Deep Reinforcement Learning for Active Wake Control*

Grigory Neustroev[1][0000−0002−7706−7778],
Sytze P. E. Andringa[1][0000−0003−4061−7104], Remco A. Verzijlbergh[1,2], and
Mathijs M. de Weerdt[1][0000−0002−0470−6241]

[1] Delft University of Technology,
P.O. Box 5031, 2600 GA, Delft, the Netherlands
{g.neustroev, s.p.e.andringa, r.a.verzijlbergh, m.m.deweerdt}@tudelft.nl
[2] Whiffle,
Molengraaffsingel 12–14, 2629 JD, Delft, the Netherlands
remco.verzijlbergh@whiffle.nl

This is an abstract of a paper published at the 21st International Conference on Autonomous Agents and Multiagent Systems [9].

While deep reinforcement learning (RL) has been successful in playing games [4, 11], its real-life applications remain limited. To bridge this gap, we discuss a problem of active wake control (AWC) in wind farms and present a simulator that can be used to address this problem using deep RL.

When a turbine extracts energy from the wind, it creates a wake behind its rotor. Wakes can extend for several kilometers and thus negatively impact other turbines in a wind farm (see Fig. 1) [14]. Field studies report 7%–13% wake losses for current wind farms [1], but potentially much higher for future wind farms larger than 1 GW [16]. Mitigation of these wake-induced losses thus becomes an increasingly important aspect of wind farm operation. A popular approach to AWC is to simulate wakes for different combinations of turbine yaws (i.e., horizontal-plane angles) [6, 5]. Optimization is then done numerically per so-called steady-state conditions in which the environment of the problem does not change. At the same time, RL is especially well-suited for AWC, as it includes stochastic transitions between states—something that steady-state models miss. Nevertheless, studies of RL for AWC remain limited [15, 12, 3].

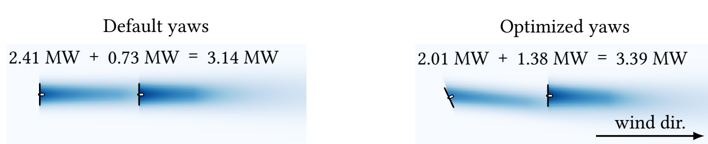

Fig. 1. An example of active wake control.

* This research received funding from the Netherlands Organization for Scientific Research (NWO).

In this paper, we present a novel dynamic wind farm simulator available at https://doi.org/10.4121/19107257. To facilitate future research in AWC from the RL community, our simulator uses the OpenAI *Gym* format [2]—a standard in RL. Our simulator is designed with real-life wind farm operation in mind and includes not only turbines, but various other equipment, such as meteorological masts and nacelle-mounted meters. This allows the simulations to better reflect the reality of wind farm operation.

RL models assume that control is performed at discrete time steps. To simulate the wakes in each time step, we use the state-of-the-art steady-state framework called FLORIS [10]. Changes of the environmental data (e.g., wind speed and direction) between steady states are simulated by a stochastic process defined by the user. We implemented one such process, namely the multivariate Ornstein–Uhlenbeck process [7, 13] estimated from publicly available meteorological observations at the *Hollandse Kust Noord* offshore wind farm zone [8].

In addition to introducing the simulator, we conducted two experiments. In the first experiment, we considered three possible ways to encode actions based on the current yaw, absolute angle, and the incoming wind direction respectively, only the first of which appeared in the previous studies of RL for AWC. We compared the performances of two state-of-the-art RL algorithms for each representation and showed that the default yaw-based encoding is outperformed by the two alternatives. These new action representations are included in our simulator and should be adopted in the future research of RL for AWC.

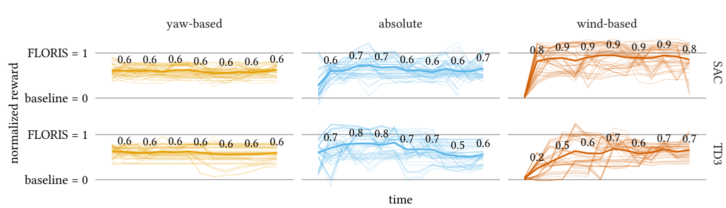

Fig. 2. Relative performance of two deep RL algorithms using different action representations. Here 0 stands for no improvement over the baseline control strategy of facing the wind with all turbines, and 1 for the maximum possible improvement.

In the second experiment, we demonstrated the benefits of RL compared to model-based control. As the scale of noise injected into the observations grows, model-based optimization struggles to outperform the baseline strategy, dropping from 9.5% improvement over the baseline to just 0.2%. While RL also suffers from the added noise, its performance improvement varies between 8.5% and 7.4%, giving a statistically significant improvement over model-based control.

Our findings show that deep RL is a promising approach to AWC. We hope that this work sparks interest of the RL community in this problem, and that our results will make it easier for other researchers to develop new RL-based methods for active wake control.

References

1. Barthelmie, R.J., Frandsen, S., Hansen, K., Schepers, J., Rados, K., Schlez, W., Neubert, A., Jensen, L., Neckelmann, S.: Modelling the impact of wakes on power output at Nysted and Horns Rev. In: European Wind Energy Conference. vol. 2, pp. 1351–1373. WindEurope, Marseille, France (2009)
2. Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., Zaremba, W.: OpenAI gym. https://gym.openai.com (2016)
3. Dong, H., Zhang, J., Zhao, X.: Intelligent wind farm control via deep reinforcement learning and high-fidelity simulations. Applied Energy **292**, 116928 (2021)
4. Hessel, M., Modayil, J., van Hasselt, H., Schaul, T., Ostrovski, G., Dabney, W., Horgan, D., Piot, B., Azar, M.G., Silver, D.: Rainbow: Combining improvements in deep reinforcement learning. Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence **32**(1), 3215–3222 (2018)
5. Houck, D.R.: Review of wake management techniques for wind turbines. Wind Energy **25**, 195–220 (2022). https://doi.org/https://doi.org/10.1002/we.2668, https://onlinelibrary.wiley.com/doi/abs/10.1002/we.2668
6. Howland, M.F., Lele, S.K., Dabiri, J.O.: Wind farm power optimization through wake steering. Proceedings of the National Academy of Sciences **116**(29), 14495–14500 (2019)
7. Meucci, A.: Review of statistical arbitrage, cointegration, and multivariate Ornstein–Uhlenbeck. https://ssrn.com/abstract=1404905 (2009)
8. Netherlands Enterprise Agency, M.o.E.A., Policy, C.: Hollandse Kust Noord (site B) dataset (Aug 2019), retrieved September 30 from https://offshorewind.rvo.nl/file/view/55040229/Processed+data+HKNB
9. Neustroev, G., Andringa, S.P.E., Verzijlbergh, R.A., de Weerdt, M.M.: Deep reinforcement learning for active wake control. In: Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems. pp. 944–953. International Foundation for Autonomous Agents and Multiagent Systems, Richland, South Carolina, USA (2022)
10. NREL: FLORIS. version 2.4 (2021), https://github.com/NREL/floris
11. Silver, D., Huang, A., Maddison, C., Guez, A., Sifre, L., Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., Hassabis, D.: Mastering the game of Go with deep neural networks and tree search. Nature **529**, 484–489 (2016)
12. Stanfel, P., Johnson, K., Bay, C.J., King, J.: Proof-of-concept of a reinforcement learning framework for wind farm energy capture maximization in time-varying wind. Journal of Renewable and Sustainable Energy **13**(4), 14 (2021)
13. Vatiwutipong, P., Phewchean, N.: Alternative way to derive the distribution of the multivariate Ornstein–Uhlenbeck process. Advances in Difference Equations **2019**(276), 7 (2019)
14. Vermeer, L.N.J., Sørensen, J., Crespo, A.: Wind turbine wake aerodynamics. Progress in Aerospace Sciences **39**(6), 467–510 (2003)
15. Verstraeten, T., Libin, P., Nowé, A.: Fleet control using coregionalized Gaussian process policy iteration. In: 24th European Conference on Artificial Intelligence. vol. 325, pp. 1571–1578. IOS Press, Santiago de Compostela, Spain (2020)
16. Verzijlbergh, R.: Atmospheric flows in large wind farms. Europhysics News **52**, 20–23 (2021). https://doi.org/10.1051/epn/2021502