

On the importance of wind predictions in wake steering optimization

Elie KADOCHE ^{1,2}, Pascal BIANCHI ², Florence CARTON ¹, Philippe CIBLAT ², and Damien ERNST ^{2,3}

¹TotalEnergies OneTech, 2 Place Jean Millier, 92400 Courbevoie, France

²Télécom Paris, 19 Place Marguerite Perey, 91120 Palaiseau, France

³Montefiore Institute, University of Liège, Belgium

Correspondence: Elie KADOCHE (elie.kadoche@totalenergies.com)

Abstract. Wake steering is a technique that optimizes the energy production of a wind farm by employing yaw control to misalign upstream turbines with the incoming wind direction. This work highlights the important dependence between wind direction variations and wake steering optimization. The problem is formalized over time as the succession of multiple steady-state yaw control problems interconnected by the rotational constraints of the turbines and the evolution of the wind. Then, this work proposes a reformulation of the yaw optimization problem of each time step by augmenting the objective function by a new heuristic based on a wind prediction. The heuristic acts as a penalization for the optimization, encouraging solutions that will guarantee future energy production. Finally, a synthetic sensitivity analysis of the wind direction variations and wake steering optimization is conducted. Because of the rotational constraints of the turbines, as the magnitude of the wind direction fluctuations increases, the importance of considering wind prediction in a steady-state optimization is empirically demonstrated. The heuristic proposed in this work greatly improves the performance of controllers and significantly reduces the complexity of the original sequential decision problem by decreasing the number of decision variables.

Nomenclature

	α_t^i	Yaw angle		$f_{\text{simulation}}$	Local velocities computation function
	β_t^i	Absolute orientation	25	f_{yaw}	Yaw angle computation function
	Δt	Time step duration		H	Horizon length
	$\epsilon_{K,t}$	Wind direction noise		K_t	Wind direction
	$\epsilon_{V,t}$	Wind speed noise		K'_t	Observed wind direction
	ν_t^i	Local wind speed		N	Number of wind turbines
	Φ	Theoretical power output function	30	P_t^i	Turbine power output
20	π	Control policy		s_t	Controller state
	f_{control}	Absolute orientation update function		u_t	All turbine yaw settings
	f_{obj}	Controller objective function		u_t^i	Yaw setting
	f_{power}	Power output function	35	V_t	Wind speed
				V'_t	Observed wind speed

1 Introduction

As global energy consumption increases, there is a strong willingness and necessity to decarbonize electricity production. Hence, renewable energies are becoming increasingly important (Chu and Majumdar, 2012). Wind energy, particularly, is the focus of considerable research and development, with turbines becoming larger and more numerous within wind farms. Assuring efficient control as wind turbines operate is necessary to maximize the benefits of wind energy.

In the context of global warming, designing more efficient wind farms is essential. Wake steering is the subject of growing interest within the community to optimize the energy production of wind farms. However, most of research regarding wind farm control technologies disregards the relevance of the wind direction variation. This work is motivated by a central question: *from what magnitude of wind direction fluctuations is it necessary to consider the wind evolution into a wake steering optimization?*

45 To answer this question, this work proposes a new controller based on wind predictions and conducts a synthetic sensitivity analysis of wake steering and wind evolution, using steady-state models and artificial wind data.

In wind farm optimization, the use of low-fidelity models (usually based on steady-state models) is favored over higher fidelity models (usually based on computational fluid dynamics and real-time wake interaction) due to the complexity and computational load associated with solving dynamic equations for every turbines in the farm. Some recent works such as
50 Janssens and Meyers (2024) explore real-time optimal control of wind farms using large-eddy simulations (LES). However, this research area is still in the early stages and for large-scale wind farm optimization, steady-state models are still widely used.

In wind farm flow control (WFFC), developing effective closed-loop controllers is essential for scaling to larger wind farms and dealing with unpredictable wind conditions. These controllers dynamically adapt their strategies in real-time, using continuous sensor feedback to guide their decisions. Model-based, closed-loop controllers, in particular, rely on simulators of the
55 environment to conduct continuous optimization while the farm is in operation. Fast and computationally efficient simulation is crucial for these controllers to quickly react to wind and turbines changes. This work focuses on the optimization process itself, adhering to community standards by using widely accepted, open-source, low-fidelity simulators.

1.1 Wake effect

60 A single wind turbine reaches its maximum power output when fully aligned with the wind. When the wind direction changes, a turbine uses its yaw to rotate its nacelle on a horizontal plane. By using active yaw control, a wind turbine can keep track of the changes in the wind direction changes and ensure maximum energy production over time by minimizing its misalignment with the wind. It corresponds to greedy control, where a wind turbine solely tries to maximize its power output (Yang et al., 2021).

65 In the space immediately behind a turbine, the wind speed is slower and more turbulent. Such a phenomenon is called the "wake effect" and is the natural consequence of wind power extraction by the machine. When a wind turbine is located in the wake of another, its power output is reduced (because of a slower wind speed) and its fatigue increased (because of the turbulence). Within a wind farm, depending on the wind direction and the farm layout, most of the turbines can be affected by the wake of others.

70 Because of wake effects, greedy control can be suboptimal within a farm. Therefore, instead of keeping every turbine aligned with the wind, yaw control can also be used to voluntary misalign some turbines in relation to the direction of the wind (Boersma et al., 2017). When a turbine is misaligned with the wind, its wake effect is steered. By intelligently yawing the turbines and steering the wake effects, the wind flow across the turbines can be optimized. Such a method is known as WFFC (Meyers et al., 2022). A simple example on a two-turbine wind farm is given in Figure 1.




Figure 1. Example of WFFC on a two-turbine wind farm with the wind coming from the west. The first (upstream) turbine is misaligned and its wake effect is steered away from the second (downstream) turbine. By letting the wind flow more freely to the second turbine, the misalignment of the first turbine increases the total power output of the farm.

75 Current implemented wake steering strategies usually involve lookup tables (LUTs) (Fleming et al., 2017; Siemens Gamesa Re-
 newable Energy, 2019). Wake steering strategies are computed for a finite set of different wind conditions prior to the farm
 operation. The yaw angles of each turbine are computed with steady-state models, regardless of the wind and turbines dynam-
 ics. Because a wake steering strategy creates misalignment with the wind, it is highly dependent on variations in the direction
 of the wind. The wind direction can change over time, and yaw control is constrained by the limited rotational speed of the
 80 nacelles. If the wind varies in directions and frequencies that the yaw actuators cannot easily track, computing adequate wake
 steering strategies over time can be a challenging task.

1.2 Wind direction dynamics

The study of wind direction dynamics is gaining interest within the research community. Wind direction dynamics can be
 broken down into large-scale drifts and small-scale fluctuations (van Doorn et al., 2000) and can be observed on different
 85 scales: the synoptic scale describes long distances and extended time periods, the mesoscale depicts the farm level and time
 periods from days to weeks, and the microscale corresponds to the turbine level and variations from seconds to minutes. The
 wind direction is fundamentally non-stationary, and there is an incomplete knowledge regarding the physical and statistical
 characteristics of wind direction fluctuations across specific length and time scales that are essential for effective WFFC (Dallas
 et al., 2023).

90 As the farm operates, the wind direction varies both in time (at the farm level) and space (at the turbine level). The authors
 von Brandis et al. (2023) found that spatial wind direction changes relevant to the operation of wind farm clusters in the German
 Bight exceed 11 degrees in 50 % of cases. In this present work, numerical simulations are run with steady-state wake models.
 Therefore, only variations in the wind at the farm level are studied. When the direction varies over time, this work considers
 that it affects the whole wind farm.

95 WFFC is most beneficial at low wind speeds because this is where small changes of the wind speeds can lead to important
 power output variations. The same wake steering strategy will lead to higher power gains at low speeds compared to higher

wind speed. Because the wind direction variability is higher for low wind speeds (von Brandis et al., 2023; van Doorn et al., 2000; Dallas et al., 2023), the study of dependence between wind direction variations and yaw control is important. Also, because the impact of climate change on wind dynamics is unknown, designing robust controllers is necessary for long-term operation.

1.3 Related works

As tracking wind direction is essential for wind turbines, the literature is rich in studies seeking better wind direction tracking mechanisms. Song et al. (2018) developed a model predictive control (MPC)-based controller on a finite control set to track the wind directions. Hure et al. (2015) designed a yaw controller based on very short-term wind predictions. But performing WFFC and wake steering is a more complex optimization problem.

LUTs can be adapted for dynamic control with different methods. Usually, a low-pass filter is used to apply control only for high variations of the direction. A sampling method can be used to adjust the yaw control frequency and hysteresis mechanisms avoid unnecessary yaw control and restrict the yaw actuators (Kanev, 2020a). Simley et al. (2021) improved a traditional LUT by anticipating the wind direction changes ahead of upstream turbines. Kanev (2020b) performed WFFC with receding horizon using gradient-based optimization and run tests in large-eddy simulations under realistic variations in wind direction and speed. But the wake steering strategies of an LUT fundamentally do not consider the wind dynamics, only their implementation does.

Regarding machine learning (ML) methods, and more particularly reinforcement learning (RL), which is becoming a source of great interest to the scientific community, wind direction variations are often overlooked. The importance of the wind direction dynamics is clearly pointed out by Saenz-Aguirre et al. (2019) and Saenz-Aguirre et al. (2020) but most of the studies carried out later only consider static or quasi-static wind directions. Some recent works have started to consider time-varying wind directions into WFFC optimization (Kadoche et al., 2023).

1.4 Contributions

The remainder of this paper is structured to mirror the three main contributions. Each contribution forms the basis of an individual Section and Section 5 concludes. The contributions, and their corresponding Sections, are as follows.

- This work proposes a discretized formalization of the WFFC problem over time as the succession of multiple steady-state optimization problems interconnected by the rotational constraints of the turbines and the evolution of the wind. Due to the discretization hypothesis and the yaw actuation constraints, the important hypotheses regarding the transition between one steady-state to the next are formulated. This formalization is conducted in Section 2.
- To develop a prediction-based controller, this work presents a reformulation of the instantaneous, steady-state, original sequential decision problem over a future time window. The default objective function is augmented by a new heuristic, computed on a prediction of the wind. The proposed heuristic acts as a penalization for the optimization without increasing its dimension and encourages solutions that will guarantee future energy production. The heuristic and the other studied controllers are detailed in Section 3.

– This work conducts a sensitivity analysis of the wind direction variations and wake steering optimization. It empirically demonstrates the importance of a wind prediction-based control when the magnitude of the wind direction fluctuations become large. The new proposed heuristic greatly improves the performance of a traditional steady-state wake steering optimization when the variations of the wind direction are important. Numerical simulations using synthetic wind data are conducted in Section 4.

2 Problem formalization

The environment is composed of a wind farm, and some exogenous variables related to the wind direction and the wind speed. The wind farm consists in N interconnected wind turbines. Over time, each turbine is controlled via its yaw. An episode consists of a succession of H time steps during which the turbines are controlled, with H the horizon length. The transition from a time step t to a time step $t+1$ corresponds to a specific time window of constant length of Δt minutes. The environment evolves from one time step t to the next time step $t+1$ in Δt minutes, with $t \in \{0, 1, \dots, H-1\}$.

2.1 Wind data

At a time step t , the exogenous variables are a global incoming wind direction $K_t \in [0, 360]$ [degrees] and a global incoming wind speed $V_t \in [\nu_{\min}, \nu_{\max}]$ [m/s]. The wind data can be measured or predicted. Then, a controller does not have access to K_t and V_t directly but to K'_t [degrees] and V'_t [m/s], with K'_t a noisy wind direction defined as $K'_t = K_t + \epsilon_{K,t}$ and V'_t is a noisy wind speed defined as $V'_t = V_t + \epsilon_{V,t}$. The random noises $\epsilon_{K,t}$ and $\epsilon_{V,t}$ can come from either measurement imprecisions or prediction errors.

Because the turbines alter the wind flow inside the farm, the wind in front of a turbine can be different from the global incoming wind. Then, at a time step t , in front of a turbine i , the local wind direction is noted κ_t^i [degrees] and the local wind speed is noted ν_t^i [m/s]. The computation of the local velocities is based on complex fluid mechanics and is subject to numerous uncertainties. This work conducts numerical experiments where such computations are done using a low-fidelity, steady-state simulator FLOW Redirection and Induction in Steady State (FLORIS) NREL (2021). The local wind directions stay equal to the global wind direction for any time steps, i.e., $\kappa_t^i = K_t \forall i \in \{0, 1, \dots, N-1\}$, as it is common for low-fidelity, steady-state simulation.

2.2 Turbines

At a time step t , a turbine i is characterized by its absolute angular position $\beta_t^i \in [0, 360]$ [degrees] and its relative orientation or yaw (often used to compute the power output) $\alpha_t^i = f_{\text{yaw}}(K_t, \beta_t^i) \in [-180, 180]$ [degrees], such that

$$f_{\text{yaw}}(K_t, \beta_t^i) = (K_t - \beta_t^i + 180) \bmod 360 - 180. \quad (1)$$

Adding and subtracting by 180 ensures that the yaw stays in the range $[-180, 180]$. As illustrated in Figure 2, the yaw corresponds to the rotational movement going from the absolute angular position β_t^i to the wind direction K_t , such that $\beta_t^i + \alpha_t^i$

mod 360 = K_t . Positive values of the yaw indicate that the turbine is rotated anti-clockwise from the wind direction and negative values of the yaw indicate that the turbine is rotated clockwise from the wind direction.

At a time step t , the yaw setting $u_t^i \in [u_{\min}, u_{\max}]$ [degrees] of a turbine i corresponds to the rotational movement of the turbine between time steps t and $t + 1$. Because of mechanical constraints related to the yaw actuator of the nacelle, the yaw setting is bounded between two consecutive time steps. As illustrated in Figure 2, the setting is used to update the orientation of the turbine $\beta_{t+1}^i = f_{\text{control}}(\beta_t^i, u_t^i)$ such that

$$f_{\text{control}}(\beta_t^i, u_t^i) = (\beta_t^i + u_t^i) \bmod 360. \quad (2)$$




Figure 2. Example of a wind turbine i seen from above at a time step t . The variables are the wind direction K_t , the absolute angular position β_t^i of the turbine, and the yaw α_t^i of the turbine. The wind direction indicates from where the wind is coming, e.g., a wind direction of 270 degrees indicates a wind coming from the west. Here, the nacelle is misaligned with the incoming wind direction: the turbine is rotated clockwise from the wind, so $\alpha_t^i < 0$. The yaw setting u_t^i gives the next orientation of the turbine at time step $t + 1$.

2.3 Power

The power curve Φ of a turbine gives the theoretical power output [megawatts (MW)] (y-axis) of the machine as a function of the wind speed ν [m/s] (x-axis), considering no yaw misalignment, such that

$$170 \quad \Phi(\nu) = \frac{1}{2} \cdot \rho \cdot A \cdot \nu^3 \cdot C_P(\nu), \quad (3)$$

with ρ [kg/m³] the air density, A [m²] the rotor blade area and C_P the power coefficient of the turbine. The theoretical power output is strictly positive if the wind speed is within certain bounds $[\nu_{\text{cut-in}}, \nu_{\text{cut-out}}]$ [m/s].

At a time step t , the power output of a turbine i considering yaw misalignment $P_t^i = f_{\text{power}}(\nu_t^i, \alpha_t^i)$ [MW] is computed from the power curve and the yaw angle, such that

$$175 \quad f_{\text{power}}(\nu_t^i, \alpha_t^i) = \Phi(\nu_t^i) \cdot \cos^p(\alpha_t^i) \cdot \mathbb{1}_{\{\alpha_{\text{cut-in}} \leq \alpha_t^i \leq \alpha_{\text{cut-out}}\}}, \quad (4)$$

with p a parameter accounting for power losses due to misalignment and $[\alpha_{\text{cut-in}}, \alpha_{\text{cut-out}}]$ [degrees] a safety bound for the yaw took into account by the indicator function. Because too much misalignment with the wind can damage the machine, if the yaw is too great, the turbine is shut down and its power output is null.

2.4 Policy

180 A policy π is a function returning the yaw settings $(u_t^0, u_t^1, \dots, u_t^{N-1})$ of all the turbines at a time step t given a state s_t . Each wind farm controller is associated with a specific policy. In this work, the state s_t may be composed of:

- the current orientations of the turbines $\{\beta_t^i\}_{i \in \{0,1,\dots,N-1\}}$;
- an observation of the current wind (K'_t, V'_t) ;
- a prediction of the wind at time step $t+1$, (K'_{t+1}, V'_{t+1}) ;
- 185 – a prediction of the wind at time step $t+2$, (K'_{t+2}, V'_{t+2}) ;
- until time step $t+L$, with L the prediction horizon.

Therefore, the general form of a state is $s_t = \{K'_t, V'_t, K'_{t+1}, V'_{t+1}, \dots, K'_{t+L}, V'_{t+L}, \{\beta_t^i\}_{i \in \{0,1,\dots,N-1\}}\}$. States can be categorized based on two distinct properties: with perfect or imperfect information and with or without foresight knowledge. Depending on the possible combinations, there are four classes of states, listed in Table 1.

	foresight	no foresight
perfect information	$\epsilon_{K,t+k} = 0$ and $\epsilon_{V,t+k} = 0, L > 0$	$\epsilon_{K,t+k} = 0$ and $\epsilon_{V,t+k} = 0, L = 0$
imperfect information	$\epsilon_{K,t+k} \neq 0$ or $\epsilon_{V,t+k} \neq 0, L > 0$	$\epsilon_{K,t+k} \neq 0$ or $\epsilon_{V,t+k} \neq 0, L = 0$

Table 1. Four different classes of states based on two distinct properties, with $t \in \{0, 1, \dots, H-1\}$ and $k \in \{0, 1, \dots, L\}$. With perfect information, the state comprises exact wind data (no noise). Without foresight, the state only comprises the current wind data (no predictions).

An episode is defined by H time steps during which turbines are controlled via their yaw. An episode is characterized by time series for the wind directions and the wind speeds and initial positions for the nacelles. During an episode, it is assumed that all the states belong to the same class (defined in Table 1) and the policy is presumed to be stationary (it does not change over time).

195 The full evolution of an episode is described in Algorithm 1. At each time step, the policy returns the yaw settings based on the current state, the system is updated and the power output of the farm is computed. The yaw setting of a turbine i at the end of time step t is indexed $t + 1$ (because it has been updated) and it is the one used for the power computation of time step t .

At a time step t , to compute the power output of each turbine, local wind velocities are needed. Such computations rely on complex fluid mechanics, depending on the incoming wind and the updated yaw angles of each turbine. For optimization, performing such complex computations is computationally expensive. Therefore, in this work, these computations are carried out by a steady-state simulator $\nu_t^i = f_{\text{simulation}}^i(K_t, V_t, \{\alpha_{t+1}^j\}_{j \in \{0,1,\dots,N-1\}})$, which is used as a substitute for real-life measurements.

The simulation is said to be steady-state because it only depends on the current global wind data and the updated yaw angles. It does not consider previous wind data, previous yaw angles or time delays in the wake propagation. The evolution of an episode over time is constrained by the rotational bounds of the turbines and the variations of the wind.

Algorithm 1 Full episode evolution over time

Input: $\{K_k, V_k\}_{k \in \{0,1,\dots,H+L-1\}}$ true wind data time series

$\{\beta_0^i\}_{i \in \{0,1,\dots,N-1\}}$ initial orientations

π yaw control policy

for $t = \{0, 1, \dots, H - 1\}$ **do**

$s_t = \{K'_t, V'_t, K'_{t+1}, V'_{t+1}, \dots, K'_{t+L}, V'_{t+L}, \{\beta_t^i\}_{i \in \{0,1,\dots,N-1\}}\}$ ▷ state

$(u_t^0, u_t^1, \dots, u_t^{N-1}) = \pi(s_t)$ ▷ control policy

$\beta_{t+1}^i = f_{\text{control}}(\beta_t^i, u_t^i), \forall i \in \{0, \dots, N - 1\}$ ▷ angular positions update

$\alpha_{t+1}^i = f_{\text{yaw}}(K_t, \beta_{t+1}^i), \forall i \in \{0, \dots, N - 1\}$ ▷ yaw angles computation

$\nu_t^i = f_{\text{simulation}}^i(K_t, V_t, \{\alpha_{t+1}^j\}_{j \in \{0,1,\dots,N-1\}}), \forall i \in \{0, \dots, N - 1\}$ ▷ local velocities computation

$P_t^i = f_{\text{power}}(\nu_t^i, \alpha_{t+1}^i), \forall i \in \{0, \dots, N - 1\}$ ▷ power outputs computation

end for

Output: $\sum_{t=0}^{H-1} \sum_{i=0}^{N-1} P_t^i$ farm power output

At each time step, during the "control policy" operation, a controller knows the evolution mechanisms of the system, i.e., it can conduct any computations with the f_{control} , f_{yaw} , $f_{\text{simulation}}^i$ and f_{power} functions but based on the wind data provided by the state. Because such data can be noisy, all the computed values can be inexact. For example, at a time step t , if a controller computes the yaw of a turbine i based on its updated orientation β_{t+1}^i , it would be equal to $\alpha_{t+1}^i = f_{\text{yaw}}(K'_t, \beta_{t+1}^i)$. Because

210 the observed wind direction K'_t can be different from the true wind direction K_t , the yaw α_{t+1}^i estimated by a controller can be different from its true value α_{t+1}^i .

2.6 Transition regime

At a time step t , for a turbine i , and due to the steady-state nature of the simulation, the WFFC problem thus formalized considers a single power output P_t^i . In reality, during a time step, the wind is time varying and a turbine takes time to rotate because of mechanical constraints. Therefore, the discretization of the continuous control problem results in the loss of some information and possibly less inaccurate power outputs. To ensure that the discretized power outputs are good approximations, from one time step to another, a turbine is supposed to rotate immediately and the wind is supposed to be quasi-constant.

220 The duration of a time step is always considered constant during an episode. At a time step t , when a setting u_t^i is applied to a turbine i , the rotational time T_r [minutes] for the turbine to go from its current orientation β_t^i to its next orientation β_{t+1}^i is always considered largely inferior to the duration of the time step, i.e., $T_r \ll \Delta t$ for all $u_t^i \in [u_{\min}, u_{\max}]$. A turbine always reaches rapidly its target position, before the end of the time step duration. But during a time step, no other control will be applied to the turbine. For this reason, the rotational constraints $[u_{\min}, u_{\max}]$ need to be consistent with the duration of a time step Δt .

225 The coherence time T_c [minutes] of a wind variable (either the direction or the speed) is the maximum duration during which the variable is quasi-constant. If the coherence time of the wind direction is strictly smaller than the time step duration, a discretized value K_t would stretch too far away from its corresponding continuous signal. The same goes for the speed. Therefore, in this work, the coherence time is always equal to the time step duration, i.e., $T_c = \Delta t$, for both the direction and the speed.

3 Controllers

230 At a time step t , the yaw settings $(u_t^0, u_t^1, \dots, u_t^{N-1})$ are noted u_t . A controller is defined by its policy $\pi(s_t)$ with the state s_t described in Subsection 2.4. This work compares a naive control where each turbine is aligned as much as possible with the wind and three optimized wake steering control strategies. In an episode, at each time step t , during the "control policy" operation of Algorithm 1, each controller computes the yaw settings such that $u_t = \pi(s_t)$, by maximizing a specific objective function $f_{\text{obj}}(s_t, u_t)$ with regard to the turbine rotational constraints, as defined below.

$$235 \quad \pi(s_t) \in \arg \max_{u_t} f_{\text{obj}}(s_t, u_t), \quad (5)$$

$$\text{subject to } u_{\min} \leq u_t^i \leq u_{\max}, \quad \forall t \in \{0, 1, \dots, H-1\}, \quad \forall i \in \{0, 1, \dots, N-1\}. \quad (6)$$

3.1 Naive controller

The naive controller always tries to keep turbines aligned with the current wind direction as much as possible. It is a weak
 240 baseline as it does not conduct any wake steering optimization. It runs with no foresight (i.e., $L = 0$), as it is only concerned
 with the current observed wind direction K'_t . Therefore, the state s_t is reduced to $\{K'_t, \{\beta_t^i\}_{i \in \{0,1,\dots,N-1\}}\}$. It consists of a
 greedy control (no wake steering) where the objective function at a time step t is minimizing the amplitude of the yaws, i.e.,

$$f_{\text{obj}}(s_t, u_t) = - \sum_{i=0}^{N-1} |\alpha_{t+1}^i|, \quad (7)$$

$$\text{with } \alpha_{t+1}^i = f_{\text{yaw}}(K'_t, \beta_{t+1}^i), \quad (8)$$

$$245 \quad \beta_{t+1}^i = f_{\text{control}}(\beta_t^i, u_t^i). \quad (9)$$

At a time step t , the rotational movement required for a turbine i to stay aligned with the observed wind direction is equal to
 $f_{\text{yaw}}(K'_t, \beta_t^i)$. Because of the rotational constraints, this movement is clipped, such that it is always an acceptable setting with
 regards to the yaw actuator, giving a closed-form expression for the solution, defined as

$$\pi_{\text{naive}} = (u_t^0, u_t^1, \dots, u_t^{N-1}), \quad (10)$$

$$250 \quad \text{such that } u_t^i = \text{clip}(f_{\text{yaw}}(K'_t, \beta_t^i), u_{\min}, u_{\max}), \forall i \in \{0, 1, \dots, N-1\}. \quad (11)$$

3.2 Wake steering

Compared to naive control, wake steering is used to optimize the power output of the farm. In this work, two distinct wake
 steering strategies are used. One is based only on the instantaneous wind data, and one is based on instantaneous and predicted
 wind data. The instantaneous controller searches for the yaw settings maximizing the instantaneous power output of the farm.
 255 The prediction-based controller maximizes the instantaneous and future power outputs. At each time step t , the same Gauss-
 Seidel (GS) method is used for both controllers, but with different objective functions. In this work, optimization is conducted
 with a GS method (described in Algorithm A1 of Appendix A). A similar approach was first proposed by Fleming et al. (2022)
 with a serial-refine algorithm.

The GS method works as follow. A first solution is initialized from the naive controller, where each initial yaw setting
 260 keeps its turbine aligned as much as possible with the wind. Then, the GS method iterates over each turbine, from upstream to
 downstream ones. At each iteration, it solves the optimization problem for the current turbine, considering the yaw settings of all
 others fixed, by conducting a grid-search over a discretized solution space $S = \{u_{\min} + l \cdot \frac{u_{\max} - u_{\min}}{n_y - 1}\}$, for all $l \in \{0, 1, \dots, n_y - 1\}$
 with n_y being a precision parameter. Once optimized, the setting of the current turbine is updated, and it goes to the next one.

3.2.1 Instantaneous controller

265 The instantaneous controller searches for the yaw settings maximizing the immediate power output of the farm. It always runs
 under no foresight (i.e., $L = 0$), as it performs wake steering for the current, observed wind data only. Therefore, the state s_t

is reduced to $\{K'_t, V'_t, \{\beta_t^i\}_{i \in \{0,1,\dots,N-1\}}\}$. It is a steady-state optimization performed on one time step where the objective function at a time step t is the immediate normalized power output, i.e.,

$$f_{\text{obj}}(s_t, u_t) = \frac{1}{N} \cdot \sum_{i=0}^{N-1} f_{\text{power}}(\nu_t^{i'}, \alpha_{t+1}^{i'}), \quad (12)$$

$$270 \quad \text{with} \quad \nu_t^{i'} = f_{\text{simulation}}^i(K'_t, V'_t, \{\alpha_{t+1}^j\}_{j \in \{0,1,\dots,N-1\}}), \quad (13)$$

$$\alpha_{t+1}^{i'} = f_{\text{yaw}}(K'_t, \beta_{t+1}^i), \quad (14)$$

$$\beta_{t+1}^i = f_{\text{control}}(\beta_t^i, u_t^i). \quad (15)$$

3.2.2 Prediction-based controller

A traditional prediction-based controller searches for the yaw settings of time steps $t, t+1, \dots, t+L$ that maximize the power output over that horizon. It always runs with foresight (i.e., $L \geq 1$). The corresponding sequential decision problem over a future time window can be stated under a form usually exploited by the MPC community, defined as

$$\max_{u_t, u_{t+1}, \dots, u_{t+L}} \frac{1}{N} \cdot \sum_{k=0}^L \sum_{i=0}^{N-1} f_{\text{power}}(\nu_{t+k}^{i'}, \alpha_{t+k+1}^{i'}), \quad (16)$$

$$\text{subject to} \quad u_{\min} \leq u_{t+k}^i \leq u_{\max}, \quad \forall t \in \{0, 1, \dots, H-1\}, \quad \forall k \in \{0, 1, \dots, L\}, \quad \forall i \in \{0, 1, \dots, N-1\}, \quad (17)$$

$$\text{with} \quad \nu_{t+k}^{i'} = f_{\text{simulation}}^i(K'_{t+k}, V'_{t+k}, \{\alpha_{t+k+1}^j\}_{j \in \{0,1,\dots,N-1\}}), \quad (18)$$

$$280 \quad \alpha_{t+k+1}^{i'} = f_{\text{yaw}}(K'_{t+k}, \beta_{t+k+1}^i), \quad (19)$$

$$\beta_{t+k+1}^i = f_{\text{control}}(\beta_{t+k}^i, u_{t+k}^i). \quad (20)$$

The optimization problem thus described multiplies the number of decision variables by $L+1$. Also, the computation of the local velocities at a given time step depends on all the previous yaw settings. Therefore, the prediction-based decision problem significantly increases the complexity, and because there is no simple solution, this work proposes a reformulation. The objective function given by Equation (16) can be split between the current time step t and the next ones, from $t+1$ to $t+L$, such that

$$\frac{1}{N} \cdot \sum_{k=0}^L \sum_{i=0}^{N-1} f_{\text{power}}(\nu_{t+k}^{i'}, \alpha_{t+k+1}^{i'}) = \underbrace{\frac{1}{N} \cdot f_{\text{power}}(\nu_t^{i'}, \alpha_{t+1}^{i'})}_{\text{first term}} + \underbrace{\frac{1}{N} \cdot \sum_{k=1}^L \sum_{i=0}^{N-1} f_{\text{power}}(\nu_{t+k}^{i'}, \alpha_{t+k+1}^{i'})}_{\text{second term}}. \quad (21)$$

The first term of Equation (21) is the normalized power output of the farm for the current time step. It corresponds to the objective function of the instantaneous controller defined in Subsubsection 3.2.1. It only depends on the current yaw settings u_t . Now, focusing on the second term, the closed-form expression of the f_{power} is written, giving

$$\frac{1}{N} \cdot \sum_{k=1}^L \sum_{i=0}^{N-1} f_{\text{power}}(\nu_{t+k}^{i'}, \alpha_{t+k+1}^{i'}) = \frac{1}{N} \cdot \sum_{k=1}^L \sum_{i=0}^{N-1} \Phi(\nu_{t+k}^{i'}) \cdot \cos^p(\alpha_{t+k+1}^{i'}) \cdot \mathbb{1}_{\{\alpha_{\text{cut-in}} \leq \alpha_{t+k+1}^{i'} \leq \alpha_{\text{cut-out}}\}}. \quad (22)$$

The complexity brought by the prediction-based controller comes from the fact that Equation (22) depends on the local velocities ν_{t+k}^i and the updated yaw angles α_{t+k+1}^i corresponding to the optimized yaw settings of each future time step. To decrease the complexity, this work proposes to modify Equation (22) in the following way.

- 295 – Each local velocity ν_{t+k}^i is replaced by the corresponding, predicted, global wind speed V'_{t+k} . It reduces the complexity coming from the steady-state simulation by removing the dependence with the updated yaw angles. While this simplification removes future local wind data, it retains some information regarding potential future energy production by relying on the predicted global wind data only.
- Each updated yaw angle α_{t+k+1}^i depending on the optimized yaw setting u_{t+k}^i is replaced by the expected yaw angle $\hat{\alpha}_{t+k+1}^i$ if a naive controller was used instead. It reduces the complexity coming from the optimization, as there is a closed-form expression for the naive controller, as provided by Equation (10). Replacing the wake steering optimization performed in the future by a naive wind tracking solution reduces the number of optimization variables of the original problem while keeping good solutions. Indeed, proper yaw settings are known to be close to the wind in average.
- 300 – The cosine function at power p of each yaw angle is replaced with a simpler penalization for yaw misalignment. The penalization chosen corresponds to one minus the normalized, absolute value of that yaw angle. It provides linearity and better interpretability.
- 305 – The indicator function is removed so that there is no discontinuity. Even if a yaw is too great, it can be of some interest for the optimization to know about the potential power output. The more a turbine is misaligned, the less likely it will be to produce energy and the more it will be penalized.
- 310 – For each time step, the overall expression is multiplied by a discounted factor $\gamma \in [0, 1]$. It gives more importance to immediate time steps. It is common practice for model-predictive based optimization.

The only variables specific to each turbine are the yaw angles updated from a naive controller, which are already normalized. Therefore, it becomes unnecessary to normalize the overall expression by N . With such modifications, Equation (22) becomes a new heuristic \mathcal{H}_t defined as

$$315 \quad \mathcal{H}_t(s_t, u_t) = \sum_{k=1}^L \gamma^{k-1} \cdot \Phi(V'_{t+k}) \cdot \left(1 - \frac{1}{N} \cdot \frac{1}{180} \cdot \sum_{i=0}^{N-1} |\hat{\alpha}_{t+k+1}^i| \right), \quad (23)$$

$$\text{with } \hat{\alpha}_{t+k+1}^i = f_{\text{yaw}}(K'_{t+k}, \hat{\beta}_{t+k+1}^i), \quad (24)$$

$$\hat{\beta}_{t+k+1}^i = f_{\text{control}}(\hat{\beta}_{t+k}^i, \hat{u}_{t+k}^i), \quad \hat{u}_{t+k}^i \text{ computed with a naive controller defined by Equation (10),} \quad (25)$$

$$\hat{\beta}_{t+1}^i = f_{\text{control}}(\beta_t^i, u_t^i), \quad u_t^i \text{ computed from a wake steering optimization.} \quad (26)$$

Because this new proposed heuristic does depend on neither the future optimized yaw settings (naive control), nor the future local velocities (no simulation), it does not increase the number of optimization variables. The heuristic is a scalar acting as a

penalization for the optimization. The final objective function of the prediction-based controller can finally be written as

$$f_{\text{obj}}(s_t, u_t) = \frac{1}{N} \cdot \sum_{i=0}^{N-1} f_{\text{power}}(\nu_t^{i'}, \alpha_{t+1}^{i'}) + \mathcal{H}_t(s_t, u_t), \quad (27)$$

$$\text{with } \nu_t^{i'} = f_{\text{simulation}}^i(K_t', V_t', \{\alpha_{t+1}^{j'}\}_{j \in \{0,1,\dots,N-1\}}), \quad (28)$$

$$\alpha_{t+1}^{i'} = f_{\text{yaw}}(K_t', \beta_{t+1}^i), \quad (29)$$

$$325 \quad \beta_{t+1}^i = f_{\text{control}}(\beta_t^i, u_t^i), \quad (30)$$

$$\mathcal{H}_t(s_t, u_t) \text{ defined by Equation (23)}. \quad (31)$$

The heuristic is the discounted, weighted sum of the future theoretical power outputs. By choosing certain optimized yaw settings u_t for the current time step, the heuristic uses a naive controller over a future time horizon of L time steps to evaluate how well the turbines will manage to stay aligned with the predicted wind directions. The higher the potential future energy production, the more critical it becomes for the current yaw settings to not rotate the turbines too far away from the predicted wind direction.

For example, if the future expected power outputs are high, the heuristic will encourage yaw settings that will put the turbines in good orientations for the future. The heuristic will penalize the objective function for yaw settings that will prevent turbines from keeping track of the wind. An illustration of the heuristic is given in Figure 3, describing the first term (wake steering optimization) and the second term (heuristic based on a wind tracking control) of Equation (27).

3.3 Upper bound

To have an upper bound in terms of performance (power output) of a wake steering strategy, the rotational constraints are relaxed. It means that in Equation (6), the variables u_{\min} and u_{\max} are equal to -180 and 180 degrees, respectively. Between two consecutive time steps, each turbine is assumed to be capable of reaching any orientation.

340 From a different point of view, the upper bound corresponds to the wake steering, instantaneous controller, but for a complete steady-state version of the system evolution presented in Algorithm 1. All time steps are entirely independent from each other, as there are no longer any rotational constraints for the turbines.

The same objective function of the instantaneous controller, presented in Subsubsection 3.2.1 is used. It always runs under no foresight (i.e., $L = 0$), as it performs wake steering for the current wind data only. Therefore, the state s_t is reduced to $\{K_t', V_t', \{\beta_t^i\}_{i \in \{0,1,\dots,N-1\}}\}$. The yaw settings computed by the upper bound would not be admissible in reality if the corresponding targeted orientations are too far away from the current ones.

4 Simulations

In Subsection 4.1 the process used to generate wind data is described and in Subsection 4.2 the experiment setting is given. Finally, the results and the empirical conclusions that can be drawn are explained in Subsection 4.3.

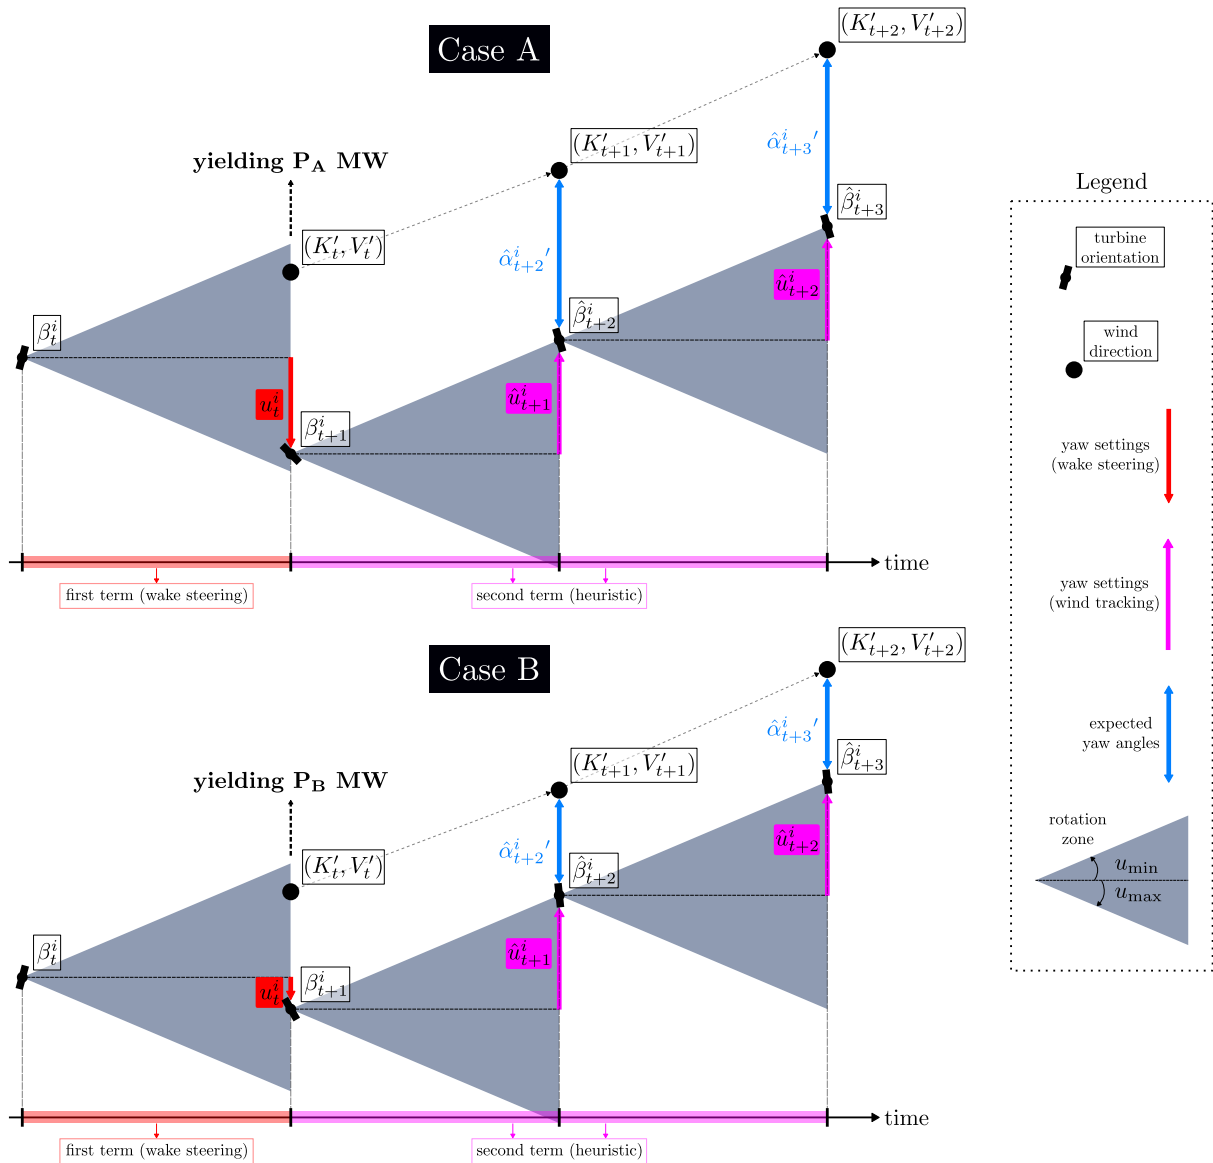


Figure 3. Illustration of the heuristic for a turbine i at time step t for two different cases. The horizon is $L = 2$ and the wind data is the same for cases A and B. The rotation zone represents the range of possible orientations a turbine can take at a given time step after being controlled. First, wake steering optimization is performed to find the setting u_t^i , which yields a power output of P_A MW for case A and P_B MW for case B at time step t . By considering $P_A > P_B$, the case A would be preferred. But then the heuristic computes the future expected yaw angles, if a naive wind tracking solution is used. From these expected yaw angles, the heuristic computes the expected power outputs based on the predicted wind speeds. Here, while the yaw setting of case A gives a better immediate solution than the one given by case B, it keeps the turbine further away (i.e., giving greater yaw angles) from the future wind. The solution of case B would then be preferable. The heuristic encourages the choice of yaw settings that may not be the best at the current time step but that ensure future power output.

350 4.1 Wind data scenario

The wind data time series are artificially generated with custom Wiener processes. The wind directions $\{K_t\}_{t \in \{0,1,\dots,H+L-1\}}$ are computed with Algorithm 2. The wind speeds $\{V_t\}_{t \in \{0,1,\dots,H+L-1\}}$ are computed with Algorithm 3. To generate the time series, an initial value is cumulatively incremented at each time step by variable m_t . Each increment m_t is independently sampled from a Normal distribution of mean 0 and standard deviation σ_t , such that

$$355 \quad \sigma_t = \tau \times \sqrt{\delta_t^X}, \quad \forall t \in \{0,1,\dots,H+L-1\} \quad (32)$$

with τ a normalization variable with regard to the number and range of the generated values and δ_t^X a variation parameter for the wind variable X (either the direction or the speed).

Algorithm 2 Wind directions generator

Input: $H + L$ number of points

K_{init} initial wind direction

$\delta_{\text{min}}^K, \delta_{\text{max}}^K$ bounds for the variation variable

$\tau = 360 / (H + L)$

for $t = \{0, 1, \dots, H + L - 1\}$ **do**

$\delta_t^K \sim U(\delta_{\text{min}}^K, \delta_{\text{max}}^K)$

$\sigma_t = \tau \times \sqrt{\delta_t^K}$

$m_t \sim N(0, \sigma_t)$

$K_t = (K_{\text{init}} + \sum_{i=0}^t m_i) \bmod 360$

end for

Output: $\{K_t\}_{t \in \{0,1,\dots,H+L-1\}}$ wind directions time series

Algorithm 3 Wind speeds generator

Input: $H + L$ number of points

V_{init} initial wind speed

$\nu_{\text{min}}, \nu_{\text{max}}$ bounds for the wind speed

$\delta_{\text{min}}^V, \delta_{\text{max}}^V$ bounds for the variation variable

$\tau = (\nu_{\text{max}} - \nu_{\text{min}}) / (H + L)$

for $t = \{0, 1, \dots, H + L - 1\}$ **do**

$\delta_t^V \sim U(\delta_{\text{min}}^V, \delta_{\text{max}}^V)$

$\sigma_t = \tau \times \sqrt{\delta_t^V}$

$m_t \sim N(0, \sigma_t)$

$V_t = \text{mirrored}(N_{\text{init}} + \sum_{i=0}^t m_i)$

end for

Output: $\{V_t\}_{t \in \{0,1,\dots,H+L-1\}}$ wind speeds time series

To maintain the wind directions in the range of valid values, i.e., $[0, 360]$ [degrees], the modulo operation is sufficient. To maintain the wind speeds in the range of valid values, i.e., $[\nu_{\min}, \nu_{\max}]$ [m/s], a mirrored function as explained in Figure 4 is proposed. The generated values inside the wind speed bounds are not modified. The generated values outside the bounds are recursively mirrored inside the bounds.




Figure 4. Toy example of the mirrored function used to keep the generated wind speeds inside specific bounds. Raw data is generated thanks to a process described by Algorithm 3. Raw data points inside the wind speed bounds are not modified: the black and red curves overlap. Data points outside the wind speed bounds are recursively mirrored inside the bounds.

The variable δ_t^X defines the level of variation of the wind variable X time series (either the direction or the speed). When equal to 0, the signal is constant. As δ_t^X increases, the absolute value of the increments increases in average. At each time step, δ_t^X is sampled from a uniform distribution defined between δ_{\min}^X and δ_{\max}^X . When δ_{\min}^X and δ_{\max}^X are equal, all increments $\{m_t\}_{t \in \{0, 1, \dots, H+L-1\}}$ are independently sampled from the same distribution: the generated time series is stationary with regard to the increments. When $\delta_{\min}^X < \delta_{\max}^X$, increments are independently sampled from different distributions: the generated time series is non-stationary with regard to the increments. In Figure 5 the impact of the δ_t^K variable is shown for the wind direction.




Figure 5. Example of different wind direction signals generated with different δ_t^K values, considering that $\delta_t^K = \delta_{\min}^K = \delta_{\max}^K$ for all $t \in \{0, 1, \dots, 49\}$ and $K_{\text{init}} = 7$ degrees. If $\delta_t^K = 0$, all the generated points are equal to K_{init} . The sine and cosine values are plotted for illustration convenience (it avoids the discontinuity issue of degrees). Note that the behavior shown in this example is the same for the wind speed, but values are in the range $[\nu_{\min}, \nu_{\max}]$.

4.2 Experimental setting

370 The function $f_{\text{simulation}}^i(K_t, V_t, \{\alpha_{t+1}^j\}_{j \in \{0, 1, \dots, N-1\}}\})$ computes the local wind speed in front of a turbine i at a time step t given wind data K_t, V_t and the yaw of each turbine $\{\alpha_{t+1}^j\}_{j \in \{0, 1, \dots, N-1\}}$. This function, introduced in Subsection 2.5, is assured by the low-fidelity, steady-state simulator FLORIS (NREL, 2021). FLORIS is used with a Gauss Curl hybrid wake model. The Gaussian velocity model is implemented based on Bastankhah and Porté-Agel (2016) and Niayifar and Porté-Agel (2016). To compute the deflection of the wakes depending on the yaws, the models described by Bastankhah and Porté-Agel (2016) and
375 King et al. (2021) are used. The turbulence model described by Crespo and Herná́ndez (1996) is used. The optional wake modelling options "secondary steering", "yaw added recovery" and "transverse velocities", provided by FLORIS and giving additional features to the $f_{\text{simulation}}^i$ function, are enabled.

A wind farm of 34 International Energy Agency (IEA) identical 15 MW (Gaertner et al., 2020) wind turbines is used. It has cut-in and cut-out speeds of $\nu_{\text{cut-in}} = 3$ m/s and $\nu_{\text{cut-out}} = 25$ m/s, respectively. Each wind turbine has a rotor diameter of
380 242.24 m, i.e., a rotor area of 46087 m². The air density is $\rho = 1.225$ kg/m³ and the tunable parameter accounting for the power losses due to misalignment is $p = 1.88$. WFFC strategies are sensible to the distances between turbines. To make the numerical simulations more robust to the distances between turbines, a diamond shape is used for the layout. With a diamond shape, there is an identical distance between each machine and its surrounding turbines. 34 machines create a sufficiently large wind farm for wake steering to be impactful, and is sufficiently small for optimization to converge quickly. A FLORIS illustration of the
385 layout used is given in Figure 6.




Figure 6. Layout in the form of a diamond shape. The farm comprises 34 identical IEA 15 MW wind turbines. There is an identical space equivalent to the diameter of four turbines between a machine and its adjacent turbines. A distance of four turbine diameters is sufficiently small to create detrimental wake effects for the farm and therefore, the optimization pertinent; and sufficient large for the design to be realistic. Here the direction is 287.4 degrees, the wind speed is 8.4 m/s and yaws are computed with the instantaneous wake steering controller.

The limits for the wind speed are $v_{\min} = 4$ m/s and $v_{\max} = 10$ m/s. The interval $[4, 10]$ m/s corresponds, approximately, to the ascending part of the power curve, where wake steering is the most beneficial for the farm. For wind speeds of between $[10, 25]$ m/s, the power output is constant; if the wind speed is reduced because of wake effects, there will be no power deficit. Because this work conducts a sensitivity analysis of yaw control, the wind speed is kept in the range of $[4, 10]$ m/s.

390 The horizon size is $H = 144$ and the length of the foresight for the prediction-based controller is $L = 10$. The initial wind values are $K_{\text{init}} = 270$ degrees and $V_{\text{init}} = 8$ m/s. The discount factor used for the heuristic \mathcal{H}_t is $\gamma = 0.99$. The precision parameter for the GS methods is $n_y = 120$, giving the grid-search method good precision.

395 More technical details regarding the simulations and numerical instabilities are given in Appendix B. The time step duration Δt is intentionally undefined, as it will be explained in Subsubsection 4.3.1. Depending on the time step duration value, different interpretations of the same results will be made. For example, if Δt corresponds to 5 minutes, then the horizon $L = 10$ means that the prediction-based controller has access to a prediction of the wind of 50 minutes. In Table 2, a summary of the experimental setting used in this work is given.

name	variable	value
number of turbines	N	34
rotor area	A	46087 m ²
low-fidelity simulator	$f_{\text{simulation}}^i$	FLORIS
air density	ρ	1.225 [kg/m ³]
cosine loss exponent yaw	p	1.88
cut-in wind speed	$\nu_{\text{cut-in}}$	3 m/s
cut-out wind speed	$\nu_{\text{cut-out}}$	25 m/s
wind speed limits	$[\nu_{\text{min}}, \nu_{\text{max}}]$	[4, 10] m/s
horizon	H	144
foresight length	L	10
initial global wind direction	K_{init}	270 degrees
initial global wind speed	V_{init}	8 m/s
discount factor	γ	0.99
precision parameter	n_y	120

Table 2. Detail of the variables and their values used across the simulations. This configuration is shared by all the numerical simulations. The foresight length is equal to 10 only for the prediction-based controller. Otherwise, it is equal to 0. The yaw rotational constraints will vary across the simulations, but $\alpha_{\text{cut-in}}$ and $\alpha_{\text{cut-out}}$ are always equal to u_{min} and u_{max} , respectively.

4.3 Results

To empirically demonstrate the importance of optimizing yaw control over a long-term time horizon, numerical simulations are performed with perfect and imperfect (noisy) wind predictions. In the graphs, for each curve, the centerline corresponds to the mean and the colored (shaded) area corresponds to the standard deviation of the results obtained through 11 Monte-Carlo trials.

For one episode, the total farm power output of a controller C given by the Algorithm 1 is denoted $P_C = \sum_{t=0}^{H-1} \sum_{i=0}^{N-1} P_t^i$. The metric to benchmark a controller C is the power gain [%] between the total farm power output of C and the total farm power output of the naive controller. The power gain is equal to $100 \cdot \frac{(P_C - P_{\text{naive}})}{P_{\text{naive}}}$.

4.3.1 Perfect predictions


The first set of simulations explores the performance of each controller over increasing variations of wind directions, using perfect predictions. Each state comprises perfect information, i.e., $\epsilon_{K,t} = 0$ and $\epsilon_{V,t} = 0$ for all $t \in \{0, 1, \dots, 153\}$. The performance of each controller presented in Section 3 is tested for increasing values of δ_t^K .

Numerical simulations are run on 21 different values of δ_t^K , with $\delta_t^K \in \{0, 1, 2, \dots, 20\}$. The wind speed is always generated with $\delta_t^V = 1$. Because this work explores the impact of wind direction on wake steering, the magnitude of the wind speed fluctuations are kept small. The wind direction and wind speed increments are stationary, $\delta_t^K = \delta_{\text{min}}^K = \delta_{\text{max}}^K$ and $\delta_t^V = \delta_{\text{min}}^V = \delta_{\text{max}}^V$ for all $t \in \{0, 1, \dots, 153\}$.

The objective here is to study the impact of the wind direction variations on yaw control. The greater the δ_t^K value, the stronger the variations. Because the nacelles have a limited rotational speed, the study of the wind direction fluctuations is crucial for yaw control. The standard deviations of the wind direction time series are related to the δ_t^K parameter in Equation (32). To better illustrate the wind direction evolution, the time series ΔK defined as

$$\Delta K = \{|f_{\text{yaw}}(K_{t+1}, K_t)|\}_{t \in \{0, 1, \dots, 153\}}, \quad (33)$$

is used. Each value of ΔK lies between $[0, 180]$ [degrees]. To study the magnitude of the variations, the absolute values are taken. Some illustrations of the ΔK time series for different values of δ_t^K are given in Figure 7.




(a) For each δ_t^K , mean values with their standard deviation of the time series ΔK .

(b) Example of time series ΔK for different values of $\delta_t^K \in \{0, 4, 8, 12, 16, 20\}$. Again, as the δ_t^K parameter increases, the magnitude of the variations of the wind direction increases.


Figure 7. Illustration of the influence of δ_t^K on the magnitude of wind direction variations ΔK . In Subfigure 7a, the mean value of wind direction variations is given as a function of δ_t^K . In Subfigure 7b, some examples are provided of the time series ΔK for different values of δ_t^K . For example, for $\delta_t^K = 5$, the mean absolute variations of the wind direction is around 6.11 degrees.

In Figure 8, the power gains of each controller compared to a naive controller are plotted. In Subfigure 8a, the yaw limits are $u_{\min} = \alpha_{\text{cut-in}} = -15$ degrees and $u_{\max} = \alpha_{\text{cut-out}} = 15$ degrees. And in Subfigure 8b, the yaw limits are $u_{\min} = \alpha_{\text{cut-in}} = -30$ degrees and $u_{\max} = \alpha_{\text{cut-out}} = 30$ degrees. These yaw constraints offer enough liberty for a wind turbine to rotate between two consecutive time steps and are small enough to limit the induced fatigue. The detailed results are given in Appendix C, in Tables C1 and C2.

As the variations of the wind direction increase, the performance of each controller diverges from each other. For small variations of the wind direction, both the instantaneous controller and the prediction-based controller give similar results. When the variations of the wind direction become large, the instantaneous controller struggles to maintain good performance. The heuristic of the prediction-based controller manages to find better yaw control strategies. The gap between the performance of the upper bound with the other controllers shows how strong wind direction variations, in relation with the rotational constraints of each machine, impact yaw control.



(a) Yaw limits are $u_{\min} = \alpha_{\text{cut-in}} = -15$ degrees and $u_{\max} = \alpha_{\text{cut-out}} = 15$ degrees. At $\delta_t^K = 6$, the prediction-based controller increases the power output of a naive approach by 6.23 %. It corresponds to absolute variations of the wind direction of 7.34 degrees as displayed in Figure 7a.



(b) Yaw limits are $u_{\min} = \alpha_{\text{cut-in}} = -30$ degrees and $u_{\max} = \alpha_{\text{cut-out}} = 30$ degrees. At $\delta_t^K = 10$, the prediction-based controller increases the power output of a naive approach by 8.93 %. It corresponds to absolute variations of the wind direction of 12.23 degrees as displayed in Figure 7a.

Figure 8. Considering future wind data in a steady-state yaw control optimization becomes mandatory when $\delta_t^k \geq 6$ for yaw constraints $[-15, 15]$ degrees and $\delta_t^k \geq 10$ for yaw constraints $[-30, 30]$. From these points, the heuristic \mathcal{H}_t provided by the prediction-based controller greatly improves the performance of a classic instantaneous steady-state optimization.

Based on the results given in Figure 8, several general statements can be drawn. As previously said, the time step duration is intentionally imprecise. The reason is that different values of Δt will lead to different interpretations. The following statement is true for any values of Δt , with respect to the hypotheses of the transition regime, described in Subsection 2.6.

- 435
- For wind turbines that can rotate from -15 to 15 degrees every Δt minutes, if the wind direction changes by more than 7.34 degrees every Δt minutes, it is important to consider future wind data in a steady-state yaw control optimization.
 - For wind turbines that can rotated from -30 to 30 degrees every Δt minutes, if the wind direction changes by more than 12.23 degrees every Δt minutes, it is important to consider future wind data in a steady-state yaw control optimization.

4.3.2 Noisy predictions

- 440 In the second set of simulations, the robustness to noisy predictions of each controller is tested. The yaw limits are $u_{\min} = \alpha_{\text{cut-in}} = -15$ degrees and $u_{\max} = \alpha_{\text{cut-out}} = 15$ degrees. The $\{K_t\}_{t \in \{0,1,\dots,153\}}$ time series are computed with $\delta_{\min}^K = 0$ and $\delta_{\max}^K = 20$. The time series $\{V_t\}_{t \in \{0,1,\dots,153\}}$ are always computed with $\delta_t^V = \delta_{\min}^V = \delta_{\max}^V = 1$. Because $\delta_{\min}^K \neq \delta_{\max}^K$, the increments are non-stationary for the wind direction. The corresponding ΔK time series is plotted in Figure 9.






Figure 9. Plot of the time series ΔK for the 11 different seeds. Wind directions are generated with $\delta_{\min}^K = 0$ and $\delta_{\max}^K = 20$. The mean is 12.53 degrees and the standard deviation is 1.12. Here, the increments vary from one time step to another because they are non-stationary.

In Figure 10a, the noise for the wind direction is increasing, i.e., $\epsilon_{K,t} \sim U(-z_K, z_K)$ with $z_K \in \{0, 1, \dots, 15\}$, for each $t \in \{0, 1, \dots, 153\}$. The noise for the wind speed is always sampled from the same distribution, i.e., $\epsilon_{V,t} \sim U(-1, 1)$. In Figure 10b, the noise for the wind speed is increasing, i.e., $\epsilon_{V,t} \sim u(-z_V, z_V)$ with $z_V \in \{0, 1, \dots, 7\}$, for each $t \in \{0, 1, \dots, 153\}$. The noise for the wind direction is always sampled from the same distribution, i.e., $\epsilon_{K,t} \sim U(-1, 1)$.

Only the noise applied to the wind directions strongly impacts the different policies. The prediction-based controller results in a poorer performance than a naive controller from a noise of 8 degrees. The wind speed noise insignificantly affects the performance of the algorithms. This corroborates the fact that yaw control mainly depends on the wind directions.



(a) The noises for the wind directions are $\epsilon_{K,t} \sim U(-z_K, z_K)$ with $z_K \in \{0, 1, \dots, 15\}$, for each $t \in \{0, 1, \dots, 153\}$. The noise for the wind speeds are always sampled from a Uniform distribution $U(-1, 1)$.



(b) The noises for the wind speeds are $\epsilon_{V,t} \sim U(-z_V, z_V)$ with $z_V \in \{0, 1, \dots, 7\}$, for each $t \in \{0, 1, \dots, 153\}$. The noise for the wind directions are always sampled from a Uniform distribution $U(-1, 1)$.

Figure 10. Only the noise applied to the wind direction strongly impacts the performance of the algorithms. This corroborates the important dependence between yaw control and the variations of the wind direction. Because the prediction-based controller uses more wind data points, it is more robust than the instantaneous controller.

5 Conclusions

As WFFC becomes more important to increase the energy production of wind farms, this work studies wake steering as a steady-state optimization problem over time. The yaw control problem is formalized as successive multiple steady-state optimization problems interconnected by the rotational constraints of the turbines and the evolution of the wind. Because the function computing the power outputs is steady-state, only the dynamics of an homogeneous, global wind and the rotational constraints of the machines are captured. Low-fidelity, steady-state simulators are used because they are not time consuming and they are suitable for optimization. But future works should perform the same studies with continuous and higher-fidelity simulators such as HAWC2Farm (Liew et al., 2023), better capturing the dynamics of the wake effects, from one time step to another. This becomes especially important when the variations of the wind direction become important.

Traditionally, yaw control is optimized in a steady-state manner. Yaw settings are computed so that they maximize the instantaneous power output of the farm. To optimize wake steering over a long-term time horizon, an MPC method is usually used. Such an approach increases the complexity of the optimization problem, making it harder to solve. To overcome such complexity, a reformulation of the steady-state optimization problem is proposed in this work to consider future wind data. The traditional objective function is augmented by a new heuristic estimating the future, expected, theoretical power outputs of the farm, weighted by how far the turbines will be from the wind if they are controlled by a naive approach. The new prediction-based controller proposed in this paper has the same number of decision variables as an instantaneous optimization.

Lastly, this work conducts a sensitivity analysis of yaw control and the variations of the wind direction. It demonstrates the importance of optimizing yaw control over future wind data when the variations of the wind directions become large. For strong wind variations, the new prediction-based controller greatly improves the performance without increasing complexity. This work shows for example that if deploying wind turbines that can rotate from -15 to 15 degrees every Δt minutes, and if the wind direction changes by more than 7.34 degrees every Δt minutes, it is important to consider future wind data in a steady-state yaw control optimization. Experiments empirically show the effectiveness of the simplifications proposed in this work in a specific experimental setting, but future research could better justify and quantify their impact on the original, prediction-based, decision problem.

This study is conducted on synthetic wind data so future works should explore the same question of dependence between the wind variations and yaw control over real wind data. Because the hypotheses regarding the transition regime may be far from reality, the proposed heuristic could be combined with low-pass filters and hysteresis mechanisms for more realistic implementations. Future works should incorporate the fatigue in the optimization process, as WFFC can have a major impact on the lifetime of each turbine. For example, the objective function of the prediction-based controller could be augmented by a heuristic taking into account the magnitude of the yaw actuations. However, the results provided by this work also suggest that with wake steering strategies more robust to wind direction variations, it would be possible to reach the same level of performance with fewer yaw actuations.

Appendix A: Gauss-Seidel method

The GS method iterates over each turbine in the direction of the wind, one by one, from upstream turbines to downstream ones.

485 The turbines' default coordinates $\{cx^i, cy^i\}_{i \in \{0,1,\dots,N-1\}}$ [m] are rotated such that the wind is coming from the west. The initial yaw settings are computed with a naive controller. By doing so, the initial solution is already a good enough solution that keeps turbines as aligned with the wind as possible. At each iteration, it solves the optimization problem by varying the yaw setting of the current turbine, considering all the others fixed. To solve each optimization problem on one variable, a grid-search approach over a discretized solution space S is used. Once solved, the setting for turbine k is fixed and optimization is

490 conducted again on turbine $k + 1$. Such approach gives good results because it exploits the sequential nature of the low-fidelity simulation.

Algorithm A1 GS method

Input: s_t input state

$cx_t^i, cy_t^i = (cx^i, cy^i) \cdot \begin{pmatrix} \cos(270-K_t^i) & -\sin(270-K_t^i) \\ \sin(270-K_t^i) & \cos(270-K_t^i) \end{pmatrix}$ ▷ rotate turbines

$R = \{i_0, i_1, \dots, i_{N-1} \mid cx_t^{i_0} \leq cx_t^{i_1} \leq \dots \leq cx_t^{i_{N-1}}\}$ ▷ order turbines

$(u_{t,0}^{i_0}, u_{t,0}^{i_1}, \dots, u_{t,0}^{i_{N-1}}) = \pi_{\text{naive}}(s_t)$ ▷ initialize yaw settings

for $i \in R$ **do**

$S = \{u_{\min} + l \cdot \frac{u_{\max} - u_{\min}}{n_y - 1}\}, \forall l \in \{0, 1, \dots, n_y - 1\}$ ▷ discretized solution space

$U = (u_{t,1}^{i_0}, u_{t,1}^{i_1}, \dots, u^i, \dots, u_{t,0}^{i_{N-2}}, u_{t,0}^{i_{N-1}})$ ▷ fix all other settings

$u_{t,1}^i \in \underset{u^i \in S}{\text{argmax}} f_{\text{obj}}(s_t, U)$ ▷ grid-search optimization

end for

Output: $(u_{t,1}^{i_0}, u_{t,1}^{i_1}, \dots, u_{t,1}^{i_{N-1}})$ yaw settings

Appendix B: Numerical instabilities

First, some modifications have been made to FLORIS in order to shut down those turbines too much misaligned with the wind during a simulation. At a time step t , for a given turbine i , all the possible yaw settings can give a similar power output. In

495 such cases, the best yaw setting is the one staying the closest to the wind direction, in order to prevent future misalignment. To incorporate such behavior in the optimization process and to make the controllers robust to numerical instabilities, the following steps are taken.

- Round out the power output computed by the simulator.
- Take the yaw setting u_t^i maximizing the objective function.

500

- Find all the yaw settings giving a performance close to the maximum.
- Among these selected settings, keep the one closest to the setting corresponding to the naive controller.

To compute the solutions of the upper-bound controller described in Subsection 3.3, a trick is used. Relaxing the rotational constraints of the turbines, i.e., making u_{\min} and u_{\max} equal to -180 and 180 degrees, respectively, increases the solution space. With the same precision parameter n_y , it reduces the precision of the grid-search method. To keep the solution space between u_{\min} and u_{\max} , and therefore, to not alter the precision of the grid-search method, the evolution of the system described in Algorithm 1 are slightly modified. At the beginning of each time step, every turbine is realigned with the wind direction. Such trick does not alter the solutions given by the upper-bound controller because good yaw settings are the one keeping turbines close to the wind direction.

Appendix C: Detailed results

δ_t^K	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20
naive	1.14	1.89	2.04	2.05	2.06	2.05	2.05	2.04	2.02	1.95	1.87	1.81	1.75	1.64	1.58	1.49	1.37	1.29	1.19	1.10	1.03
instantaneous	1.49	2.12	2.20	2.20	2.21	2.18	2.17	2.14	2.10	2.00	1.89	1.85	1.78	1.67	1.60	1.49	1.40	1.32	1.25	1.16	1.08
prediction_based	1.49	2.12	2.20	2.20	2.20	2.18	2.18	2.17	2.15	2.09	2.03	1.98	1.92	1.82	1.77	1.68	1.61	1.52	1.43	1.35	1.27
upper_bound	1.50	2.13	2.22	2.23	2.24	2.23	2.24	2.24	2.24	2.23	2.22	2.23	2.24	2.23	2.25	2.23	2.23	2.25	2.24	2.25	2.23

Table C1. Detailed results of the simulations conducted on perfect predictions in Subsubsection 4.3.1. Yaw limits are $u_{\min} = \alpha_{\text{cut-in}} = -15$ degrees and $u_{\max} = \alpha_{\text{cut-out}} = 15$ degrees. For each δ_t^K are given the total power output of the farm in 10^4 MW for each controller.

δ_t^K	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20
naive	1.14	1.89	2.04	2.05	2.06	2.05	2.06	2.06	2.07	2.05	2.04	2.04	2.06	2.04	2.05	2.01	1.99	2.00	1.97	1.94	1.87
instantaneous	1.65	2.18	2.25	2.25	2.26	2.25	2.25	2.25	2.25	2.24	2.22	2.21	2.21	2.19	2.20	2.15	2.11	2.10	2.07	2.04	1.96
prediction_based	1.65	2.18	2.25	2.25	2.26	2.25	2.25	2.25	2.25	2.24	2.22	2.22	2.22	2.20	2.22	2.18	2.17	2.16	2.13	2.12	2.07
upper_bound	1.65	2.19	2.26	2.27	2.27	2.26	2.27	2.27	2.28	2.27	2.26	2.26	2.27	2.27	2.28	2.27	2.27	2.28	2.28	2.28	2.27

Table C2. Detailed results of the simulations conducted on perfect predictions in Subsubsection 4.3.1. Yaw limits are $u_{\min} = \alpha_{\text{cut-in}} = -30$ degrees and $u_{\max} = \alpha_{\text{cut-out}} = 30$ degrees. For each δ_t^K are given the total power output of the farm in 10^4 MW for each controller.

510 *Author contributions.* EK contributed to the original idea of the heuristic, implemented the codes, conducted all the simulations and wrote most of the paper. PB, FC, PC and DE contributed to the writing and review of the paper.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. The authors are grateful for the support of the R&D Wind program of TotalEnergies OneTech, especially Cédric ENEAU for his support.

515 **Acronyms**

FLORIS FLOW Redirection and Induction in Steady State. 5, 17, 19, 24

GS Gauss-Seidel. 10, 18, 24

IEA International Energy Agency. 17, 18

LES large-eddy simulations. 2

520 **LUT** lookup table. 3, 4

ML machine learning. 4

MPC model predictive control. 4, 11, 23

MW megawatts. 6, 7, 14, 17, 18, 25

RL reinforcement learning. 4

525 **WFFC** wind farm flow control. 2–4, 9, 17, 23

References

- Bastankhah, M. and Porté-Agel, F.: Experimental and theoretical study of wind turbine wakes in yawed conditions, *Journal of Fluid Mechanics*, 806, 506–541, <https://doi.org/10.1017/jfm.2016.595>, 2016.
- Boersma, S., Doekemeijer, B. M., Gebraad, P. M. O., Fleming, P. A., Annoni, J., Scholbroek, A. K., Frederik, J. A., and van Wingerden, J.: A tutorial on control-oriented modeling and control of wind farms, in: 2017 American Control Conference (ACC), pp. 1–18, <https://doi.org/10.23919/ACC.2017.7962923>, 2017.
- Chu, S. and Majumdar, A.: Opportunities and challenges for a sustainable energy future, *Nature*, 488, 294–303, <https://doi.org/10.1038/nature11475>, 2012.
- Crespo, A. and HERNÁNDEZ, J.: Turbulence characteristics in wind-turbine wakes, *Journal of Wind Engineering and Industrial Aerodynamics*, 61, 71–85, [https://doi.org/10.1016/0167-6105\(95\)00033-X](https://doi.org/10.1016/0167-6105(95)00033-X), 1996.
- Dallas, S., Stock, A., and Hart, E.: Control-Oriented Modelling of Wind Direction Variability, *Wind Energy Science Discussions*, 2023, 1–42, <https://doi.org/10.5194/wes-2023-93>, 2023.
- Fleming, P., Annoni, J., Shah, J. J., Wang, L., Ananthan, S., Zhang, Z., Hutchings, K., Wang, P., Chen, W., and Chen, L.: Field test of wake steering at an offshore wind farm, *Wind Energy Science*, 2, 229–239, <https://doi.org/10.5194/wes-2-229-2017>, 2017.
- Fleming, P. A., Stanley, A. P. J., Bay, C. J., King, J., Simley, E., Doekemeijer, B. M., and Mudafort, R.: Serial-Refine Method for Fast Wake-Steering Yaw Optimization, *Journal of Physics: Conference Series*, 2265, 032 109, <https://doi.org/10.1088/1742-6596/2265/3/032109>, 2022.
- Gaertner, E., Rinker, J., Sethuraman, L., Zahle, F., Anderson, B., Barter, G. E., Abbas, N. J., Meng, F., Bortolotti, P., Skrzypinski, W., Scott, G. N., Feil, R., Bredmose, H., Dykes, K., Shields, M., Allen, C., and Viselli, A.: IEA Wind TCP Task 37: Definition of the IEA 15-Megawatt Offshore Reference Wind Turbine, <https://doi.org/10.2172/1603478>, 2020.
- Hure, N., Turnar, R., Vašák, M., and Benčić, G.: Optimal wind turbine yaw control supported with very short-term wind predictions, in: 2015 IEEE International Conference on Industrial Technology (ICIT), pp. 385–391, <https://doi.org/10.1109/ICIT.2015.7125129>, 2015.
- Janssens, N. and Meyers, J.: Towards real-time optimal control of wind farms using large-eddy simulations, *Wind Energy Science*, 9, 65–95, <https://doi.org/10.5194/wes-9-65-2024>, 2024.
- Kadoche, E., Gourvénéec, S., Pallud, M., and Levent, T.: MARLYC: Multi-Agent Reinforcement Learning Yaw Control, *Renewable Energy*, 217, 119 129, <https://doi.org/10.1016/j.renene.2023.119129>, 2023.
- Kanev, S.: Dynamic wake steering and its impact on wind farm power production and yaw actuator duty, *Renewable Energy*, 146, 9–15, <https://doi.org/10.1016/j.renene.2019.06.122>, 2020a.
- Kanev, S.: Dynamic wake steering and its impact on wind farm power production and yaw actuator duty, *Renewable Energy*, 146, 9–15, <https://doi.org/10.1016/j.renene.2019.06.122>, 2020b.
- King, J., Fleming, P., King, R., Martínez-Tossas, L. A., Bay, C. J., Mudafort, R., and Simley, E.: Control-oriented model for secondary effects of wake steering, *Wind Energy Science*, 6, 701–714, <https://doi.org/10.5194/wes-6-701-2021>, 2021.
- Liew, J., Göçmen, T., Lio, A. W. H., and Larsen, G. C.: Extending the Dynamic Wake Meandering Model in HAWC2Farm: Lillgrund Wind Farm Case Study and Validation, *Wind Energy Science Discussions*, 2023, 1–25, <https://doi.org/10.5194/wes-2023-14>, 2023.
- Meyers, J., Bottasso, C., Dykes, K., Fleming, P., Gebraad, P., Giebel, G., Göçmen, T., and van Wingerden, J.-W.: Wind farm flow control: prospects and challenges, *Wind Energy Science*, 7, 2271–2306, <https://doi.org/10.5194/wes-7-2271-2022>, 2022.

- Niayifar, A. and Porté-Agel, F.: Analytical Modeling of Wind Farms: A New Approach for Power Prediction, *Energies*, 9, <https://doi.org/10.3390/en9090741>, 2016.
- NREL: FLORIS. Version 3.4.1, <https://github.com/NREL/floris>, 2021.
- 565 Saenz-Aguirre, A., Zulueta, E., Fernandez-Gamiz, U., Lozano, J., and Lopez-Guede, J. M.: Artificial Neural Network Based Reinforcement Learning for Wind Turbine Yaw Control, *Energies*, 12, <https://doi.org/10.3390/en12030436>, 2019.
- Saenz-Aguirre, A., Zulueta, E., Fernandez-Gamiz, U., Ulazia, A., and Teso-Fz-Betono, D.: Performance enhancement of the artificial neural network–based reinforcement learning for wind turbine yaw control, *Wind Energy*, 23, 676–690, <https://doi.org/https://doi.org/10.1002/we.2451>, 2020.
- 570 Siemens Gamesa Renewable Energy: Siemens Gamesa now able to actively dictate wind flow at offshore wind locations, <https://www.siemensgamesa.com/en-int/newsroom/2019/11/191126-siemens-gamesa-wake-adapt-en>, last accessed on 15 May 2023, 2019.
- Simley, E., Fleming, P., King, J., and Sinner, M.: Wake Steering Wind Farm Control With Preview Wind Direction Information: Preprint, <https://www.osti.gov/biblio/1778702>, 2021.
- Song, D., Yang, J., Fan, X., Liu, Y., Liu, A., Chen, G., and Joo, Y. H.: Maximum power extraction for wind turbines through a novel yaw control solution using predicted wind directions, *Energy Conversion and Management*, 157, 587–599, <https://doi.org/https://doi.org/10.1016/j.enconman.2017.12.019>, 2018.
- 575 van Doorn, E., Dhruva, B., Sreenivasan, K. R., and Cassella, V.: Statistics of wind direction and its increments, *Physics of Fluids*, 12, 1529–1534, <https://doi.org/10.1063/1.870401>, 2000.
- von Brandis, A., Centurelli, G., Schmidt, J., Vollmer, L., Djath, B., and Dörenkämper, M.: An investigation of spatial wind direction variability and its consideration in engineering models, *Wind Energy Science*, 8, 589–606, <https://doi.org/10.5194/wes-8-589-2023>, 2023.
- 580 Yang, J., Fang, L., Song, D., Su, M., Yang, X., Huang, L., and Joo, Y. H.: Review of control strategy of large horizontal-axis wind turbines yaw system, *Wind Energy*, 24, 97–115, <https://doi.org/https://doi.org/10.1002/we.2564>, 2021.