

ANAC 2018 Diplomacy Challenge

Official Results

Dave de Jonge

July 19, 2018

Goal of the Challenge

The ANAC 2018 Diplomacy Challenge consisted of two rounds. In order to win the Challenge an agent had to outperform the non-negotiating D-Brane in Round 1, as well as beat all the other agents in Round 2.

Submissions

We received the following submissions:

- **CoalitionBot**, by Ido Westler, Yehuda Callen, Moche Uzan, Arie Cattan, and 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, DDAgent2 was too slow to participate in the competition, because in many rounds it was not able to submit its orders within the deadline.

Results Round 1

Each agent played 100 games with 4 instances of itself and 3 instances of the non-negotiating D-Brane agent. We counted the average number of Supply Centers conquered per agent instance per game. We say an agent *passes Round 1* if it scores a statistically significant higher average than D-Brane.

To test for significance we applied a one-sided Student-t test with respect to the null hypothesis that the submitted agent has a true mean score of $\frac{34}{7} = 4.857$ Supply Centers per game.

	Average	Std. Error	Result
CoalitionBot	5.528	0.110	PASS
D-Brane	3.963	0.146	
Gunma	4.95	0.128	FAIL
D-Brane	4.733	0.171	
D-Brane	4.93	0.164	
M@sterMind	4.803	0.123	FAIL
D-Brane	5.44	0.184	
GamlBot	4.42	0.138	FAIL

Table 1: Results of Round 1

Only CoalitionBot and Gunma were able to outperform the D-Brane. However, the results of Gunma were not significant (p-value 0.23). Therefore, CoalitionBot was the only agent to pass Round 1 (p-value $9.7 \cdot 10^{-9}$).

Results Round 2

In Round 2 we played 100 games with 1 instance of each submitted agent together with 3 instances of the non-negotiating D-Brane. We see that when negotiating with other agents the CoalitionBot is not able to significantly outperform the other agents.

	Average	Std. Error
1. Gunma	5.69	0.3
2. GamlBot	5.31	0.334
3. CoalitionBot	4.94	0.289
D-Brane	4.54	0.168
4. M@sterMind	4.44	0.29

Table 2: Results of Round 2

We see that Gunma performed best in Round 2, but the difference between the first three agents is non-significant.

Conclusions

The CoalitionBot was the only agent that passed Round 1, but it did not manage to beat all the other agents in Round 2. Therefore, we conclude that:

The ANAC 2018 Diplomacy Challenge ends with NO WINNER