Attila: a Negotiating Agent for the Game of Diplomacy, Based on Purely Symbolic A.I.

Extended Abstract

Dave de Jonge
IIIA-CSIC
Bellaterra (Barcelona), Spain
davedejonge@iiia.csic.es

Laura Rodriguez Cima
IIIA-CSIC
Bellaterra (Barcelona), Spain
laura.rodriguez@iiia.csic.es

ABSTRACT

The board game of Diplomacy is considered one of the most challenging test cases for automated negotiation. While many bots have been developed for this game, very few of them are able to negotiate successfully, and the ones that do have been trained on large data sets of human example games. This makes it hard to apply the same techniques to other games or negotiation scenarios for which no human knowledge is (yet) available. Furthermore, since those bots were trained using deep learning, they are essentially black-boxes for which it is hard to understand how they work. So, these bots do not help us much in gaining a better understanding of strong negotiation techniques. We therefore present a new Diplomacy bot, called Attila, that is purely based on symbolic A.I. techniques. It makes use of an existing oracle for the tactical part of the game, called the 'D-Brane Tactical Module' (DBTM). We explain how the DBTM can be converted into a search algorithm for automated negotiation and we present experiments that show that Attila strongly outperforms several state-of-the-art Diplomacy bots.

KEYWORDS

Automated Negotiation; Diplomacy

ACM Reference Format:
Dave de Jonge and Laura Rodriguez Cima. 2024. Attila: a Negotiating Agent for the Game of Diplomacy, Based on Purely Symbolic A.I.: Extended Abstract. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 3 pages.

1 INTRODUCTION

Diplomacy is a board game for seven players, which need to form coalitions and negotiate with each other. Each player has a number of armies that are placed on a map of Europe, and the goal is to conquer the so-called *Supply Centers* (SC) on this map. One can say that Diplomacy consists of two 'layers'; a *tactical layer*, involving the players moving their pieces over the board, and a *negotiation layer*, involving the players negotiating with each other about those moves. Good players need to master both aspects of the game.

Although Diplomacy has been studied by A.I. researchers for a long time [13, 14], most bots that have been developed for this game are not able to negotiate [11, 16, 18, 19], or have only been shown to outperform a very simple bot called DumbBot [7, 10, 17, 20]. One of the strongest bots that have been developed is called D-Brane [6]. This bot is much stronger than DumbBot and it was shown that, even if D-Brane did not negotiate, it would still outperform some of the above mentioned negotiating bots. However, it was also shown that with negotiations D-Brane only performed marginally better than without negotiations.

D-Brane consists of two separate modules, namely a tactical module and a negotiation module, that each focus on the two respective layers of the game. Its tactical module (the DBTM) was made publicly available to allow other researchers to implement new negotiation algorithms on top of it. This allows researchers to focus purely on the development of negotiation algorithms without having to care about the tactical aspects of the game.

Since then, many bots have been implemented by adding a new negotiation algorithm on top of the DBTM, but none of them resulted in a clear increase of performance with respect to the non-negotiating D-Brane. One such agent, called AlphaDip [15], did show some improvement over D-Brane, but the authors still concluded that this improvement was very small. To increase the attention of the automated negotiations community to the game of Diplomacy, the annual Automated Negotiating Agents Competition (ANAC) hosted a special Diplomacy league from 2017 till 2019. However, none of the agents that were submitted to this competition managed to clearly improve over the non-negotiating D-Brane [1, 3].

A major breakthrough, however, was made recently by a team of researchers from Meta, who implemented a negotiating Diplomacy bot that was able to beat human players [8]. Their approach used a deep-learning model that was trained on a database of games played by humans. Similarly, a team from DeepMind also implemented a successfully negotiating bot, trained on human data [12].

While the key challenge of Diplomacy has now been tackled using deep learning, this approach has two downsides. Firstly, the aforementioned bots are essentially black-boxes, which makes it very hard to tell *why* they are successful. They therefore do not help us much in gaining more insight into the topic of automated negotiation.[1] Secondly, the fact that both of these bots required a large corpus of human data to be trained means that one cannot follow the same approach to apply negotiation to other games or use cases for which no such data is (yet) available.

We therefore present a new negotiating Diplomacy player, called Attila (Advanced negoTiaTing dIpLomAcy bot), that is based on

[1] One could try to gain such insight using techniques from Explainable A.I., but to the best of our knowledge this has not been done for those bots.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand.

purely symbolic A.I. techniques, and we show that it outperforms several state-of-the-art Diplomacy bots. The source code of Attila has been made publicly available at:

https://www.iiia.csic.es/~davedejonge/downloads.

2 ATTILA

For the implementation Attila we followed an approach that was suggested in [3]. That is, we applied a trick that allowed us to use the DBTM not only for the tactical layer of the game, but also for the negotiation layer. This trick consists in pretending that we are controlling the armies of two different players (the one played by our agent, plus its coalition partner) and asking the DBTM to find the best moves for this combined player. These moves can then be proposed to the coalition partner as a potential deal. This idea may seem simple, but, to the best of our knowledge, no one else has tried this approach in the game of Diplomacy before. Given the proposal found by the DBTM we then make small random mutations to it, to generate more alternative proposals. In order to determine when to propose which proposal, Attila applies a combination of a time-based strategy [9] and the MiCRO strategy [2].

While the implementation of our bot is entirely specific to Diplomacy, we think that the underlying ideas can be generalized to other games or negotiation scenarios as long as one has access to some algorithm to determine good actions for an individual non-negotiating agent. For example, in [4, 5] a scenario was described in which logistics companies negotiated to exchange their truck loads. In this scenario one can use the same approach by pretending all truck loads belong to the same company, and then finding the optimal solution for that one company.

3 EXPERIMENTS

Our experiments were implemented on the BANDANA framework [6], and were conducted on a laptop with 11th Gen Intel Core i7-1165G7@2.80GHz CPU, 16 GB RAM, and Windows 11. We made sure that each game was automatically stopped after 80 rounds (i.e. after the Winter 1940 phase) and solo-victories were disabled (so a game would continue even if one of the players had 18 or more SCs). Furthermore, any agreement between players was considered officially binding, so players always had to obey their agreements. Deadlines were set to 3 seconds per round. We verified that all conclusions were statistically significant (with $p < 0.01$).

Our first experiment involved 4200 games with two instances of Attila and five instances of the non-negotiating D-Brane and we observed how many SCs each agent captured. The results are displayed in the right-hand column of Table 1. We see that the two Attilas each score an average of 8.40 SCs per game, while the non-negotiating D-Branes each only score 3.44 SCs. So, we conclude that Attila is much better than D-Brane.

Next, we repeated the experiment, but this time the Attilas did not negotiate. Instead, we only configured them to not attack each other. The results of this experiment are displayed in the center column of Table 1. We see that in this case the two Attilas each score an average of 5.86 SCs per game, while the non-negotiating D-Branes each only score 4.46 SCs. While this is still an improvement, it is a lot smaller than the improvement achieved by the fully negotiating Attilas. From this we conclude that the improvement of Attila over D-Brane is mainly due to the negotiation among the Attilas, rather than simply due to the fact that they do not attack each other.

Agent	SCs (Peace)	SCs (Negotiation)
Attila	5.86 ± 0.07	8.40 ± 0.08
D-Brane	4.46 ± 0.03	3.44 ± 0.03

Table 1: The number of SCs scored by each agent (averaged per game and agent instance) in two settings with two instances of Attila against five instances of D-Brane.

Agent	SCs
Attila	6.80 ± 0.20
D-Brane	4.55 ± 0.08
Gunma	3.80 ± 0.12

Table 2: The number of SCs scored by each agent (averaged per game and agent instance), in an experiment with one instance of Attila, one instance of Gunma, and five instances of D-Brane.

In our second experiment we played 420 games with one instance of Attila, one instance of another negotiating agent called Gunma, which participated in the ANAC Diplomacy Leage of 2018, and five instances of the non-negotiating D-Brane. We observed the average number of SCs captured by each of these agents (per game and per agent instance) and displayed the results in Table 2.

Finally, we ran an experiment with two instances of Attila, two instances of Gunma, two instances of Saitama (another participant of the ANAC Diplomacy League) and one instance of D-Brane. It should be stressed that the agents were anonymous, so Attila did not know which opponent was also a copy of Attila. The results are displayed in Table 3.

From these last two experiments we see that Attila clearly outperforms the other agents negotiating agents.

Agent	SCs
Attila	8.0 ± 0.72
Gunma	4.6 ± 0.68
D-Brane	4.1 ± 1.08
Saitama	2.3 ± 0.31

Table 3: The number of SCs scored by each agent (averaged per game and per agent instance), in an experiment with two instances of Attila, two of Gunma, two of Saitama, and one D-Brane.

ACKNOWLEDGMENTS

This work was supported by a Juan de la Cierva research grant (IJC2018-036443-I) from the Spanish Ministry of Science and Innovation, by a JAE-INTRO-ICU grant (JAEIntroICU-2021-IIIA-09) funded by CSIC, and by Grant no. TED2021-131295B-C31 funded by MCIN/AEI/10.13039/501100011033 and by the European Union NextGenerationEU/PRTR.

REFERENCES

[1] Reyhan Aydoğan, Tim Baarslag, Katsuhide Fujita, Johnathan Mell, Jonathan Gratch, Dave de Jonge, Yasser Mohammad, Shinji Nakadai, Satoshi Morinaga, Hirotaka Osawa, Claus Aranha, and Catholijn M. Jonker. 2020. Challenges and Main Results of the Automated Negotiating Agents Competition (ANAC) 2019. In *Multi-Agent Systems and Agreement Technologies - 17th European Conference, EUMAS 2020, and 7th International Conference, AT 2020, Thessaloniki, Greece, September 14-15, 2020, Revised Selected Papers*, Nick Bassiliades, Georgios Chalkiadakis, and Dave de Jonge (Eds.). Springer International Publishing, Cham, 366–381.

[2] Dave de Jonge. 2022. An Analysis of the Linear Bilateral ANAC Domains Using the MiCRO Benchmark Strategy. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, Luc De Raedt (Ed.). ijcai.org, 223–229. https://doi.org/10.24963/ijcai.2022/32

[3] Dave de Jonge, Tim Baarslag, Reyhan Aydoğan, Catholijn Jonker, Katsuhide Fujita, and Takayuki Ito. 2019. The Challenge of Negotiation in the Game of Diplomacy. In *Agreement Technologies, 6th International Conference, AT 2018, Bergen, Norway, December 6-7, 2018, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 11327)*, Marin Lujak (Ed.). Springer International Publishing, Cham, 100–114. https://doi.org/10.1007/978-3-030-17294-7_8

[4] Dave de Jonge, Filippo Bistaffa, and Jordi Levy. 2021. A Heuristic Algorithm for Multi-Agent Vehicle Routing with Automated Negotiation. In *AAMAS '21: 20th International Conference on Autonomous Agents and Multiagent Systems, Virtual Event, United Kingdom, May 3-7, 2021*, Frank Dignum, Alessio Lomuscio, Ulle Endriss, and Ann Nowé (Eds.). ACM, 404–412. https://dl.acm.org/doi/10.5555/3463952.3464004

[5] Dave de Jonge, Filippo Bistaffa, and Jordi Levy. 2022. Multi-Objective Vehicle Routing with Automated Negotiation. *Applied Intelligence* 52, 14 (01 Nov 2022), 16916–16939. https://doi.org/10.1007/s10489-022-03329-2

[6] Dave de Jonge and Carles Sierra. 2017. D-Brane: a Diplomacy Playing Agent for Automated Negotiations Research. *Applied Intelligence* 47, 1 (2017), 158–177. https://doi.org/10.1007/s10489-017-0919-y

[7] Angela Fabregues. 2012. *Facing the Challenge of Human-aware Negotiation.* Ph.D. Dissertation. Universitat Autònoma de Barcelona.

[8] Meta Fundamental AI Research Diplomacy Team (FAIR)†, Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, Athul Paul Jacob, Mojtaba Komeili, Karthik Konath, Minae Kwon, Adam Lerer, Mike Lewis, Alexander H. Miller, Sasha Mitts, Adithya Renduchintala, Stephen Roller, Dirk Rowe, Weiyan Shi, Joe Spisak, Alexander Wei, David Wu, Hugh Zhang, and Markus Zijlstra. 2022. Human-level play in the game of Diplomacy by combining language models with strategic reasoning. *Science* 378, 6624 (2022), 1067–1074. https://doi.org/10.1126/science.ade9097 arXiv:https://www.science.org/doi/pdf/10.1126/science.ade9097

[9] Peyman Faratin, Carles Sierra, and Nicholas R. Jennings. 1998. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems* 24, 3-4 (1998), 159 – 182. https://doi.org/10.1016/S0921-8890(98)00029-3 Multi-Agent Rationality.

[10] André Ferreira, Henrique Lopes Cardoso, and Luís Paulo Reis. 2015. DipBlue: 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.* SciTePress, 54–65.

[11] Michael R Hall and Daniel E Loeb. 1992. Thoughts on programming a diplomat. *Heuristic Programming in Artificial Intelligence* 3, 9 (1992), 123–145.

[12] János Kramár, Tom Eccles, Ian Gemp, Andrea Tacchetti, Kevin R McKee, Mateusz Malinowski, Thore Graepel, and Yoram Bachrach. 2022. Negotiation and honesty in artificial intelligence methods for the board game of Diplomacy. *Nature Communications* 13, 1 (2022), 7214.

[13] Sarit Kraus, Daniel Lehman, and Eithan Ephrati. 1989. An automated diplomacy player. In *Heuristic Programming in Artificial Intelligence: The 1st Computer Olympia*, D. Levy and D. Beal (Eds.). Ellis Horwood Limited, 134–153.

[14] Sarit Kraus and Daniel Lehmann. 1995. Designing and building a negotiating automated agent. *Computational Intelligence* 11 (1995), 132–171.

[15] João Marinheiro and Henrique Lopes Cardoso. 2018. Towards General Cooperative Game Playing. In *Transactions on Computational Collective Intelligence XXVIII*, Ngoc Thanh Nguyen, Ryszard Kowalczyk, Jaap van den Herik, Ana Paula Rocha, and Joaquim Filipe (Eds.). Springer, Cham, 164–192.

[16] João Ribeiro, Pedro Mariano, and Luís Seabra Lopes. 2009. DarkBlade: A Program That Plays Diplomacy. In *Progress in Artificial Intelligence, 14th Portuguese Conference on Artificial Intelligence, EPIA 2009, Aveiro, Portugal, October 12-15, 2009. Proceedings (Lecture Notes in Computer Science, Vol. 5816)*, Luís Seabra Lopes, Nuno Lau, Pedro Mariano, and Luis M. Rocha (Eds.). Springer, 485–496. https://doi.org/10.1007/978-3-642-04686-5_40

[17] Jaspreet Shaheed. 2004. Creating a diplomat. *Master's thesis, Department of Computing, Imperial College Of Science, Technology and Medicine* 180 (2004).

[18] Ari Shapiro, Gil Fuchs, and Robert Levinson. 2003. Learning a game strategy using pattern-weights and self-play. In *Computers and Games: Third International Conference, CG 2002, Edmonton, Canada, July 25-27, 2002. Revised Papers 3*. Springer, 42–60.

[19] Alexios Theodoridis and Georgios Chalkiadakis. 2020. Monte Carlo Tree Search for the Game of Diplomacy. In *SETN 2020: 11th Hellenic Conference on Artificial Intelligence, Athens, Greece, September 2-4, 2020*, Constantine D. Spyropoulos, Iraklis Varlamis, Ion Androutsopoulos, and Prodromos Malakasiotis (Eds.). ACM, 16–25. https://doi.org/10.1145/3411408.3411413

[20] Adam Webb, Jason Chin, Thomas Wilkins, John Payce, and Vincent Dedoyard. 2008. Automated negotiation in the game of diplomacy.