

The Challenge of Negotiation in the Game of Diplomacy

Dave de Jonge^{1,2}, Tim Baarslag³, Reyhan Aydoğan⁴, Catholijn Jonker⁵,
Katsuhide Fujita⁶, and Takayuki Ito⁷

¹ IIIA-CSIC, Bellaterra, Spain

² Western Sydney University, Australia

³ Centrum Wiskunde & Informatica, Amsterdam, The Netherlands

⁴ Özyeğin University, Istanbul, Turkey

⁵ Delft University of Technology, The Netherlands

⁶ Tokyo University of Agriculture and Technology, Japan

⁷ Nagoya Institute of Technology, Japan

Abstract. The game of Diplomacy has been used as a test case for complex automated negotiations for a long time, but to date very few successful negotiation algorithms have been implemented for this game. We have therefore decided to include a Diplomacy tournament within the annual Automated Negotiating Agents Competition (ANAC). In this paper we present the setup and the results of the ANAC 2017 Diplomacy Competition and the ANAC 2018 Diplomacy Challenge. We observe that none of the negotiation algorithms submitted to these two editions have been able to significantly improve the performance over a non-negotiating baseline agent. We analyze these algorithms and discuss why it is so hard to write successful negotiation algorithms for Diplomacy. Finally, we provide experimental evidence that, despite these results, coalition formation and coordination do form essential elements of the game.

1 Introduction

Automated negotiations have been studied extensively, but traditionally most work has focused on the strategy to determine which deals to propose *given* the utility values of those deals. A point that has received less attention is the fact that in many real-world negotiation settings, for any given proposal, a negotiator would need to spend considerable effort on estimating its value. Only recently, more attention has been given in the literature to negotiation domains where the calculation of utility is a highly non-trivial and time-consuming task. For example, [9] treated a problem in which determining the value of a deal was NP-hard and in [11] an algorithm was presented for negotiations applied to non-zero-sum General Game Playing.

The Automated Negotiating Agents Competition (ANAC) is an annually returning competition that aims to improve the state-of-the-art in automated negotiations [3]. It was first held in 2010 and has been steadily growing in popularity. The setup of this competition has been updated each year to reflect the

advancements made in the field of research. While ANAC started with small contract spaces and linear utility functions [4], it has featured increasingly complex scenarios, involving very large agreement spaces [8], multilateral negotiations [1], human-agent interactions [13], and non-linear utility functions [8, 2].

However, in all of these editions, the process of evaluating a proposal was abstracted away. The agents would know the value of any potential proposal almost instantaneously because it could be calculated with a simple linear formula (this was true even in the editions with non-linear utility functions). Furthermore, the agents were not required to have any background knowledge of the negotiation domains and did not need to apply any form of reasoning to obtain the utility value of a proposal. The utility functions of the agents' opponents, on the other hand, were assumed to be completely unknown.

We argue that in real negotiations it is important to have knowledge of the domain and one should be able to reason about it. One cannot, for example, expect to make profitable deals in the antique business without having any knowledge of antique, no matter how good one is at bargaining. Moreover, a good negotiator should also be able to reason about the desires of its opponents. A good car salesman, for example, would try to find out what type of car best suits his client's needs to increase the chances of making a profitable deal. Therefore, we envisioned a need to add a new league to ANAC that does involve this kind of complex reasoning.

The game of Diplomacy forms an excellent test case for this type of complex negotiations, as it is a game that includes many of the difficulties one would also have to face in real-life negotiations [6]. It involves constraint satisfaction, coalition formation, game theory, trust, and even psychology. Now that modern Chess and Go computers are already far superior to any human player [15], we expect that Diplomacy will start to draw more attention as the next big challenge for computer science.

Although the game of Diplomacy has already been under attention of the Automated Negotiations community for a long time, to date very few successful negotiating Diplomacy players have been developed. Some of the earliest work on this game, for example, was presented in [14], but they only managed to play a very small number of games, because they had to play them with humans.

An informal online community called DAIDE exists which is dedicated to the development of Diplomacy playing agents.⁸ Many agents have been developed by this community but only very few are capable of negotiation. One of the main non-negotiating bots developed on this framework is called the DumbBot.

In [6] a new platform called DipGame was introduced to make the development of Diplomacy agents easier for scientific research. This platform was later extended into the BANDANA platform [10]. Several negotiating agents have been developed using DipGame such as DipBlue [7] which consists of a negotiation algorithm built on top of the DumbBot. Unfortunately, its negotiation algorithm did not result in a very strong increase in performance with respect to the non-negotiating DumbBot. An entirely new agent was presented in [10],

⁸ <http://www.daide.org.uk>

called D-Brane, which can play with or without negotiations. Again, it turned out that when applying negotiations it is only slightly stronger than when it plays without negotiating. In 2015 the non-negotiating version of D-Brane won the Computer Diplomacy Challenge⁹ which was organized as part of the ICGA Computer Olympiad.

On the other hand, another negotiation algorithm was implemented on top of DumbBot [5], which did strongly outperform the non-negotiating DumbBot. Unfortunately, this agent required a supercomputer to run.

Another negotiating agent, called AlphaDip, was presented in [12], which was largely based on D-Brane. Although it did improve over D-Brane, the authors still concluded that adding negotiations to their agent only had a very small influence on its overall performance.

This paper presents the setup and results of the ANAC 2017 Diplomacy Competition and the ANAC 2018 Diplomacy Challenge and provides an analysis of the proposed negotiation strategies for Diplomacy. The rest of the paper is organized as follows: Section 2 introduces the game of Diplomacy while Section 3 explains the negotiation protocol used in this game. Sections 4 and 5 present the setup and results of the 2017 and 2018 editions respectively. In Section 6 we present an experiment we conducted to show the importance of cooperation in Diplomacy. Finally, in Section 7, we conclude the paper with the lessons learned.

2 Diplomacy

Diplomacy is a widely played game for seven players. Just like chess it is completely deterministic (i.e. there are no dice, cards, or any other source of randomness) and there is no hidden information.¹⁰ Players make their moves simultaneously. It is designed in such a way that each player needs to negotiate with the other players in order to have a chance of winning. It can be played as a classical board game, or it can be played online.¹¹

The game takes place on a map of Europe in the year 1901, which is divided into 75 *Provinces*. Each player plays one of the seven great *Powers* of that time: *Austria* (AUS), *England* (ENG), *France* (FRA), *Germany* (GER), *Italy* (ITA), *Russia* (RUS) and *Turkey* (TUR) and each player starts with three or four units (armies or fleets) which are placed in fixed initial positions on the map. In each round of the game, each player must ‘*submit an order*’ for each of its units, which tells those units how to move around the map and allows them to conquer the map’s provinces.

Some of the Provinces are so-called *Supply Centers* and the goal for the players is to conquer those Supply Centers. A player is eliminated when he or

⁹ https://icga.leidenuniv.nl/?page_id=987

¹⁰ One might argue that Diplomacy does have hidden information, because players make secret agreements. However, these agreements have no formal meaning, and form part of the players’ strategies rather than of the rules of the game. Therefore, *formally* speaking there is no hidden information.

¹¹ <http://www.playdiplomacy.com/>

she loses all his or her Supply Centers and a player wins the game when he or she has conquered 18 or more of the 34 Supply Centers (a *Solo Victory*). However, the game may also end when all surviving players agree to a draw.

The game iterates through five types of rounds (or ‘phases’), in the following order: *Spring, Summer, Fall, Autumn, Winter*. The first round of the game is referred to as Spring 1901, followed by Summer 1901, etcetera. After Winter 1901 follows Spring 1902, Summer 1902, and so on.

The main difference between Diplomacy and other deterministic games like Chess and Go, is that in Diplomacy players are allowed to negotiate with each other and form coalitions. At each round, before the players submit their orders, the players are given time to negotiate with each other and make agreements about the orders they will submit. Negotiations take place in private, and each agreement that is made is only known to the players involved in that agreement.

Typically, players may agree not to invade certain provinces, or they may agree that one player will help the other player to invade a certain province. In this way, players essentially form coalitions. These coalitions are not given beforehand. Instead, during the course of the game players may form and break coalitions as they like.

3 The Negotiation Protocol

In a real Diplomacy game there are no formal rules for the negotiations. Players are allowed to negotiate anything and there is no guarantee that players will obey their agreements. However, for our competition we needed to establish a well-defined negotiation language and protocol so that the agents could understand each other. Furthermore, in order to simplify the game and eliminate the issue of trust, we imposed the rule that the players are always obliged to obey their agreements. This means that our negotiation language needed to have well-defined formal semantics, which are explained below.

As the negotiation protocol, we used the Unstructured Negotiation Protocol [9], because it most closely resembles how negotiations in real games of Diplomacy take place. In this protocol, the agents do not take turns, but instead are allowed to propose or accept a deal whenever they want. A deal may involve any number of agents. Once all players involved in the deal have accepted it, a special *Notary* agent checks whether it is consistent with earlier made agreements. If this is indeed the case then the Notary will send a confirmation message to all agents involved in the deal. Once the Notary has sent this confirmation message the deal is considered officially binding. Players may propose and accept as many deals as they wish and negotiations continue after a deal has been confirmed.

If an agent has proposed or accepted a deal, but then changes its mind, and the deal has not yet been confirmed by the Notary, it can send a reject message to withdraw from the proposal and hence prevent it from becoming confirmed. However, once the deal is confirmed by the Notary the agents involved must always obey it.

Since each proposal is only sent to those players that are involved in it the other players will never be aware that this deal was proposed. Also, the Notary sends its confirmation message only to the players involved in the deal, so the agreement remains secret.

3.1 Allowed Proposals

In this section we define the set of deals that agents may propose to each other. A deal may consist of any number of *Order Commitments* and any number of *Demilitarized Zones*.

Definition 1. An **Order Commitment** oc is a tuple: $oc = (y, \phi, o)$, where y is a ‘year’ (an integer greater than 1900), $\phi \in \{Spring, Fall\}$ is a ‘phase’ and o is any legal order for any unit.

An Order Commitment represents a promise that a power will submit a certain order during a certain phase and year. For example: “*In the Spring of 1902 the army in Holland will move to Belgium*”. Formally, an Order Commitment (y, ϕ, o) is obeyed if Power P submits the order o during phase ϕ of year y , where P is the owner of the unit defined by the details of the order o .

Definition 2. A **Demilitarized Zone** dmz is a tuple: $dmz = (y, \phi, A, B)$ with y and ϕ as in Def.1, A is a nonempty set of Powers and B is a nonempty set of Provinces.

A Demilitarized Zone is an agreement between the specified Powers that none of them will invade (or stay inside) any of the specified Provinces during the specified phase and year. For example, the Demilitarized Zone

$$(1903, Fall, \{FRA, GER, ENG\}, \{NTH, ECH\})$$

has the interpretation “*In the Fall of 1903 France, Germany, and England will keep out of the North Sea and the English Channel*”. Formally, a Demilitarized Zone is obeyed if none of the powers in A submits any order during phase ϕ of year y to move any unit into any of the provinces in B .

Definition 3. A **Deal** d is a non-empty set:

$$d = \{oc_1, \dots, oc_n, dmz_1, \dots, dmz_m\}$$

where each oc_i is an Order Commitment, each dmz_i is a Demilitarized Zone, and where n and m can be any non-negative integers.

When a deal is confirmed by the Notary it means that all Order Commitments and all Demilitarized Zones in it must be obeyed.

A proposed deal can only be accepted or rejected in its entirety. If an agent wishes to accept only a part of the deal, it can simply propose a new deal which only consists of the subset of Order Commitments and Demilitarized Zones it desires.

Apart from proposing this type of deals, agents are also allowed to propose a draw to all other players. The game ends in a draw if all agents that have not been eliminated propose a draw in the same round of the game.

4 The ANAC 2017 Diplomacy Competition

4.1 Submission Rules and Tournament setup

The assignment for the participants was to implement a negotiation algorithm using the BANDANA framework. This negotiation algorithm would then be combined with the tactical module of D-Brane to form a complete agent. This tactical module would then choose which moves the agent makes, while obeying the agreements made by the negotiation algorithm. The participants were not allowed to implement a complete Diplomacy playing agent from scratch. They were only allowed to implement a negotiation algorithm so that the competition focused purely on the negotiation aspect of Diplomacy.

In order to determine whether to accept a proposal or not, the participants' negotiation algorithms had the possibility to consult D-Brane's Tactical Module to see which moves would be played if that proposal was accepted.

The tournament was run using the Parlance game server.¹² We let all agents participating in the competition play 110 games together. Since a game requires 7 players and we only had 4 participants, we supplemented the agents with 3 instances of the non-negotiating D-Brane. In each game the players were randomly assigned to the 7 Powers.¹³ Every round of each game had a deadline of 30 seconds. In order to prevent the games from continuing forever a draw was declared automatically in any game that advanced to the Winter 1920 phase. The agents' overall score for the tournament was determined by the number of Supply Centers they conquered.

4.2 Submissions

We received the following submissions:

- **Frigate**, by Ryohei Kawata and Katsuhide Fujita, *Tokyo University of Agriculture and Technology*, Japan
- **Agent Madoff**, by Tan Hao Hao, *Nanyang Technological University*, Singapore
- **DDAgent**, by Daichi Shibata, *Nagoya Institute of Technology*, Japan
- **NaiveThinkerG**, by Giancarlo Nicolo, *Universitat Politècnica de València*, Spain

Due to lack of space we cannot give a description of all of these agents. Therefore, we will only discuss the winner and the runner-up.

¹² <https://pypi.python.org/pypi/Parlance/1.4.1>

¹³ It would have been better to assign each agent to each Power an equal number of times, because some Powers are stronger than others. Unfortunately, however, the Parlance game server does not provide this option.

Frigate Frigate only proposes bilateral deals, and only to Powers that own at least 3 and at most 10 Supply Centers. Furthermore, it does not deal with any Power that forms a direct threat to any of Frigate’s own Supply Centers. For each Power that does qualify Frigate constructs a proposal by consulting the D-Brane Tactical Module to find the best plans for itself and the other agent, under the restriction that they do not invade each others’ Supply Centers. The proposal will then consist of the union of these plans.

Frigate randomly chooses a deal from the proposals it found, where the probability depends on the strength of the other agent (the weaker the agent, the higher the probability) and the number of Supply Centers that Frigate expects to gain from it. Furthermore, the probability is multiplied by 5 if the other agent is considered an ally. An agent is considered an ally if it was involved in the last confirmed deal that Frigate was involved in.

Although Frigate does implement an acceptance strategy, due to a bug in the code, it never accepts any incoming proposals.

Agent Madoff In order to generate proposals Agent Madoff first tries to predict the opponents’ orders using the D-Brane Tactical Module under the assumption that the opponents have not made any agreements. Then, it identifies which orders are in conflict with its own interests, namely orders for units to invade any of Agent Madoff’s own Home Supply Centers, or any province that Agent Madoff is also trying to invade. It then tries to find alternative orders for such units and proposes them. If it cannot find any suitable alternative order then Agent Madoff will try to ask a third party for support to defend or attack the province in question.

Agent Madoff does not really apply a coalition formation strategy. However, it does keep track of each opponent’s ‘hostility’. Initially, it assigns to each Power has a *hostility value* of 0. This value is decreased whenever a Power steals a Supply Center from Agent Madoff, and is increased whenever a Power agrees to give support to Agent Madoff. This value is then used by Agent Madoff’s acceptance strategy. The higher this value, the more likely it is that Agent Madoff will accept a proposal from this opponent.

When Agent Madoff receives a proposal it calculates for each component of this deal a value between 0 and 1 which depends on various heuristics, such as the value of the province that is the destination of the order (in case of a move order commitment), or the hostility of the supported power (in case of a support order commitment). It then calculates the average value over these components. The higher this average value, the higher the probability that Agent Madoff will accept it.

4.3 Results

Initially, we ran the competition according to the setup announced to the participants. Unfortunately, no agent performed significantly better than the non-negotiating D-Brane, which means that the ability to negotiate did not really

improve the results of the agents. We then played 50 games with 4 instances of each agent versus 3 instances of D-Brane. The idea behind this was that it might be easier for the agents to negotiate with a copy of themselves, rather than with a different agent. Unfortunately, this setup also did not result in any of the players significantly outperforming the others.

Therefore, to decide a winner, we counted the number of proposals made by each agent that were accepted by every other agent involved in them, and considered that value as the final score of each agent. The idea being that if an agent’s proposals are accepted by the other agents, this can be seen as a measure of quality, even though the agreement did not in the end result in a higher number of Supply Centers. The results are displayed in Table 1. We see that Frigate was proclaimed the winner of the competition and Agent Madoff was awarded the second prize.

	<u>Confirmed proposals</u>
Frigate	372
Agent Madoff	170
DDAgent	61
NaiveThinkerG	30

Table 1. Final results of the 2017 Diplomacy Competition. We counted the number of proposals made by each agent that were eventually accepted by all the other agents involved in it.

5 The 2018 Diplomacy Challenge

Because the 2017 Diplomacy Competition did not end with one agent being significantly better than any of the other agents, or even better than the non-negotiating agent, we decided to change the setup for 2018. Instead of a ‘Competition’ we turned it into a ‘Challenge’, meaning that a winner would only be proclaimed if its results are significant.

5.1 Tournament Setup

Most of the setup for 2018 was identical to the setup of 2017. We used exactly the same negotiation protocol, and the participants were again required to implement a negotiation algorithm on top of D-Brane. The main difference was that the 2018 Challenge consisted of two rounds.

In the first round for each agent we ran 100 games with 4 instances of that agent against 3 instances of the non-negotiating D-Brane agent. We say an agent *passed* the first round if the instances of that agent conquered a statistically significant higher number of Supply Centers on average than the D-Branes. The agents that did not pass the first round were eliminated from the Challenge.

For the second round we then let all agents that passed the first round play together. Since it was likely that there would be less than 7 such agents, we stated the rule that the field would be supplemented with as many agents that did not pass the first round as necessary, even though such agents were not eligible to win the challenge. Furthermore, if there still would not be enough agents, we would supplement the field with instances of the non-negotiating D-Brane agent. We played 100 games and the agent that conquered the highest number of Supply Centers would be the winner of Round 2.

In order to win the Challenge an agent had to win the second round, as well as pass the first round. This means that if the winner of the second round did not pass the first round there would be no winner at all.

5.2 Motivation

The motivation behind this setup is that in Round 2 the real negotiation skills of the agents are tested. In theory, if an agent makes purely selfish proposals, it will not succeed, because its proposals will not be accepted by the other agents. On the other hand, if it makes purely altruistic proposals or accepts any proposal it receives, it will not succeed either, because it will be exploited by its opponents. In practice, however, a bad negotiator could still be able to win Round 2, because its opponents are not perfect either and therefore it might purely benefit from bad proposals made by the other agents. In order to prevent such ‘freeloading’ behavior we demanded that each agent was also able to successfully negotiate with only copies of itself. For this reason we have included Round 1 in this challenge. One could also roughly say that Round 1 tests the agents’ ‘proposing strategy’, while Round 2 tests their ‘acceptance strategy’.

5.3 Submissions

We received the following submissions:

- **CoalitionBot**, by Ido Westler, Yehuda Callen, Moche Uzan, Arie Cattan, Avishay Zagury *Bar Ilan University*, Israel
- **M@sterMind**, by Jonathan Ng, *Nanyang Technological University*, Singapore
- **Gunma**, by Ryohei Kawata and Katsuhide Fujita, *Tokyo University of Agriculture and Technology*, Japan
- **GamlBot**, by Michael Vassernis, *Bar Ilan University*, Israel
- **DDAgent2**, by Daichi Shibata, *Nagoya Institute of Technology*, Japan

Unfortunately, it turned out that DDAgent2 was too slow to participate, because in many rounds it was not able to submit its orders before the deadline.

Again, due to space constraints we will only discuss the two best agents of the two respective rounds.

CoalitionBot CoalitionBot is a very passive player. It only proposes demilitarized zones and it accepts any incoming proposal. In the first turn, it proposes a bilateral deal to every other Power. This deal proposes that the other Power will not invade any of the CoalitionBot’s own supply Centers during the current turn, and in return the CoalitionBot will not invade the other power’s Supply Centers during the same turn. Any agent that accepts this proposal will be considered an ally for the rest of the game. In all other turns, CoalitionBot proposes to all its allies that they will not attack each others’ Supply Centers.

We will see below that CoalitionBot was able to perform strongly in the first round, but not in the second round. This is not surprising, given that it always accepts any incoming proposal and does not try to exploit its opponents. Its implementation seems to be based on the idea that it can always completely trust its opponents. Clearly, this works well when playing against copies of itself, but not when playing against less altruistic opponents.

Gunma Gunma proposes two types of deals, which the authors call a ‘Mutual Support’ and a ‘Combined Attack’, respectively. A Mutual Support is a deal in which one unit of Gunma supports an opponent’s unit to hold, and the opponent’s unit supports Gunma’s unit to hold in return. A Combined Attack is a deal in which one of Gunma’s units attacks a province owned by an enemy, with support from as many units from allies as possible. Whenever Gunma can find a Combined Attack, it will propose it. On the other hand, it will only propose a Mutual Support if it finds one for which it is sure it can gain a Supply Center.

For any received proposal Gunma predicts how many Supply Centers it would gain from it. It accepts the deal that yields the highest gain, but if there are multiple such deals, it uses the current number of supply centers owned by the proposer as a tie-breaker. In that case it will accept the deal from the currently weakest Power.

Note that Gunma’s proposing strategy is rather greedy. It only proposes deals that yield benefit to himself, and never considers the needs of its negotiation partners. When it comes to accepting, however, it is less selfish. If no deal yields any gain, than it is willing to accept a deal that does not cause Gunma to lose any Supply Centers.

5.4 Results

The Results of Round 1 are displayed in Table 2. We see that only CoalitionBot and Gunma were able to outperform D-Brane. However, a one-sided Student-t test¹⁴ revealed that the results of Gunma were not significant (p-value 0.23). Therefore, only CoalitionBot managed to pass Round 1 (p-value $9.7 \cdot 10^{-9}$).

The results of Round 2 are shown in Table 3. As explained above, we needed to include all the agents in this round, as well as 3 instances of D-Brane, in

¹⁴ With respect to the null-hypothesis that each agent has a mean score of $\frac{34}{7}$ Supply Centers per game.

	Sup. Centers	Result
CoalitionBot	5.528 ± 0.110	PASS
D-Brane	3.963 ± 0.146	
Gunma	4.950 ± 0.128	FAIL
D-Brane	4.733 ± 0.171	
D-Brane	4.930 ± 0.164	
M@sterMind	4.803 ± 0.123	FAIL
D-Brane	5.440 ± 0.184	
GamlBot	4.420 ± 0.138	FAIL

Table 2. Results of the 2018 Diplomacy Challenge, Round 1. Displayed are the average number of conquered supply centers per game, with their standard errors.

	Supply Centers
Gunma	5.69 ± 0.300
GamlBot	5.31 ± 0.334
CoalitionBot	4.94 ± 0.289
D-Brane	4.54 ± 0.157
M@sterMind	4.44 ± 0.290

Table 3. Results of the 2018 Diplomacy Challenge, Round 2. Gunma scores highest, but the results are not significant.

order to have 7 players, even though CoalitionBot was the only agent that passed Round 1 and therefore the only candidate to win the challenge.

We see that Gunma performed best, although the difference between the first three agents is non-significant. Since the CoalitionBot did not beat the other agents in Round 2, and it was not able to clearly outperform the D-Brane in this round either, the 2018 Diplomacy Challenge ended with no winner.

6 Is Cooperation even Possible?

One question that may come to mind when looking at the results, is whether it is really possible at all to improve performance by means of negotiation. Any experienced Diplomacy player would answer this question with a definite ‘yes’, but we would like to back this claim up with scientific evidence.

The question is then how we could show that it is possible to negotiate successfully, without having any algorithm that can do this to our disposal. Fortunately, we have managed to design an experiment that allows us to show the benefit of cooperation, without actually using a negotiation algorithm.

It worked as follows. We first let 7 instances of the non-negotiating D-Brane play 200 games and recorded how many Supply Centers each Power conquered on average. The results are displayed in Table 4. Next, we repeated this experiment, but with only 6 instances of D-Brane while one of those agents was playing two Powers at the same time. For each possible combination of two Powers we played 200 games (there are $\binom{7}{2} = 21$ such combinations, so we played $21 \cdot 200 = 4200$ games) and recorded the number of Supply Centers conquered by the agent playing two Powers.

In this way we have been able to show that if one agent plays the role of two Powers at once, it scores more Supply Centers than if two agents individually play the same two Powers. In other words, when two Powers work together as a team, they have a clear advantage. These results are displayed in Table 5.

For example, in the first row we see that when AUS and ENG are played by one agent, then that agent scores on average 6.99 Supply Centers. However, we see in Table 4 that when these Powers are played by individual agents, they only score 1.60 and 4.39 Supply Centers respectively, yielding a total score of $1.60 + 4.39 = 5.99$, which is also displayed in the first row of Table 5.

The combination of AUS and ENG only yields a small advantage, but for many other coalitions we see much stronger synergy effects. For example, FRA and GER together score 22.1 Supply Centers when played by a single agent, while when playing individually they only score¹⁵ $4.98 + 4.11 = 9.09$.

In general, we see a clear advantage in 12 out of the 21 possible combinations (more than 4 Supply Centers difference, indicated with ++) and a small advantage in 4 of those combinations (indicated with +). When we calculate the average over all combinations we find that the agent playing two Powers scores around 14 Supply Centers, which is clearly more than the $\frac{2}{7} \cdot 34 = 9.71$ Supply Centers that two individual agents would conquer on average.

Although it is clear that players have an advantage when cooperating, we also conclude that this highly depends on *which* two Powers are forming a coalition. FRA and GER, for example, form a much stronger coalition than AUS and ENG. This is an important observation, because this may also explain why it is hard for the submitted agents to negotiate successfully. In many games the negotiating agents may be assigned to Powers that do not form strong combinations, making it hard to benefit from negotiation.

The fact that some coalitions are stronger than others is well-known among experienced Diplomacy players, and is a consequence of the topology of the map. For example, Russia and Turkey are two bordering Powers, which means that if they form a coalition at the beginning of the game then each of them does not have to worry about being attacked by the other, and can therefore completely focus on its other direct neighbors. Furthermore, the fact that they are located next to each other means they can easily give support to one another.

On the other hand, Turkey and England form a weak coalition because they are positioned on opposite ends of the map, so they cannot attack each other in early stages of the game, which means they would not benefit from any mutual peace agreement, and they cannot give each other support either.

For some coalitions we even see a detrimental effect. Although in most cases they are relatively small, they cannot be attributed to statistical fluctuations. We suspect that this results from the fact that they play different opening moves when playing together, which coincidentally happen to be worse.

Finally, we should note that the difference in strength between the various coalitions may not only be caused by the topology of the map, but may also partially be a consequence of the strategy applied by D-Brane. Therefore, we expect these results to be different, but not radically different, if we repeated this experiment with a different agent.

¹⁵ Table 5 shows a value of 9.08 instead of 9.09. This difference is due to rounding errors.

We conclude from these experiments that it should definitely be possible for two agents to benefit from negotiations. Interestingly, these results also suggest *how* such a negotiation algorithm could be implemented. The idea is that if our agent is playing, for example, FRA, then it could consult the D-Brane Tactical module to ask which moves it should play if it were playing as both FRA and GER. Then, it could propose those moves to GER.

Power	Supply Centers	Power	Supply Centers
AUS	1.60 ± 0.16	ITA	2.41 ± 0.16
ENG	4.39 ± 0.17	RUS	10.44 ± 0.42
FRA	4.98 ± 0.20	TUR	6.09 ± 0.17
GER	4.11 ± 0.24		

Table 4. The number of Supply Centers conquered by each Power, when 7 instances of D-Brane are playing without negotiations

Coalition of 2 Powers	Score by 1 agent	Score by 2 agents	Coalition of 2 Powers	Score by 1 agent	Score by 2 agents
AUS + ENG	6.99	5.99	FRA + GER	22.1	9.08
AUS + FRA	9.91	6.57	FRA + ITA	13.43	7.39
AUS + GER	4.11	5.7	FRA + RUS	9.91	15.41
AUS + ITA	12.91	4.01	FRA + TUR	8.71	11.07
AUS + RUS	17.61	12.03	GER + ITA	11.7	6.52
AUS + TUR	17.95	7.69	GER + RUS	21.27	14.54
ENG + FRA	17.95	9.37	GER + TUR	8.99	10.20
ENG + GER	17.78	8.50	ITA + RUS	19.21	12.85
ENG + ITA	8.88	6.8	ITA + TUR	11.92	8.5
ENG + RUS	20.37	14.83	RUS + TUR	24.34	16.53
ENG + TUR	8.82	10.48	Overall	14.04	9.71

Table 5. The number of Supply Centers conquered by each combination of two Powers played by one agent, compared to their score when played by two agents. Differences greater than 4 Supply Center are indicated with ++ or --, while smaller differences are indicated with + or -. In all cases except AUS+ENG, AUS+GER and GER+TUR the p -value was smaller than 10^{-4} .

7 Discussion and Conclusions

From these two competitions we have learned that it is still very hard for the Automated Negotiations community to implement algorithms for domains as complex as Diplomacy. So far, no submission has been able to significantly outperform a non-negotiating agent, even though we have experimentally shown that it is definitely possible for agents to benefit from cooperation.

However, it is important to understand that we are not expecting the Diplomacy Challenge to have a winner any time soon. We regard it as a long term challenge which might take several years to tackle. After all, in the cases of Chess, Go, and Poker it also took many years to develop strong programs.

Diplomacy is a very complex game and it is hard for participants to write a strong algorithm in the few months they have between the call for participation and the submission deadline. Before they could even start implementing they first needed to learn the rules of the game (which are fairly complex), learn the rules of the competition, and learn to work with the BANDANA framework. After that, they needed to come up with a smart algorithm, implement it, debug it, and optimize it.

Studying the source codes of the agents, we made two important observations:

1. Most agents never make any proposals for any of the future turns. They only make proposals for the current turn.
2. Many of the agents seem to have bugs in their code.

We think that both of these observations play an important role in the reason why the agents fail to negotiate successfully.

Any experienced Diplomacy player would agree that it is essential to plan several steps ahead. An important reason for this is that one does not often encounter a situation in which two players can both directly benefit from cooperation. Although it often happens that one player can give support to another player, it may then take several turns before a situation occurs in which the other player can return the favor. Therefore, it is essential that, in the short term, players are not purely selfish. They should be willing to help another player, while only expecting the favor to be returned at a later stage. Currently, none of the submitted agents seem to exhibit this kind of long term negotiation strategy.

Similarly, we think that the second observation is a very important one. As explained, the participants only have a limited amount of time to implement their agents, so perhaps we can only expect any participant to win the challenge after participating for several years. We noticed, for example, that due to a bug Frigate never accepted any proposals, even though it did implement an acceptance strategy. Also, Agent Madoff was more likely to accept a proposal if it involved a unit invading a province currently occupied by a Power that is considered a friend. We think that this is an error and that the author intended the opposite. Luckily, we see that two participants from 2017 have continued to participate in 2018, so the necessary drive seems to exist to commit to this long-term challenge.

In future editions of the Diplomacy Challenge, whenever negotiating agents play together with non-negotiating agents, we may need to make sure the negotiating agents play Powers that are more likely to form successful coalitions, as indicated by our experiments in Section 6.

8 Acknowledgments

This work is part of the Veni research programme with project number 639.021.751, which is financed by the Netherlands Organisation for Scientific Research (NWO), and project LOGISTAR, funded by the E.U. Horizon 2020 research and innovation programme, Grant Agreement No. 769142

References

1. Reyhan Aydođan, Katsuhide Fujita, Tim Baarlag, Catholijn M. Jonker, and Takayuki Ito. ANAC 2017: Repeated multilateral negotiation league. In *The 11th International Workshop on Automated Negotiation, ACAN 2018*, 2018.
2. Reyhan Aydođan, Catholijn M. Jonker, Katsuhide Fujita, Tim Baarlag, Takayuki Ito, Rafik Hadfi, and Kohei Hayakawa. *A Baseline for Nonlinear Bilateral Negotiations: The full results of the agents competing in ANAC 2014*, pages 96–122. Bentham Science Publishers, 2017.
3. Tim Baarlag, Reyhan Aydođan, Koen V. Hindriks, Katsuhide Fujita, Takayuki Ito, and Catholijn M. Jonker. The automated negotiating agents competition, 2010-2015. *AI Magazine*, 36(4):115–118, 12/2015 2015.
4. Tim Baarlag, Koen Hindriks, Catholijn M. Jonker, Sarit Kraus, and Raz Lin. The first automated negotiating agents competition (ANAC 2010). In *New Trends in Agent-based Complex Automated Negotiations, Series of Studies in Computational Intelligence*. Springer-Verlag, 2010.
5. Angela Fabregues. *Facing the Challenge of Human-aware Negotiation*. PhD thesis, Universitat Autònoma de Barcelona, 2012.
6. Angela Fabregues and Carles Sierra. Dipgame: a challenging negotiation testbed. *Engineering Applications of Artificial Intelligence*, 24(7):1137–1146, 2011.
7. André Ferreira, Henrique Lopes Cardoso, and Luís Paulo Reis. Diplblue: A diplomacy agent with strategic and trust reasoning. In *ICAART 2015 - Proceedings of the International Conference on Agents and Artificial Intelligence, Volume 1, Lisbon, Portugal, 10-12 January, 2015.*, pages 54–65. SciTePress, 2015.
8. Katsuhide Fujita, Reyhan Aydođan, Tim Baarlag, Takayuki Ito, and Catholijn M. Jonker. The fifth automated negotiating agents competition (ANAC 2014). In *Recent Advances in Agent-based Complex Automated Negotiation .*, volume 638 of *Studies in Computational Intelligence*, pages 211–224. Springer, 2014.
9. Dave de Jonge and Carles Sierra. NB3: a multilateral negotiation algorithm for large, non-linear agreement spaces with limited time. *Autonomous Agents and Multi-Agent Systems*, 29(5):896–942, 2015.
10. Dave de Jonge and Carles Sierra. D-brane: a diplomacy playing agent for automated negotiations research. *Applied Intelligence*, 47(1):158–177, 2017.
11. Dave de Jonge and Dongmo Zhang. Automated negotiations for general game playing. In *Proc. of the 16th Conf. on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017*, pages 371–379. ACM, 2017.
12. João Marinheiro and Henrique Lopes Cardoso. Towards general cooperative game playing. In Ngoc Thanh Nguyen, Ryszard Kowalczyk, Jaap van den Herik, Ana Paula Rocha, and Joaquim Filipe, editors, *Transactions on Computational Collective Intelligence XXVIII*, pages 164–192, Cham, 2018. Springer.
13. J. Mell, J. Gratch, T. Baarlag, R. Aydođan, and C. Jonker. Results of the first annual human-agent league of the automated negotiating agents competition. In *Proceedings of the 2018 Int. Conference on Intelligent Virtual Agents*, 2018.
14. E. Ephrati S. Kraus, D. Lehman. An automated diplomacy player. In D. Levy and D. Beal, editors, *Heuristic Programming in Artificial Intelligence: The 1st Computer Olympiad*, pages 134–153. Ellis Horwood Limited, 1989.
15. D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. P. Lillicrap, M. Leach, K. Kavukcuoglu, Th. Graepel, and D. Hassabis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.