
WFCRL: A Multi-Agent Reinforcement Learning Benchmark for Wind Farm Control

Claire Bizon Monroc

Inria and DI ENS, École Normale Supérieure, PSL Research University, Paris, France
IFP Energies nouvelles
claire.bizon-monroc@inria.fr

Ana Bušić

Inria and DI ENS, École Normale Supérieure, PSL Research University
Paris, France

Donatien Dubuc

IFP Energies nouvelles
Solaize, France

Jiamin Zhu

IFP Energies nouvelles
Rueil-Malmaison, France

Abstract

1 The wind farm control problem is challenging, since conventional model-based
2 control strategies require tractable models of complex aerodynamical interactions
3 between the turbines and suffer from the curse of dimension when the number of
4 turbines increases. Recently, model-free and multi-agent reinforcement learning
5 approaches have been used to address this challenge. In this article, we introduce
6 WFCRL (Wind Farm Control with Reinforcement Learning), the first suite of multi-
7 agent reinforcement learning environments for the wind farm control problem.
8 WFCRL frames a cooperative Multi-Agent Reinforcement Learning (MARL)
9 problem: each turbine is an agent and can learn to adjust its yaw, pitch or torque
10 to maximize the common objective (e.g. the total power production of the farm).
11 WFCRL also offers turbine load observations that will allow to optimize the
12 farm performance while limiting turbine structural damages. Interfaces with two
13 state-of-the-art farm simulators are implemented in WFCRL: a static simulator
14 (FLORIS) and a dynamic simulator (FAST.Farm). For each simulator, 10 wind
15 layouts are provided, including 5 real wind farms. Two state-of-the-art online
16 MARL algorithms are implemented to illustrate the scaling challenges. As learning
17 online on FAST.Farm is highly time-consuming, WFCRL offers the possibility of
18 designing transfer learning strategies from FLORIS to FAST.Farm.

19 1 Introduction

20 The development of wind energy plays a crucial part in the global transition away from fossil energies,
21 and it is driven by the deployment of very large offshore wind farms [38, 28]. Significant gains
22 in wind energy production can be made by increasing the amount of wind power captured by the
23 farms [28]. The power production of a wind farm is greatly influenced by wake effects: an operating
24 upstream turbine causes a decrease in wind velocity and an increase in wind turbulence behind its
25 rotor, which creates sub-optimal wind conditions for other wind turbines downstream. An illustration




Figure 1: Left: Wake effects in the offshore wind farm of Horns Rev 1 - Vattenfall. Right: Schema of a wind turbine [6]. The *pitch*, *yaw* or *torque* can be controlled.

26 of this phenomenon can be seen on Figure 1. Wake effects are a major cause of power loss in wind
 27 farms, with the decrease in power output estimated to be between 10% and 20% in large offshore
 28 wind farms [4]. Higher turbulence in wakes also increases fatigue load on the downstream turbines
 29 by 5% to 15%, which can shorten their lifespans [37].

30 The wind farm control problem is challenging. Conventional model-based control strategies require
 31 tractable models of complex dynamic interactions between turbines, and suffer from the curse
 32 of dimensionality when the number of turbines increases. Moreover, optimal strategies differ
 33 significantly with modeling choices. Reinforcement Learning (RL) provides a model-free, data-based
 34 alternative, and recent work applying RL algorithms to wind farm control has yielded promising
 35 results (see e.g. [1, 25]). Single agent approaches, where a single RL controller must learn a
 36 centralized policy, encounter scaling challenges [10], are slow to converge under dynamic conditions
 37 [23] and do not explore the graph structure of the problem induced by local perturbations. Several
 38 multi-agent RL approaches have been proposed to tackle this issue, relying on both centralized
 39 critics [9, 10, 27] and independent learning approaches [5, 21, 33]. Different authors have published
 40 code relative to their specific applications, but there is to the best of our knowledge no open-source
 41 reinforcement learning environment for the general wind farm control problem.

42 In this article, we propose WFCRL, the first suite of reinforcement learning environments for the
 43 wind farm control problem. WFCRL is highly customizable, allowing researchers to design and run
 44 their own environments for both centralized and multi-agent RL.

45 Wind turbines can be controlled in several ways. A turbine can adjust its *yaw* (defined as the angle
 46 between the rotor and the wind direction) to deflect its wake, increase its *pitch* (the angle between the
 47 turbine blades and the incoming wind) to decrease its wind energy production, or directly control the
 48 *torque* of its rotor. WFCRL makes it possible to control yaw, pitch or torque, and a schema of these
 49 different control variables can be found in fig. 1. WFCRL offers a large set of observations including
 50 local wind statistics, power production, and fatigue loads for each turbine. This makes it possible to
 51 consider different objective, including the maximization of the total production, the minimization of
 52 loads to reduce maintenance costs over the wind turbine life-cycle [22], or, as wind energy becomes a
 53 larger part of the energy mix, the tracking of power or frequency targets that will allow operators to
 54 offer ancillary services for grid integration [24].

55 In WFCRL, interfaces with two state-of-the-art farm simulators are implemented : a static simulator
 56 FLORIS [12] and a dynamic simulator FAST.Farm [20]. Indeed, the choice of a static or dynamic
 57 model is particularly important: the overwhelming majority of proposed approaches are evaluated on
 58 static models, but it was shown in [34] that successful learning approaches under static conditions

59 generally do not adapt to dynamic ones. However, online learning from scratch with dynamic
60 simulators is often too slow, making transfer learning from static to dynamic simulators of great
61 interest. From the broader literature on transfer learning and learning from simulators we know that
62 it is challenging to train policies that can improve on previously learned behavior when deployed
63 on new environments with unseen dynamics [40, 14]. In spite of this problem, to the best of our
64 knowledge, most approaches so far have been trained and evaluated on the same environment, and
65 it is therefore not clear whether the policies learned with simulators are robust enough to be useful,
66 or even safe, when deployed on real wind farms. With two simulators of different model-fidelity
67 (referring to how closely the model represents the real system), WFCRL offers the possibility of
68 designing transfer learning strategies between these simulators.

69 Contributions of the paper

- 70 • We introduce WFCRL, the **first reinforcement learning** suite of environments for **wind**
71 **farm control**. WFCRL is highly customizable, allowing researchers to design and run their
72 own environments for both centralized and multi-agent RL. It includes a default suite of
73 wind farm layouts to be used in benchmark cases.
- 74 • We interface all our wind farm layouts with **two different wind farm simulators**: a static
75 simulator FLORIS [12] and a dynamic simulator FAST.Farm [20]. They can be used to
76 **design transfer learning strategies**, with the goal to **learn robust policies that can adapt**
77 **to unseen dynamics**.
- 78 • We include two implementations of PPO-based state-of-the-art MARL algorithms, IPPO
79 and MAPPO [39], adapted to our environments.
- 80 • We propose a benchmark example for **wind power maximization** with two wind condition
81 scenarios. It takes into account the costs induced by wind turbine fatigue.

82 The paper is organized as follows. In Section 2, we introduce the WFCRL environment suite. First in
83 Section 2.1 we introduce the simulators, the specifications of the simulated wind farms and turbines
84 and the wind conditions scenarios we consider. We then lay out in Section 2.2 the cooperative MARL
85 framework for the wind farm control problem, and finally detail the learning tasks and algorithms
86 available with the suite in Section 2.3. In the second part Section 3, we illustrate the possibilities
87 of the WFCRL environment suite by introducing a benchmark example: the maximization of total
88 power production with fatigue-induced costs. In Section 3.1, we explicit the actions, observations
89 and rewards used in this problem, then in Section 3.2, we present and discuss the results of the IPPO
90 and MAPPO on our benchmark tasks. In Section 4, we discuss perspectives and limitations, and we
91 conclude in Section 5

92 2 WFCRL environments suite

93 In this section, we present our WFCRL environments suite. We first present the simulators interfaced
94 in WFCRL (FLORIS and FAST.Farm), several pre-defined layouts and wind condition scenarios.
95 Note again that having two simulation environments with different model-fidelity offers the possibility
96 of designing transfer learning strategies between simulation environments. Then, we describe briefly
97 the MARL framework for the wind farm control problem. More precisely, we consider a wind farm
98 with M turbines, which operate in the same wind field and create turbulence that propagates across
99 the farm. In our multi-agent environment, each turbine is considered as an agent receiving local
100 observations, and all cooperate to maximize a common objective.

101 2.1 The simulation environments

102 In WFCRL, users can choose one of the two state-of-the-art wind farm simulators (FLORIS or
103 FAST.Farm), select a pre-defined wind farm layout or define a custom one, and choose one of the
104 implemented wind conditions. Though designed for the MARL framework, we note that it is also
105 possible to apply single-agent RL algorithms by considering global observations and actions.

| WFCRL environment | Real wind farm |
|-------------------|--|
| Ablaincourt | Ablaincourt Energies onshore wind farm, Somme, France |
| Ormonde | Ormonde Offshore Wind Farm, Irish Sea, UK |
| WMR | Westermost Rough Wind Farm is an offshore wind farm, North Sea, UK |
| HornsRev1 | Horns Rev 1 Offshore Wind Farm, North Sea, Denmark |
| HornsRev2 | Horns Rev 2 Offshore Wind Farm, North Sea, Denmark |

Table 1: Correspondences between WFCRL environments and real wind farms.

106 **FLORIS environments** The wind farm simulator FLORIS implements static wind farm models,
107 which predict the locations of wake centers and velocities at each turbine in the steady state: the
108 dynamic propagation of wakes are neglected. The yaws of all wind turbines can be controlled, and the
109 power production of the wind farm is then a function of all yaw angles and the so-called free-stream
110 wind conditions: wind measurements - e.g. velocity and direction - taken at the entrance of the farm.
111 FLORIS has been released as an open-source Python software tool¹. In WFCRL environments built
112 on FLORIS, global and local states contain time-averaged, steady-state wind and production statistics
113 for both global and local observations.

114 The models used by FLORIS do not compute any estimate of fatigues on wind turbines, and we
115 propose to use local wind statistics to compute proxy for load estimates indeed. We detail this when
116 introducing our benchmark example in Section 3.1.

117 **FAST.Farm environments** Unlike FLORIS, FAST.Farm is a dynamic simulator that produces
118 time-dependent wind fields that take into account the dynamics of wake propagation [20]: wakes in
119 wind farms tend to meander, and the wakes of different turbines interact and eventually merge as
120 they propagate in the farms. One consequence is that under dynamic conditions there is a significant
121 delay between the time agents take an action and the time this action finally impacts the turbines
122 downstream.

123 FAST.Farm is built on wind turbine simulation tool OpenFAST [26] which computes an estimate of
124 the strength of the bending moment on each turbine blades. This reflects the structural loads induced
125 on turbine blades, and thus can be used to design rewards in RL problems to reduce or avoid physical
126 damages to turbines.

127 FAST.Farm is coded in Fortran. To allow for integration with the large ecosystem libraries and RL
128 research practices developed in Python, we implement an interface between the simulator and the
129 Python wind farm environment via MPI communication channels. The details of the interfacing
130 infrastructure are reported in Appendix B.

131 **Wind farm layouts** Any custom layout - the arrangement of the wind turbines in the farm - can be
132 used in WFCRL. We also propose several pre-defined wind farm layouts for use in benchmark cases.
133 The coordinates of the wind turbines of 5 real wind farms with 7 to 92 wind turbines are obtained
134 from [2]. A complete list of all correspondences between wind farms inspired by real environments
135 and their locations is in Table 1, and a list of all available environments can be found in Appendix C.
136 We also include in WFCRL several toy layouts, including a simple row of 3 turbines (the *Turb3_Row1*
137 environment) for validation purpose and the 32 turbines layout of the *FarmConnors* benchmark [13].
138 A visual representation of the layouts can be found in Appendix G.

139 For all cases, we simulate instances of the NREL Reference 5MW wind turbines, whose specifications
140 have been made public by the National Renewable Energy Laboratory (NREL) [19]. It has become
141 standard reference for wind energy research and is used by the majority of proposed evaluations of
142 RL methods [1].

143 **Wind condition scenarios** For all environments, we distinguish two scenarios.

¹<https://nrel.github.io/floris/>

144 *Wind scenario I:* In this scenario, all trajectories in a given environment are run under the prevailing
 145 wind velocity and direction at the location.

146 *Wind scenario II:* In this scenario, we let the wind farm be subject to variations in wind change, and
 147 sample new free-stream wind conditions u_∞, ψ_∞ at the beginning of each episode:

$$u_\infty \sim \mathcal{W}(\bar{u}, \lambda) \quad \psi_\infty \sim \mathcal{N}(\bar{\psi}, \sigma_\psi) \quad (1)$$

148 where \mathcal{W} is a Weibull distribution modeling wind speed with shape λ and scale \bar{u} , and \mathcal{N} is a Normal
 149 distribution with $\bar{\psi}$ being the dominant wind direction for a given farm.

150 By default, at the beginning of each simulation, all wind turbines have the yaw angle zero. This
 151 corresponds to the so-called *greedy* case, the strategy that would allow each of them to maximize its
 152 production in un-waked conditions.

153 2.2 The MARL framework for the wind farm control problem

154 A Decentralized Partially Observable Markov Decision Process (Dec-POMDP) with M interacting
 155 agents is a tuple $\{M, S, O_1, \dots, O_M, A_1, \dots, A_m, P, r\}$. S is the full state space of the system,
 156 while for any $i \in \{1, \dots, M\}$, O_i is the observation space of the i th agent, with the mapping from
 157 the full state to the local observation defined by a function $o^i : S \rightarrow O_i$. A_i is the local action
 158 space of the agent, and the global action space is the product of all local action spaces $A = \times_i^M A_i$.
 159 At each iteration, all agents observe their local information, chose an action and receive a reward
 160 $r : S \times A \times S \rightarrow \mathbb{R}$. The system then moves to a new state, the transition kernel $P : S \times A \rightarrow S$ gives
 161 the probability of transition from a state $s \in S$ to $s' \in S$ when agents have taken global action $a \in A$.
 162 We call π_1, \dots, π_M the policies followed by each agent, where $\pi_i(a_i|o_i)$, defined the probability for
 163 agent i to chose action a_i when observing o_i . The corresponding global policy $\pi = (\pi_1, \dots, \pi_M)$
 164 simply concatenates the outputs of all local policies.

165 **Objective** The MARL problem is to find a policy π^* that maximizes the expectation of the
 166 discounted sum of rewards collected over a finite or infinite sequence of time-steps

$$\max_{\pi} \mathbb{E}_{s_0, a_0, s_1, \dots} [J], \quad J := \sum_{k=0}^T \beta^k r_k \quad (2)$$

167 with $0 < \beta < 1$ the discount factor and T the number of steps in the environment, or the length of an
 168 episode. For the wind farm control problem, possible rewards include the total production of the farm
 169 or a distance to a target production. As the fatigue load measurements are also available, rewards can
 170 also be designed to encourage actions that preserve the turbine structure.

171 **State and Observation** As the production of each turbine is a function of the local wind conditions
 172 at its rotor, a Markovian description of the full state of the system should contain the whole wind
 173 velocity field of the entire farm. This is so far impossible to know in practice. We rather assume
 174 that local measurements of wind speed and direction are available at each wind turbine, and that an
 175 estimate of the free-stream wind speed and direction can be accessed, but might not necessarily be sent
 176 to the turbines in real time. Our environments therefore distinguish between the local observations o_i
 177 for each $i \in \{1, \dots, M\}$ and a global observation o_g . Each $o_i = (u_i, \phi_i, \theta_i)$ contains a local measure
 178 of the wind velocity u_i and direction ψ_i , as well as the last target value sent to each actuator θ_i . The
 179 global observation $o_g = (o_1, \dots, o_M, u_\infty, \psi_\infty)$ contains the concatenation of all local states, as well
 180 as the free-stream measure of the wind u_∞, ψ_∞ . Table 2 summarizes all observations and actions
 181 available with the two simulators.

182 **Actions** WFCRL offers several ways to control wind turbines: the *yaw*, the *pitch* or *torque*. Yaw
 183 control is available on the FLORIS environments, and all three can be controlled on FAST.Farm
 184 environments. The yaw is the angle between a wind turbine’s rotor and the wind direction: turbines
 185 facing the wind have a yaw of 0° which maximizes their individual power output. Increasing the yaw
 186 can deflect the wake away from downstream turbines, which may increase the total production of the

| | FLORIS | FAST.Farm |
|---------------------------|--|--|
| Local Observations o_i | u_i, ϕ_i (steady-state), y_i | u_i, ϕ_i (time-dependent), y_i, p_i, τ_i |
| Global Observations o_g | $o_1, \dots, o_M, u_\infty, \phi_\infty$ | |
| Actions | Δy_i | $\Delta y_i, \Delta p_i, \Delta \tau_i$ |

Table 2: Observations (global and local) and actions available for an agent i in FLORIS and FAST.Farm environments. y_i, p_i, τ_i refer respectively to the yaw, pitch and torque of the turbine.

187 wind farm. The pitch is the angle of the attack of the rotor blades with respect to the incoming wind,
 188 while the torque of the turbine’s rotor directly controls the rotation speed. Increasing the blade pitch
 189 or decreasing the torque target both decrease the fraction of the power in the wind extracted by the
 190 turbine, and therefore decrease the turbulence in its wake. To reflect the fact that the actuation rate of
 191 the wind turbines is limited by physical constraints, we conceive actions as increases or decreases in
 192 the actuator target value rather than absolute values, with the limits being implemented by the upper
 193 and lower bounds of a continuous action space.

194 2.3 Learning in WFCRL

195 All environments are implemented with known RL and MARL Python interfaces Gymnasium [7]
 196 and PettingZoo [36]. The code is open-sourced under the Apache-2.0 license.

197 2.3.1 Online Learning

198 Environments implemented on both FLORIS and FAST.Farm can be used in an episodic learning
 199 approach. This is the traditional setting of the RL problem, and we will refer to it as the *Online*
 200 *Learning Task*. In Wind scenario I, where agents learn to cooperate against a single set of wind
 201 conditions, we look at the evolution of the sum of rewards collected over an episode. In Wind scenario
 202 II, where a different set of wind conditions is sampled at each episode, we evaluate the policies on a
 203 predefined set of wind conditions and use a weighted average as the final score. This gives us our
 204 evaluation score:

$$\text{score}(\pi_1, \dots, \pi_M) = \sum_{j=1}^{n_w} \rho_j \sum_{k=0}^T r_k \quad (3)$$

205 where T is the length of the episode, n_w is the number of wind conditions considered and the
 206 $\rho_1, \dots, \rho_{n_w}$ are the weights on each conditions, with for all j , $0 < \rho_j < 1$ and $\sum_j^{n_w} \rho_j = 1$. The
 207 wind conditions distributions on which policies are evaluated need not be identical to the one from
 208 which conditions were sampled during training.

209 2.3.2 Transfer

210 Exploration on real wind farms is costly: as prototype models are typically not available for large
 211 wind farms, adjusting to the real dynamics of the system will require exploring in real time on an
 212 operating wind farm. Every move of exploring in a suboptimal direction is a cost for the farm
 213 operator. Learning efficient policies offline that can quickly adapt to the real system is therefore
 214 critical. Since the dynamic FAST.Farm simulator is considered a higher fidelity version of the static
 215 simulator FLORIS, we propose to use the former as a proxy of a real wind farm to evaluate the
 216 robustness of policies learned on the latter, and their ability to adjust to the real dynamics of a farm.
 217 We will refer to this as the *Transfer Task*.

218 2.3.3 Algorithms

219 We consider two state-of-the-art algorithms IPPO (Independent PPO) and MAPPO (Multi-Agent
 220 PPO) introduced in [8, 39]. Both are based on the on-policy PPO algorithm [31]. In [39], it was found
 221 that PPO-based methods can perform very well when extended to cooperative multi-agent tasks,
 222 outperforming algorithms specifically designed for cooperative problems like QMIX [30]. Following
 223 an approach called independent learning, IPPO builds on PPO by allowing every agent to run a PPO
 224 algorithm in parallel. On the other hand MAPPO maintains both M agent policies taking actions
 225 based on local information and a shared critic, which estimates the value of a global observation. The

226 choice of the global observation fed to the critic is an important factor influencing the performance of
 227 the algorithm [39]. We follow the recommendations of [39] to adapt PPO to the multi-agent case.
 228 They suggest to include both local and global observation features to the value function input. We
 229 therefore feed to our shared critic network the full global observation introduced in Section 2.2, that
 230 is both the of concatenation of all local observations and the free-stream wind velocity

231 For implementation, we adapt the CleanRL² [17] baseline implementations of PPO to our multi-agent
 232 Petting Zoo environments. Since given a local observation the optimal policies are not identical
 233 for all turbines, we do not implement weight sharing between different agents.

234 3 Benchmark example: the maximization of the total power production

235 We consider the problem of finding the optimal yaws to maximize the total power production under
 236 a set of wind conditions, and taking into account the costs induced by turbine fatigue load. This
 237 problem is known as the *wake steering* problem, and is an active area of research in the wind energy
 238 literature [15, 16].

239 3.1 Problem formulation

240 **Actions and observations** Local observations include the local yaw and local wind statistics. The
 241 concatenation of all local observations along with free-stream wind statistics in the global observation
 242 is as described in Section 2.2. Recall that actions are defined as increase or decrease in the actuator
 243 target value. In this problem, all agents control their yaws, and we define the continuous action space
 244 $[-5, 5]$, defining changes in yaw angle expressed in degrees. To constraint the load on the turbines
 245 caused by the control strategies and reduce its impact on the lifetime of the turbines, the time each
 246 turbine spends actuating is limited. We choose the upper bound of 10% of the time, which is the
 247 same upper bound value discussed in [29]. At every iteration, the time needed to change the state of
 248 the actuator is computed, and any action violating this condition is not allowed.

249 **Rewards** At each iteration k , all agents receive a reward r_k^P which is the currently measured
 250 production of the wind farm in kW divided by the number of agents and normalized by the free-
 251 stream wind velocity:

$$r_k^P = \frac{1}{M} \sum_i^M \frac{\hat{P}_k^i}{(u_{\infty,k})^3} \quad (4)$$

252 where \hat{P}_k^i is the measured power production and $u_{\infty,k}$ the free-stream wind velocity at time-step k .
 253 To discourage agents from taking risky policies damaging the turbines, we also return a load penalty
 254 r_k^L which increases with the sum of loads on all the turbine blades.

255 FLORIS does not provide estimates of the loads on structures. Instead, we evaluate the impact of
 256 actuations on loads with a proxy based on local estimates of turbulence and velocities on the surface
 257 of the rotor planes. Our proxy takes into account 2 factors increasing stress on wind turbine structures
 258 as noted in [35]: first, the turbulence of the wind and second, the variation of velocities on the turbine
 259 rotor. We therefore define the load penalty in FLORIS environments as

$$r_{k,S}^L = \frac{1}{M} \sum_i^M \left(\sum_j^9 TI_k[x_{i,j}, y_{i,j}] + \sigma(u_k) + \sigma(v_k) + \sigma(w_k) \right) \quad (5)$$

260 where TI_k is the turbulence field at time-step k , u_k , v_k and w_k are respectively the x , y and z
 261 components of the velocity field at time-step k , and the $x_{i,j}$ define the coordinates of the $9 \times M$ grid
 262 points at which these values are computed for the M rotor planes. σ denotes the standard-deviation.
 263 For FAST.Farm, we use the estimates of the the blades' bending moment strength as a proxy for the

²<https://github.com/vwxyzjn/cleanrl>

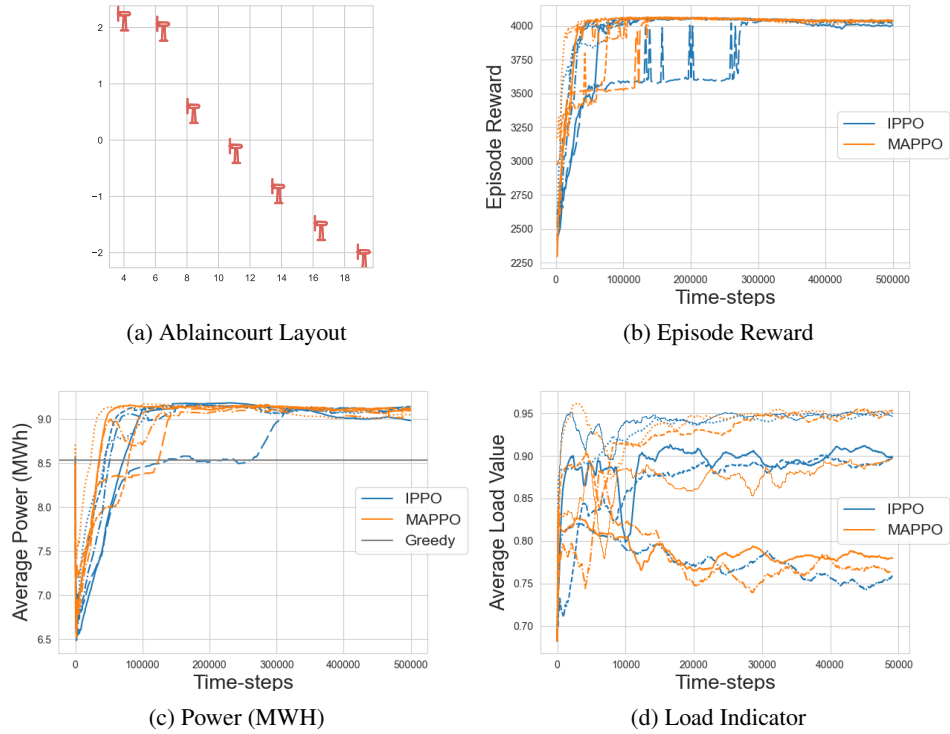


Figure 2: The evolution of episode reward, average power output and average load on the *Ablaincourt* environment, simulated with FLORIS. A visual representation of the layout is in (a), the evolution of the episode reward is reported (b), the power output averaged on an episode length (here $T=2048$) is reported on the (c) and the loading indicator is on (d). The curves are plotted for all 5 seeds.

264 structural loads induced on the turbines, and define the load penalty as

$$r_{k,D}^L = \frac{1}{M} \sum_i^M \left(\sum_j^3 |Mop_k[i,j]| + \sum_j^3 |Mip_k[i,j]| \right) \quad (6)$$

265 where Mop_k is the $M \times 3$ matrix of out-of-plane bending moments for the 3 blades of every turbine
 266 at time-step k , and Mip_k is the corresponding matrix of in-plane bending moments.

267 Both rewards are common to all turbines, and all must therefore maximize (2) with $r_k = (r_k^P - 0.1r_k^L)$.
 268 We downscale the load penalty to account for the difference in magnitude between current production
 269 energy and load-induced maintenance cost for wind energy projects.

270 **Wind conditions** To evaluate the algorithms with score (3), we need to choose weights ρ_j . We
 271 use data acquired during the SmartEole project at the location of the Ablaincourt wind farm [11].
 272 It consists of estimates of free-stream wind direction and velocity computed from measures taken
 273 during a 3 months field campaign every 10 min. Since in real conditions wind velocity and wind
 274 direction are correlated, we compute the bi-dimensional histogram for the two variables, taking 5 bins
 275 for each dimension. We obtain a set of 25 wind condition rectangle. The wind condition w_1, \dots, w_j
 276 are the center of each rectangle, and the corresponding weights ρ_j are defined as the frequencies at
 277 which wind conditions in the time series appeared in the rectangle.

278 3.2 Results

279 We apply algorithms IPPO and MAPPO available in WFCRL to our benchmark example. We
 280 distinguish two scenarios, the first (resp. second) is trained with the wind scenario I (resp. II)

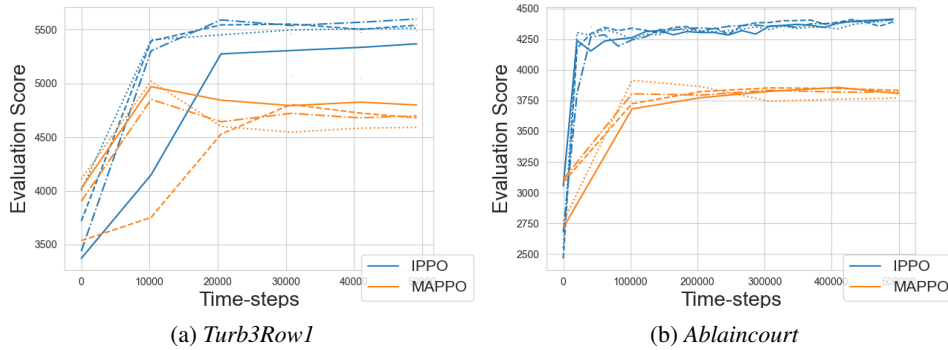


Figure 3: Evolution of the evaluation score, defined in (3), during the training of IPPO and MAPPO on the two environments *Turb3Row1* (left) and *Ablaincourt* (right).

281 described in section 2.1. Both are learned with the static simulator FLORIS. For the first scenario,
 282 the score is reward obtained on a single policy rollout of 2048 steps in the environment. Results with
 283 on the *Ablaincourt* layout are given in Fig. 2. For the second scenario, the score is the one defined in
 284 (3). The training curves for the the *Ablaincourt* and the *Turb3Row1* layouts are illustrated in Fig. 3.
 285 A table detailing training scores at convergence is available in Appendix F and hyper-parameters are
 286 given in Appendix D

287 To illustrate the *Transfer* case, we then deploy the learned IPPO policies on a *Turb3Row1* on 900
 288 steps (45 minutes in simulated time) in the corresponding FAST.Farm environment. We report in
 289 Appendix F.1 the average percentage increase of power output and load compared to the greedy case,
 290 and compare it to a naive deployment of strategies learned online. Our results illustrate the difficulty
 291 of adapting learned policies to unseen dynamics.

292 4 Limitations

293 The choice of the wind farm simulators included in WFCRL relies on three criteria: the trade-off
 294 between fidelity and computation time, the popularity of the simulators in the wind farm energy
 295 community, and its open-source availability. Both FLORIS and FAST.Farm are developed and
 296 actively maintained by the US-based National Renewable Energy Laboratory³, and have a large user
 297 base among wind energy researchers. FAST.Farm was explicitly designed to provide good fidelity
 298 at a limited computation cost [20]. Despite this, dynamic wind farm simulators remain slow. The
 299 development and open-sourcing of faster dynamic simulators will be critical. Machine-learning
 300 accelerated simulators could be an important step in that direction.

301 5 Conclusion

302 We have introduced WFCRL, the first reinforcement learning suite of environments for wind farm
 303 control. WFCRL is highly customizable, allowing researchers to design and run their own envi-
 304 ronments for both centralized and multi-agent RL. It is interfaced with two different wind farm
 305 simulators: a static simulator FLORIS and a dynamic simulator FAST.Farm. They can be used to
 306 design transfer learning strategies with the goal to learn robust policies that can adapt to unseen
 307 dynamics. We have proposed a benchmark example for wind power maximization with two wind
 308 condition scenarios that take into account the costs induced by wind turbine fatigue. We hope that
 309 WFCRL will help building a bridge between the RL and wind energy research communities.

³<https://www.nrel.gov/>

References

- 310
- 311 [1] Mahdi Abkar, Navid Zehtabiyani-Rezaie, and Alexandros Iosifidis. Reinforcement learning for
312 wind-farm flow control: Current state and future actions. *Theoretical and Applied Mechanics*
313 *Letters*, 13(6):100475, 2023.
- 314 [2] Ramon Abritta. Wind power plants layouts according to arbitrary reference points, thanet, west
315 of duddon sands, ormonde, westernmost rough, horns rev 1 & 2, anholt, and london array [data
316 set]. zenodo. <https://zenodo.org/records/10927983>, 2023.
- 317 [3] Cristina L. Archer, Ahmadreza Vassel-Be-Hagh, Chi Yan, Sicheng Wu, Yang Pan, Joseph F.
318 Brodie, and A. Eoghan Maguire. Review and evaluation of wake loss models for wind energy
319 applications. *Applied Energy*, 226:1187–1207, 9 2018.
- 320 [4] R. J. Barthelmie, S. C. Pryor, S. T. Frandsen, K. S. Hansen, J. G. Schepers, K. Rados, W. Schlez,
321 A. Neubert, L. E. Jensen, and S. Neckelmann. Quantifying the impact of wind turbine wakes
322 on power output at offshore wind farms. *Journal of Atmospheric and Oceanic Technology*,
323 27(8):1302 – 1317, 2010.
- 324 [5] Claire Bizon Monroc, Ana Bušić, Donatien Dubuc, and Jiamin Zhu. Actor critic agents for
325 wind farm control. In *2023 American Control Conference (ACC)*, pages 177–183, 2023.
- 326 [6] Sjoerd Boersma, Bart M Doekemeijer, Pieter MO Gebraad, Paul A Fleming, Jennifer Annoni,
327 Andrew K Scholbrock, Joeri Alexis Frederik, and Jan-Willem van Wingerden. A tutorial on
328 control-oriented modeling and control of wind farms. In *2017 American control conference*
329 *(ACC)*, pages 1–18. IEEE, 2017.
- 330 [7] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang,
331 and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- 332 [8] Christian Schroeder de Witt, Tarun Gupta, Denys Makoviichuk, Viktor Makoviychuk, Philip HS
333 Torr, Mingfei Sun, and Shimon Whiteson. Is independent learning all you need in the starcraft
334 multi-agent challenge? 2020.
- 335 [9] Zhiwen Deng, Chang Xu, Xingxing Han, Zhe Cheng, and Feifei Xue. Decentralized yaw
336 optimization for maximizing wind farm production based on deep reinforcement learning.
337 *Energy Conversion and Management*, 286:117031, 2023.
- 338 [10] Hongyang Dong and Xiaowei Zhao. Reinforcement learning-based wind farm control: Toward
339 large farm applications via automatic grouping and transfer learning. *IEEE Transactions on*
340 *Industrial Informatics*, 19(12):11833–11845, 2023.
- 341 [11] Thomas Duc, Olivier Coupiac, Nicolas Girard, Gregor Giebel, and Tuhfe Göçmen. Local
342 turbulence parameterization improves the jensen wake model and its implementation for power
343 optimization of an operating wind farm. *Wind Energy Science*, 4(2):287–302, 5 2019.
- 344 [12] PMO Gebraad, FW Teeuwisse, JW van Wingerden, PA Fleming, SD Ruben, JR Marden, and
345 L.Y Pao. Wind plant power optimization through yaw control using a parametric model for
346 wake effects - a cfd simulation study. *Wind Energy*, 19(1):95 – 114, 2016.
- 347 [13] T. Göçmen, F. Campagnolo, T. Duc, I. Eguinoa, S. J. Andersen, V. Petrović, L. Imširović,
348 R. Braunbehrens, J. Liew, M. Baungaard, M. P. van der Laan, G. Qian, M. Aparicio-Sanchez,
349 R. González-Lope, V. V. Dighe, M. Becker, M. J. van den Broek, J.-W. van Wingerden, A. Stock,
350 M. Cole, R. Ruisi, E. Bossanyi, N. Requate, S. Strnad, J. Schmidt, L. Vollmer, I. Sood, and
351 J. Meyers. Farmconners wind farm flow control benchmark – part 1: Blind test results. *Wind*
352 *Energy Science*, 7(5):1791–1825, 2022.
- 353 [14] Sebastian Höfer, Kostas Bekris, Ankur Handa, Juan Camilo Gamboa, Melissa Mozifian, Florian
354 Golemo, Chris Atkeson, Dieter Fox, Ken Goldberg, John Leonard, et al. Sim2real in robotics
355 and automation: Applications and challenges. *IEEE transactions on automation science and*
356 *engineering*, 18(2):398–400, 2021.
- 357 [15] Daniel R Houck. Review of wake management techniques for wind turbines. *Wind Energy*,
358 25(2):195–220, 2022.


- 359 [16] Michael F. Howland, Sanjiva K. Lele, and John O. Dabiri. Wind farm power optimization
360 through wake steering. *Proceedings of the National Academy of Sciences*, 116(29):14495–14500,
361 2019.
- 362 [17] Shengyi Huang, Rousslan Fernand Julien Dossa, Chang Ye, Jeff Braga, Dipam Chakraborty,
363 Kinal Mehta, and João G.M. Araújo. Cleanrl: High-quality single-file implementations of
364 deep reinforcement learning algorithms. *Journal of Machine Learning Research*, 23(274):1–18,
365 2022.
- 366 [18] J Jonkman, P Doubrawa, N Hamilton, J Annoni, and P Fleming. Validation of FAST.farm
367 against large-eddy simulations. *Journal of Physics: Conference Series*, 1037:062005, 6 2018.
- 368 [19] Jason Jonkman, Sandy Butterfield, Walter Musial, and George Scott. Definition of a 5-mw
369 reference wind turbine for offshore system development. Technical report, National Renewable
370 Energy Lab.(NREL), Golden, CO (United States), 2009.
- 371 [20] Jason M Jonkman, Jennifer Annoni, Greg Hayman, Bonnie Jonkman, and Avi Purkayastha.
372 Development of fast. farm: A new multi-physics engineering tool for wind-farm design and
373 analysis. In *35th wind energy symposium*, page 0454, 2017.
- 374 [21] Elie Kadoche, Sébastien Gourvéne, Maxime Pallud, and Tanguy Levent. Marlyc: Multi-agent
375 reinforcement learning yaw control. *Renewable Energy*, 217:119129, 2023.
- 376 [22] Ali C. Kheirabadi and Ryoza Nagamune. A quantitative review of wind farm control with
377 the objective of wind farm power maximization. *Journal of Wind Engineering and Industrial
378 Aerodynamics*, 192:45–73, 2019.
- 379 [23] Jaime Liew, Tuhfe Göçmen, Wai Hou Lio, and Gunner Chr. Larsen. Model-free closed-loop
380 wind farm control using reinforcement learning with recursive least squares. *Wind Energy*,
381 2023.
- 382 [24] Nicholas W. Miller and Kara Clark. Advanced controls enable wind plants to provide ancillary
383 services. In *IEEE PES General Meeting*, pages 1–6, 2010.
- 384 [25] Grigory Neustroev, Sytze PE Andringa, Remco A Verzijlbergh, and Mathijs M De Weerd.
385 Deep reinforcement learning for active wake control. In *Proceedings of the 21st International
386 Conference on Autonomous Agents and Multiagent Systems*, pages 944–953, 2022.
- 387 [26] NREL. Openfast documentation, 2022.
- 388 [27] Venkata Ramakrishna Padullaparathi, Srinarayana Nagarathinam, Arunchandar Vasan, Vishnu
389 Menon, and Depak Sudarsanam. Falcon-farm level control for wind turbines using multi-agent
390 deep reinforcement learning. *Renewable Energy*, 181:445–456, 2022.
- 391 [28] Sara C Pryor, Rebecca J Barthelmie, and Tristan J Shepherd. Wind power production from very
392 large offshore wind farms. *Joule*, 5(10):2663–2686, 2021.
- 393 [29] Alban Puech and Jesse Read. An improved yaw control algorithm for wind turbines via
394 reinforcement learning. In *Joint European Conference on Machine Learning and Knowledge
395 Discovery in Databases*, pages 614–630. Springer, 2022.
- 396 [30] Tabish Rashid, Mikayel Samvelyan, Christian Schroeder De Witt, Gregory Farquhar, Jakob
397 Foerster, and Shimon Whiteson. Monotonic value function factorisation for deep multi-agent
398 reinforcement learning. *Journal of Machine Learning Research*, 21(178):1–51, 2020.
- 399 [31] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal
400 policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- 401 [32] Coen-Jan Smits, Jean Gonzalez Silva, Valentin Chabaud, and Riccardo Ferrari. A fast.farm and
402 matlab/simulink interface for wind farm control design. *Journal of Physics: Conference Series*,
403 2626(1):012069, oct 2023.
- 404 [33] Paul Stanfel, Kathryn Johnson, Christopher J. Bay, and Jennifer King. A distributed reinforce-
405 ment learning yaw control approach for wind farm energy capture maximization. In *2020
406 American Control Conference (ACC)*, pages 4065–4070, 2020.

- 407 [34] Paul Stanfel, Kathryn Johnson, Christopher J. Bay, and Jennifer King. Proof-of-concept of a
408 reinforcement learning framework for wind farm energy capture maximization in time-varying
409 wind. *Journal of Renewable and Sustainable Energy*, 13(4), 8 2021.
- 410 [35] A. P. J. Stanley, J. King, C. Bay, and A. Ning. A model to calculate fatigue damage caused by
411 partial waking during wind farm optimization. *Wind Energy Science*, 7(1):433–454, 2022.
- 412 [36] J Terry, Benjamin Black, Nathaniel Grammel, Mario Jayakumar, Ananth Hari, Ryan Sullivan,
413 Luis S Santos, Clemens Dieffendahl, Caroline Horsch, Rodrigo Perez-Vicente, et al. Pettingzoo:
414 Gym for multi-agent reinforcement learning. *Advances in Neural Information Processing*
415 *Systems*, 34:15032–15043, 2021.
- 416 [37] Kenneth Thomsen and Poul Sørensen. Fatigue loads for wind turbines operating in wakes.
417 *Journal of Wind Engineering and Industrial Aerodynamics*, 80(1):121–136, 1999.
- 418 [38] P. Veers, K. Dykes, S. Basu, A. Bianchini, A. Clifton, P. Green, H. Holttinen, L. Kitzing,
419 B. Kosovic, J. K. Lundquist, J. Meyers, M. O’Malley, W. J. Shaw, and B. Straw. Grand
420 challenges: wind energy research needs for a global energy transition. *Wind Energy Science*,
421 7(6):2491–2496, 2022.
- 422 [39] Chao Yu, Akash Velu, Eugene Vinitsky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and Yi Wu.
423 The surprising effectiveness of ppo in cooperative multi-agent games. *Advances in Neural*
424 *Information Processing Systems*, 35:24611–24624, 2022.
- 425 [40] Zhuangdi Zhu, Kaixiang Lin, Anil K Jain, and Jiayu Zhou. Transfer learning in deep reinforce-
426 ment learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*,
427 2023.


428 **Checklist**

- 429 1. For all authors...
- 430 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s
431 contributions and scope? [Yes]
- 432 (b) Did you describe the limitations of your work? [Yes] See Section 3.2 for a Discussion
433 of the limitations of our work
- 434 (c) Did you discuss any potential negative societal impacts of your work? [N/A] To the
435 best of our knowledge our work does not have any potential negative societal impacts.
- 436 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
437 them? [Yes] Yes, we believe our paper conform to the ethics review guidelines.
- 438 2. If you are including theoretical results...
- 439 (a) Did you state the full set of assumptions of all theoretical results? [N/A] We have not
440 included theoretical results
- 441 (b) Did you include complete proofs of all theoretical results? [N/A] We have not included
442 theoretical result
- 443 3. If you ran experiments (e.g. for benchmarks)...
- 444 (a) Did you include the code, data, and instructions needed to reproduce the main ex-
445 perimental results (either in the supplemental material or as a URL)? [Yes] Yes, all
446 code and instructions needed to reproduce the experimental results are included in the
447 supplementary material in Appendix D, along an URL to both the environment suite
448 and the training and evaluation scripts
- 449 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
450 were chosen)? [Yes] Yes, see Appendix D as well as Section 3.2
- 451 (c) Did you report error bars (e.g., with respect to the random seed after running exper-
452 iments multiple times)? [Yes] Yes, all our results figures and tables either plot the
453 curves for all seeds like in Appendix F or report error bars like in Section 3.2
- 454 (d) Did you include the total amount of compute and the type of resources used (e.g., type
455 of GPUs, internal cluster, or cloud provider)? [Yes] See Appendix D
- 456 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 457 (a) If your work uses existing assets, did you cite the creators? [Yes] Of course, for both
458 simulators in Section 2.1 and base RL algorithms implementations in Section 2.3.3
- 459 (b) Did you mention the license of the assets? [Yes] See Appendix H.1
- 460 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
461 See Appendix C
- 462 (d) Did you discuss whether and how consent was obtained from people whose data you’re
463 using/curating? [Yes] See Appendix H.1, we only use open-source wind data and
464 software dependencies
- 465 (e) Did you discuss whether the data you are using/curating contains personally identifiable
466 information or offensive content? [N/A] This is not relevant to our work.
- 467 5. If you used crowdsourcing or conducted research with human subjects...
- 468 (a) Did you include the full text of instructions given to participants and screenshots, if
469 applicable? [N/A] We have not used crowdsourcing or conducted research with human
470 subjects
- 471 (b) Did you describe any potential participant risks, with links to Institutional Review
472 Board (IRB) approvals, if applicable? [N/A] We have not used crowdsourcing or
473 conducted research with human subjects
- 474 (c) Did you include the estimated hourly wage paid to participants and the total amount
475 spent on participant compensation? [N/A] We have not used crowdsourcing or con-
476 ducted research with human subjects

477 **A Difference between FLORIS and FAST.Farm**



(a) FLORIS



(b) FAST.Farm

Figure 4: Wind velocity field for the simulation of our 3-turbines layout on the 2 simulators: FLORIS and FAST.Farm.

478 Wind farm models serve two main purposes, in the broader literature as in WFCRL. First, when
479 experiments on real wind farms or tunnel experiments on scaled farms are not possible, they are the
480 only way to evaluate and compare control strategies by predicting their impact on the total power
481 output of a farm. For that purpose, the value of models lie in their accuracy, and the best results are
482 achieved by complex dynamic models involving costly computations. Secondly, a model of the farm
483 can be used to estimate an optimal command. Here, accuracy must be balanced by tractability, and a
484 constraint on computation time arises for real-time optimization. Of course, evaluating the command
485 on the model used to derive it will likely overestimate its performance: it should rather be evaluated
486 on an other, higher fidelity model which will serve as a substitute for the real farm.

487 Static models estimate the time-averaged features of the wind flow while ignoring the dynamics
488 of short-term effects, including wake propagation time and wake meandering. They rely on the
489 design of an analytical solution to predict wind speed deficit at a downwind turbine with respect to an
490 upstream turbine [3]. This gives them the advantage of a very low computation time, as they usually
491 return a solution instantaneously. The wind farm simulation software FLORIS (NREL 2021), created
492 and maintained by the National Renewable Energy Laboratory (NREL), proposes a variety of such
493 models in a single Python framework, and has become a reference for wind farm control engineering.
494 These parametric models combine several components to estimate the effects of turbine yaws on both
495 the redirection of the wake behind the turbine and the velocity in the wake. An example of such a
496 simulation can be found on Figure 4a.

497 At higher accuracy and higher computational complexity is FAST.Farm [20]. It relies on OpenFAST
498 (NREL 2022) to model the dynamics of each individual turbine, but considers additional physics to
499 account for farm-wide ambient wind, as well as wake deficits, propagation dynamics and interactions
500 between different wakes. It has been shown to be of similar accuracy with high-fidelity large-eddy
501 simulations with much less computational expense [18], and it supports the implementation of
502 controllers tracking a received yaw reference for each turbine. Figure 4b provides an example of
503 these realistic wind fields.

504 **B Details of the FAST.Farm interface**

505 The Python-FAST.Farm interfacing tool relies on two interfaces with RECEIVE and SEND functions.
 506 Following [32] in which the authors designed an interface between FAST.Farm and Matlab based on
 507 the MPI communication protocol, we rely on an MPI communication channel between the Python
 508 and FAST.Farm processes. We choose to let the Python process spawn a new child process to launch
 509 the FAST.Farm simulation in the background, allowing the user to only interface with Python. The
 510 architecture of the interfacing tool is illustrated in Figure 5.

511 At every iteration, the FAST.Farm interface retrieves 12 measures per turbine:

- 512 • 2 wind measurements: wind velocity and direction at the entrance of the farm. The wind
 513 direction is estimated by subtracting the yaw estimation error from the current yaw measure.
- 514 • The current output power of the turbine
- 515 • The yaw of the turbine
- 516 • The pitch of the turbine
- 517 • The torque of the turbine
- 518 • 6 measures of blade loads: the out-of-plane bending moment estimate for each blade, and
 519 the in-plane bending moment estimate on each blade

520 and sends the 3 control targets - yaw, pitch, torque - to each local turbine controller. Other actuators
 521 that are not controlled by the RL algorithm are controlled by the default naive FAST.Farm controllers.

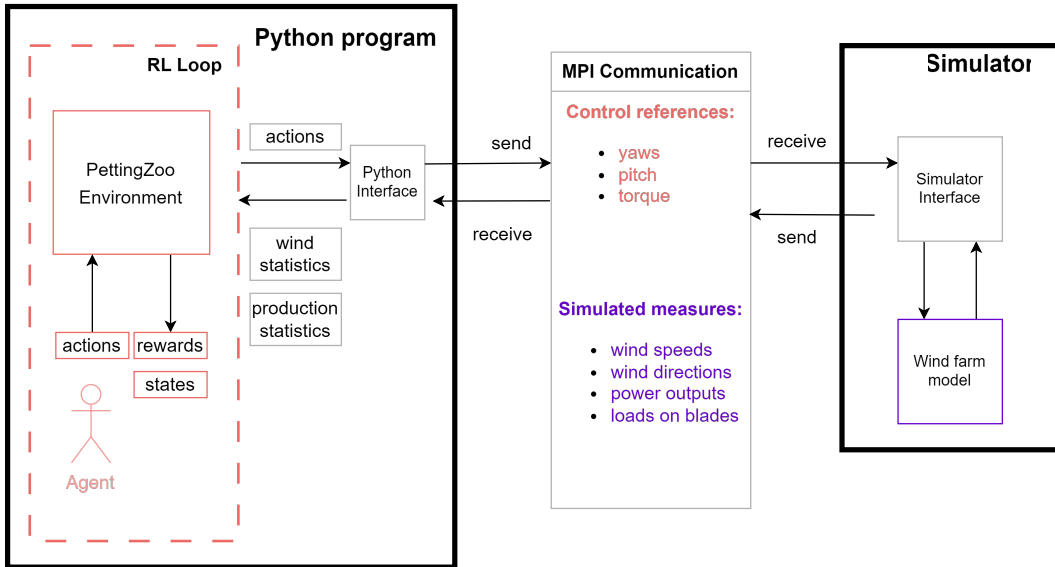


Figure 5: Schema: interfacing infrastructure between FAST.Farm and Python

522 **C Characteristic of all environments**

| | Centralized Control | Decentralized Control |
|------------------|----------------------------|--------------------------------|
| Floris | <i>LayoutName_Floris</i> | <i>Dec_LayoutName_Floris</i> |
| FAST.Farm | <i>LayoutName_Fastfarm</i> | <i>Dec_LayoutName_Fastfarm</i> |

Table 3: Creating environment IDs: Prefix, Root, Suffix

523 For preregistered layouts, every environment is characterized by a tuple of 3 options, and every
 524 environment ID is a combination of the corresponding 3 parts: a prefix, a root, and a suffix.

- 525 • The choice to formalize it as a centralized or decentralized control problem. Environments
 526 with centralized control are *Gymnasium* environments and expect global actions, i.e. vectors
 527 concatenating all actions, and have no prefix. Environments with decentralized control are
 528 *PettingZoo* environments and expect local actions sent by each agent. They have the prefix
 529 *Dec*.
- 530 • The choice of the layout, i.e. the arrangement of wind turbines in the field. A list of all
 531 layouts is given in Table 4, and a visual overview of them in Appendix G. The name of the
 532 layout is the root of the environment ID.
- 533 • The choice of a simulator. Two simulators are for now implemented in WFCRL: the static
 534 FLORIS and the dynamic FAST.Farm. The corresponding suffix *Floris* or *Fastfarm* is
 535 appended to the environment ID.

| Layout Name | # Agents | Description |
|-----------------------------|-----------------|--|
| Ablaincourt | 7 | Inspired by layout of the Ablaincourt farm in France, (Duc et al, 2019) |
| Turb16_TCRWP | 16 | Layout of the Total Control Reference Wind Power Plant (TC RWP) (the first 16 turbines) |
| Turb6_Row2 | 6 | Custom case - 2 rows of 6 turbine |
| Turb16_Row5 | 16 | Layout of the first 16 turbines in the CL-Windcon project as implemented in WFSim |
| Turb32_Row5 | 32 | Layout of the farm used in the CL-Windcon project as implemented in WFSim |
| TurbX_Row1 for X in [1, 12] | X | Procedurally generated single row layout with X turbines, spaced by 4D with the D the diameter of the turbine. |
| Ormonde | 31 | Layout of the Ormonde Offshore Wind Farm |
| WMR | 36 | Layout of the Westernmost Rough Offshore Wind Farm |
| HornsRev1 | 76 | Layout of the Horns Rev 1 Offshore Wind Farm |
| HornsRev2 | 92 | Layout of the Horns Rev 2 Offshore Wind Farm |

Table 4: Preregistered layouts: name, number of agents, and description

536 **D Environment and Training procedure details**

537 The source code for the WFCRL package is open-sourced under the license Apache v2, and publicly
 538 released here www.github.com/ifpen/wfcrl-env, along with notebook tutorials and documenta-
 539 tion. An example of code snippet allowing the creation of the Floris Ablaincourt environment with
 540 decentralized control is given below:

```
541     from wfcrl import environments as envs
542     env = envs.make("Dec_Ablaincourt_Floris")
```

543 The code to reproduce all experiments is available here www.github.com/ifpen/wfcrl-benchmark. We report in Table 5 the hyper-parameters used for both algorithms.


| Parameter | Value |
|-----------------|----------|
| learning_rate | 0.0003 |
| gamma | 0.99 |
| gae_lambda | 0.95 |
| num_minibatches | 32 |
| update_epochs | 10 |
| norm_adv | True |
| clip_coef | 0.2 |
| clip_vloss | True |
| ent_coef | 0.0 |
| vf_coef | 0.5 |
| max_grad_norm | 0.5 |
| target_kl | None |
| kl_coef | 0.0 |
| hidden_layer_nn | (64, 64) |
| num_steps | 2048 |
| anneal_lr | True |
| batch_size | 2048 |
| minibatch_size | 64 |
| num_iterations | 24 |

Table 5: Experiment Hyperparameters


545
 546 The experiments were run on 3 different computers. The first computer, which has no GPU and a
 547 Intel Xeon Gold 6240Y processor, was used to train IPPO and MAPPO on *Wind Scenario I* during
 548 1 week. On the second computer, an internal cluster with a GPU Quadro RTX 6000 24Go, 1 week
 549 of compute was used to to train experiments of *Wind Scenario II*. The last computer which has a
 550 Intel Xeon Gold 6240Y processor and a GPU Quadro RTX 6000 24Go was used for training models
 551 during 3 days of compute on *Wind Scenario I*, and for evaluation purposes.

552 **E Score: wind rose and weights**

553 In this section we illustrate the use of wind statistics from the SMARTEOLE dataset to extract
 554 wind conditions weights ρ of the evaluation score (3). In Figure 6a, we report the distribution of
 555 wind velocity and direction in the SMARTEOLE dataset. In Figure 6b, we show the corresponding
 556 extracted weights ρ for the 25 corresponding wind conditions.



(a) Wind conditions in SMARTEOLE



(b) ρ_i extracted from SMARTEOLE

Figure 6: Extraction of the ρ_i weights from the SMARTEOLE dataset. The empirical distribution of wind speed and direction in the data represented as a windrose is in (a), and the corresponding extracted weights ρ_i given to each of the 25 wind conditions are in (b).

557 **F More benchmark results**

558 **F.1 Evaluation and transfer on FAST.Farm**

559 We evaluate the agents trained on the FLORIS environments by rolling out their determinist policies
 560 in this new environment (they always pick the likeliest action under their policy functions). On
 561 this task, we simulate a 900 steps episode of the *Turb3_Row1* layout on FAST.Farm (environment
 562 *Dec_Turb3_Row1_Fastfarm*). The average rewards collected during the episode are in Table 6.

563 For the *Transfer* task, we pursue the training in the new environments, and report the percentage
 change in power and load compared to the *Greedy* case in Figure 7 for the agents trained under IPPO.




Figure 7: Evaluation and transfer on FAST.Farm: evolution of power increase (left) and load increase (Right) on the *Dec_Turb3_Row1_Fastfarm* environment with respect to the Greedy case. Standard deviations are re

| | IPPO | MAPPO |
|--------------------------|-----------|------------|
| <i>Turb3Row1 (Sc. 1)</i> | 1238 ± 24 | 1369 ± 41 |
| <i>Turb3Row1 (Sc. 2)</i> | 1607 ± 41 | 1369 ± 124 |

Table 6: FAST.Farm evaluation task

564

565 **F.2 Some more training results**

566 In this section we report more benchmark results. The training curves of IPPO and MAPPO under
 567 *Wind Scenario I* on the *Turb3_Row1* layout are in Figure 8. Table 7 summarizes the results at
 568 convergence: both on the total score and the increase or decrease in average power of load compared
 569 to the *Greedy* baseline, for both the *Turb3_Row1* and *Ablaincourt* layouts.

| | IPPO | | | MAPPO | | |
|-----------------------|-------------------|-----------|----------|-----------------|-----------|----------|
| | Score | Power (%) | Load (%) | Score | Power (%) | Load (%) |
| <i>Turb3. (Sc. 1)</i> | 3431 ± 138 | +18 ± 5 | +30 ± 11 | 3362 ± 135 | +15 ± 5 | +27 ± 12 |
| <i>Turb3. (Sc. 2)</i> | 5501 ± 86 | - | - | 4757 ± 164 | - | - |
| <i>Abl. (Sc. 1)</i> | 3968 ± 29 | +5 ± 1 | -15 ± 2 | 4035 ± 7 | +7 ± 0.3 | -16 ± 3 |
| <i>Abl. (Sc. 2)</i> | 4430 ± 22 | - | - | 3808 ± 275 | - | - |

Table 7: Results at the end of training IPPO and MAPPO, on 50k and 500k time-steps for *Turb3Row1* (*Turb3.* in the table) and *Ablaincourt* (*Abl.* in the table) respectively. *Sc. 1* (resp *Sc. 2*) corresponds to the firts Wind Scenario I (resp. II).

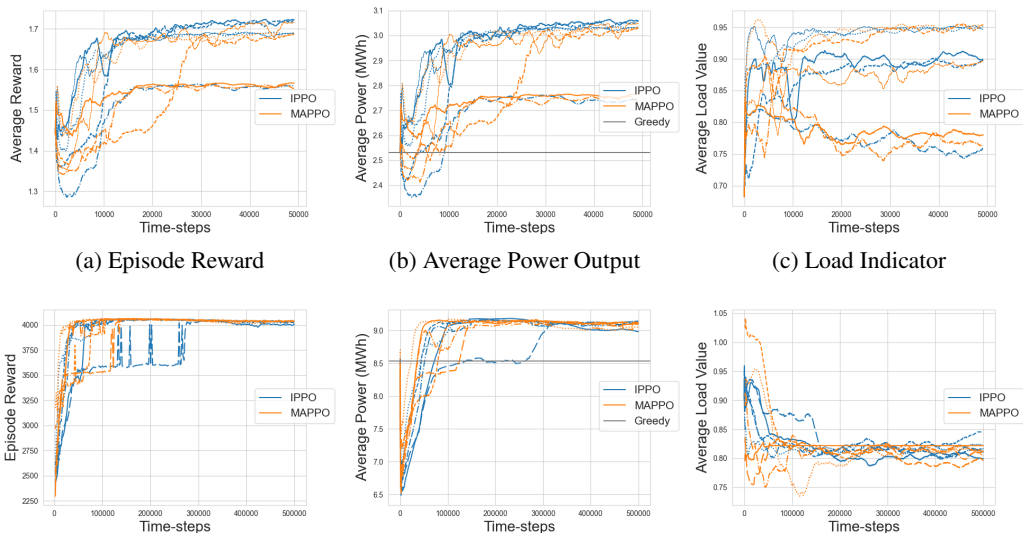




Figure 8: Evolution of episode reward, average power output and average load on the layout *Turb3Row1* (top) and *Ablaincourt* (down) simulated with FLORIS. The evolution of the episode reward is reported on the first column (a), the power output averaged on an episode length (here $T=2048$) is reported on the second column (b) and the loading indicator is on column (c).


570 **G Visual Overview of Layouts**




(a) Turb6_Row2: 6 turbines




(b) Ablaincourt: 7 turbines




(c) Turb16_Row5: 16 turbines




(d) Turb_TCRWP: 32 turbines




(e) Ormonde: 31 turbines




(f) Turb32_Row5: 32 turbines




(g) WMR: 36 turbines




(h) HornsRev1: 76 turbines




(i) HornsRev2: 92 turbines



(j) TurbX_Row1. X = 3



(k) X=6



(l) X = 12

Figure 9: Coordinates of each wind turbine for the pre-registered layouts in WFCRL. Distances are in turbine diameters ($126m$ for the NREL 5MW Reference turbine). The *TurbX_Row1* toy layouts are procedurally generated for any value of X between 1 and 12.

571 **H Additional information on WFCRL**

572 **H.1 List of dependencies**

573 We report in the table below the list of open-source Python packages and other open-source software
574 that WCFRL relies on.

| Software | License | License Link |
|----------------------|--------------------------------|---|
| numpy | <i>Custom</i> | https://numpy.org/doc/stable/license.html |
| Gymnasium | MIT | https://github.com/Farama-Foundation/Gymnasium/blob/main/LICENSE |
| PettingZoo | MIT | https://github.com/Farama-Foundation/PettingZoo/blob/master/LICENSE |
| Floris | Apache v2.0 | https://github.com/NREL/floris/blob/main/LICENSE.txt |
| FAST.Farm (OpenFAST) | Apache v2.0 | https://github.com/OpenFAST/openfast/blob/main/LICENSE |
| mpi4py | <i>Custom</i> | https://github.com/erdc/mpi4py/blob/master/LICENSE.txt |
| Microsoft-MPI | MIT | https://github.com/microsoft/Microsoft-MPI/blob/master/LICENSE.txt |
| Open MPI | BSD 3-Clause | https://www.open-mpi.org/community/license.php |
| Seaborn | BSD 3-Clause | https://github.com/mwaskom/seaborn/blob/master/LICENSE.md |
| Matplotlib | <i>Custom - BSD-compatible</i> | https://matplotlib.org/stable/project/license.html |
| PyYAML | MIT | https://github.com/yaml/pyyaml/blob/main/LICENSE |
| Pandas | BSD 3-Clause | https://github.com/pandas-dev/pandas/blob/main/LICENSE |

575 **H.2 Licence**

576 The WFCRL package is licensed under the Apache v2 license. The text of the license can be found
577 here: <https://github.com/ifpen/wfcrl-env/blob/main/LICENSE>.

578 **H.3 Responsibility**

579 The authors bear all responsibility in case of violation of rights.