agent planning,” *Journal of Artificial Intelligence Research*, vol. 51, pp. 293–332, 2014.
- [80] L. Ruan, J. Chen, Q. Guo, H. Jiang, Y. Zhang, and D. Liu, “A coalition formation game approach for efficient cooperative multi-uav deployment,” *Applied Sciences*, vol. 8, no. 12, p. 2427, 2018.
- [81] L. Ilany and Y. Gal, “Algorithm selection in bilateral negotiation,” *Autonomous Agents and Multi-Agent Systems*, vol. 30, no. 4, pp. 697–723, 2016.
- [82] T. D. Güneş, E. Ardit, and R. Aydoğan, “Collective voice of experts in multilateral negotiation,” in *International Conference on Principles and Practice of Multi-Agent Systems*, pp. 450–458, Springer, 2017.
- [83] K. Fujita, “Efficient strategy adaptation for complex multi-times bilateral negotiations,” *Proceedings - IEEE 7th International Conference on Service-Oriented Computing and Applications, SOCA 2014*, pp. 207–214, 12 2014.
- [84] P. Li, X. Luo, X. Meng, C. Miao, M. He, and X. Guo, “A two-stage win-win multiattribute negotiation model: Optimization and then concession,” *Computational Intelligence*, vol. 29, 11 2013.
- [85] S. Liu, S. Okuhara, A. Moustafa, and T. Ito, “An automated negotiating agent that searches the bids around nash bargaining solution to obtain high joint utilities,” in *Annual Conference of the Japanese Society for Artificial Intelligence*, pp. 66–76, 02 2020.
- [86] M. Amini, M. Fathian, and M. Ghazanfari, “A boa-based adaptive strategy with multi-party perspective for automated multilateral negotiations,” *Applied Intelligence*, vol. 50, 09 2020.
- [87] A. Azaria, Y. Gal, S. Kraus, and C. V. Goldman, “Strategic advice provision in repeated human-agent interactions,” *Autonomous Agents and Multi-Agent Systems*, vol. 30, no. 1, pp. 4–29, 2016.
- [88] H. H. Bui, D. Kieronska, and S. Venkatesh, “Learning other agents’ preferences in multiagent negotiation,” in *Proc. of the national conference on artificial intelligence*, pp. 114–119, 1996.

- [89] S. Buffett and B. Spencer, “Learning opponents’ preferences in multi-object automated negotiation,” in *Proc. of the 7th Int. Conf. on Electronic commerce*, pp. 300–305, 2005.
- [90] M. Ikarashi and K. Fujita, “Compromising strategy using weighted counting in multi-times negotiations,” in *Proc. of the 3rd Int. Conf. on Advanced Applied Informatics*, pp. 453–458, 2014.
- [91] R. Aydoğ̃an and P. Yolum, “Ontology-based learning for negotiation,” in *IEEE/WIC/ACM Int. Conf. on Intelligent Agent Technology*, pp. 177–184, 2009.
- [92] R. M. Coehoorn and N. R. Jennings, “Learning an opponent’s preferences to make effective multi-issue negotiation tradeoffs,” in *The 6th Int. Conf. on E-Commerce*, pp. 59–68, 2004.
- [93] C. Williams, V. Robu, E. Gerding, and N. Jennings, “Using gaussian processes to optimise concession in complex negotiations against unknown opponents,” *IJCAI International Joint Conference on Artificial Intelligence*, 01 2011.

VITA

Education

- Master of Science , Computer Engineering, 2016 – 2021,
Özyeğin University
- Bachelor of Science , Economics, 2011 – 2014,
Bahçeşehir University
- Minor, Energy Management, Energy Systems Eng., 2010 – 2014,
Bahçeşehir University
- Bachelor of Science, Electrical and Electronics Eng., 2008 – 2014,
Bahçeşehir University

Academic Deployment

- Teaching Assistant, 2016 – 2018, Özyeğin University.
Assistant of Artificial Intelligence, Engineering Computation, Object Oriented Programming, Computer Programming

Publications

- Rethinking frequency opponent modeling in automated negotiation
O Tunalı, R Aydoğan, V Sanchez-Anguix. International Conference on Principles and Practice of Multi-Agent Systems, 7, 2017
- Can social agents efficiently perform in automated negotiation?
Victor Sanchez-Anguix, Okan Tunalı and Reyhan Aydoğan, and Vicente Julian, Journal of Applied Sciences, 6, 2021