Treball Final de

Màster en Modelització per a la Ciència i l'Enginyeria

Advancing Automated Negotiations in Diplomacy: Design and Implementation of an Effective Negotiation Bot

Laura Rodriguez Cima

Director
Dave De Jonge

Any
2022/23

Convocatòria
Setembre

To my colleagues,

Abstract

This thesis focuses on the field of automated negotiation, a significant research area encompassing different fields such as computer science, artificial intelligence, economics and business, motivated by the increasing complexity of modern transactions, the growing number of participants involved in negotiations, and the need for more efficient and effective decision-making. The strategic game Diplomacy provides a complex and realistic platform for examining automated negotiation techniques, as it involves multiple agents with conflicting objectives and requires both cooperative and competitive interactions. However, existing negotiation bots for Diplomacy have not demonstrated significant improvement compared to non-negotiating bots, except those utilizing NLP and reinforcement learning techniques. The present study aims to fill this gap and is a direct continuation of the paper "The Challenge of Negotiation in the Game of Diplomacy", by De Jonge *et al*, implementing and testing some of the ideas presented in there. It begins with empirical testing to confirm that negotiations enhance gameplay outcomes and proceeds to develop a proposing agent using Pareto optimal points for mutually beneficial deals, as well as different acceptor agents to test different acceptance strategies. The study culminates with the development of a comprehensive negotiation bot incorporating enhanced proposal and acceptance strategies, Attila, which is the first agent ever for the game of Diplomacy capable of negotiating successfully without being trained on human data, and distinctly outperforming non-negotiating agents.

Resum

Aquesta tesi se centra en el camp de la negociació automatitzada, una àrea de recerca significativa que abasta diversos camps com la informàtica, la intel·ligència artificial i l'economia, motivada per la creixent complexitat de les transaccions modernes, el creixent nombre de participants implicats en les negociacions, i la necessitat d'un procés presa de decisions eficient. El joc estratègic Diplomacy proporciona una plataforma complexa i realista per examinar tècniques de negociació automatitzades, ja que implica múltiples agents amb objectius contradictoris i requereix tant interaccions cooperatives com competitives. No obstant això, els bots de negociació existents per aquest joc no han demostrat una millora significativa en comparació amb els bots no negociants, excepte els que utilitzen el processament de llenguatge natural i les tècniques d'aprenentatge automàtic. L'estudi actual pretén omplir aquest buit i és una continuació directa del document "The Challenge of Negotiation in the Game of Diplomacy", de De Jonge *et al*, implementant i provant algunes de les idees presentades allà. Comença amb les proves empíriques per confirmar que les negociacions milloren els resultats dels participants i procedeix a desenvolupar un agent "proposador" que utilitza punts òptims de Pareto per a trobar acords mútuament beneficiosos, així com diferents agents "acceptadors" per provar diferents estratègies d'acceptació. L'estudi culmina amb el desenvolupament d'un bot de negociació integral que incorpora estratègies de proposta i acceptació millorades, Attila, que és el primer bot presentat per al joc Diplomacy capaç de negociar amb èxit i superar clarament als agents no negociants, sense utilittzar tècniques d'aprenentatge automàtic.

Acknowledgements

"If I have seen further, it is by standing on the shoulders of giants."

— Isaac Newton, *in a letter to Robert Hooke dated February 5, 1676, to express Newton's recognition of the contributions of previous scientists and scholars whose work paved the way for his own discoveries.*

I would like to express my sincere gratitude to all those who have supported and contributed to the completion of this master's thesis.

First and foremost, I extend my heartfelt appreciation to my thesis supervisor, Dr. Dave de Jonge, for their invaluable guidance, expertise, and continuous support throughout this research journey. Their insightful feedback, constructive criticism, and encouragement have been instrumental in shaping the direction and quality of this work.

Also, I would like to extend my thanks to the IIIA (Insitut d'Investigació en Intel·ligència Artificial) for providing the necessary resources, facilities, and research environment that have been vital in carrying out this study. The opportunity to collaborate closely with esteemed researchers has been enriching and inspiring.

I am also grateful to the developers and contributors of the Bandana framework, particularly Dr. Dave de Jonge, for their efforts in creating a powerful and accessible platform for the development of automated agents in the game of Diplomacy. The availability of this framework has been instrumental in the successful implementation and evaluation of the negotiation bot presented in this thesis.

Finally, I want to express my deepest gratitude to my family and friends for their unwavering support, love, and belief in me. Their encouragement and understanding during the challenges and demands of this thesis have been invaluable.

This research would not have been possible without the collective contributions and support of all those mentioned above and many others who have, in various ways, contributed to this work. Thank you all for being an essential part of this journey.

LRC

Contents

Abstract **ii**

Acknowledgements **v**

Contents **vii**

1 Introduction **1**
1.1 Context and justification . . . 1
1.2 Goals . . . 1
1.3 Structure of the document . . . 2

2 Background **3**
2.1 Automated Negotiations . . . 3
2.1.1 The negotiation model . . . 3
2.1.2 Knowledge About Opponent Utility . . . 5
2.1.3 Knowledge About Own Utility . . . 5
2.2 Diplomacy Basics . . . 6
2.3 D-Brane and Bandana Framework . . . 7
2.4 The Negotiation Protocol . . . 8

3 Related Work **10**

4 Approach **13**
4.1 General Description . . . 13
4.2 Experimental Evaluation . . . 13
4.3 Negotiation Domain . . . 13
4.4 General Set-up . . . 14

5 Methods and Experiments **15**
5.1 Experiment 1: Establishing a Baseline with Non-negotiating Agents . . . 15
5.1.1 Agents and Set-up . . . 15
5.1.2 Results . . . 15
5.2 Experiment 2: Proving the Power of Negotiation with Simple negotiating agents . 16
5.2.1 Agents . . . 16
5.2.2 Set-up . . . 16
5.2.3 Results . . . 17
5.3 Experiment 3: Enhancing Bidding and Accepting Strategies with Smart Negotiating Agents . . . 17
5.3.1 Agents . . . 18
5.3.2 Set-up . . . 20
5.3.3 Results . . . 21
5.4 Experiment 4: ATTILA, a Fully-Implemented Bot . . . 22
5.4.1 Introduction . . . 22
5.4.2 Attila's Accepting Strategy . . . 22
5.4.3 Attila's Bidding Strategy . . . 23

5.4.4 Set-up 24
5.4.5 Results 24

6 Conclusions **30**
6.1 Conclusions 30
6.2 Future Work 30

7 References **31**

Chapter 1

Introduction

1.1 Context and justification

Automated negotiation is the process of negotiating between two or more parties using computational methods rather than face-to-face communication. The ability to negotiate automatically has become increasingly important due to the growing complexity of modern transactions, the increasing number of participants involved in negotiations, and the need for more efficient and effective decision-making.

Diplomacy is a strategic board game that has gained popularity in the field of artificial intelligence (AI) and automated negotiation. The game involves seven players representing different European nations during World War I (Russia, Turkey, France, Germany, Italy, Austria-Hungary and England), who use diplomacy and negotiation to achieve their objectives and, ultimately, win the game. Each player must form alliances, negotiate with their opponents, and strategically move their military forces to gain control of key territories.

The game of Diplomacy is considered important for automated negotiation research because it provides a complex and realistic negotiation scenario that involves multiple agents with conflicting objectives. The game involves both cooperative and competitive interactions between the players, which requires sophisticated negotiation strategies that can be automated using AI techniques. Diplomacy is also interesting because it is a fully observable, turn-based game with perfect information, making it a suitable domain for AI research. This means that researchers can accurately model the game and test their algorithms in a controlled environment.

Despite the efforts the research community has made to advance automated negotiations in Diplomacy, only a few proficient automated Diplomacy players have emerged so far. A Diplomacy competition was held during some years as part of the Automated Negotiating Agents Competition (ANAC), and served as a crucial platform for testing and comparing Diplomacy automated agents. However, results from these competitions revealed that negotiating agents frequently failed to surpass their non-negotiating counterparts.

In their paper "The Challenge of Negotiation in the Game of Diplomacy", De Jonge *et al.* present the setup and the results of the ANAC 2017 Diplomacy Competition and ANAC 2018 Diplomacy Challenge, analyzing the participating algorithms and discussing why it is so hard to write successful negotiation algorithms for Diplomacy. This thesis is a direct continuation of this paper, actively implementing some of the ideas the authors proposed in there.

1.2 Goals

The primary goal of this study is to **design and implement a negotiation bot for Diplomacy that effectively enhances the results achieved by non-negotiating bots**.

To accomplish this main goal, we identified the following objectives:

1. Design a proposition strategy capable of determining optimal moves for the agents involved in the negotiation and proposing them.

2. Design an acceptance strategy capable of evaluating received proposals, discerning their benefits, and accepting or rejecting them accordingly.
3. Implement a bot that incorporates both the bidding and acceptance strategies, and conduct tests with a fixed coalition of two participating agents, letting the bot play as both.
4. Compare the results with those achieved by each country without negotiation, as documented in the existing literature.
5. Make necessary adjustments and repeat the testing process until the desired results are attained.

In addition to the research aims, this project also seeks to fulfill the following personal objectives:

1. Apply the knowledge acquired during my bachelor and master's degrees to practical applications.
2. Enhance my programming skills in Java, a language I have had limited experience with but find highly intriguing to explore.
3. Collaborate closely with researchers from a prestigious research institution, engaging in a project that enhances my research abilities, scientific writing skills, and other related competencies.

1.3 Structure of the document

In the subsequent chapter, *Chapter 2: Background*, we provide a comprehensive introduction to automated negotiations, the game of Diplomacy and the employed frameworks, highlighting key concepts and procedures.

Then, in *Chapter 3: Related Work*, we review the state of the art in automated negotiations, with a particular focus on the applications in the game of Diplomacy.

In *Chapter 4: Approach*, we present a general explanation of the approach followed during our research process, the implied negotiation domain and some general experimental set-up specifications.

After that, *Chapter 5: Methods and Experiments* details the agents developed, experiments conducted, and their corresponding results. This chapter concludes with the presentation of our final bot: Attila, and its obtained results.

Finally, conclusions and suggestions for future work are discussed in *Chapter 6: Conclusions*.

Chapter 2

Background

2.1 Automated Negotiations

2.1.1 The negotiation model

Negotiations are a crucial process for forging alliances and establishing trade agreements that has long been a subject of interdisciplinary research, drawing insights from economics, social science, game theory, and artificial intelligence ([1],[2]).

While traditionally negotiations have been reliant on human intuition and interpersonal skills, the introduction of automated negotiations offers a new path in decision-making processes, operating on quantifiable metrics and strategic algorithms to determine outcomes.

Central to these mechanisms lies the *negotiation model*, which delineates a set of possible outcomes, representing the various agreements or resolutions that can be achieved between the negotiating agents. To provide a quantitative measure of the desirability of each outcome, each agent's preferences are modeled by means of a *utility function* that maps every possible outcome to a numerical value, signifying its relative benefit to the agent. The higher the *utility value*, the more preferred the outcome is for the agent.

Within the negotiation algorithms, each agent, driven by self-interest, collaborates with others to attain mutually beneficial outcomes. Each agent proposes potential solutions, accepting or rejecting them based on the individual utilities. Crucially, these utilities are only gained if the solution is accepted by all agents involved. In the case that the agents do not come to an agreement before the deadline, their utility defaults to a baseline value called the *reservation value*. An intelligent agent will only agree to a proposal if the utility it gains from it matches or surpasses its reservation value. This is because, even without striking any deal, the agent is assured of this gain. Hence, in automated negotiations, the primary focus is on solutions where each agent's derived utility is at least equal to this minimum, which are termed as *individually rational*. In mathematical terms, this can be expressed as, $\forall\, a_i$:

$$u_i(s) \geq rv_i \tag{2.1}$$

Here, $u_i(s)$ is the utility that player a_i receives for the agreement s, and rv_i is the minimum payoff that agent a_i would accept. If the utility that the agent receives from the agreement s is less than this reservation value, then the agent would prefer not to participate on the negotiation.

Also, a solution is *Pareto efficient* if there is no other solution that makes at least one agent better off without making any other agent worse off. Pareto efficiency is mathematically represented when there is no other agreement s' such that:

$$u_i(s') \geq u_i(s) \text{ for every } a_i \text{ and } u_j(s') > u_j(s) \text{ for some } a_j \tag{2.2}$$

In this equation, s and s' represent agreements, u_i and u_j represent the utilities for agents a_i and a_j, respectively [3]. To illustrate what a Pareto optimal point is, Figure 2.1 shows a set of points where Pareto optimal points are bigger and highlighted in red. In this example, (3,6) and (4,5) are Pareto optimal points as no other points improve them in one dimension without being

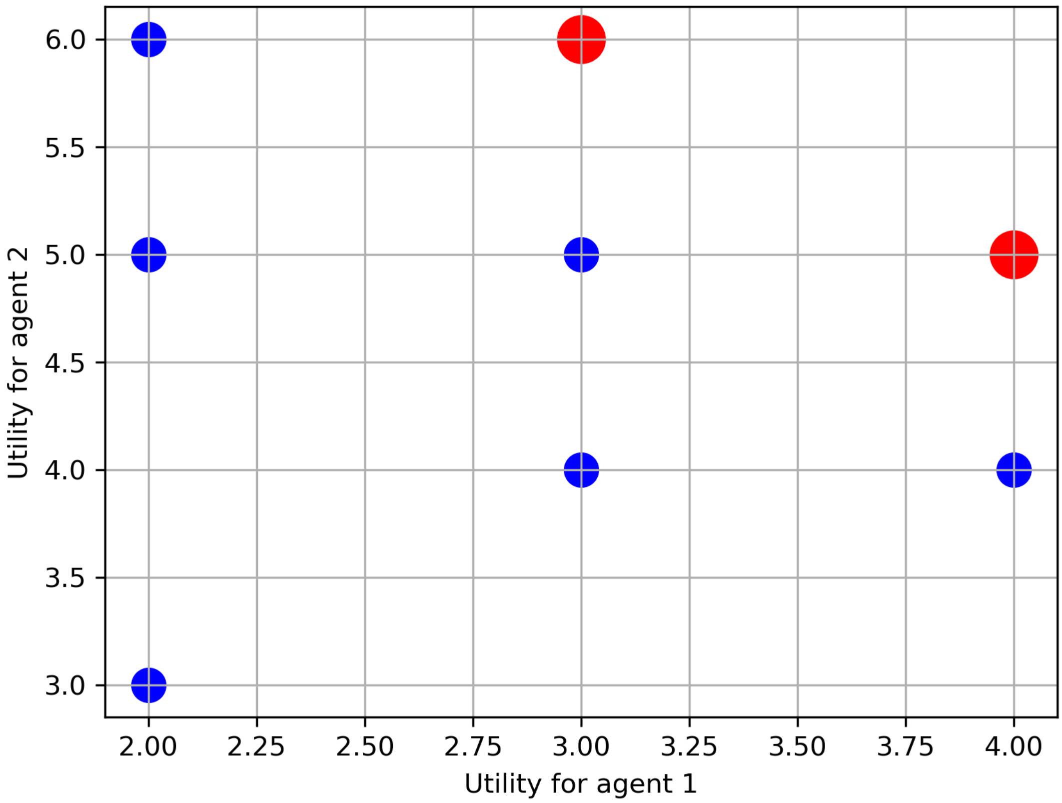

Figure 2.1: Example of a set of points highlighting Pareto optimal points.

worse in the other dimension. The other points in the set don't meet this condition: for example, (3,4) and (3,5) are dominated by (3,6), as (3,6) increases the y-coordinate without decreasing the x-coordinate. Similarly, (2,5) and (3,5) are dominated by (4,5), as (4,5) increases the x-coordinate without decreasing the y-coordinate.

The performance of the agents is guided by a defined *negotiation protocol*, which outlines the rules of how and when proposals can be exchanged. The Alternating Offers Protocol is the predominant negotiation protocol in the literature [4]. According to this protocol, one of the agents starts the negotiation by making an offer and the other agent has the choice to accept or reject the given proposal. By accepting the offer, the negotiation ends with an agreement. When rejecting the offer, the other party can either terminate the negotiation or propose a counter offer. This process proceeds with each party taking turns.

Typically, agents start negotiations with proposals that favor their interests. Over time, the proposals evolve, balancing personal interests with the need for mutual acceptance. This evolutionary aspect necessitates agents having a diverse range of proposal options, reflecting varying degrees of self-interest.

Formally, a problem instance in the field of automated negotiations (a negotiation domain) is defined as follows:

- A specified group of **agents**, such as $a_1, a_2, \ldots, a_m$.
- An **agreement space**, Ω, representing potential proposals.
- A set of **utility functions**, $U_1, U_2, \ldots, U_m$, each mapping the agreement space to real numbers, essentially $U_i : \Omega \to \mathbb{R}$.
- A set of **reservation values**, one for each agent, exemplified by $rv_1, rv_2, \ldots, rv_m \in \mathbb{R}$.

A common example of a negotiation domain involves a car sale between a car seller and a customer, bargaining over the price of a car. Here, there are two agents (the customer and the salesperson), and the agreement space include all possible combinations of a specific car c and its associated price p. Typically, the car seller would start the negotiation by making an offer with a high price, whereas the customer would begin with making an offer with a low price. They then take turns making offers, progressing until they meet somewhere in the middle, where one's offer aligns with the other's acceptance criteria.

2.1.2 Knowledge About Opponent Utility

Utility functions play a pivotal role in automated negotiations, which can be characterized by various scenarios, depending on the extent of an agent's understanding of the opponent's utility function:

1. **Ignorance Scenario:** The agent is completely unaware of the opponent's utility function.

2. **Complete Knowledge Scenario:** The agent has full knowledge of the opponent's utility function.

3. **Partial Knowledge Scenario:** The agent has fragmented or imperfect knowledge about the opponent's utility function. This can manifest in two distinct ways:

    - **3A. Insufficient Knowledge:** The agent possesses some information but lacks a holistic understanding. A good example for this can be found in the car sale example, where the car seller knows that the buyer prefers to pay a lower price, but does not know exactly which price is willing to pay.
    - **3B. Computational Complexity:** The agent theoretically has access to all the necessary information to compute the utility value. However, it is computationally too expensive.

The first scenario is the one that is most commonly studied in the literature, characterized by ignorance about the opponent's utility function. Negotiation strategies often comprise three key components, referred to as the BOA model (Bidding, Opponent-Modeling, and Acceptance [5]). While Bidding and Acceptance components are relevant to all kinds of negotiation, the Opponent-Modeling component is specific for this scenario. Central to it, there is the principle of learning through observation. As agents receive incoming proposals, they analyze them to infer and learn about the opponent's utility function, assuming it is easy to calculate through a linear function.

As De Jonge *et al.* [6], [7], we argue that Scenario 3A, or a combination of 3A and 3B, are the most realistic. Often, in real world, negotiators know only a bit about what the other side wants. They might have some clues or hints, but they usually don't see the whole picture. Also, sometimes even when they know a lot, figuring out how to use that knowledge can be very hard.

The interesting aspect of Diplomacy as a test case for Automated Negotiations is that it falls into the third scenario (particularly 3B): despite the absence of hidden information, there is no explicit formula to calculate utility values. Similar to games like Chess or Go, the complexity makes it nearly impossible to precisely compute them.

2.1.3 Knowledge About Own Utility

Another common assumption in automated negotiations is that all possible agreements and the agent's *own* utility values are given in advance. This assumption can simplify the problem of automated negotiation, but it is not always realistic. The actual complexity of real-world negotiations often lies in accurately assessing the value of a proposal, an aspect that has traditionally received less attention. In many real-world settings, for any given proposal, negotiators may need to invest considerable effort to estimate its utility values, which in several scenarios can be a highly non-trivial and time-consuming task.

2.2 Diplomacy Basics

Diplomacy is a strategic board game that involves seven players and is played on a map of Europe set in the year 1901. Similar to chess, Diplomacy is entirely deterministic, meaning there are no elements of chance or hidden information involved. The players control one of the seven great *Powers* of that time: Austria (AUS), England (ENG), France (FRA), Germany (GER), Italy (ITA), Russia (RUS), and Turkey (TUR). Each player starts with three or four units (armies or fleets) placed in fixed initial positions on the map. Figure 2.2 shows the initial configuration of a standard game.

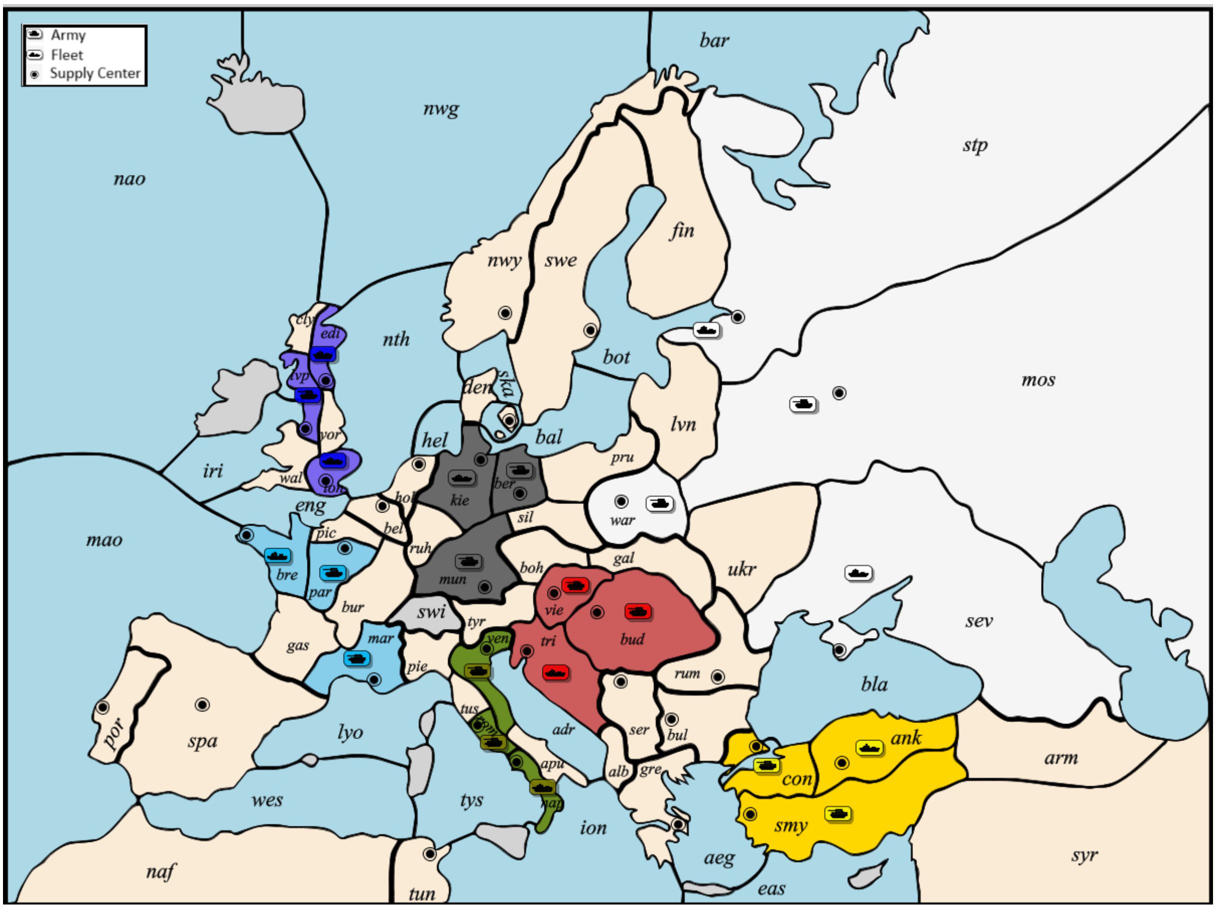

Figure 2.2: Initial configuration of a standard Diplomacy game

In the map of the game, every Power's territory is divided into regions known as *Provinces*. Some of these Provinces are designated as *Supply Centers*. The objective of the game is to conquer these Supply Centers, as players must vie for control over them. When a player loses all their Supply Centers, they are eliminated from the game. Conversely, if a player conquers 18 or more of the 34 Supply Centers, they win the game, and this is called a "Solo Victory." Alternatively, the game may end in a *draw* if all surviving players mutually agree to end it without a clear winner.

The game is divided into five types of rounds or *Phases*: Spring, Summer, Fall, Autumn, and Winter. These five phases together are called a *Year*. The phases are labeled sequentially, starting with Spring 1901, followed by Summer 1901, Fall 1901, Autumn 1901, Winter 1901, Spring 1902, and so on, indicating the game's progression.

Each turn, a player's unit can only move to an adjacent province. But movement isn't the only action units can take. They can also choose to hold or to give support to another unit that's trying to move. Therefore, units can receive support from other units, enhancing their chances of success when moving into a new province. A move is only successful if the unit, along with its supporting units, possesses greater strength than any defending unit already stationed in the target province.

Interestingly, the concept of support is not restricted to a player's own units. One player's unit can give support to another player's unit, fostering cooperative tactics and strategic alliances.

All players are required to submit their moves simultaneously, preserving the game's uncertainty. This implies players must strategize without prior knowledge of their opponents' intended moves.

In each round, before the players submit their orders, a designated time allows them to engage in private conversations so they can form alliances, coordinate movements and discuss strategies. Typical agreements arising from these conversations include mutual support pledges or designating certain provinces as *demilitarized zones*, ensuring everyone avoids invading those provinces.

De Jonge and Sierra provide a simple example in their paper [8]:

Let us focus on the three players: ENG, FRA and GER, and suppose that ENG and FRA together form a coalition. These players submit the following orders, which are all illustrated in Figure 2.3 and individually highlighted in Figure 2.4:

1. ENG moves his unit in the North Sea to Holland (Figure 2.4a).
2. FRA's unit in Belgium supports ENG's unit in the North Sea (Figure 2.4b).
3. GER moves his unit in Kiel (Figure 2.4e)to Holland (Figure 2.4c).
4. FRA moves his unit in Burgundy to Munich (Figure 2.4d).
5. GER holds with his unit in Munich (Figure 2.4e).
6. GER's unit in Silesia supports GER's own unit in Munich (Figure 2.4f).

We see here there are two battles going on: a battle for Holland and a Battle for Munich. The first two orders together form a battle plan of the coalition ENG,FRA to conquer Holland and the third order is the battle plan of GER to conquer Holland. The fourth order is a battle plan of FRA to conquer Munich, and the fifth and sixth orders form GER's battle plan to defend Munich. Although ENG and GER both try to move to Holland, only ENG will succeed, because ENG's unit has support from FRA. Furthermore, FRA is trying to expel GER 's unit from Munich, but fails, because FRA 's unit does not have any support, while GER's unit in Munich does have support (in fact, even without support GER's unit would not be expelled from Munich, because FRA and GER would have equal strength).

2.3 D-Brane and Bandana Framework

In order to develop and test our agents, we used the Bandana framework [9]. Bandana is a Java framework specifically designed for the development of automated agents that participate in the game of Diplomacy. It serves as an extension of the existing DipGame framework [10] and introduces a new game server and negotiation server, along with a simplified negotiation language. The source code for Bandana is publicly available at: https://gitlab.iiia.csic.es/davedejonge/bandana.

A key feature of Bandana is its ability to integrate with D-Brane. D-Brane itself is a highly proficient Diplomacy player that emerged as the victor of the Computer Diplomacy Challenge during the 2015 ICGA Computer Olympiad [8]. The core component of D-Brane is its *tactical module*; when provided with a proposal including a list of deals, the D-Brane tactical module attempts to identify the most favorable list of moves for the agent's units. The tactical module does not only return a set of moves but also a number that represents an (approximate) minimum number of supply centers the agent will own after executing those moves. While there is no absolute guarantee that the list returned by the D-Brane tactical module will be the theoretically optimal set of orders, in general, it delivers very good results; players can rely on its assessments to make well-informed decisions and gain a significant advantage in the game.

The Bandana framework's integration with D-Brane's tactical module allows for concentrated research on negotiations while leveraging a reliable and successful baseline player. This approach ensures that the research outcomes are not influenced by the tactical components of the game, enabling a focused analysis of the negotiation strategies implemented in the designed bot.

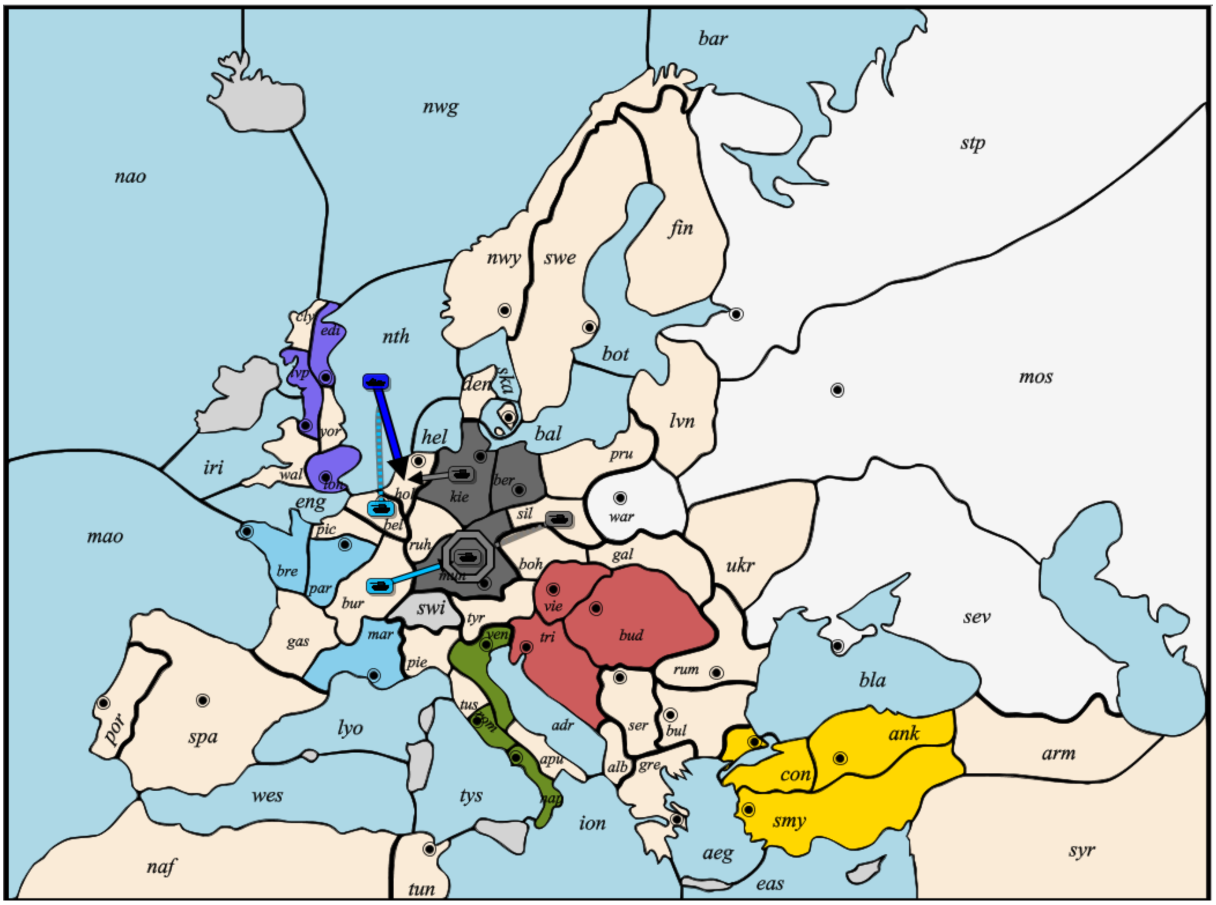

Figure 2.3: Illustration of the example given in Section 2.2, including the implied units and orders. 'Move' orders are represented by solid lines, 'Support' orders by dotted lines, and 'Hold' orders by an external octagon around the unit.

2.4 The Negotiation Protocol

The negotiation protocol used in the Bandana Framework is the Unstructured Negotiation Protocol [7], which closely resembles how negotiations take place in real Diplomacy games. In this protocol, agents are not restricted to taking turns; they can propose or accept deals at any time. A deal may involve any number of agents. Once all players involved in the deal accept it, a special Notary agent checks whether the deal is consistent with earlier agreements. If it is, the Notary sends a confirmation message to all agents involved, making the deal officially binding. Players can propose and accept multiple deals, and negotiations continue even after a deal is confirmed.

If an agent proposes or accepts a deal but later changes their mind before the Notary confirms it, they can send a reject message to withdraw from the proposal and prevent it from becoming binding. However, once the Notary confirms the deal, all agents involved must always honor it.

Each proposal is only sent to the players involved in it, ensuring that other players are unaware of the deal's details. Additionally, the Notary sends its confirmation message only to the players involved in the deal, keeping the agreement secret.

The allowed proposals in this negotiation language consist of deals comprising *Order Commitments* and *Demilitarized Zones*. An *Order Commitment* represents a promise by a power to submit a specific order during a particular phase and year. On the other hand, a *Demilitarized Zone* is an agreement among specified powers to refrain from invading or staying inside certain provinces during a designated phase and year.

A *Deal* is a non-empty set consisting of Order Commitments and Demilitarized Zones. When a deal is confirmed by the Notary, all Order Commitments and Demilitarized Zones within it must be obeyed. In addition to proposing deals with Order Commitments and Demilitarized Zones, agents can propose a "draw" to all other players. The game ends in a draw if all non-eliminated agents propose a draw in the same round of the game.

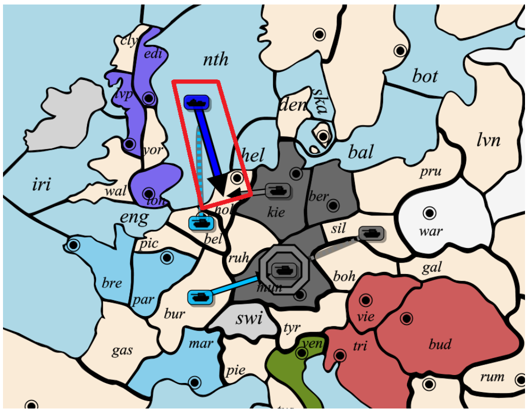

(a) ENG moves his unit in the North Sea to Holland.

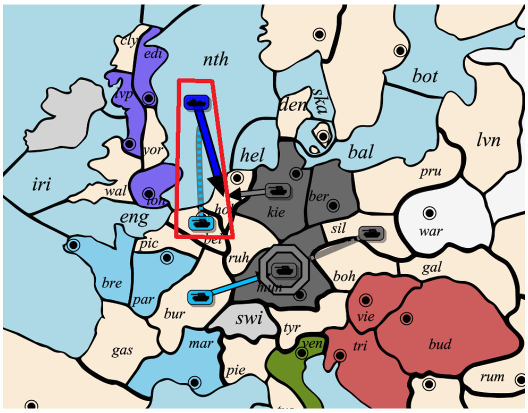

(b) FRA's unit in Belgium supports ENG's unit in the North Sea.

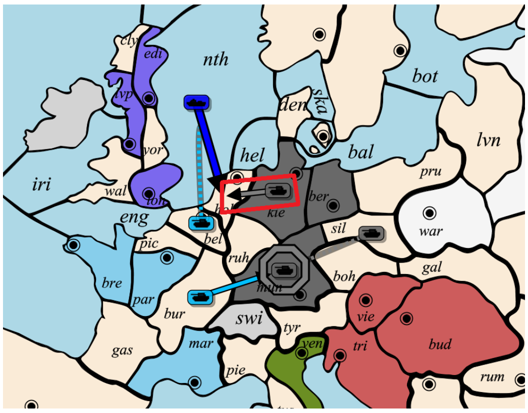

(c) GER moves his unit in Kiel.

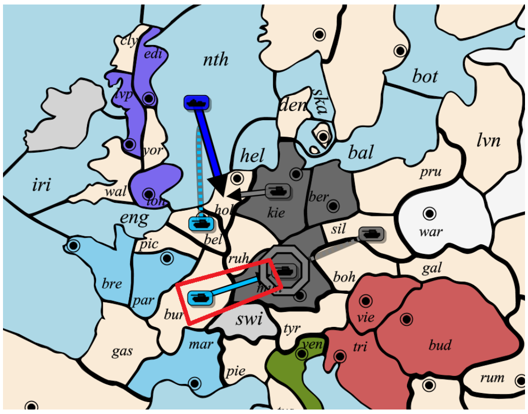

(d) FRA moves his unit in Burgundy to Munich.

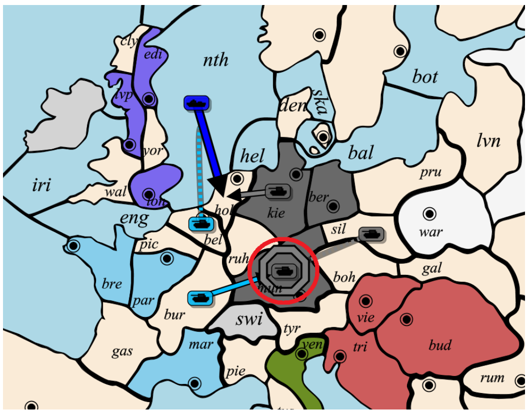

(e) GER holds with his unit in Munich.

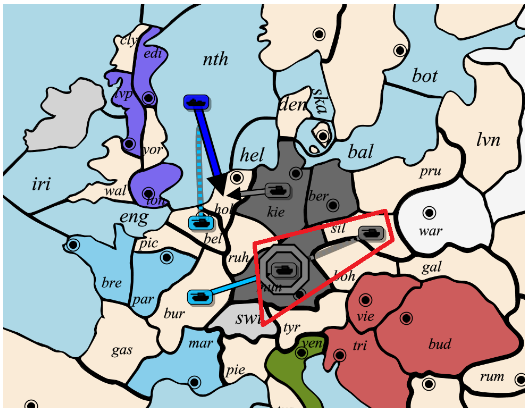

(f) GER's unit in Silesia supports GER's own unit in Munich.

Figure 2.4: Detail of the orders included in the example given in Section 2.2. 'Move' orders are represented by solid lines, 'Support' orders by dotted lines, and 'Hold' orders by an external octagon around the unit.

Chapter 3

Related Work

Automated negotiation has been widely studied and applied in various fields such as e-commerce, supply chain management, resource allocation, etc. In the early work by Jennings *et al.* [11], the authors delve into the realm of negotiation opportunities for autonomous agents. They identify, evaluate, and discuss key techniques within this domain, while also outlining major challenges for future automated negotiation research. Within their analysis, the authors also identify several properties that an ideal negotiation should meet under game theoretic techniques, such as Pareto efficiency.

Different methods have been developed for automated negotiations. For example, Buffet and Spencer [12] ventured into the application of Bayesian methods to understand an opponent's preferences during a bilateral multi-issue negotiation. Their model demonstrates how early work in the field can be built upon, marking a significant improvement in strategies that focus on understanding an opponent's preferences.

In the realm of e-commerce, automated negotiation holds a significant position, particularly in dynamic trading. While much of the existing research has concentrated on the design of negotiation protocols and strategies, Cao *et al.* [13] point out a notable gap: the insufficient attention given to multi-strategy selection in human-computer negotiation, a critical factor for improved online negotiation outcomes. To address this gap, they introduce a multi-strategy negotiating agent system, underpinned by a integration of time-dependent and behavior-dependent tactics.

In 2018, Weiss and Chen [14] delivered an extensive analysis of practical applications and approaches in automated negotiation. They underscored the challenges of large-scale, multi-issue negotiations and proposed a classification of state-of-the-art negotiation agents. Their study demonstrates how the field has evolved, leveraging the power of machine learning to manage negotiation complexities.

In a more recent practical application, Marsico *et al.* [15] applied automated negotiation concepts to networking and telecommunications. Their work introduced an approach to manage shortage scarcity through automated negotiation. This evolution of the field illustrates the real-world applicability and utility of automated negotiation.

Although automated negotiations have been widely studied, most prior research has focused on creating strategies for deciding proposals using their utility values. Yet, in real-world negotiations, the challenge often lies in accurately evaluating these values, which hasn't received as much attention. In many real situations, negotiators must invest significant effort to compute the utility values of proposals, which can be a tough and time-consuming task.

Recent literature has started acknowledging and addressing this complexity. For instance, De Jonge and Sierra [7] challenge the common assumption in automated negotiations that all possible agreements and their utility values are given in advance. The authors present NB3, a new family of negotiation algorithms designed for scenarios with many agents, large agreement spaces, non-linear utility functions, and limited time. The algorithm incorporates a heuristic Branch & Bound search to identify suitable proposals and does not rely on a mediator, assuming agents are self-interested and potentially unreliable.

Similarly, De Jonge and Zhang [16] present a new algorithm for negotiations in non-zero-sum games. Challenging the traditional non-communication assumption in game-playing algorithms, the authors argue that many real-world problems can be modeled as non-zero-sum games where players can mutually benefit from action coordination, which requires negotiation. Their proposed algorithm is completely domain-independent and can be applied to any non-zero-sum game whose rules are described in the Game Description Language. This marks a significant step towards the integration of automated negotiations into broader, more complex gaming contexts.

Transitioning to automated negotiation in Diplomacy, despite being a focal point of automated negotiations for some time, only a few proficient automated Diplomacy players have emerged so far. An online community known as DAIDE has made efforts in this direction, but the majority of developed agents lack negotiation capabilities.

Researchers have explored various approaches to developing automated negotiation agents for the game, ranging from dynamic programming to reinforcement learning. Polberg *et al.* [17] introduced an architecture of a bot designed for the Diplomacy game within the 'dip' framework [10], a testbed for multi-agent negotiations. Their proposed SillyNegoBot was an extension of an earlier bot which didn't incorporate negotiation, SillyBot, demonstrating a simple but effective method of automated negotiation. The authors report early successes with the bot, though note that it was often overly trustful and generous towards other players, highlighting the challenge of building a negotiation bot that can robustly handle adversarial and cooperative strategies in games like Diplomacy.

A few years later, DipBlue was presented as a more advanced bot that capitalized on strategic opportunities and trust reasoning [18]. Unlike previous bots that relied primarily on solution search and complex heuristics, DipBlue strategically negotiated with opponents to secure peace treaties, form alliances, and suggest actions to allies. It even incorporated a basic trust assessment mechanism to detect and respond to potential betrayals. The field continued to evolve with other bots such as M@sterMind [19]. Its negotiating algorithm, divided into deal consideration and proposal stages, involved the agent evaluating each deal and returning an acceptance probability. Subsequently, it proposed deals that were Pareto Optimal, a method that bears similarities to the approach we adopted in our own research. Despite these approaches, DipBlue's and M@stermind's results demonstrated only limited improvement when compared with non-negotiation agents.

The field continued to evolve with the introduction of other bots such as M@sterMind. Its negotiating algorithm, divided into deal consideration and proposal stages, involved the agent evaluating each component of a deal and providing an acceptance probability. Subsequently, it would propose deals that are Pareto Optimal, a method that bears similarities to the approach we adopted in our own research. Despite these advanced strategies, the results from both DipBlue and M@sterMind showed only limited improvement when compared with non-negotiating agents.

More recently, the Meta Fundamental AI Research Diplomacy Team represented a significant advancement in the field [20] using language models and strategic reasoning to create an agent capable of human-level play in Diplomacy. Notably, this bot could understand other players' motivations and perspectives and use natural language to negotiate complex shared plans, a vast improvement over previous agents. This work underscores the potential of incorporating natural language understanding and strategic reasoning in automated negotiation agents; however, it is important to note that these recent breakthroughs are based on the application of machine learning techniques to user data, while our work takes a different approach using the Bandana framework, and does not make use of any human data.

Some of the developed Diplomacy negotiation agents participated in the different editions of ANAC Diplomacy League. The Automated Negotiating Agents Competition (ANAC) is a yearly event dedicated to advancing the field of automated negotiations, having taken place annually since 2010. This event has consistently embraced increasingly complex scenarios, with problems encompassing vast agreement spaces, multilateral negotiations, human-agent interactions, and non-linear utility functions. The growing recognition of the importance of domain knowledge and reasoning ability in real-world negotiations led to the establishment of a new league within ANAC focusing on these complex negotiation elements, and the game of Diplomacy was chosen as the test case. The Diplomacy League provided a robust platform for cross-evaluation, enabling participating agents to be rigorously tested not only against non-negotiating entities, but also against each other.

De Jonge *et al.* present and analyze the results for the 2017 and 2018 editions of Diplomacy League [6], which witnessed no agent able to outperform a non-negotiating agent. As explained by Aydogan *et al.* [21], this inspired the decision in the 2019 competition to assign negotiating agents to 'Powers' that are known to collaborate well together, while maintaining the competition structure identical to previous years. Participants were required to implement a negotiation algorithm on top of the existing non-negotiating D-Brane agent, and the competition was organized into two rounds. Unfortunately, none of the five submissions could surpass the performance of D-Brane in Round 1, suggesting a lack of cooperative strategy. However, in Round 2, an agent named Oslo A, developed by Liora Zaidner *et al.*, outperformed all other agents, marking a first in the history of the ANAC Diplomacy League. Despite this achievement, the competition rules led to the declaration of Ryohei Kawata's Saitama as the winner. Until the present moment, no further editions of the ANAC Diplomacy League have been held, with 2019 being the last year in which the tournament took place.

Overall, a few attempts to build negotiating agents, such as SillyNegoBot, DipBlue, M@stermind and other participants of the ANAC Diplomacy competitions, showed limited improvement when negotiation algorithms were incorporated. Furthermore, while some agents have shown superior performance, they required substantial computational resources, making them impractical for regular use. These results underscore the pressing need to design better bots for the Diplomacy game. While progress is noticeable, there is much potential for future research to further improve upon automated negotiation strategies in this complex domain, and our work aims to make advancements in this direction.

Chapter 4

Approach

4.1 General Description

This master thesis is a direct continuation of the paper "The Challenge of Negotiation in the Game of Diplomacy" [6]. Recognizing the valuable insights and propositions from this paper, our work seeks to implement the proposed ideas and evaluate its performance.

The main idea of our approach is to use D-Brane tactical module to find proposals that are beneficial for the two agents involved in a negotiation. We do this by "pretending" to play two powers at once, and asking the D-Brane Tactical module to find the best set of moves for this "combined power".

However, while the moves returned by the D-Brane Tactical module are (ideally) pareto-optimal, they are not necessarily individually rational. This means that it is possible for one of the agents to be worse off after accepting a proposal than they would be if they did not accept it. Therefore, the main challenge of this thesis is to find proposals that are beneficial to both agents individually, by using bidding and accepting strategies that take individual rationality into account.

4.2 Experimental Evaluation

Before diving deep into negotiation strategies, we first run an experiment to evaluate the performance of the different powers when played by non-negotiation agents. The results of this experiment served as a benchmark which our negotiation agents aimed to surpass.

After that, as in the mentioned paper, we conducted an experiment to empirically validate that players in Diplomacy can indeed benefit from negotiations. Although this served as a preliminary exercise to establish an upper bound on what realistic negotiating agents could achieve, in this scenario neither of the two agents took individual rationality into account. Therefore, the results showed how negotiation improved performance for both powers as a coalition, but this improvement was not always seen in both agents.

We continued to refine our bidding and acceptance strategies, developing and testing different agents. The culmination of our efforts is our final agent: Attila. Attila ensures that no deal with a utility value below the reservation value is either proposed or accepted, and that the only deals that can be proposed or accepted are the ones that are both Pareto efficient and individually rational.

4.3 Negotiation Domain

The negotiation domain in our research is characterized by the following:

- A group of **two agents** confirming the coalition, a_1, a_2.
- An **agreement space**, Ω, representing potential proposals that can be made on each turn, consistent with the rules of the Diplomacy game and the negotiation protocol.

- **Utility values**, u_1, u_2 are approximated by the outputs from the D-Brane tactical module, as the true utilities are not known.
- A set of two **reservation values**, one for each agent, exemplified by $rv_1, rv_2 \in \mathbb{R}$, representing the utility the agents would obtain without negotiation, as determined by the D-Brane tactical module.

4.4 General Set-up

To ensure a controlled testing environment that facilitates performance comparison of various agents and negotiation strategies, all our experiments were structured around a specific coalition of two countries, which could be manually set when initiating a tournament. This structure ensured a stable baseline for performance assessment, enabling a focused analysis of the agents' behavior and negotiation tactics within a consistent context.

We also established a consistent experimental framework to ensure the reliability of our results. Each experiment comprised a *tournament*, and each tournament comprised 200 games of 40 'years' of duration (recall that in Diplomacy the five phases of the game are together called a 'year'). A negotiation deadline of 15 seconds was imposed for every phase within the game. The key metric of success in our experiments was the number of Supply Centers captured by each Power. At the conclusion of all rounds, we recorded the average number and standard error (4.1) of these Supply Centers acquired by each Power across the 200 games.

$$\text{SE} = \frac{\sigma}{\sqrt{n}} \tag{4.1}$$

In Equation (4.1), SE represents the standard error, calculated as the standard deviation (σ) divided by the square root of the sample size (n). The standard deviation, σ, can be computed using the following equation:

$$\sigma = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x})^2} \tag{4.2}$$

Here, x_i is an individual data point, $\bar{x}$ is the sample mean, and n is the sample size.

To extend the concept of standard error to coalition or a group of m agents, the collective standard error (SE_{group}) can be calculated using the following equation:

$$\text{SE}_{\text{group}} = \sqrt{(SE_{a_1})^2 + (SE_{a_2})^2 + \ldots + (SE_{a_m})^2} \tag{4.3}$$

In Equation (4.3), $SE_{a_1}, SE_{a_2}, \ldots, SE_{a_m}$ represent the standard errors for individual agents $a_1, a_2, \ldots, a_m$, respectively. These values are computed using Equation (4.1).

Finally, it is important to note that there is a total of 7 powers in Diplomacy, resulting in 21 potential coalitions between pairs of countries, as given by Equation (4.4):

$$C(7,2) = \frac{7!}{2! \times (7-2)!} = 21 \tag{4.4}$$

By using the specified set-up, each coalition test required approximately 17 hours, which made it impossible to conduct all the experiments with all the coalitions. Therefore, in some cases, we chose to run experiments only for the most important coalitions.

Chapter 5

Methods and Experiments

5.1 Experiment 1: Establishing a Baseline with Non-negotiating Agents

To establish a benchmark for evaluating the efficacy of our negotiation algorithms, we began our research with a baseline experiment in which we assessed the performance of each Power when playing without negotiation.

The results provided us with a standard performance level that our negotiation algorithms would have to surpass. It let us measure how good our negotiating agents were when compared to agents that didn't negotiate. So, our aim was not just to get a good result from negotiation, but to make sure this result was better than this baseline set by non-negotiating agents.

5.1.1 Agents and Set-up

In order to test the individual performance of the Powers without negotiation, we used a non-negotiating agent provided in the Bandana Framework.

In this setup, we let seven instances of the non-negotiating agent compete against each other in a big tournament comprising 2000 games. At the conclusion, we calculated the average number of Supply Centers (SC) and the standard error for each Power.

5.1.2 Results

The results from Experiment 1 are detailed in Table 5.1. Each entry shows the mean number of Supply Centers (SC) secured by a specific Power across all tournaments. It's evident that certain Powers tend to outperform others. For instance, the Power "RUS" achieved the highest average with 8.4 SCs, while "ENG" recorded the lowest average at 2.0 SCs.

Power	**SC (avg $\pm$ SE)**
RUS	8.4 ± 0.23
TUR	6.9 ± 0.14
GER	6.1 ± 0.19
FRA	4.4 ± 0.15
AUS	3.6 ± 0.14
ITA	2.5 ± 0.09
ENG	2.0 ± 0.05

Table 5.1: Results for Experiment 1 in terms of average ($\pm$ standard error) SC per Power

5.2 Experiment 2: Proving the Power of Negotiation with Simple negotiating agents

Upon reviewing the outcomes of several Diplomacy Competitions and observing the unsuccessful attempts to design a bot capable of surpassing non-negotiation agents, one might question the true value of negotiation for Diplomacy players. While human players naturally perceive the benefits of negotiation, we sought to substantiate this claim with robust scientific evidence.

In their paper, De Jonge *et al.* performed an experiment in which one single agent was used to play two powers at a time [6]. Here, we try to emulate a comparable experiment using two dedicated components: the SimpleProposer, which is responsible for making the proposals, and the SimpleAcceptor, that simply and blindly accepts any received proposal. Therefore, the SimpleAcceptor itself does not really form an autonomous agent, and they both play together as one single agent.

Much like the previous experiment provided us with a lower bound that our negotiating agents aim to surpass, the results of this experiment served to establish an upper bound to what realistic negotiating agents could achieve.

5.2.1 Agents

SimpleProposer

The SimpleProposer focuses on the bidding strategy. It applies the main idea of our research in its simplest form: it "pretends" to play the two powers of the coalition as one, by simulating an altered game state where both powers are merged into a "combined power" that controls all the units. Then, it uses D-Brane tactical module to find the best moves for this combined power. Finally, the agent proposes the set of moves returned by the D-Brane tactical module to the SimpleAcceptor.

SimpleAcceptor

In contrast to the SimpleProposer, the SimpleAcceptor bot is designed with minimalistic principles. Its primary function is to serve as a passive receiver, accepting any incoming proposal without making any of its own.

5.2.2 Set-up

In Experiment 2 we paired SimpleProposer with SimpleAcceptor in a coalition, while the remaining five powers were controlled by the non-negotiation agent utilized in Experiment 1.

The procedure followed SimpleProposer identifying and proposing the optimal set of moves for them both as a coalition, and SimpleAcceptor accepting all the proposed deals. If the coalition's cumulative score (in terms of average number of SC) exceeds the sum of their individual scores, it would provide empirical evidence supporting the beneficial role of negotiation for players in Diplomacy.

Considering time constraints and the impracticality of testing all 21 possible coalition combinations, we strategically focused on three specific coalitions: *TUR + RUS*, *FRA + GER* and *GER + RUS*. This coalitions were chosen because of the notable cooperation benefits observed in De Jonge et al.'s experiment [6].

Coalition	**SimpleProposer**	**SimpleAcceptor**
TUR + RUS	TUR	RUS
FRA + GER	FRA	GER
GER + RUS	GER	RUS

Table 5.2: Configuration of Experiment 2

5.2.3 Results

Coalition	Exp. 1	Exp. 2
TUR + RUS	**15.3 ± 0.27**	**20.8 ± 0.6**
TUR	6.9 ± 0.14	16.9 ±0.6
RUS	8.4 ± 0.23	3.9 ± 0.3
FRA + GER	**10.5 ± 0.24**	**13.2 ± 0.6**
FRA	4.4 ± 0.15	6.7 ± 0.4
GER	6.1 ± 0.19	6.5 ± 0.5
GER + RUS	**14.5 ± 0.30**	**18.0 ± 0.6**
GER	6.1 ± 0.19	6.1 ± 0.4
RUS	8.4 ± 0.23	11.9 ± 0.4

Table 5.3: Results for Exps 1 and 2, as a coalition and individually, in terms of SC (avg±SE)

Table 5.3 summarizes performance from Experiments 1 and 2, facilitating a direct comparison between non-negotiating agents (Exp 1) and basic negotiating agents (Exp 2). Green highlights indicate cases where negotiation outperformed the absence of negotiation, while red highlights indicate cases where negotiation did not surpass non-negotiation results.

The key insights derived from this results are as follows:

1. **Negotiation Impact**: Comparing Experiment 1 (baseline) to Experiment 2, it is evident that the introduction of negotiation produced a notable improvement in the coalition's performance. The average SC for every coalition exhibited a positive bump, attesting to the merit of cooperative play. For instance, the TUR + RUS coalition surged from 15.3 to 20.8 SC, manifesting the potential of cooperation.

2. **Imbalanced Gains**: However, a closer scrutiny reveals disparities within these coalitions. Take TUR + RUS in Experiment 2: while Turkey's individual score almost doubled from its baseline, Russia's score substantially reduced. This indicates that while the coalition as a whole benefit, individual gains were not always equitable, showcasing that simplistic negotiation strategies might cause imbalances in benefits.

It's important to note that, although this experiment provided us with interesting results for our research, this set-up is not realistic because it assumes that the acceptor unconditionally trusts the proposer, accepting everything he proposes without analyzing if the suggested moves are aligned with its interests. Moreover, there is another important limitation: the moves returned by D-Brane tactical module are (ideally) Pareto efficient, but there is no guarantee that they are individually rational. Therefore, in this set-up, neither of the two agents was taking individual rationality into consideration.

5.3 Experiment 3: Enhancing Bidding and Accepting Strategies with Smart Negotiating Agents

Upon establishing the advantages of negotiation in the previous experiment, we advanced to the next phase of our investigation by starting to explore the efficiency of agents powered with more sophisticated bidding and accepting strategies.

SimpleProposer's approach offers a viable starting point for proposing strategies. However, it has limitations in its current form. It presents only a single deal: the best deal as determined by the D-Brane tactical module for the given coalition. Once this proposal is made, regardless of acceptance or rejection, no further proposals are forthcoming. This rigid approach becomes a drawback when interacting with agents other than the SimpleAcceptor, as the acceptance criteria of the agent may lead to rejections of the deals and, consequently, a premature termination of negotiations.

In response to these limitations, we developed the SmartProposer agent, which has the capability to propose a new deal if the event of a rejection, and also ensures that proposed agreements are mutually beneficial by employing Pareto optimal points. Through the introduction of this agent, we aimed to create a dynamic and adaptable bidding strategy that could navigate rejections and continue to work towards mutually beneficial outcomes.

Similarly, while the SimpleAcceptor agent serves as a great test subject for examining the quality of the proposals of other agents, it is important to note that it's behaviour is not realistic. Unlike a real player who would assess proposed deals for personal benefit, the SimpleAcceptor unquestioningly trusts the proposer. In real negotiations, participants would evaluate offers based on their own interests, accepting or rejecting them accordingly. Therefore, when designing our final bot, it becomes necessary to develop an accepting strategy that aligns with these more nuanced and practical negotiation considerations. To address this need, we developed two different accepting agents, SmartAcceptor1 and SmartAcceptor2, each adhering to its own distinct yet simple strategy, to start exploring the influence of acceptance criteria on negotiation outcomes. These agents do not make any proposals of their own; they solely evaluate incoming proposals and decide whether to accept or reject them.

Both SmartAcceptor1 and SmartAcceptor2 take into account some considerations that are common in the negotiation strategies within this domain:

1. Accepting a deal from a stronger opponent might give them even more advantage in the game, increasing their probability of winning the game —the goal that all players share.

2. If the opponent becomes too strong, their need for collaboration with SmartAcceptor decreases, reducing their motivation to extend support in return.

By developing SmartAcceptor1 and SmartAcceptor2 and combining them with a proposer agent, we were able to track the impact of the different acceptance strategies. These were the first steps to equip our final agent with a robust accepting strategy that was designed to protect its interests while fostering beneficial negotiations in the game.

5.3.1 Agents

SmartProposer

SmartProposer expands the bidding strategy through the following steps:

1. SmartProposer, like its predecessor, simulates a new game state that combines both powers into a single entity. Using D- Brane Tactical Module, it identifies an optimal plan of moves in this adjusted context.

2. It begins computing alternative plans through systematic move elimination (detailed explanation below).

3. It computes the benefits for itself and its coalition ally under all the plans and selects combinations of benefits that represent Pareto optimal points.

4. It organizes these plans according to its own benefit (from most to least beneficial).

5. It starts with proposing the most beneficial plan for itself and continues to propose other plans in descending order of benefit if the previous proposal is rejected.

To illustrate how SmartProposer computes the alternative plans (Step 2), consider a hypothetical scenario where the initial optimal plan entails five moves: A, B, C, D and E. Moves A, B, C consist of orders to move units to an adjacent country, while move D consist of a supporting order to move A, and move E consist of a supporting order to move C.

In this scenario, SmartProposer iterates through all the moves in the initial optimal plan, as returned by the D-Brane tactical module. During each iteration, the algorithm generates a copy of this initial optimal plan and removes a particular move, along with its supporting orders. For example, the first iteration would commence by creating a copy of the initial optimal plan, then removing move A and its supporting move, D. Following their removal, SmartProposer constructs a new plan based on the remaining moves—B, C, and E in this case. Once this new plan has been formulated, the algorithm advances to the next iteration. Here, another move and its corresponding supporting orders are removed for subsequent plan re-evaluation. This iterative process continues until all possible combinations of moves have been exhaustively explored. The pseudocode illustrating this process is presented in Algorithm 1.

Algorithm 1 Generate Alternative Plans with SmartProposer

Set initialOptimalPlan to $[A, B, C, D, E]$
Set supportingMoves to $\{A : D, C : E\}$
procedure GENERATEALTERNATIVEPLANS(initialOptimalPlan, supportingMoves)
    alternativePlans ← []
    for each move in initialOptimalPlan **do**
        copyOfInitialPlan ← COPY(initialOptimalPlan)
        REMOVE move FROM copyOfInitialPlan
        if move in supportingMoves **then**
            REMOVE supportingMoves[move] FROM copyOfInitialPlan
        end if
        APPEND copyOfInitialPlan TO alternativePlans
    end for
    return alternativePlans
end procedure=0

Therefore, for the given example, the returned *alternativePlans* would include these plans:

1. [B, C, E]
2. [A, C, D, E]
3. [A, B, D]
4. [A, B, C, E]
5. [A, B, C, D]

SmartAcceptor1

SmartAcceptor1 evaluates incoming proposals and accepts them based on the following criteria: if, before the players make their moves, he owns at least 5 more Supply Centers (SCs) than the proposing agent, it automatically accepts the proposal. This reflects a strategic decision to cooperate when in a position of relative strength. Conversely, if the proposing agent is too strong, SmartAcceptor1 exercises more caution in accepting the proposal. This caution exists because of the considerations mentioned before.

Nevertheless, SmartAcceptor1 may still choose to collaborate with a stronger proposing agent under certain conditions. Specifically, it will accept the proposal only if it has received an equivalent or higher level of support from the proposing agent, taking into account the game history or the specifics of the current deal. This mechanism ensures that acceptance is predicated on fair reciprocity. SmartAcceptor1's acceptance strategy is illustrated in Figure 5.1, which presents a decision tree mapping the sequence of decisions made by this bot.

SmartAcceptor2

SmartAcceptor2 operates with a slightly different logic compared to SmartAcceptor1. Initially, it evaluates whether the proposed deal will result in a gain in Supply Centers (SCs) for itself. If the proposal meets this criterion, SmartAcceptor2 is inclined to accept it. This aligns with a strategy focused on ensure immediate gains. However, for the same reasons as before, caution is exercised if the proposing agent owns more than three times the SCs that SmartAcceptor2 holds.

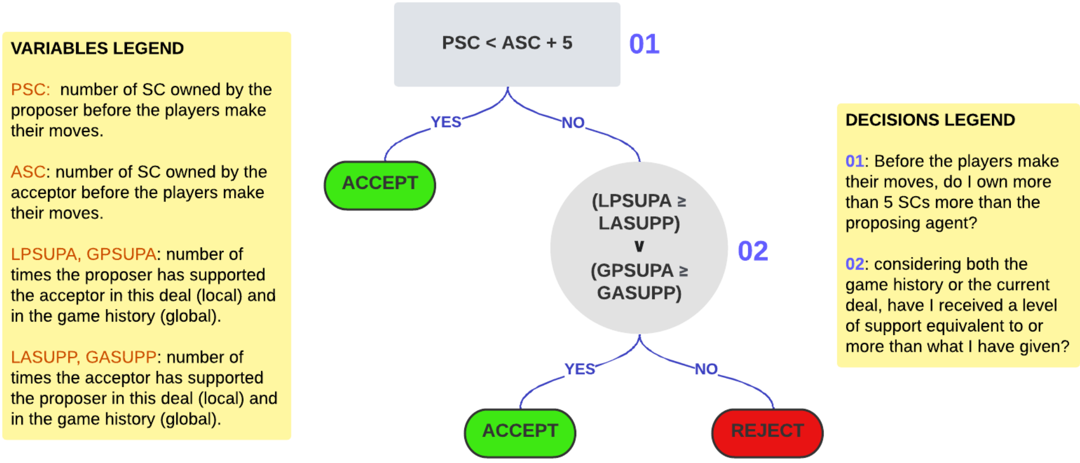

Figure 5.1: SmartAcceptor1's acceptance strategy

When the proposed deal does not result in a gain of SCs, SmartAcceptor2 becomes more selective. It will only accept the proposal if the proposing agent does not own more than three times the SCs it does and if there is a balanced exchange of support. A before, this mechanism ensures that acceptance is predicated on fair reciprocity. SmartAcceptor2's acceptance strategy is elaborated in Figure 5.2.

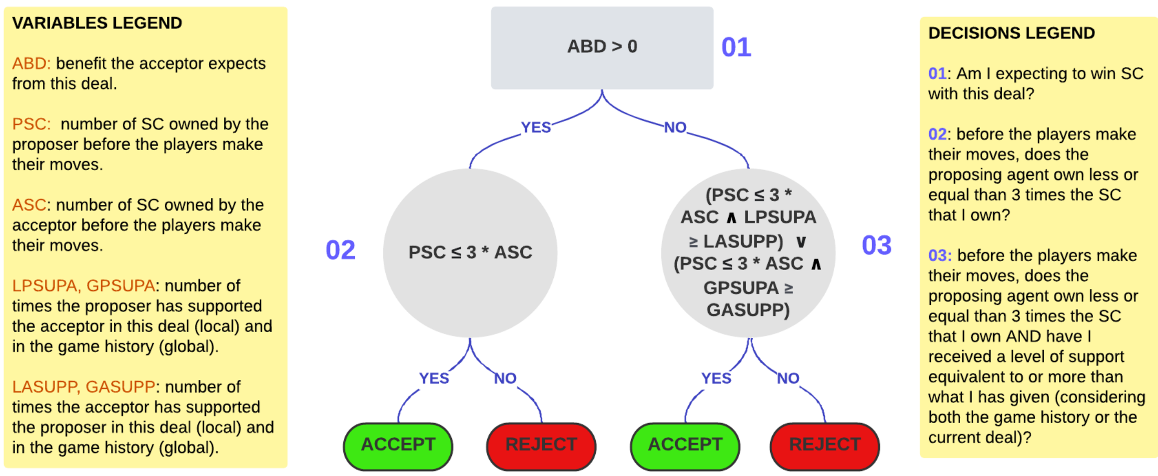

Figure 5.2: SmartAcceptor2's acceptance strategy

5.3.2 Set-up

Experiment 3.1

In the initial stage of Experiment 3, we formed a coalition comprising SmartProposer and SmartAcceptor1, while the remaining five powers were controlled by the non-negotiation agent. This combination was evaluated in three tournaments, across the same three coalitions we used in Experiment 2.

Coalition	SmartProposer	SmartAcceptor1
TUR + RUS	TUR	RUS
FRA + GER	FRA	GER
GER + RUS	GER	RUS

Table 5.4: Configuration of Experiment 3.1

Experiment 3.2

Following the preceding phase, we paired the SmartProposer with SmartAcceptor2 for another set of evaluations, also across the same three coalitions and with the remaining five powers controlled by the non-negotiation agent.

Coalition	SmartProposer	SmartAcceptor2
TUR + RUS	TUR	RUS
FRA + GER	FRA	GER
GER + RUS	GER	RUS

Table 5.5: Configuration of Experiment 3.2

5.3.3 Results

As before, results of Experiments 3.1 and 3.2 are shown next to previous results of Experiments 1 and 2 (Table 5.6), enabling a direct comparison between non-negotiation (Exp 1), basic negotiation (Exp 2), and enhanced negotiation strategies (Exp 3). Green highlights indicate cases where negotiation outperformed the absence of negotiation, while red highlights indicate cases where negotiation did not surpass non-negotiation results.

Coalition	Exp. 1	Exp. 2	Exp. 3.1	Exp. 3.2
TUR + RUS	**15.3 ± 0.27**	**20.8 ± 0.6**	**20.3 ± 0.6**	**20.4 ± 0.5**
TUR	6.9 ± 0.14	16.9 ±0.6	16.2 ± 0.5	7.6 ± 0.3
RUS	8.4 ± 0.23	3.9 ± 0.3	4.1 ±0.3	12.8 ± 0.4
FRA + GER	**10.5 ± 0.24**	**13.2 ± 0.6**	**13.0 ± 0.6**	**11.7 ± 0.6**
FRA	4.4 ± 0.15	6.7 ± 0.4	7.6 ± 0.4	6.3 ± 0.3
GER	6.1 ± 0.19	6.5 ± 0.5	5.4 ± 0.5	5.4 ± 0.5
GER + RUS	**14.5 ± 0.30**	**18.0 ± 0.6**	**20.2 ± 0.5**	**20.2 ± 0.5**
GER	6.1 ± 0.19	6.1 ± 0.4	8.4 ± 0.4	8.2 ± 0.4
RUS	8.4 ± 0.23	11.9 ± 0.4	11.8 ± 0.4	12.0 ± 0.4

Table 5.6: Results for Experiments 1, 2 and 3, as a coalition and individually, in terms of SC (avg)

When we proceed to Experiments 3.1 and 3.2, it can be seen that our slightly improved bidding and accepting strategies yield mixed results. Overall, they seem to offer a marginal or no improvement in total SC for coalitions. However, in some scenarios, such as in the coalition TUR + RUS, they help harmonize the distribution of gains. Russia's individual score increased, reflecting a more balanced negotiation outcome, which suggests that more advanced strategies can, at times, provide equitable benefits.

Despite these potential improvements shown in Experiment 3, the advantages of our improved strategies weren't consistently significant for all coalitions and players. This was due to the fact that while the bidding strategy had been enhanced to propose alternative deals when receiving a rejection, it still didn't guarantee individually rational proposals.

5.4 Experiment 4: ATTILA, a Fully-Implemented Bot

5.4.1 Introduction

Following our exploration of negotiation enhancements and the potential of various strategies, we developed our culminating agent: Attila. This bot incorporates both bidding and acceptance mechanisms, surpassing its predecessors with a more sophisticated negotiation strategy that ensures that proposed and accepted deals are both Pareto efficient and individually rational. Our methodical experimentation with diverse strategies and parameters led us to the optimal approach that consistently delivered superior results.

Each turn, before delving into proposing or accepting deals, Attila uses the D-Brane tactical module to compute two pivotal values to inform its negotiation strategy:

- The anticipated outcome (expressed in terms of final SC count) if Attila chooses to abstain from negotiations. This serves as the reservation value.
- The maximum attainable outcome (again in final SC count) for the current turn through negotiation.

5.4.2 Attila's Accepting Strategy

Central to Attila's negotiation strategy is the boolean method *isDealAcceptable()*, dictating whether a proposed deal should be accepted or not. Its decision tree is shown in Figure 5.3.

One of the key-points in Attila's acceptance strategy is that, contrary to the previous SmartAcceptors, it considers the reservation values and employs them to refine its strategy. These reservation values are simulated by consulting the D-Brane tactical module to determine the best plan and outcome in the current state of the game if Attila chooses not to negotiate.

"One of the key points in Attila's acceptance strategy is that, contrary to previous SmartAcceptors, it considers reservation values and employs them to refine its strategy. These reservation values are simulated by consulting the D-Brane tactical module to determine the best plan and outcome in the current state of the game, should Attila choose not to negotiate."

Initially, Attila assesses if the deal's output is below the reservation value. If so, the deal is instantly rejected. However, if the outcome matches or surpasses the reservation value and Attila predicts a post-deal position that's either equal to or stronger than the proposing agent's, the deal is accepted.

When Attila anticipates a post-deal disadvantage, it employs a more meticulous scrutiny. It directly rejects deals that don't offer the maximum attainable outcome for that turn. If a deal does match this maximum, Attila accepts it only if the outcome exceeds the reservation value. However, if the outcome just meets the reservation value, Attila agrees to the deal only if the other party doesn't end up with 15 or more SCs more than Attila.

The driving logic behind Attila's strategies mirrors many complexities inherent to human negotiations. Just as individuals in real-world negotiations balance the desire to form mutually beneficial agreements with the necessity of guarding their interests, so does Attila. While maximizing the number of agreements takes advantage of the proven benefits of negotiation, it's essential not to compromise one's position, especially considering partners may have conflicting end goals.

Hence, deals that perform worse than going solo (without negotiations) are immediately rejected. If Attila projects strength equal to or surpassing its ally, any deal surpassing the initial threshold is accepted, emphasizing the value of continued collaboration. However, when weaker, Attila exercises prudence, balancing its gains against those of its counterpart and ensuring its strategic interests remain safeguarded. This can lead Attila to reject deals even when their outcomes match or exceed the reservation value, especially if the proposer is much stronger than him. This mechanism serves to force the opponent to offer more equitable deals. For instance, consider a scenario where two agents must decide how to divide 10$, knowing that failure to agree results in the loss of the money (the reservation value equals 0 for both agents). If the first agent proposes keeping 9$ and giving 1$ to the second, the proposal may exceed the second agent's reservation value, but still lack fairness. For the second agent, accepting this proposal would mean missing out on the chance for more beneficial proposals.

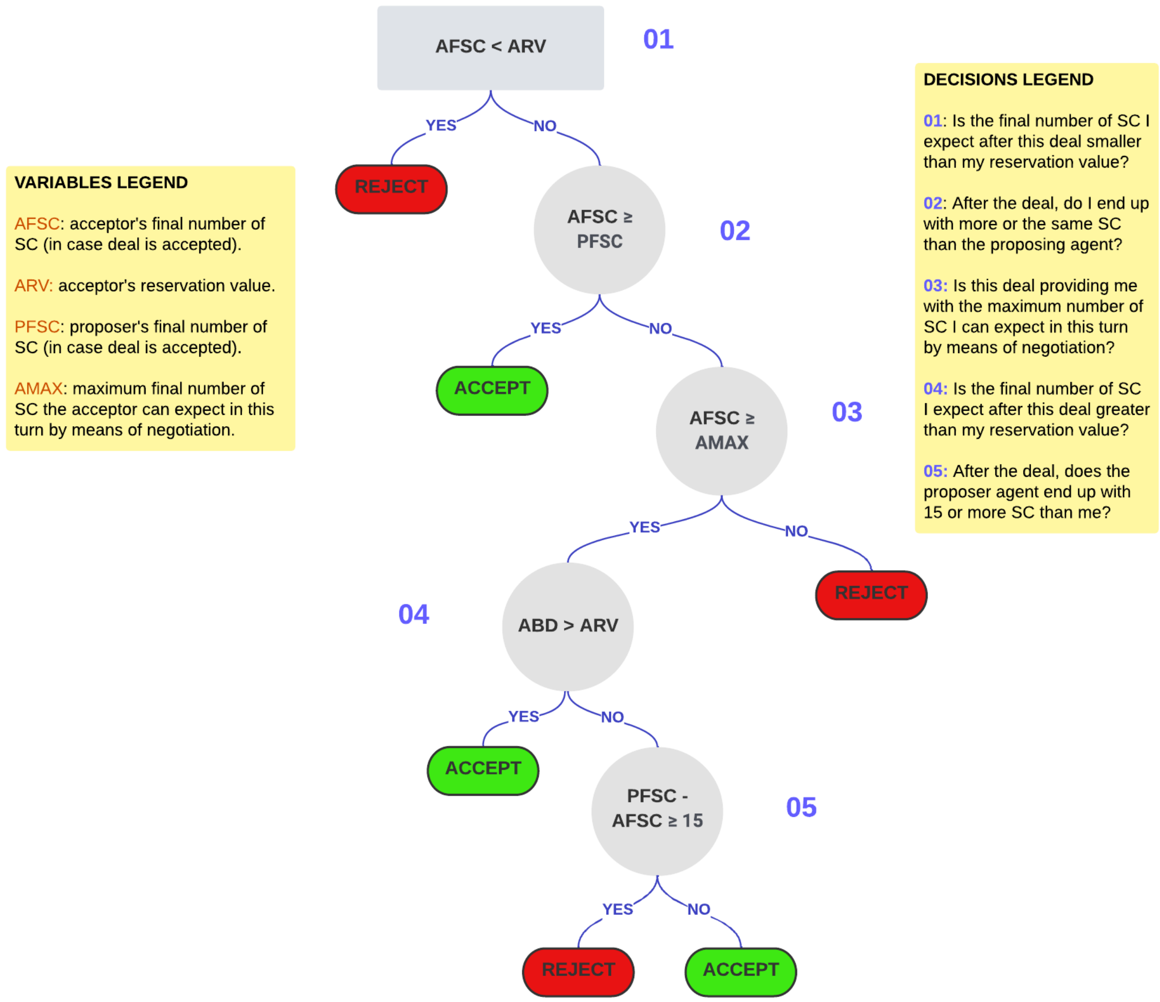

Figure 5.3: Attila's acceptance strategy

5.4.3 Attila's Bidding Strategy

Attila's bidding strategy is heavily influenced by the SmartProposer agent. As before, it starts by performing the following steps:

1. It uses the method *generateNewGame()* to simulate a joint game state, merging both powers.
2. It uses the method *searchForNewDealsToPropose()* to identify the best deals in this new context, and sorts them based on the benefit to itself.

However, after that, Attila differentiates itself in a significant manner. It employs the aforementioned *isDealAcceptable()* method to evaluate the potential deals (those returned by the method *searchForNewDealToPropose()*). Instead of proposing all identified deals, Attila uses *isDealAcceptable()* to assess the desirability of deals within the current context of the game. This avoids proposing deals that, even being the best available in a given turn, might not align with Attila's current interests. This introduced step is pivotal; it would be counter-intuitive for Attila to suggest deals that are not inside its acceptance criteria, even if they are the best it can find on this context, for example when they don't meet the reservation value and fail to outperform non-negotiated outcomes.

By incorporating the *isDealAcceptable()* function to guide both deal proposals and acceptances, Attila guarantees that all proposed and accepted deals are both Pareto efficient and individually rational.

5.4.4 Set-up

Experiment 4.1

Our research journey proceeded to our main experiment, the combination of two instances of Attila as a coalition, with the remaining five powers controlled by the non-negotiation agent.

In this case, we evaluated the performance across all 21 potential coalitions, running a tournament for each one, to ensure the reliability of our findings and conclusions.

Experiment 4.2

Finally, we ran a last experiment motivated by a question that came to our mind. As Attila makes decisions based on the information from the D-Brane tactical module when asked to find the best moves for both itself and its ally (coalition partner), we started wondering about Attila's success against non-negotiating agents. We questioned whether this success could be attributed mainly to the coalition partners avoiding to attack each other. Therefore, we wanted to ensure that our negotiation strategies genuinely contributed to better results, beyond the advantage of coalition partnership avoiding to conquer other's territories.

With this aim, we conducted an experiment which involved two instances of Attila that were playing without negotiation (we added an option to disable negotiation when running a tournament if desired), but which were taking its coalition partner as an ally when looking for the possible own moves, avoiding to conquer its territories. As before, the remaining five powers were controlled by the non-negotiation agent and this was evaluated across all 21 potential coalitions.

5.4.5 Results

Table 5.7 presents the results for all 21 coalitions from experiments with both Negotiating Attila (Exp 4.1) and Non-negotiating Attila (Exp 4.2), alongside the results obtained by Non-negotiating agents (Exp 1), for an easy comparison. Results surpassing those achieved by non-negotiating agents are highlighted in green, with a distinction made between the outcomes: lighter green indicates inferior results between Negotiating and Non-Negotiating Attila, while darker green signifies superior performance. Conversely, results falling short of those achieved by non-negotiating agents are marked in red.

To facilitate the visualization of the results, these outcomes are also shown in Figure 5.4, which presents a plot with the score obtained in the Experiments 1, 4.1 and 4.2 for each of 21 coalitions.

Additionally, Table 5.8 expands upon these results, displaying not only the outcomes for each coalition but also the individual results for each Power within the coalition. This provides insight into how the improvements are balanced and whether both Powers within the coalition are benefiting from negotiation.

In a detailed examination of the results, several patterns emerge that substantiate the effectiveness of Negotiating Attila over Non-negotiating agents. With an overall score of 16.4 ± 1.9 compared to 9.7 ± 0.8 for the non-negotiating agents, it's clear that Attila's negotiation strategies enhanced agents' performance. In all 21 coalitions, results over Non-negotiating agents were clearly improved, and only in 3 out of 21 coalitions, Non-negotiating Attila outperformed Negotiating Attila.

For instance, in the coalition of AUS + ITA, Negotiating Attila achieved an outstanding score of 22.6 ± 0.6, substantially surpassing Non-Negotiating Attila's 18.1 ± 0.5 and regular non-negotiating agents' 6.1 ± 0.17. Such findings are particularly interesting given that in the individual Powers of AUS and ITA, the scores were 9.6 ± 0.4 and 13 ± 0.5 respectively, implying that the negotiation strategies were exceptionally effective in this coalition, both at the coalition and individual levels. In other cases such the coalition AUS + RUS, while Negotiating Attila still outperformed with a score of 16.5 ± 0.6, the margin of superiority over Non-Negotiating Attila (14.9 ± 0.4) was less pronounced.

Only in 3 out of the 21 coalitions, non-negotiating Attila outperformed Negotiating Attila. For example, in the case of FRA + TUR, Non-Negotiating Attila actually surpassed with a score of 13.8 ± 0.6 compared to 11.9 ± 0.5 for Negotiating Attila. This variability showcases that the specific nature of coalitions also plays a role in determining the likelihood of a Power/Coalition to perform better.

The fact that in most of the cases the results improved both for the coalition and for the individual players, supports our idea that the right mix of accepting and proposing deals, along with considering the reservation and utility values, helps achieve results that are fair and make sense for everyone involved.

The fact that the results obtained with Negotiating Attila surpass those of Non-Negotiating Attila in most of the coalitions, underscores that the improved outcomes are not merely a consequence of avoiding to conquer allies' territory but indeed arises from an intelligent, balanced, and context-aware negotiation process.

Coalition	**Experiment 1**	**Negotiating Attila (4,1)**	**Non-Negotiating Attila (4,2)**
AUS + ENG	5.6 ± 0.15	6.3 ± 0.4	8.5 ± 0.4
AUS + FRA	8.0 ± 0.20	11.8 ± 0.6	9.6 ± 0.5
AUS + GER	9.7 ± 0.24	14.0 ± 0.7	13 ± 0.6
AUS + ITA	6.1 ± 0.17	22.6 ± 0.6	18.1 ± 0.5
AUS + RUS	12.0 ± 0.27	16.5 ± 0.6	14.9 ± 0.4
AUS + TUR	10.5 ± 0.20	19.6 ± 0.6	17.3 ± 0.4
ENG + FRA	6.4 ± 0.16	18.7 ± 0.5	11.9 ± 0.4
ENG + GER	8.1 ± 0.19	19 ± 0.5	14.0 ± 0.5
ENG + ITA	4.5 ± 0.10	7.5 ± 0.4	6.2 ± 0.3
ENG + RUS	10.4 ± 0.24	15.8 ± 0.5	12.6 ± 0.4
ENG + TUR	8.9 ± 0.15	13.5 ± 0.4	11.8 ± 0.4
FRA + GER	10.6 ± 0.24	23.2 ± 0.6	15.9 ± 0.6
FRA + ITA	6.9 ± 0.17	16.0 ± 0.5	14.3 ± 0.6
FRA + RUS	12.8 ± 0.27	18.2 ± 0.6	17.2 ± 0.6
FRA + TUR	11.2 ± 0.20	11.9 ± 0.5	13.8 ± 0.6
GER + ITA	8.6 ± 0.21	13.7 ± 0.6	11 ± 0.6
GER + RUS	14.5 ± 0.30	25.9 ± 0.5	19.4 ± 0.5
GER + TUR	12.9 ± 0.24	13.2 ± 0.6	11.5 ± 0.5
ITA + RUS	10.8 ± 0.25	19.2 ± 0.7	13.0 ± 0.7
ITA + TUR	9.3 ± 0.17	13 ± 0.5	13.2 ± 0.4
RUS + TUR	15.3 ± 0.20	25.2 ± 0.6	19.6 ± 0.6
Overall	**9.7 ± 0.8**	**16.4 ± 1.9**	**13.7 ± 1.7**

Table 5.7: Comparison of the average number of SCs (± standard error) achieved across all 21 coalitions in experiments with Negotiating Attila (Exp 4.1), Non-negotiating Attila (Exp 4.2), and Non-negotiating agents (Exp 1)

Coalition	**Experiment 1**	**Negotiating Attila (4.1)**	**Non-Negotiating Attila (4.2)**
AUS + ENG	5.6 ± 0.15	6.3 ± 0.4	8.5 ± 0.4
AUS	3.6 ± 0.14	2.7 ± 0.3	4.4 ± 0.4
ENG	2.0 ± 0.05	3.6 ± 0.2	4.1 ± 0.2
AUS + FRA	8.0 ± 0.20	11.8 ± 0.6	9.6 ± 0.5
AUS	3.6 ± 0.14	4.4 ± 0.4	4.4 ± 0.3

Coalition	**Experiment 1**	**Negotiating Attila (4.1)**	**Non-Negotiating Attila (4.2)**
FRA	4.4 ± 0.15	7.4 ± 0.5	5.2 ± 0.4
AUS + GER	9.7 ± 0.24	14.0 ± 0.7	13 ± 0.6
AUS	3.6 ± 0.14	6.8 ± 0.5	5.3 ± 0.3
GER	6.1 ± 0.19	7.2 ± 0.5	7.7 ± 0.5
AUS + ITA	6.1 ± 0.17	22.6 ± 0.6	18.1 ± 0.5
AUS	3.6 ± 0.14	9.6 ± 0.4	8.7 ± 0.3
ITA	2.5 ± 0.09	13.0 ± 0.5	9.4 ± 0.4
AUS + RUS	12.0 ± 0.27	16.5 ± 0.6	14.9 ± 0.4
AUS	3.6 ± 0.14	7.0 ± 0.4	6.7 ± 0.2
RUS	8.4 ± 0.23	9.5 ± 0.5	8.2 ± 0.4
AUS + TUR	10.5 ± 0.20	19.6 ± 0.6	17.3 ± 0.4
AUS	3.6 ± 0.14	9.3 ± 0.3	8.3 ± 0.3
TUR	6.9 ± 0.14	10.3 ± 0.3	9.0 ± 0.2
ENG + FRA	6.4 ± 0.16	18.7 ± 0.5	11.9 ± 0.4
ENG	2.0 ± 0.05	6.3 ± 0.2	4.9 ± 0.2
FRA	4.4 ± 0.15	12.4 ± 0.5	7.0 ± 0.4
ENG + GER	8.1 ± 0.19	19 ± 0.5	14.0 ± 0.5
ENG	2.0 ± 0.05	7.6 ± 0.2	5.3 ± 0.2
GER	6.1 ± 0.19	11.4 ± 0.5	8.7 ± 0.5
ENG + ITA	4.5 ± 0.10	7.5 ± 0.4	6.2 ± 0.3
ENG	2.0 ± 0.05	4.7 ± 0.3	4.5 ± 0.2
ITA	2.5 ± 0.09	2.8 ± 0.2	1.7 ± 0.2
ENG + RUS	10.4 ± 0.24	15.8 ± 0.5	12.6 ± 0.4
ENG	2.0 ± 0.05	6.2 ± 0.3	4.8 ± 0.2
RUS	8.4 ± 0.23	9.6 ± 0.4	7.8 ± 0.4
ENG + TUR	8.9 ± 0.15	13.5 ± 0.4	11.8 ± 0.4
ENG	2.0 ± 0.05	6.1 ± 0.2	5.3 ± 0.2
TUR	6.9 ± 0.14	7.4 ± 0.3	6.5 ± 0.3
FRA + GER	10.6 ± 0.24	23.2 ± 0.6	15.9 ± 0.6
FRA	4.4 ± 0.15	15.1 ± 0.5	7.0 ± 0.3
GER	6.1 ± 0.19	8.1 ± 0.4	8.9 ± 0.5
FRA + ITA	6.9 ± 0.17	16.0 ± 0.5	14.3 ± 0.6
FRA	4.4 ± 0.15	11.5 ± 0.5	11.0 ± 0.6
ITA	2.5 ± 0.09	4.5 ± 0.2	3.3 ± 0.2
FRA + RUS	12.8 ± 0.27	18.2 ± 0.6	17.2 ± 0.6
FRA	4.4 ± 0.15	9.7 ± 0.5	6.3 ± 0.4
RUS	8.4 ± 0.23	8.5 ± 0.4	10.9 ± 0.5
FRA + TUR	11.2 ± 0.20	11.9 ± 0.5	13.8 ± 0.6
FRA	4.4 ± 0.15	5.6 ± 0.4	5.4 ± 0.5
TUR	6.9 ± 0.14	6.3 ± 0.3	8.4 ± 0.4
GER + ITA	8.6 ± 0.21	13.7 ± 0.6	11 ± 0.6
GER	6.1 ± 0.19	8.1 ± 0.5	7.1 ± 0.5
ITA	2.5 ± 0.09	5.6 ± 0.4	3.9 ± 0.3
GER + RUS	14.5 ± 0.30	25.9 ± 0.5	19.4 ± 0.5
GER	6.1 ± 0.19	9.6 ± 0.3	7.0 ± 0.3

Coalition	Experiment 1	Negotiating Attila (4.1)	Non-Negotiating Attila (4.2)
RUS	8.4 ± 0.23	16.3 ± 0.4	12.4 ± 0.4
GER + TUR	12.9 ± 0.24	13.2 ± 0.6	11.5 ± 0.5
GER	6.1 ± 0.19	6.2 ± 0.5	5.9 ± 0.5
TUR	6.9 ± 0.14	7 ± 0.3	5.6 ± 0.2
ITA + RUS	10.8 ± 0.25	19.2 ± 0.7	13.0 ± 0.7
ITA	2.5 ± 0.09	6.6 ± 0.4	3.3 ± 0.3
RUS	8.4 ± 0.23	12.6± 0.6	9.7 ± 0.6
ITA + TUR	9.3 ± 0.17	13.0 ± 0.5	13.2 ± 0.4
ITA	2.5 ± 0.09	6.0 ± 0.4	5.8± 0.3
TUR	6.9 ± 0.14	7.0 ± 0.3	7.4 ± 0.3
RUS + TUR	15.3 ± 0.20	25.2 ± 0.6	19.6 ± 0.6
RUS	8.4 ± 0.23	8.5 ± 0.4	12.0 ± 0.6
TUR	6.9 ± 0.14	16.6 ± 0.5	7.6 ± 0.2
Overall	**9.7 ± 0.8**	**16.4 ± 1.9**	**13.7 ± 1.7**

Table 5.8: Detailed breakdown of the average number of SCs (± standard error) for each coalition and individual Power within the coalition.

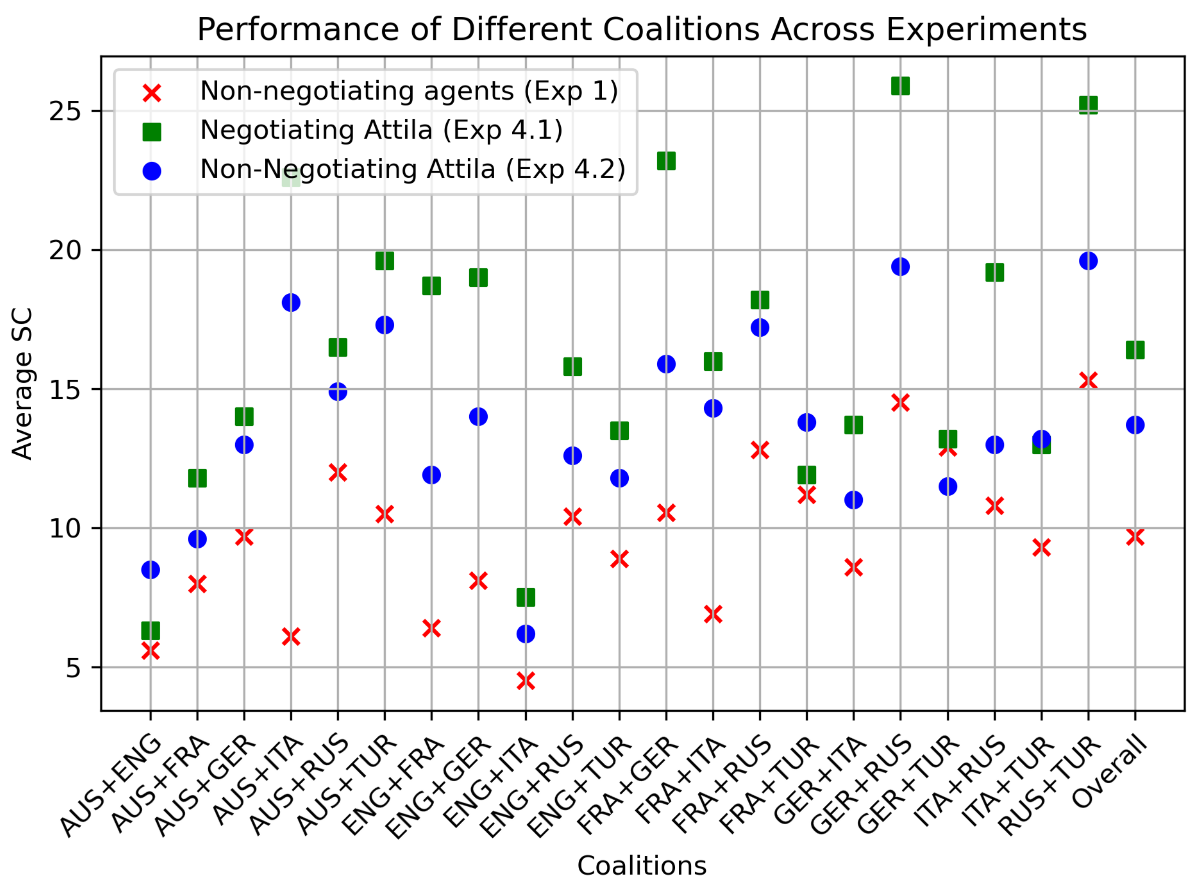

Figure 5.4: Plot showing results of Table 5.7.

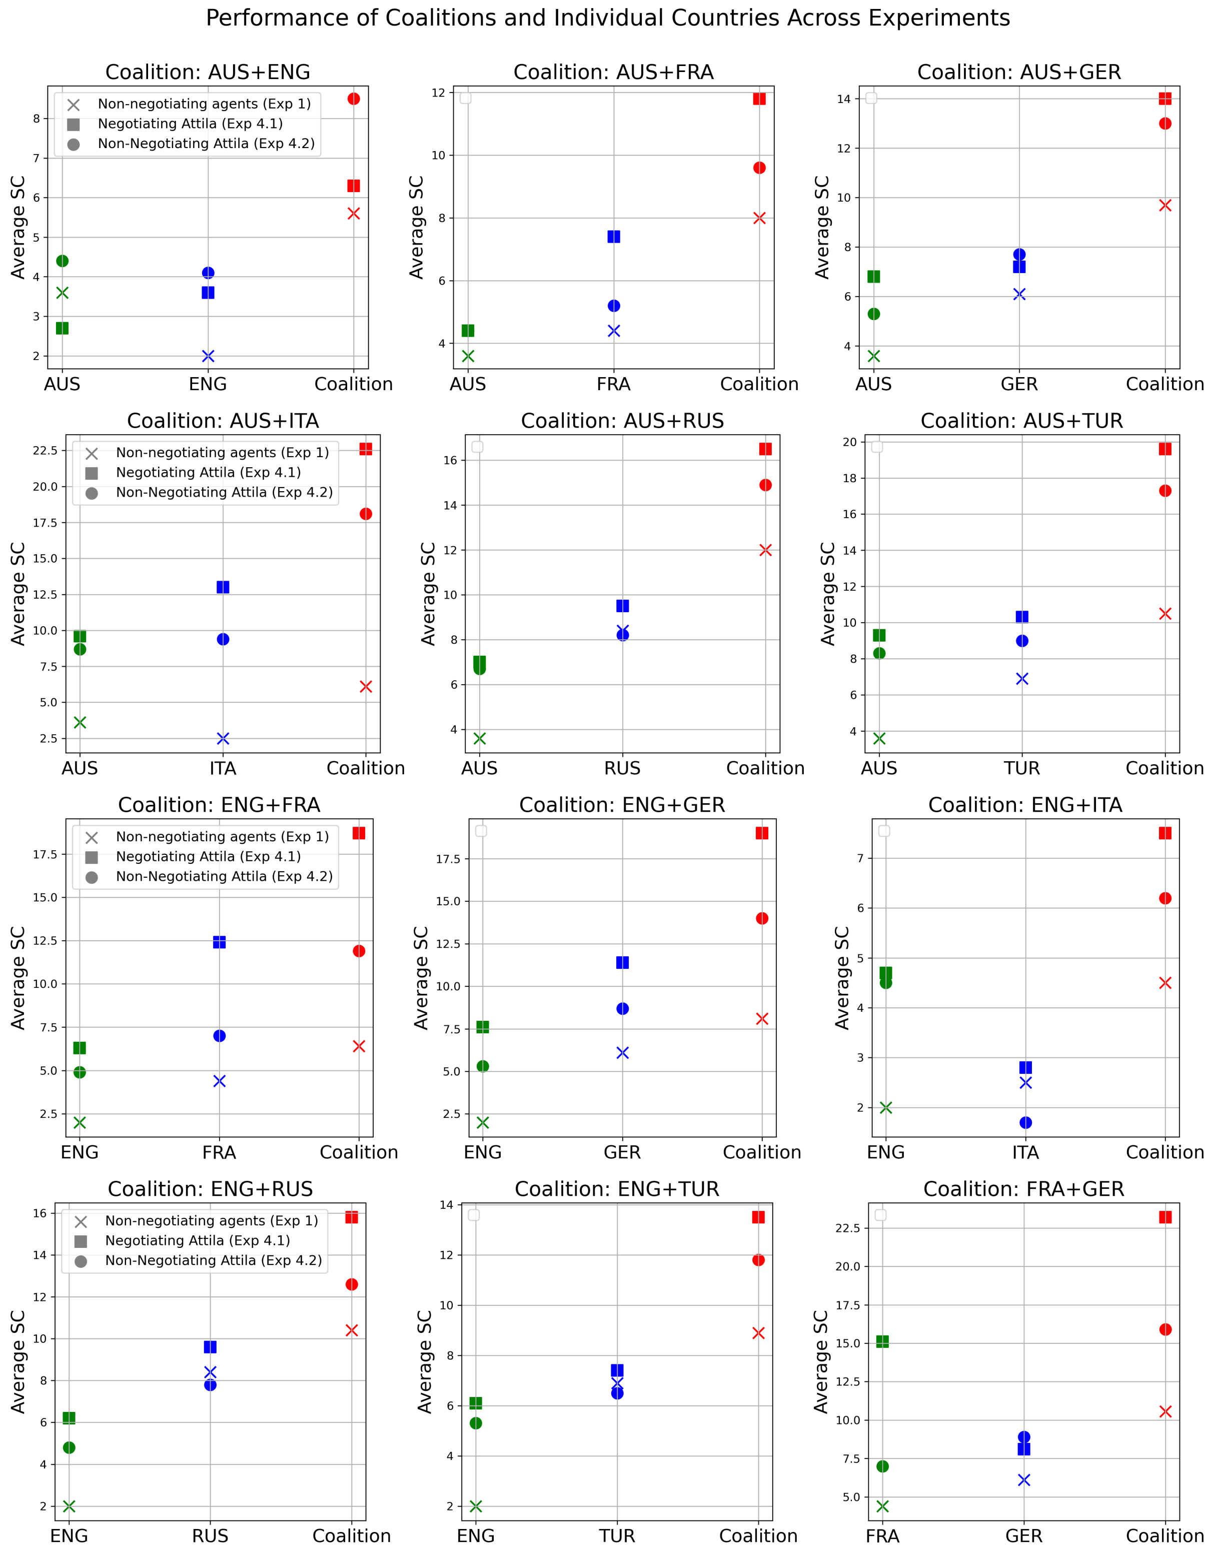
Performance of Coalitions and Individual Countries Across Experiments
Coalition: AUS+ENG
Coalition: AUS+FRA
Coalition: AUS+GER
Coalition: AUS+ITA
Coalition: AUS+RUS
Coalition: AUS+TUR
Coalition: ENG+FRA
Coalition: ENG+GER
Coalition: ENG+ITA
Coalition: ENG+RUS
Coalition: ENG+TUR
Coalition: FRA+GER
Non-negotiating agents (Exp 1)
Negotiating Attila (Exp 4.1)
Non-Negotiating Attila (Exp 4.2)
Average SC
AUS
ENG
FRA
GER
ITA
RUS
TUR
Coalition

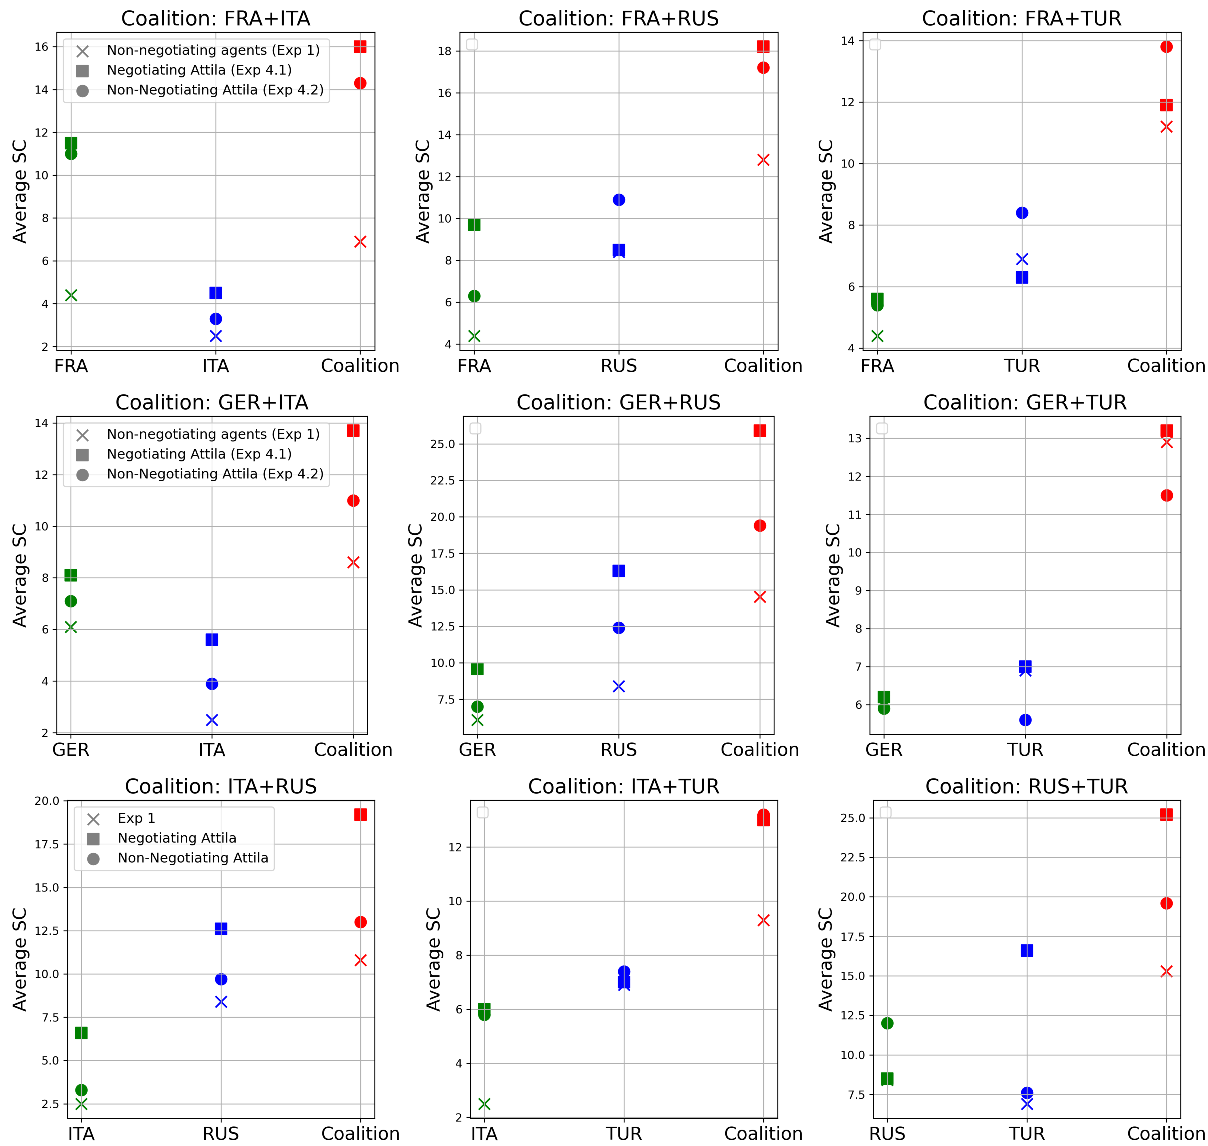

Figure 5.6: Plot showing results of Table 5.8.

Chapter 6

Conclusions

6.1 Conclusions

This thesis embarked on the task of designing and implementing a negotiation bot, aimed at enhancing the performance outcomes of non-negotiating agents in the game of Diplomacy.

Empirical results across various experiments underscore the success of this endeavor. Our designed bot, Attila, demonstrated a marked improvement over non-negotiating agents, providing evidence that effective negotiation strategies can indeed yield superior outcomes in this domain. Specifically, the introduction of Negotiating Attila led to an average score increase across multiple coalitions from 9.7 ± 0.8 to 16.4 ± 1.9.

Therefore, this research marks an important advancement in the domain of automated negotiation bots for the game of Diplomacy by presenting *the first agent ever for the game of Diplomacy capable of negotiating successfully without being trained on human data, and distinctly outperforming non-negotiating agents*. Furthermore, we have demonstrated that these outcomes are not merely because the agents avoid attacking their coalition partners, but arise from a well-structured and fair negotiation process. This emphasizes that the smart design of the algorithms, ensuring that all deals are both Pareto efficient and individually rational, contributes to fair and effective negotiations.

While the data attests to the success of Negotiating Attila, there is still much work that can be done to refine the agent.

6.2 Future Work

While Attila has demonstrated significant advancements, this work also open avenues for future refinement and exploration. One such avenue involves providing Attila with the capability to autonomously identify and establish its coalition partners, which would add a layer of dynamism and adaptability to the negotiation process.

Moreover, it would be instructive to test Attila in competitive scenarios involving other agents, particularly those that have participated in previous Diplomacy Competitions. These experiments would provide valuable insights into the relative efficacy and limitations of Attila's negotiation algorithms.

Another interesting research avenue is exploring the behaviour of Attila in larger coalitions (more than two Powers). As larger coalitions bring more complexity, this would be good for testing the bot's versatility and efficacy in different negotiation scenarios.

Finally, we think that Attila's performance can also be improved by making it capable to make and accept proposals for future turns. At this moment, it is only capable to make and accept proposals for the current turn, however, as De Jonge *et al.* explain in their paper [6], any experienced Diplomacy player would agree that it is essential to plan several steps ahead. This is especially important because it's rare for two players to both benefit immediately when they cooperate. Usually, one player helps another, expecting that the favor will be returned later on in the game.

Chapter 7

References

[1] R. J. Aumann and S. Hart. *Handbook of Game Theory with Economic Applications*, volume 1. Elsevier, 1 edition, 1992.

[2] P. Faratin, C. Sierra, and N. Jennings. Using similarity criteria to make issue trade-offs in automated negotiation. *Artificial Intelligence*, 142:205–237, 12 2002. doi: 10.1016/S0004-3702(02)00290-4.

[3] J. Watson. *Strategy: An Introduction to Game Theory.* W. W. Norton and Company, 3 edition, 2013.

[4] J.S. Rosenschein and G. Zlotkin. *Rules of Encounter.* The MIT Press, Cambridge, USA, 1994.

[5] T. Baarslag, K. Hindriks, M. Hendrikx, A. Dirkzwager, and C. Jonker. *Decoupling Negotiating Agents to Explore the Space of Negotiation Strategies*, volume 535, pages 61–83. 01 2014. ISBN 9784431547570. doi: 10.1007/978-4-431-54758-7_4.

[6] Dave Jonge, Tim Baarslag, Reyhan Aydogan, Catholijn Jonker, Katsuhide Fujita, and Takayuki Ito. *The Challenge of Negotiation in the Game of Diplomacy*, pages 100–114. 04 2019. ISBN 978-3-030-17293-0. doi: 10.1007/978-3-030-17294-7_8.

[7] D. De Jonge and C. Sierra. Nb3: A multilateral negotiation algorithm for large, non-linear agreement spaces with limited time. *Autonomous Agents and Multi-Agent Systems*, 29, 2015. doi: 10.1007/s10458-014-9271-3.

[8] D. De Jonge and C. Sierra. D-brane: A diplomacy playing agent for automated negotiations research. *Applied Intelligence*, 47:158–, 2017. doi: 10.1007/s10489-017-0919-y.

[9] D. De Jonge. *The BANDANA Framework v2.0, User's Manual*, 2022. URL `https://www.iiia.csic.es/~davedejonge/bandana/download.php`.

[10] A. Fabregues and C. Sierra. Dipgame: A challenging negotiation testbed. *Engineering Applications of Artificial Intelligence*, 24(7):1137–1146, 2011. doi: https://doi.org/10.1016/j.engappai.2011.06.005. URL `https://www.sciencedirect.com/science/article/pii/S0952197611001059`.

[11] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, M. Wooldridge, and C. Sierra. Automated negotiation: Prospects, methods and challenges. *Group Decision and Negotiation*, 10:199–215, 2001. doi: 10.1023/A:1008746126376.

[12] S. Buffett and B. Spencer. A bayesian classifier for learning opponents' preferences in multi-object automated negotiation. *Electronic Commerce Research and Applications*, 6(3):274–284, 2007. doi: https://doi.org/10.1016/j.elerap.2006.06.008. URL `https://www.sciencedirect.com/science/article/pii/S1567422306000445`.

[13] 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*, 73:1–14, 2015. doi: https://doi.org/10.1016/j.dss.2015.02.012. URL `https://www.sciencedirect.com/science/article/pii/S0167923615000329`.

[14] G. Weiss and S. Chen. *Automated Negotiation: An Efficient Approach to Interaction Among Agents*, pages 149–177. World Scientific Publishing Company, 2018.

[15] A. Marsico, M. Savi, D. Siracusa, and E. Salvadori. An automated negotiation framework for application-aware transport network services. *Optical Switching and Networking*, 38:100571, 2020. doi: https://doi.org/10.1016/j.osn.2020.100571. URL `https://www.sciencedirect.com/science/article/pii/S1573427720300163`.

[16] D. De Jonge and D. Zhang. Automated negotiations for general game playing, 2017. URL `https://digital.csic.es/handle/10261/241155`.

[17] S. Polberg, M. Paprzycki, and M. Ganzha. Developing intelligent bots for the diplomacy game. In *2011 Federated Conference on Computer Science and Information Systems (FedCSIS)*, pages 589–596, 2011.

[18] A. Ferreira, H. Lopes Cardoso, and L. Reis. Dipblue: A diplomacy agent with strategic and trust reasoning. *ICAART 2015 - 7th International Conference on Agents and Artificial Intelligence, Proceedings*, 1:54–65, 2015. doi: 10.5220/0005205400540065.

[19] J. S. G. Ng. M@stermind: A diplomacy game negotiating agent, 2018. URL `http://hdl.handle.net/10356/76131`.

[20] Meta Fundamental AI Research Diplomacy Team (FAIR)†, A. Bakhtin, N. Brown, E. Dinan, G. Farina, C. Flaherty, D. Fried, A. Goff, J. Gray, H. Hu, A. P. Jacob, M. Komeili, K. Konath, M. Kwon, A. Lerer, M. Lewis, A. H. Miller, S. Mitts, A. Renduchintala, S. Roller, D. Rowe, W. Shi, J. Spisak, A. Wei, D. Wu, H. Zhang, and M. Zijlstra. Human-level play in the game of diplomacy by combining language models with strategic reasoning. *Science*, 378(6624): 1067–1074, 2022. doi: 10.1126/science.ade9097. URL `https://science.sciencemag.org/content/378/6624/1067`.

[21] Reyhan Aydogan, Tim Baarslag, Katsuhide Fujita, Johnathan Mell, Jonathan Gratch, Dave Jonge, Yasser Mohammad, Shinji Nakadai, Satoshi Morinaga, Hirotaka Osawa, Claus Aranha, and Catholijn Jonker. *Challenges and Main Results of the Automated Negotiating Agents Competition (ANAC) 2019*, pages 366–381. 01 2020. ISBN 978-3-030-66411-4. doi: 10.1007/978-3-030-66412-1_23.