# Risk-averse wake steering optimization for energy and power maximization under uncertain wind direction changes

Journal of Physics: Conference Series **3224** (2026) 032124 doi:10.1088/1742-6596/3224/3/032124

# Risk-averse wake steering optimization for energy and power maximization under uncertain wind direction changes

**Marcus Becker**[1] **and Jan-Willem van Wingerden**[1]

[1]Delft Center for Systems and Control, Delft University of Technology, Delft, Netherlands

E-mail: marcus.becker@tudelft.nl

**Abstract.** Wind farm flow control strategies aim to manipulate the flow between the turbines to achieve a farm-wide goal. Wake steering is one such strategy, typically employed to increase the yield of a wind farm by redirecting upstream turbines' wakes away from downstream ones. This control approach is sensitive to the wind direction, and frequently varying wind directions create the need for robust yaw steering control setpoints. In the past, this has been achieved in the steady-state domain, as dynamic simulations are typically deemed too expensive to perform a large grid-search for optimal setpoints. This paper utilizes a computationally cheap dynamic wake model and explores how robust control setpoints can be derived in the time domain. To this end, the paper presents a methodology for generating synthetic wind direction changes with a prescribed variation in wind direction. These wind direction time series are then used to create a database of a two-turbine wind farm, which allows the exploration of different cost functions in time. The database provides both the expected value and the uncertainty of both power and energy. The obtained data is then used to explore four cost functions to derive robust setpoints. Comparing energy and power performance, we define useful quantities of interest to connect the two and to highlight necessary assumptions made when using steady-state setpoints. The paper concludes by applying the resulting look-up table controllers in a ten-turbine wind farm. The performance shows that maximizing for the expected power is the best approach to increase the farm efficiency. The results also show that an alternative cost function, which avoids losses, does lead to similar but smaller gains at a much lower yaw angle investment.

## 1. Introduction

Wind turbines in farms are subject to wake losses, induced by upstream turbines extracting energy from the flow, thereby creating a wake that reduces the power generated by downstream turbines. While wake losses can typically not be entirely avoided, they can be partially mitigated by wind farm flow control (WFFC) strategies [1]. Wake steering is one WFFC strategy, where the rotor of the upstream wind turbine is intentionally misaligned by the angle $\gamma$ to redirect its wake away from the downstream turbine. This results in a slight loss at the upstream turbine and a larger gain at the downstream turbine. The resulting optimization problem is typically solved in steady-state, where, for example, wind direction changes are neglected. To relax the assumption, previous literature included uncertainty in the wind direction by considering its variance $\sigma_\varphi$ [2, 3]. The optimal $\gamma^*$ is then determined by its performance in multiple Gaussian

weighted steady-state simulations in discrete wind directions around a given mean wind direction $\bar{\varphi}$.

This approach is necessary as only steady-state simulations have been deemed fast enough and efficient enough to optimize yaw angles under arbitrary conditions. In recent years, however, models have been developed to simulate wake dynamics at a low computational cost, e.g., [4–6], which begs the question of how the outlined derivation of robust yaw angles can be translated to the time domain.

To this end, three questions are discussed in this paper: (i) How can realistic synthetic wind direction changes be obtained? (ii) What are suitable robust cost functions? And (iii), under which conditions is it valid to choose optimal steady-state setpoints rather than dynamically simulating the wind farm? Related to these questions, the contributions of this paper are as follows:

(i) We provide the dynamic pendant to the steady-state optimization of robust yaw angles under uncertain wind direction changes. To this end, we utilize a dynamic engineering model to generate a dataset that enables us to incorporate the costs associated with transitioning from a starting angle to a new one, and to calculate the resulting expected power and energy, along with their associated uncertainty.

(ii) We investigate four cost functions, all of which aim to return robust control setpoints using different strategies.

(iii) We introduce quantities of interest like the amortization time and the uniform convergence time to assess under which conditions it is valid to use steady-state set points and when it is necessary to optimize over a time horizon.

The remainder of the paper is structured as follows: Section 2 introduces the dataset created for this study, as well as the cost functions applied to it. Section 3 discusses the resulting optimization problem and the relation between power and energy setpoints. Successively, Section 4 applies the resulting controllers to a ten-turbine wind farm, and Section 5 draws the overall conclusions from the presented paper.

## 2. Methodology

The methodology for the simulations consists of two parts: Sec. 2.1 describes the generation of synthetic but realistic wind direction changes with a prescribed variation, as well as the applied wind farm simulation, and Sec. 2.2 introduces the formulation of the cost functions.

### *2.1. Wind direction signal and farm simulation*

To generate wind direction signals, we utilize vertical-axis LiDAR data recorded offshore [7] and filter after cleaning for at least 1 h long time series with turbulence intensity levels below 10 %, and a mean wind speed within the rated region of the IEA 10 MW turbine [8]. The resulting 66 time series were used to calculate the normalized power spectral density (NPSD) of the wind direction signal. Following [3], we fit a low-pass function to the data to isolate the low-frequency wind direction signal NPSD below 1/270 Hz from the high-frequency components caused by turbulence. Subsequently, a random wind direction time series is created by summing frequencies with the amplitudes given by the low-frequency NPSD and random phases. The signals are normalized by their standard deviation and multiplied by the prescribed one. For each $\sigma_\varphi$ value 15 random wind direction signals are generated and mirrored to avoid bias due to under-sampling, leading to 30 wind direction offset signals, each 30 min long. The free wind speed is set to 8 $\mathrm{ms}^{-1}$.

To simulate a two-turbine wind farm, we employ FLORIDyn implemented in the OFF toolbox in tandem with the Gauss Curl Hybrid Model [5, 9, 10]. The turbines are spaced 5 diameters (D) apart and are subject to 17 initial wind directions; from a view angle $\psi$ of $-20$ deg to a view angle

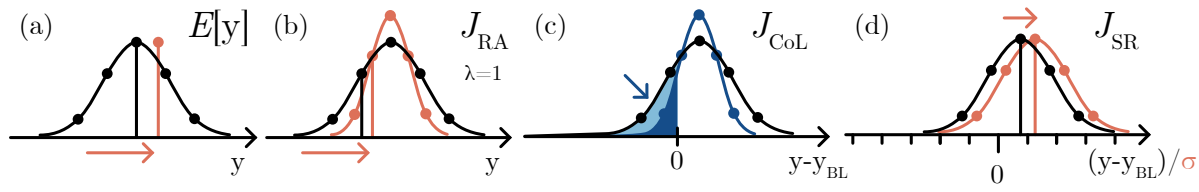


**Figure 1.** Cost function mechanism comparison. The black curve depicts a normal distribution with dots at the expected value, as well as $\pm 1\sigma$ and $\pm 2\sigma$. The second, colored curve is the optimized version. (a) The expected value cost function does maximize the mean of the distribution, with disregard of the resulting spread, (b) The risk averse (RA) cost function maximizes the value of a point on the distribution below the expected value, (c) the Chance-of-Loss (CoL) minimizes the area of the distribution associated with loss, and (d) the Sharpe-Ratio (SR) maximizes the expected difference on a scale normalized by the standard deviation - a smaller $\sigma$ therefore leads to a higher expected value in the normalized space.

of +20 deg in 2.5 deg steps. For each mean wind direction 25 starting and final yaw misalignment angles of T1 are tested, ranging from −30 to 30 deg in 2.5 deg resolution. The wind direction variation is set to $\sigma_\varphi \in [3, 5, 7]$ deg. The resulting database consists of $\approx 9.6 \cdot 10^5$ simulations.

The angles involved can be observed in two frames of reference: (i) in the world coordinate system and (ii) in the relative coordinate system to the direction in which the downstream turbine is located. The world coordinate system consists of the wind direction $\varphi$, the turbine orientation $\gamma_o$, and the turbine misalignment $\gamma$. The relative coordinate frame looks at the view angle $\psi$, which is 0 deg for full alignment between the turbines, and the turbine view angle orientation $\psi_o$. To translate between them, the following set of equations is used:

$$\gamma_o = \varphi - \gamma\,, \quad \psi = \angle T - \varphi\,, \quad \psi_o = \angle T - \gamma_o = \psi + \gamma\,, \tag{1}$$

where $\angle T$ is the angle of the vector pointing from the upstream turbine to the downstream one. We further utilize the initial misalignment $\gamma_0$ of the upstream turbine.

### *2.2. Optimization formulation*

Traditional wind farm performance cost functions focus on the expected value of the steady-state wind farm power $p$, or of the generated energy $e$ over a prediction horizon $\tau_{\text{ph}}$. A fundamental difference between maximizing power and energy is that the latter takes transition costs from a starting yaw angle to a final yaw angle into account. This paper explores three additional cost functions that incorporate uncertainty beyond the expected value. The presented cost functions are primarily used in economics and finance, where investment decisions are weighed against one another. Either put money into low-but-guaranteed-gain options or invest in volatile options with the possibility of high returns. The link to wake steering is evident: here, the question is whether to invest in a high or low yaw misalignment of the upstream turbine for an uncertain return at the downstream turbine.

The presented formulations hold for both power and energy, which we substitute with $y$ and $\sigma_y$. The covariance between $y$ and a baseline (BL) $y_{\text{BL}}$ is denoted as $\sigma^2_{yy}$. The first function is the risk-averse (RA) cost function [11], which maximizes the expected value minus $\lambda\sigma_y$:

$$\mathcal{J}_{\text{RA}} : \gamma^*_{\text{RA}} = \text{argmax}_\gamma \; E\left[y(\gamma)\right] - \lambda\sigma_y(\gamma)\,. \quad \lambda \in \mathbb{R}^+ \tag{2}$$

Fig. 1 (b) depicts a schematic of the mechanism, where both a smaller $\sigma_y$ and a larger $E[y]$ lead to a better value. Increasing $\lambda$ accentuates the importance of the variance. In this work,

we select a common value of $\lambda = 1$ [11]. Alternatively, the chance of loss (CoL), see Fig. 1 (c), minimizes the integral of the distribution function associated with a loss with respect to a BL:

$$\mathcal{J}_{\mathrm{CoL}} : \gamma^*_{\mathrm{CoL}} = \mathrm{argmin}_\gamma \; \Phi\left(0,\, E\left[y(\gamma)\right] - E\left[y(\gamma_{\mathrm{BL}})\right],\; \sqrt{\sigma_y^2(\gamma) + \sigma_y^2(\gamma_{\mathrm{BL}}) - 2 \cdot \sigma_{yy}^2(\gamma, \gamma_{\mathrm{BL}})}\right). \quad (3)$$

An issue can be that the cost function becomes less sensitive to changes once the distribution is sufficiently positive. Lastly, the Sharpe Ratio [12] normalizes the difference with the uncertainty:

$$\mathcal{J}_{\mathrm{SR}} : \gamma^*_{\mathrm{SR}} = \mathrm{argmax}_\gamma \left[E\left[y(\gamma)\right] - E\left[y(\gamma_{\mathrm{BL}})\right]\right] \,/\, \left[\sqrt{\sigma_y^2(\gamma) + \sigma_y^2(\gamma_{\mathrm{BL}}) - 2 \cdot \sigma_{yy}^2(\gamma, \gamma_{\mathrm{BL}})}\right]. \quad (4)$$

One way to think about the mechanism of this cost function is that it improves by increasing the expected value and by reducing the axis spacing of the expected difference, as shown in Fig. 1 (d). This cost function can become highly sensitive when uncertainty is minimal.

## 3. Results

The shape of the cost functions introduced in Section 2.2, based on the power generated, is discussed in Section 3.1. Section 3.2 then discusses the optimal setpoints for energy and their relation to the power setpoints.

### *3.1. Power*

We present a selection of the results in Fig. 2, analyzed with respect to the wind farm power and its uncertainty. The landscape of $E[p]$ is composed of (i) the maxima on either end of the view angle range, where both turbines operate in their optimum, (ii) the yawing losses once T1 diverts from $\gamma = 0$, and (iii) the wake losses for small view angles. Notable is the diagonal loss induced by the wake of the upstream turbine, which is steered into the downstream turbine. Right and left of the certain-loss area, $\sigma[p]$ increases as small wind direction changes cause the wake to be temporarily displaced into the downstream turbine, which causes fluctuations. The covariance in the third column shows that some of these fluctuations are unavoidable and affect both the controlled and baseline cases with $\gamma = 0$ deg. Based on these fluctuations, the risk-averse cost function proposes different optimal setpoints: Since it punishes $\sigma[p]$, it leads to more extreme absolute $\gamma$ values, ensuring that the wake is consistently steered away from the second turbine. Fig. 2 (a,b) shows how $\sigma[p(\gamma)]$ develops in the neighborhood of the optimal $E[p(\gamma)]$ values: As $\gamma$ is pushed to higher absolute values, $\sigma[p(\gamma)]$ decreases. The wake is more consistently steered away, leading to fewer fluctuations, but at the cost of lower $E[p(\gamma)]$ values. On the other side, $\mathcal{J}_{\mathrm{CoL}}$ and $\mathcal{J}_{\mathrm{SR}}$ base their cost on the difference to the baseline. Both CoL and SR prefer a high expected difference and a low related standard deviation. The standard deviation of the difference is 0 for $\gamma = 0$ deg, and very low for small $|\gamma|$ values, making this area of the cost function particularly attractive. This only changes for view angles close to 0 deg, where the expected difference is large enough to justify a yaw misalignment despite higher related standard deviation.

### *3.2. Energy*

As the integral of power over time, the energy-related optimal control setpoints are time-dependent. To this end, two additional quantities of interest are introduced: the amortization time $t_a$ (Sec. 3.2.1) and the uniform convergence time $t_{uc}$ (Sec. 3.2.2). Subsequently, Sec. 3.2.3 comments on the resulting optimal control setpoints. Time units are given in convection time scales $t_c$, based on the turbine distance and the free stream velocity: $t_c = 5D/u_\infty = 123.75$ s.

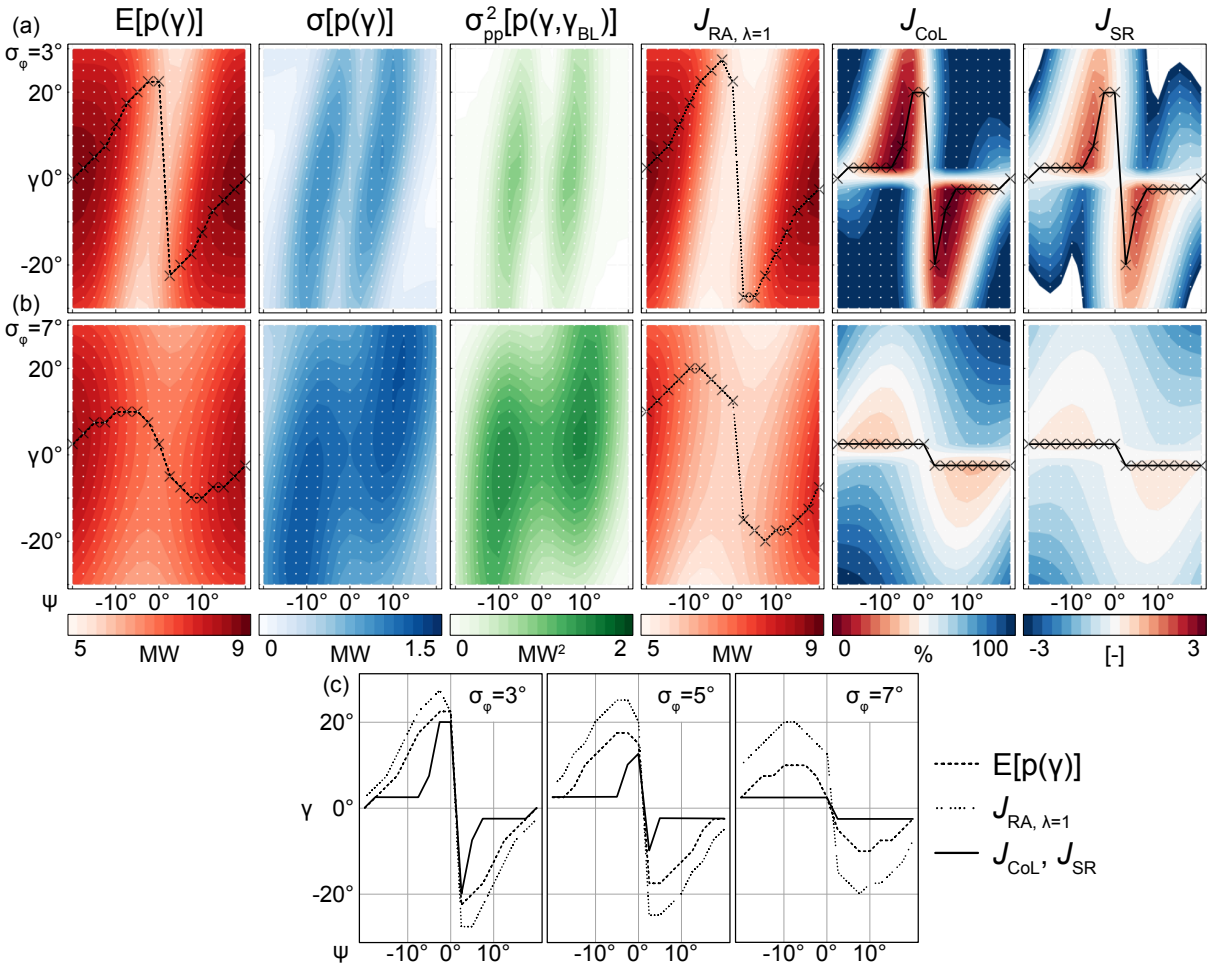


**Figure 2.** (a,b) The columns depict the expected power generated, along with its standard deviation and covariance to the baseline ($\gamma = 0$ deg), followed by the cost function landscape related to Eq. (2 - 4) for $\sigma_\varphi \in [3, 7]$. The comparison of the optimal setpoints for $E\left[p(\gamma)\right]$ and the cost-functions is given in (c) for $\sigma_\varphi \in [3, 5, 7]$. Note that in general the optimal solutions of $\mathcal{J}_{\mathrm{CoL}}$ do not have to be equal to the solutions of $\mathcal{J}_{\mathrm{SR}}$.

*3.2.1. Amortization time* The amortization time $t_a$ is defined as the time it takes to recoup the energy loss due to actuating the upstream turbine. Mathematically, this value is the infimum of the time at which the integral over the difference between the power generated by the controlled wind farm and the baseline wind farm becomes positive:

$$t_a(\gamma, \gamma_0) = \inf \left\{ \int_{t_0}^{t_a} \sum_{i \in [1, n_T]} p_{i,\mathrm{Con}}(t, \gamma, \gamma_0) - \sum_{i \in [1, n_T]} p_{i,\mathrm{Bl}}(t)\, \mathrm{d}t > 0 \right\}, \tag{5}$$

where $p_{i,\mathrm{Con}}$ and $p_{i,\mathrm{Bl}}$ denote the power generated by turbine $i$ in the controlled and baseline case, respectively. The metric overlooks temporary gains that may result from reducing the upstream turbine misalignment, but it focuses solely on the time scales of the long-term gain. Figure 3 (a) depicts an example of the amortization time. Note that for any gain to manifest, $t_a$ has to be greater than 1. For more extreme misalignment angles, we see $t_a$ approach values around 2.5 $t_c$, meaning that $\approx 5$ minutes pass without outperforming the baseline. The amortization time is

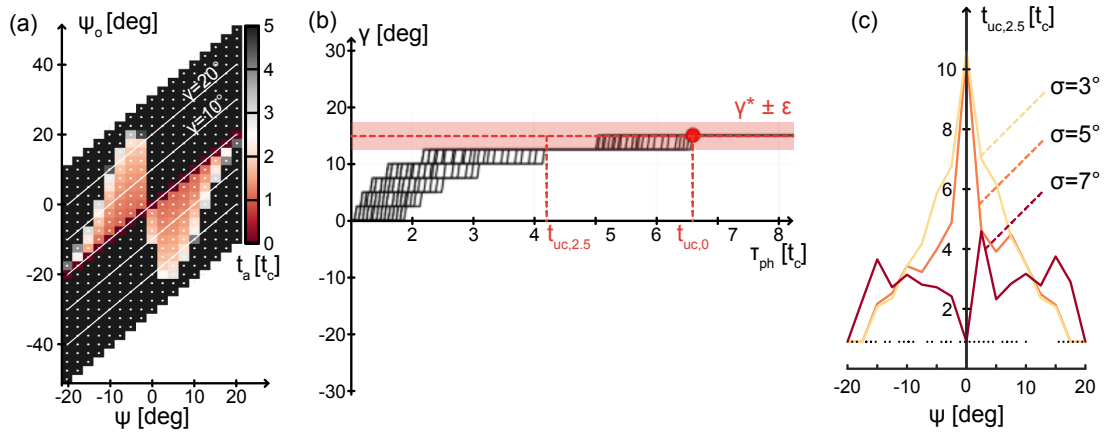


**Figure 3.** (a) Amortization time for $\gamma_0 = 0\,\text{deg}$, $\sigma_\varphi = 5\,\text{deg}$ relative to maintaining $\gamma_{\text{Bl}} = 0\,\text{deg}$, (b) one example of convergence of $\gamma^*(\tau_{\text{ph}})$ for the maximum expected energy given different initial $\gamma_0$ values. This example is given for $\psi = -7.5\,\text{deg}$, and $\sigma_\varphi = 5\,\text{deg}$, (c) $t_{uc,\,2.5}$ over the view angles and for different wind direction variations. All numbers are derived for the cost function, which maximizes the expected energy.

sensitive to the starting state - if the turbine is already at a favorable orientation, $t_a$ will shrink to a value closer to, but not below 1, as realigning the turbine with the flow always yields a short-term improvement. In this work, we further exclude effects related to different maximum yawing rates, the energy consumption of the yawing actuators, and possible imposed limits on the number of yaw-drive activations, all of which may also impact $t_a$.

*3.2.2. Uniform convergence time* The dataset contains different starting orientations, allowing us to study how long it takes to arrive at a starting-state-independent optimum. This is solely the case for energy-based cost functions, as maximizing power does not account for time. The uniform convergence time $t_{uc,\epsilon}$ is defined as follows: first, we define the convergence time $t_c(\gamma_0, \epsilon)$, as the time it takes until the difference between the global optimum and the time-dependent optimum for a given starting orientation is less or equal than an $\epsilon \geq 0\,\forall t \geq t_c$. Successively $t_{uc,\epsilon}$ is achieved by maximizing $t_c(\gamma_0, \epsilon)$ across all possible starting states:

$$|\gamma^* - \gamma^*(t, \gamma_0)| \leq \epsilon \quad \forall t \geq t_c(\gamma_0, \epsilon)\,, \tag{6}$$

$$t_{uc} = \max_{\gamma_0} t_c(\gamma_0, \epsilon)\,. \tag{7}$$

This formulation assumes that a global optimum $\gamma^*$ exists, which depends on the cost function employed to determine $\gamma^*$. Figure 3 (b) depicts one example of the convergence of $\gamma^*(\tau_{\text{ph}})$ for different initial turbine orientations and the resulting $t_{uc}$ for an $\varepsilon = 2.5\,\text{deg}$. In this case, the optimal $\gamma$ converges towards $\gamma^* = 15$ deg. Based on $\gamma_0$, this convergence can take longer or shorter, but after $t_{uc,\,2.5} \approx 4.1\,t_c$ all optimizations are within $\pm 2.5\,\text{deg}$ of $\gamma^*$. The resulting $t_{uc,\,2.5}$ values for the expected energy across all investigated $\psi$ and $\sigma_\varphi$ are shown in Fig. 3 (c).

*3.2.3. Optimal energy setpoints* Fig. 4 depicts the cost function evolution in time for increasingly long prediction horizons for $\sigma_\varphi = 5\,\text{deg}, \gamma_0 = 0\,\text{deg}$. As expected, all cost-functions prefer $\gamma = 0$ at $\tau_{\text{ph}} = 1\,t_c$, after which the positive impacts of the $\gamma \neq 0$ angles start to manifest. Similarly to the results presented in Fig. 2, the expected energy cost function results in moderate yaw misalignment, while the Risk-averse cost function leads to more aggressive misalignment

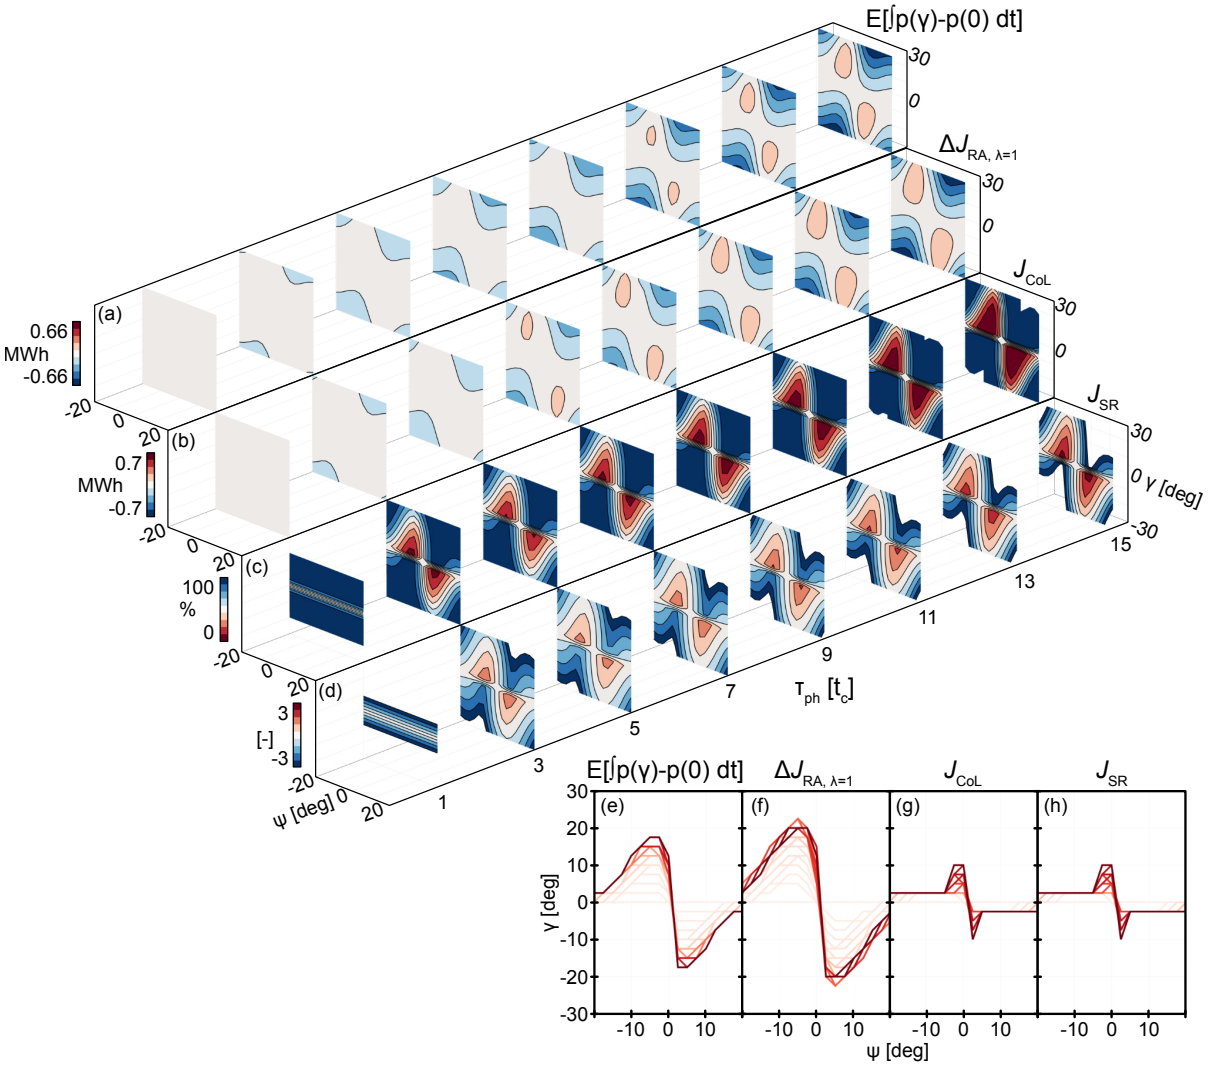


**Figure 4.** Energy related cost functions over time for $\sigma_\varphi = 5\,\text{deg}, \gamma_0 = 0\,\text{deg}$: (a) Expected difference to baseline ($\gamma = 0$), (b) Risk averse cost, subtracted by $\gamma = 0$ values, (c) the chance of loss, and (d) the Sharpe ratio. The respective optimal $\gamma$ angles are given in (e-h), where darker lines indicate larger $\tau_{\text{ph}}$ values.

angles. Both $\mathcal{J}_{\text{CoL}}$ and $\mathcal{J}_{\text{SR}}$ prefer small misalignment angles over large ones, see Fig. 4 (e-h). A consistent theme is the convergence of the optimal setpoints to larger values for a longer $\tau_{\text{ph}}$. In this context Fig. 3 (c) depicts $t_{uc,\,2.5}$, after which the optimal settings are within $\pm 2.5\,\text{deg}$ of the $\tau_{\text{ph}} = \infty$ solution. As a result, it is viable to simplify the optimization and to choose the optimal power setpoints instead, as long as it is justified to choose a $\tau_{\text{ph}} \geq t_{uc,\,2.5}$. However, how to pick $\tau_{\text{ph}}$ is not obvious and remains an unanswered question.

## 4. Wind farm control

Section 4.1 describes the control laws derived from the cost functions. The wind farm simulation is described in Section 4.2, followed by a discussion of the performance in Section 4.3.

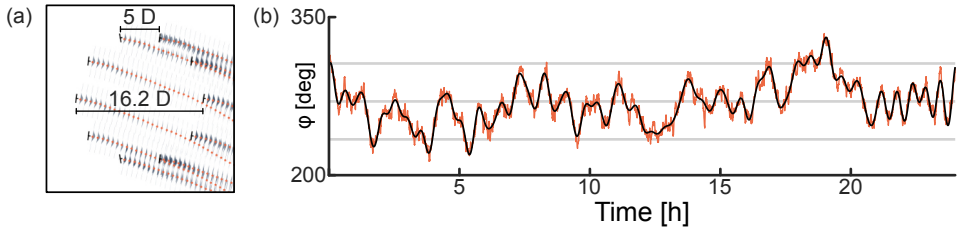


**Figure 5.** (a) Circular 10 turbine wind farm layout [15] simulated in OFF, (b) 24 h wind direction signal with low and high frequency components. The high-frequency components are scaled with $\sigma_\varphi$ (here 5 deg), the low-frequency ones are scaled to utilize the periodicity of the wind farm layout. Grey lines indicate wind directions for which two pairs of neighboring turbines are fully aligned (every 38 deg). The wind speed is fixed to $u_\infty = 8$ ms$^{-1}$ with 6% TI.

### *4.1. Controllers*

In order to apply the cost functions introduced in Sec. 2.2, we limit ourselves to the optimal setpoints derived for power, which are the same or close to the optimal points for energy given $\tau_{\text{ph}} \geq t_{uc,\,2.5}$. This removes the dependency on $\tau_{\text{ph}}$. We formulate Look-up Table (LuT) control laws $\psi_o^* = \mathcal{C}(\sigma_\varphi, \psi, \psi_o)$, which accept the wind direction variation $\sigma_\varphi$, the current wind direction in the form of the view angle $\psi$, and the turbine orientation $\psi_o$ in the view angle frame of reference. The control laws return the optimal orientation of the upstream turbine $\psi_o^*$. Note that the given controller operates under the assumption of a turbine spacing of 5 D and a free wind speed of 8 ms$^{-1}$. Given that yaw steering control laws have a predictable geometric pattern [13, 14], extensions of laws are imaginable but outside of the scope of this paper. To continuously apply a control law based on the presented data set, linear interpolation is used between the $(\sigma_\varphi, \psi, \psi_o, \psi_o^*)$ data points[1]. A yaw rate limit of 0.3 deg s$^{-1}$ is employed.

### *4.2. Application in a wind farm*

We adapt the circular ten-turbine wind farm used by [15] with 5 D spacing between the neighboring turbines, depicted in Fig. 5 (a). The advantage of the layout is that it is periodic with respect to the wind direction in $\Delta\varphi = 36$ deg bins, allowing a focus on the performance of controllers designed for 5 D, while also containing out-of-design spacings, such as 9.5, 13.1, 15.4, and 16.2 D. Contrary to, e.g., [16], we apply the turbine orientation control continuously to avoid an additional degree of freedom in the shape of a dead-band. However, the turbines only follow the wind direction signal below 1/1800 Hz; everything above is assumed to be captured by $\sigma_\varphi$. The wind farm layout, along with the low-frequency wind direction signal, is used to precompute each turbine's $\psi$ relative to its two direct neighbors. If $-20\,\text{deg} \leq \psi \leq 20\,\text{deg}$, the controller is applied.

### *4.3. Simulation results*

Figure 6 (a-d) depicts the view-angle binned cumulative density function of the controller performance, going from underperforming to overperforming compared to the power generated by the baseline. The plot displays the median performance for each view angle, along with the spread. The median of each bin is plotted in Figure 6 (e-h), respectively, with ±25% bounds of the data. A key takeaway is that all controllers can generally enhance the performance of the wind farm for most cases, particularly for small view angles. The expected power cost function translates large yaw angles into high gains around $\psi = 0$ deg, but also suffers occasional losses.

[1] The derived LuTs are available under `doi.org/10.4121/50497d89-3745-4b29-a048-9bf3a61fbbea` . Further data is available upon request.

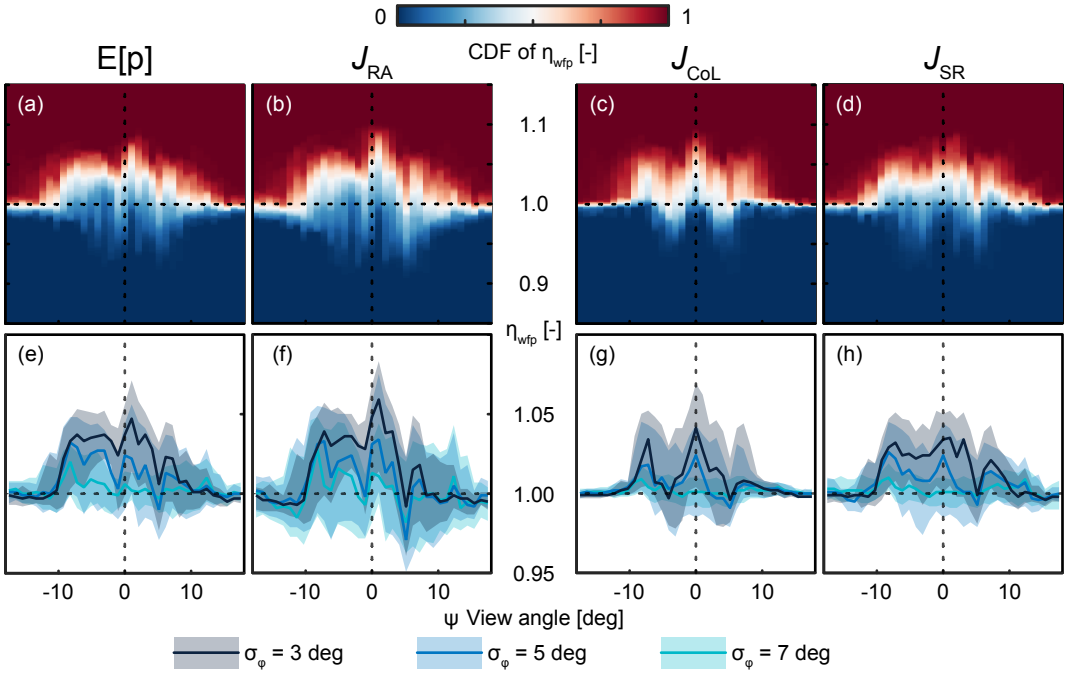


**Figure 6.** (a-d) Cumulative density function of the wind farm power efficiency per 1 deg view angle bin. The white area in each bin indicates that 50% of the performance data in this bin is equal to this performance or worse. All data is shown for the $\sigma_\varphi = 3\,\text{deg}$ case and normalized by the baseline with $\gamma = 0$. (e-f) Median performance for all $\sigma_\varphi \in [3, 5, 7]\,\text{deg}$, as well as bands to indicate $\pm 25\%$ of the data.

Similarly, the RA cost function generates even higher gains but also deeper losses, increasing the overall spread. The CoL controller performance is an outlier compared to the rest; the effect of yawing is limited to $|\psi| \leq 10\,\text{deg}$, with an overall smaller spread. This leads to small gains and small losses. Lastly, the SR controller delivers a gain similar to the expected power cost function, but with a lower overall magnitude. The difference to the CoL controller may be unexpected, given the equal setpoints presented in Figure 2 (c). The difference lies in the sensitivity to the starting misalignment angle $\gamma_0$, which leads to different behavior of the control strategies. The interested reader is referred to the LuT dataset linked in Sec. 4.1. A notable similarity between most controllers is the consistently low efficiency for $|\psi| > 13\,\text{deg}$. This is due to the steered wake now decreasing the performance of the further downstream turbine at a distance of 9.5 D. The interaction falls outside of the cost function design and is therefore not taken into account during the optimization of the LuTs.

Beyond power efficiency, Tables 1 and 2 investigate the energy efficiency over the simulated 24 hours and the median absolute misalignment angles as a measure of investment. To calculate the latter, the low frequent wind direction signal was used and angles with $|\gamma| \leq 0.1\,\text{deg}$ were excluded. Looking at the efficiency of the energy generated in Table 1, $E[p]$ consistently outperforms the other controllers for all tested $\sigma_\varphi$. The least yaw steering investment is shown by the $\mathcal{J}_{\text{CoL}}$ controller, with a consistent $\text{Med}|\gamma| = 2.5\,\text{deg}$. The consistency stems from the combination of yaw-angle step width considered in the study and the loss-avoiding nature of the cost-function, which only leads to larger $|\gamma|$ values for $\psi$ close to 0 deg. This strategy results in the best overall ratio of $\eta_e/\text{Med}|\gamma|$, showing that a large share of the gains can be achieved with minimal yaw misalignment effort. Looking at Table 2, $\mathcal{J}_{\text{RA}}$ consistently performs the worst, showing that while the approach can lead to occasional high gains, the related effort does

| | $\eta_e$ | | | | Med$\lvert\gamma\rvert$ | | | |
|---|---|---|---|---|---|---|---|---|
| $\sigma_\varphi$ | $E[p]$ | $\mathcal{J}_{\mathrm{RA}}$ | $\mathcal{J}_{\mathrm{CoL}}$ | $\mathcal{J}_{\mathrm{SR}}$ | $E[p]$ | $\mathcal{J}_{\mathrm{RA}}$ | $\mathcal{J}_{\mathrm{CoL}}$ | $\mathcal{J}_{\mathrm{SR}}$ |
| 3 deg | **+1.37**% | +1.13% | +1.11% | +1.23% | 12.11 deg | 16.38 deg | **2.50** deg | 10.00 deg |
| 5 deg | **+0.85**% | +0.56% | +0.64% | +0.70% | 12.08 deg | 18.37 deg | **2.50** deg | 9.32 deg |
| 7 deg | **+0.35**% | +0.13% | +0.21% | +0.23% | 7.50 deg | 16.30 deg | **2.50** deg | 5.00 deg |

**Table 1.** Gain in energy efficiency over the baseline after 24 hours of simulated time, as well as median absolute yaw angle, excluding values $\leq 0.1$ deg. The best performances in terms of largest gain and smallest yawing effort are marked in bold.

| | $\eta_e$/Med$\lvert\gamma\rvert$ [%/ deg] | | | |
|---|---|---|---|---|
| $\sigma_\varphi$ | $E[p]$ | $\mathcal{J}_{\mathrm{RA}}$ | $\mathcal{J}_{\mathrm{CoL}}$ | $\mathcal{J}_{\mathrm{SR}}$ |
| 3 deg | 0.113 | 0.068 | **0.444** | 0.123 |
| 5 deg | 0.070 | 0.030 | **0.256** | 0.075 |
| 7 deg | 0.047 | 0.008 | **0.084** | 0.046 |

**Table 2.** Gain in energy efficiency per invested median absolute yaw misalignment angle. The best performance is marked in bold.

not justify the approach. The $\mathcal{J}_{\mathrm{SR}}$ cost function overall performs similar compared to $E[p]$, highlighting the scaled-down version of the performance already observed in Fig. 6 (a) and (d).

In conclusion, the showcase demonstrates that maximizing for the expected power remains a valid approach, and no immediate improvement can be shown by also including the variance alongside the expected value. However, the results also suggest that cost functions like $\mathcal{J}_{\mathrm{CoL}}$ can lead to better gain-to-investment ratio while utilizing lower absolute yaw angles overall.

## 5. Conclusion

This paper lays the groundwork for designing robust cost functions for yaw steering controllers using dynamic wake models. To this end, optimization problems that maximize power and energy are considered. An extensive database for two turbine performance is generated, providing insight into farm output statistics and the impact of state transitions under an uncertain, time-varying wind direction signal. Four cost functions are presented, which aim to provide robust control setpoints that maximize the expected value while considering the connected uncertainty. The paper introduces the amortization time and the uniform convergence time as metrics to parametrize the impact of control actions in time and to compare optimal setpoints for power and energy. The look-up tables resulting from the cost functions are tested in a ten-turbine wind farm over a 24-hour period. The results suggest that maximizing for the expected power remains a valid approach, and no immediate improvement can be shown by also including the variance alongside the expected value. However, the results also show that alternative cost functions can lead to a better return on investment ratio with much lower absolute yaw angles for slightly smaller gains.

## Acknowledgments

This work is part of Hollandse Kust Noord wind farm innovation program where CrossWind C.V., Shell, Eneco, Grow and Siemens Gamesa are teaming up; funding for the PhD's and PostDocs was provided by CrossWind C.V. and Siemens Gamesa.

## References

[1] Meyers J, Bottasso C, Dykes K, Fleming P, Gebraad P, Giebel G, Göçmen T and van Wingerden J W 2022 Wind farm flow control: Prospects and challenges *Wind Energy*

*Science* **7** 2271–2306

[2] Rott A, Doekemeijer B, Seifert J K, van Wingerden J W and Kühn M 2018 Robust active wake control in consideration of wind direction variability and uncertainty *Wind Energy Science* **3** 869–882

[3] Simley E, Fleming P and King J 2020 Design and analysis of a wake steering controller with wind direction variability *Wind Energy Science* **5** 451–468

[4] Lejeune M, Moens M and Chatelain P 2022 A Meandering-Capturing Wake Model Coupled to Rotor-Based Flow-Sensing for Operational Wind Farm Flow Prediction *Frontiers in Energy Research* **10** 884068 ISSN 2296-598X

[5] Becker M, Allaerts D and van Wingerden J W 2022 FLORIDyn - A dynamic and flexible framework for real-time wind farm control *Journal of Physics: Conference Series* **2265** 032103

[6] van den Broek M J, De Tavernier D, Hulsman P, van der Hoek D, Sanderse B and van Wingerden J W 2023 Free-vortex models for wind turbine wakes under yaw misalignment – a validation study on far-wake effects *Wind Energy Science* **8** 1909–1925

[7] Knoop S 2019 Wind - lidar wind profiles measured at North Sea wind farm TenneT platforms 1 second raw data

[8] Bortolotti P, Tarres H, Dykes K, Merz K, Sethuraman L, Verelst D and Zahle F 2019 IEA Wind TCP Task 37: Systems Engineering in Wind Energy - WP2.1 Reference Wind Turbines Tech. Rep. NREL/TP–5000-73492, 1529216

[9] Becker M, Lejeune M, Chatelain P, Allaerts D, Mudafort R and van Wingerden J W 2025 A dynamic open-source model to investigate wake dynamics in response to wind farm flow control strategies *Wind Energy Science* **10** 1055–1075

[10] Bay C J, Fleming P, Doekemeijer B, King J, Churchfield M and Mudafort R 2023 Addressing deep array effects and impacts to wake steering with the cumulative-curl wake model *Wind Energy Science* **8** 401–419

[11] Shapiro A, Dentcheva D and Ruszczynski A 2021 *Lectures on stochastic programming: modeling and theory* (SIAM)

[12] Sharpe W F 1998 The sharpe ratio *Streetwise–the Best of the Journal of Portfolio Management* **3** 169–85

[13] Stanley A P J, Bay C J and Fleming P 2023 Enabling control co-design of the next generation of wind power plants *Wind Energy Science* **8** 1341–1350

[14] Taschner E, Becker M, Verzijlbergh R and van Wingerden J 2024 Comparison of helix and wake steering control for varying turbine spacing and wind direction *Journal of Physics: Conference Series* **2767** 032023

[15] Jensen N O 1983 *A Note on Wind Generator Interaction* (Roskilde, Denmark: Risø National Laboratory)

[16] Kanev S 2020 Dynamic wake steering and its impact on wind farm power production and yaw actuator duty *Renewable Energy* **146** 9–15