MADDPG Reinforcement Learning for Active Wake Control

Guus van der Schaaf (g.vanderschaaf@student.tudelft.nl) - Supervisors: Mathijs de Weerdt, Greg Neustroev - EEMCS, Delft University of Technology, The Netherlands

1. Introduction and Background

- Wind turbines in wind farms generate a turbulent air stream behind them, the so called wake.
- This wake can negatively affect the next turbine in the row.
- By steering the turbine this wake field can be redirected.
- This decreases the power output of the single turbine. but can increase the total power output of the whole farm.
- A optimal solution is a solution where the total output power of the whole farm is maximized This optimal solution can be derrived by a numerical algorithm, named Floris [1].
- Using Floris is not feasible for larger problems (larger than 6 turbines), since the action space grows exponentially.

A traditional Deep Reinforcement Learning algorithm has also been used in the past [2], this booked great results, but also suffered from the exponential growth of larger farms.

My research question is:

"Can we apply a Multi Agent Deep Deterministic Policy Gradient algorithm to the Active Wake Control Problem successfully?"

2. Methodology

- I will use a Multi Agent Deep Deterministic Policy Gradient algorithm.
- This algorithm uses a seperate agent per wind turbine.
- Each agent has 2 neural networks, the actor and the critic.
- The actor uses the current observation (o) of the environment to choose an appropriate action (a).
- This action is executed and a reward is returned from the environment.
- The critic uses the current state and the taken actions and approximates the reward (Q). This is used to update the actor model.

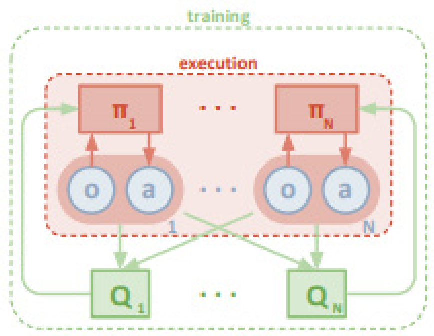

Figure 1. Overview of a multi-agent actor critic approach [3]

3. Results

- We ran MADDPG on the real Princess Amalia Windfarm
- This wind farm consists of 60 wind turbines
- It can be observed that MADDPG learns in less training steps
- These training steps do take significantly longer

Algorithm	3 turbine	4x4 grid	Amalia
TD3	11ms	31ms	217ms
MADDPG	82ms	418ms	1680ms
TD3 Corrected	4ms	2ms	4ms
MADDPG Corrected	27ms	26ms	28ms

Table 1. Time per training step

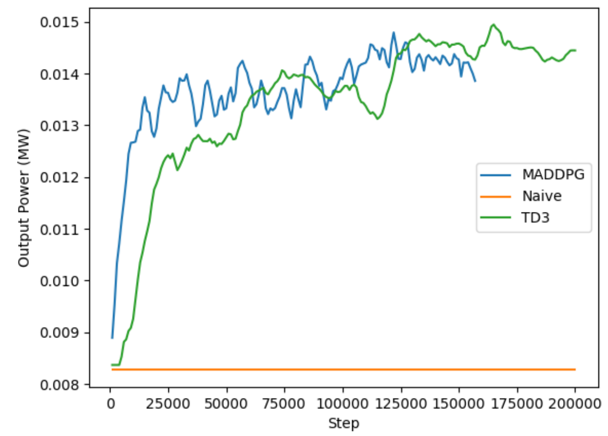

Figure 2. A graph of the mean reward per training step

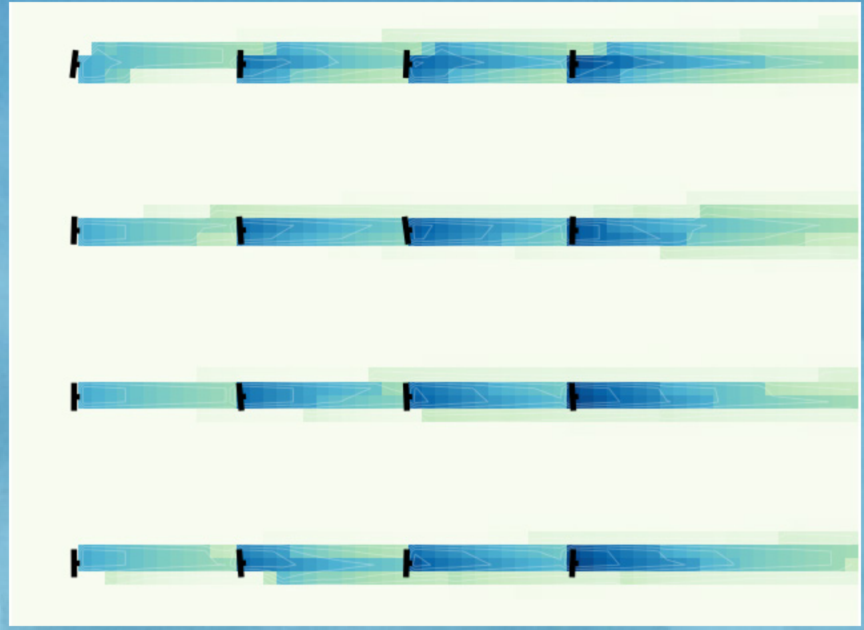

Figure 3. An example of a 4x4 windfarm with wakes (generated by FLORIS visualizer)

4. Conclusion and Future work

MADDPG seems to be working quite well on the problem. The next steps are:

- Running on the bigger (and more interesting models). The 4x4 layout in figure 3 and as a final on the Princess Amalia windfarm.

References
[1] NREL. 2021. FLORIS. Version 2.4. https://github.com/NREL/floris
[2] Grigory Neustroev, Sytze P. E. Andringa, Remco A. Verzijlbergh, and Mathijs M. de Weerdt. 2022. Deep Reinforcement Learning for Active Wake Control
[3] Ryan Lowe et al. Multi-agent actor-critic for mixed cooperative-competitive environments. In: Advances in neural information processing systems 30 (2017) page 4
[4] "Will wind-wake slow industry's ambitions offshore?," Recharge | Latest Renewable Energy News, Nov. 02, 2019. https://www.rechargenews.com/wind/will-wind-wake-slow-industrys-ambitions-offshore-/2-1-699430