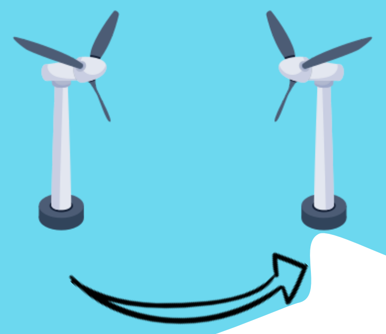

Mean-Field Multi Agent Reinforcement Learning for Active Wake Control

Ion Plămădeală
i.plamadeala@student.tudelft.nl
Supervised by:
Mathijs de Weerdt
Greg Neustroev

1. Active Wake Control

Turbines **placed behind one another** are affected by the **wake** of the upstream turbine, which causes a reduction in power output and an increase in fatigue loads. This can be mitigated by **turning the upstream turbine slightly out of the wind**, which is called active wake control. The total power output is overall increased, and the fatigue loads are reduced.[1]

2. Reinforcement Learning

Reinforcement Learning is an ML technique that learns the best actions that an agent (**turbine/wind farm**) can take in a given situation (**wind speed, wind direction, neighbouring turbines, etc.**) to maximize a reward (**power output**). The single agent solution is effective for a small number of turbines, due to the combination explosion of possible states and actions.

Mean Field MARL in a Diagram

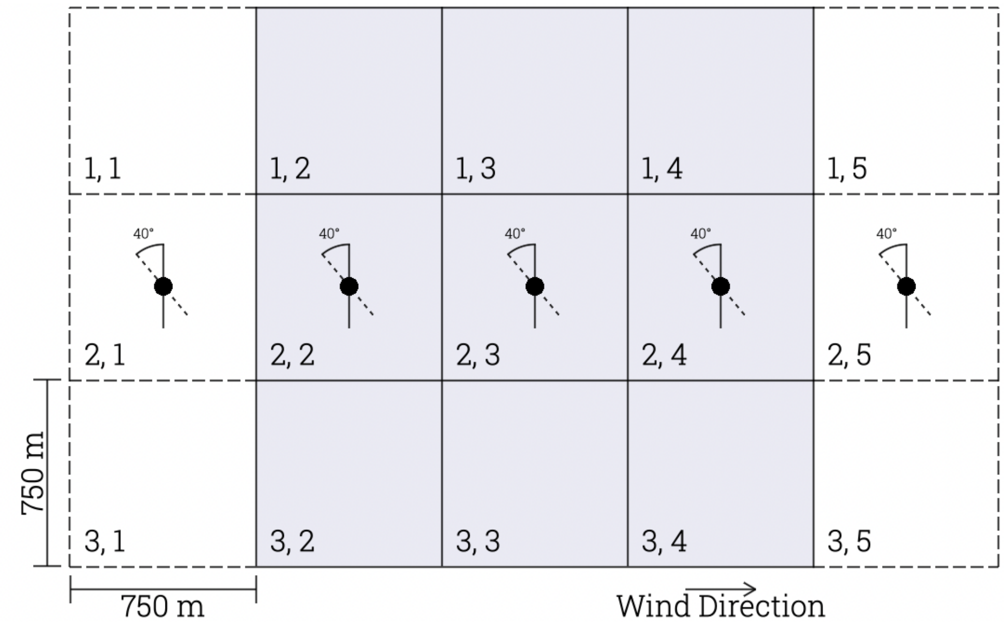

3. Mean-Field MARL

Mean-Field MARL is a technique that reduces the complexity by including only a subset of neighbors in each turbine's computation. It has been used for the Ising Model, and a 64 vs 64 battle agent.[2] The question for the Research Project is: **Can Mean-Field MARL be used for Active Wake Control?**, and **How does it compare against TD3?**. I've also researched how the **reward and view space** affect the performance (training time and power output) of the model.

4. Methodology

The research uses FLORIS, a wind wake simulator. In the experiments below, I used a wind tunnel with 3 turbines, and 4 wind tunnels of 4 turbines in parallel. I run each experiment for 3000-4000 episodes (between 2h and 12h), with around 150 steps per episode. I varied the reward in different experiments, from the total power output to global delta between power output for each step, and, in the mean-field spirit, a limited delta.

Results

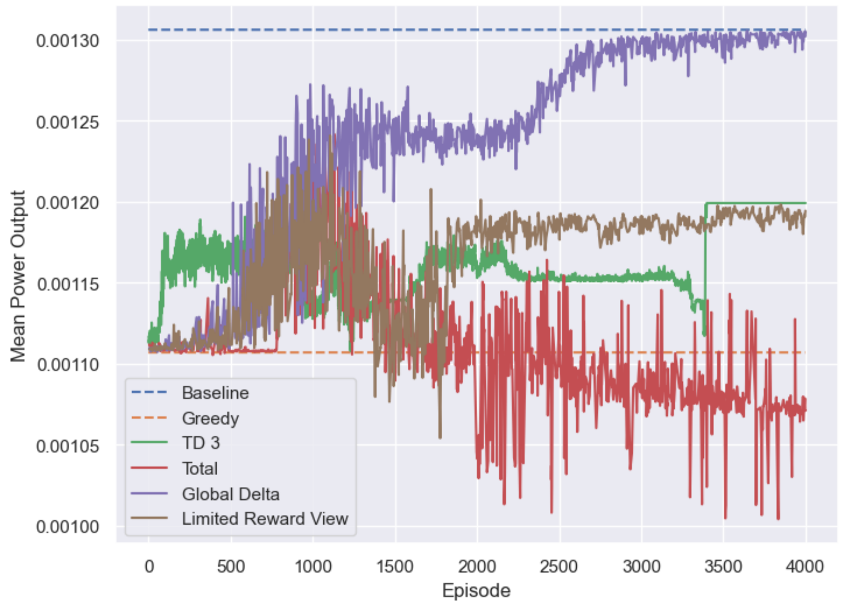

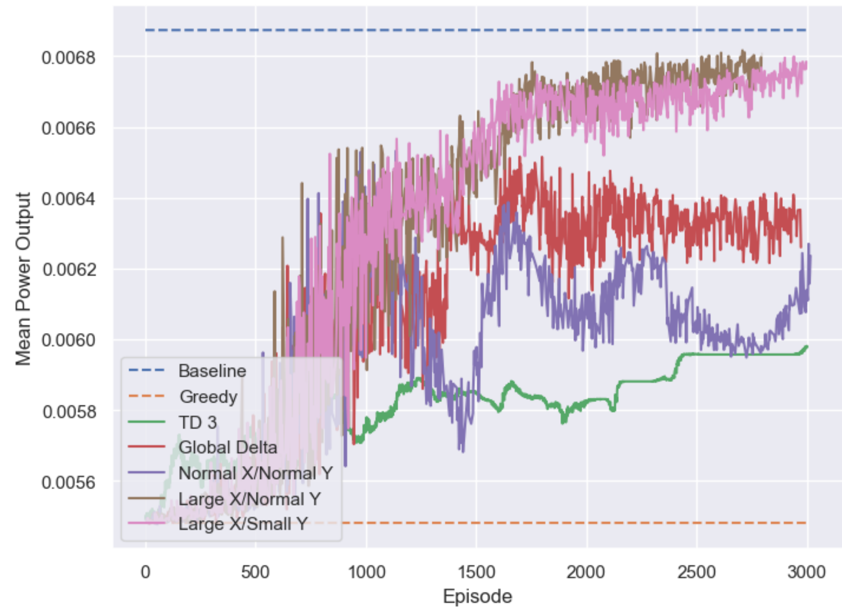

5. Conclusion

After 3000-4000 episodes, the model converges to near-optimal yaw configuration. While these results are promising, they are limited to parallel wind tunnels. The next step is researching real-world layouts and changing wind directions.

[1] Kanev, S. K., Savenije, F. J., & Engels, W. P. (2018). Active wake control: An approach to optimize the lifetime operation of wind farms. Wind Energy, 21(7), 488-501.
[2]Yang, Y., Luo, R., Li, M., Zhou, M., Zhang, W., & Wang, J. (2018, July). Mean field multi-agent reinforcement learning. In International conference on machine learning (pp. 5571-5580). PMLR.