https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

Integrated Wind Farm Layout Optimization Accounting for Wake-Induced Blade Fatigue and Wake-Steering Effects on LCOE

Patrick De Smet[1], Georg Jacobs[1], Dustin Frings[1], Leon Schenke[1], Stefan Witter[1], Thorsten Reichartz[1], and Martin Knops[1]

[1]Chair for Wind Power Drives, RWTH Aachen University, Germany

Correspondence: Patrick De Smet (patrick.de.smet@rwth-aachen.de)

Abstract. While contemporary layout optimization considers losses in energy production due to wake effects, the effects of blade fatigue remain mostly underrepresented. However, evaluating this coupled mechanism within the design loop is critical to accurately assess the levelized cost of energy (LCOE) under more realistic assumptions. In this paper, a layout optimization method that evaluates energy production costs based on turbine lifetime expectancy and accounts for wake-induced blade fatigue is presented. The method is used to study the effects on optimal layouts and park economics. Turbine lifetime expectancy is estimated using Miner's cumulative rule, including wake-induced turbulence, asymmetry introduced by wake steering, and azimuth-dependent variations in blade loading, to determine individual remaining lifetimes. This work compares three optimization scenarios that differ in whether the objective accounts for lifetime degradation and quantifies the resulting differences in the levelized cost of energy and optimal turbine placement. We find that the annual-energy-production (AEP) landscape of the studied site contains multiple near-degenerate optima, while blade fatigue varies steeply across these layouts, so a lifetime-aware objective helps in selecting between layouts that a pure AEP objective cannot distinguish. Under a demonstration reference damage equivalent load (DEL), anchored such that the most wake-exposed turbine of the AEP-optimal layout retains half of its assumed design life, the AEP-optimal layout carries a hidden lifetime-aware LCOE penalty of 7% that is invisible to a standard fixed-lifetime objective, and a layout optimized directly against the lifetime-aware LCOE removes this penalty at no cost in energy yield. The fatigue model is calibrated against single-turbine OpenFAST and two-turbine FAST.Farm wake simulations. Applying a power-maximizing wake-steering pass to each optimized layout further increases AEP and, despite the added once-per-revolution fatigue loading of the yawed rotors, moderately extends the remaining life of the most-exposed turbine. The results show that for the site assumed in this work, lifetime-aware LCOE optimization identifies the best available layout and lowers the worst-case fatigue loading.

1 Introduction

By the end of 2024, global wind energy capacity exceeded 1136 GW (Global Wind Energy Council, 2025). To meet the low-carbon energy goals, further expansion is needed. Developers are under increasing pressure to lower the levelized cost of energy (LCOE). Achieving this means optimizing both annual energy production (AEP) and managing capital expenditures

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

(CAPEX) and operational expenditures (OPEX) throughout the life of a project (Fingersh et al., 2006). Layout optimization is a common approach to reduce wake losses and increase AEP (Mosetti et al., 1994; Grady et al., 2005).

However, the long-term economic performance of a wind farm is determined by factors that extend far beyond the first-year energy yield captured by AEP. For example, wake effects increase the intensity of the turbulence (TI) downstream of the turbines, increasing the blade-root loads relative to the free-stream conditions (Barthelmie et al., 2009; Frandsen et al., 2006). The Palmgren-Miner linear accumulated damage rule predicts that turbines operating in high-TI wake conditions accumulate structural fatigue damage at a rate that is proportional to the damage equivalent load (DEL) raised to the Wöhler exponent of the material (Miner, 1945; Palmgren, 1924). Due to this damage, the turbines could reach the end of their life well before the assumed project duration.

Standard LCOE evaluations treat the design lifetime as a fixed project parameter, which is typically 20 years in accordance with IEC 61400-1 (International Electrotechnical Commission, 2019). This practice neglects variations in actual service life among turbines due to factors such as varying wake exposure. A turbine positioned in the wake of several upstream turbines accumulates fatigue damage at an elevated rate and potentially reaches the end-of-life well before the design horizon. Neglecting this variation in turbine lifetime can lead to an overestimation of the economic value of closely spaced layouts. The objective of this study is therefore to quantify how much of this hidden fatigue cost can be recovered by carrying turbine-specific lifetimes directly into the layout-optimization objective.

Several studies have examined wake-induced turbulence intensity, the resulting blade fatigue and its relationship to farm layout. Frandsen (2007) derived analytical expressions for wake-added turbulence intensity, and Frandsen et al. (2006) modeled the associated wind speed deficit. These directly link inter-turbine spacing to effective turbulence loading. Stanley et al. (2022) demonstrated a fast reduced-order blade-fatigue model for layout optimization which primarily captures fatigue effects from partial waking. Mendez Reyes et al. (2019) validated a model that predicts fatigue from active wake control. Wake steering yaws upstream turbines away from the prevailing wind direction to deflect their wakes laterally and reduce the velocity deficit that downstream turbines experience. The method has been established as a standard power-maximizing control strategy in field trials and can be simulated and optimized with frameworks such as FLORIS (Mudafort et al., 2025; Jiménez et al., 2010; Fleming et al., 2019; Gebraad et al., 2016). The resulting yaw misalignment introduces an additional once-per-revolution fatigue loading on the upstream turbine. The coupling between wake-induced fatigue accumulation, fatigue from wake-steering, and a lifetime-adjusted LCOE objective in a layout optimizer has not been fully investigated in the literature.

This paper presents an integrated wind farm model that evaluates a per-turbine remaining design life from Miner's accumulated damage rule applied to blade fatigue, and computes an LCOE in which each turbine contributes to the revenue stream only up to its individually predicted lifetime. We compare three optimization scenarios:

a) AEP maximization without degradation modeling,

b) LCOE minimization with fixed lifetimes,

c) LCOE minimization with lifetime-aware fatigue.

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

Comparing these three scenarios isolates how much of the economic gap between an AEP-driven layout and a degradation-aware one is recoverable through layout choice alone.

2 Methodology

The model links four parts of the layout evaluation in order. First, a wake model calculates the wake-reduced wind speed v_{loc} and the effective turbulence intensity I_{eff} at each turbine for all wind directions and speeds. Second, a load model converts these into blade-root moment ranges (ΔS_{turb}, ΔS_{1P}) and accumulates them into a damage equivalent load (DEL). Third, the model translates the accumulated damage into each turbine's remaining lifetime. This lifetime is then used in a lifetime-aware LCOE formula, which is evaluated within a genetic-algorithm layout optimizer. Figure 1 gives an overview of the complete pipeline and the next sections describe each component in turn.

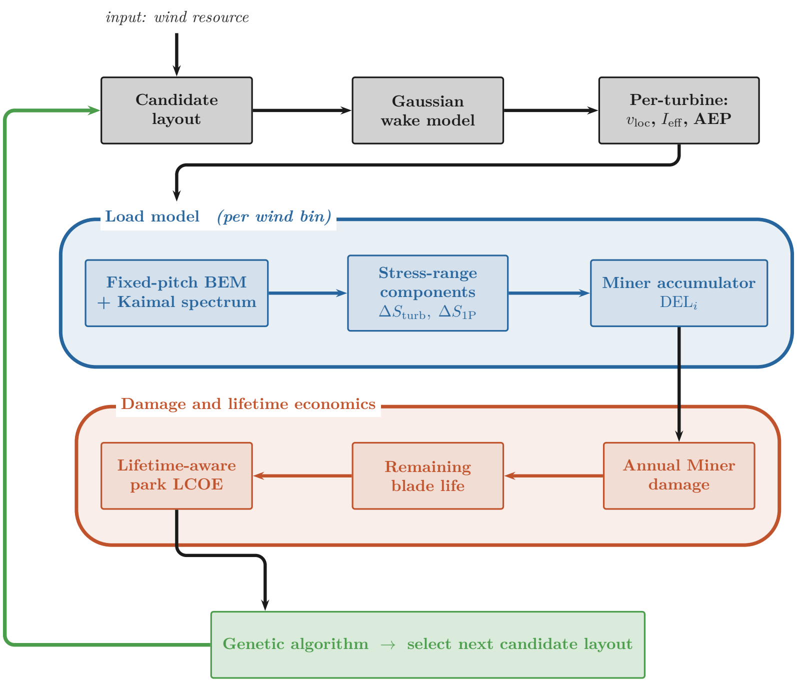

Figure 1. Overview of the lifetime-aware wind-farm layout optimization pipeline. The wake model and load model provide per-turbine power and fatigue estimates for each candidate layout. The genetic algorithm iterates over candidate layouts until the park LCOE converges.

2.1 Wake model and effective turbulence intensity

The wake effects of the wind farm layouts are modeled using the Bastankhah and Porté-Agel Gaussian model (2014), with wake decay parameter $k^* = 0.022$ and near-wake correction factor $\varepsilon = 0.2\sqrt{\beta}$, where C_T denotes the thrust coefficient of the

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

wake-generating turbine and

$$\beta = \frac{1}{2}\frac{1+\sqrt{1-C_T}}{\sqrt{1-C_T}}. \tag{1}$$

To combine the effects of multiple upstream turbines, we use a linear superposition of velocity deficits. The normalized wake velocity deficit $\Delta u_{ij}/U_\infty$ at the hub of turbine j due to upstream turbine i, with the free-stream wind speed U_∞, is:

$$\frac{\Delta u_{ij}}{U_\infty} = \left(1-\sqrt{1-\frac{C_{T,i}}{8\sigma_{ij}^2/D^2}}\right)\exp\left(-\frac{r_{ij}^2}{2\sigma_{ij}^2}\right), \tag{2}$$

where $\sigma_{ij} = k^* x_{ij} + \varepsilon D$ is the Gaussian wake width, x_{ij} is the streamwise separation, r_{ij} is the lateral offset, and D is the diameter of the rotor. The wind resource is discretized into direction bins of 10° and wind-speed bins of $1\ \mathrm{m\,s^{-1}}$ for all AEP and fatigue evaluations (the wind rose of the study site is described in Sect. 2.5). The effective turbulence intensity at each turbine position combines ambient and wake-added turbulence in quadrature. The wake-added term follows Frandsen's (2007) fitted maximum-added-wake-turbulence model:

$$I_{\mathrm{eff}} = \sqrt{I_0^2 + I_{\mathrm{wake}}^2}, \quad I_{\mathrm{wake}} = \frac{1}{1.5+0.3\,s\,\sqrt{v}}, \quad s = \frac{x}{D}, \tag{3}$$

where I_0 is the ambient turbulence intensity, x is the downstream distance to the wake-generating turbine, D is its rotor diameter, and v is the hub-height wind speed in $\mathrm{m\,s^{-1}}$.

The same directional and speed discretization is used to compute the AEP. For each wind direction of the wind rose, the wake deficits described above are evaluated for the candidate layout, giving a wake-reduced wind speed $v_{\mathrm{loc},i}$ at every turbine i. The park AEP is the frequency-weighted sum of the power curve P evaluated at these speeds,

$$\mathrm{AEP} = T_{\mathrm{yr}} \sum_i \sum_\theta \sum_k f_\theta\, p_k\, P(v_{\mathrm{loc},i}(\theta, v_k)), \tag{4}$$

where $T_{\mathrm{yr}} = 8760$ h, f_θ is the annual frequency of direction sector θ, and p_k the probability of wind-speed bin v_k within that sector. The effective TI per-turbine and direction computed at this stage are stored and reused directly by the fatigue model below, so that both quantities (AEP and DEL) are evaluated from the same underlying flow field rather than from separate assumptions.

2.2 Blade fatigue model

We evaluate blade root fatigue via a damage equivalent load (DEL) model based on the Palmgren-Miner linear accumulated damage rule (Miner, 1945; Palmgren, 1924). The annual moment-range spectrum includes five physically distinct contributions:

1. a once-per-revolution (1P) component from azimuthal induction asymmetry in turbulent inflow, scaled by an internal calibration factor k_{asym} against the effective turbulence intensity ($\Delta S_{\mathrm{1P,TI}}$)
2. a 1P component from yaw misalignment, dominated by the in-plane wind component that alternately advances and retreats against the blade's rotational velocity ($\Delta S_{\mathrm{1P,yaw}}$)

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

3. a 1P component from the blade's own weight, projected into the flapwise direction by the blade pitch angle ($\Delta S_{\mathrm{1P,grav}}$)
4. a 1P component from the vertical wind-speed gradient across the rotor disk ($\Delta S_{\mathrm{1P,shear}}$)
5. a broadband turbulence component (ΔS_{turb}).

Throughout, ΔS denotes the peak-to-peak blade-root flapwise bending-moment range, carried in kNm. The 1P components combine into a single moment range. The yaw and gravity terms both peak when the blade is horizontal (the 3 and 9 o'clock positions) and therefore add coherently, while the shear term peaks at 12 o'clock, 90° out of phase, and the turbulence-driven term carries no fixed phase.

$$\Delta S_{\mathrm{1P}} = \sqrt{\Delta S^2_{\mathrm{1P,TI}} + (\Delta S_{\mathrm{1P,yaw}} + \Delta S_{\mathrm{1P,grav}})^2 + \Delta S^2_{\mathrm{1P,shear}}}. \tag{5}$$

Together with the broadband component ΔS_{turb}, this range enters the per-condition contribution to the fatigue accumulator over a wind condition bin of duration t_s, with cycle counts n_{1P} and n_{turb}.

$$\sum_j n_j \, \Delta S_j^m = n_{\mathrm{1P}} \, \Delta S^m_{\mathrm{1P}} + n_{\mathrm{turb}} \, \Delta S^m_{\mathrm{turb}}. \tag{6}$$

The damage equivalent load and the annual Miner damage D_{ann} follow with:

$$\mathrm{DEL} = \left[\frac{\sum_j n_j \, \Delta S_j^m}{N_{\mathrm{eq}}}\right]^{1/m}, \quad D_{\mathrm{ann}} = \left(\frac{\mathrm{DEL}}{\mathrm{DEL_{ref}}}\right)^m \frac{1}{T_d}. \tag{7}$$

Here $m = 10$ is the Wöhler exponent for composite blade materials (International Electrotechnical Commission, 2019), $N_{\mathrm{eq}} = 10^7$ is the reference cycle count, $\mathrm{DEL_{ref}}$ is the design-basis reference DEL at which a turbine consumes its fatigue budget in exactly the nominal design life, and $T_d = 30$ yr is our assumed unreduced nominal design life. We adopt 30 yr rather than the 20-yr minimum because modern multi-megawatt platforms are increasingly designed and certified for longer operation. A shorter T_d would raise all LCOE levels through the annuity factor but leaves the comparison between scenarios structurally unchanged, since every scenario shares the same T_d. $\mathrm{DEL_{ref}}$ is a demonstration parameter rather than a calibrated aeroelastic load. The actual reference DEL depends on the full aeroelastic load spectrum and varies with turbine type, rotor class, and site wind conditions. We fix it by a simple convention, $\mathrm{DEL_{ref}} = \mathrm{DEL_{max,A}} \, (15/30)^{1/m}$, such that the most wake-exposed turbine of a reference layout optimized for only AEP retains exactly half of its 30-year design life. This reference layout is defined formally as Scenario A in Sect. 2.5, and its numerical anchor value is reported in Sect. 3.5. The sensitivity of all results to this choice is examined in Sect. 3.8. In principle, this parameter could be chosen to reflect the actual design DEL of the turbine type, based on the full aeroelastic load spectrum and fatigue measurements.

Equations (5) through (7) are the complete moment-range and damage model. The remainder of this section derives the reference moments, the calibration constants, and the cycle counts n_{1P} and n_{turb} that enter them, in that order.

2.2.1 Load-amplitude reference moments

Each load-amplitude component scales with a primary amplitude moment reference $M_{\mathrm{blade,amp}}$. We obtain it by differentiating the rotor thrust at the steady-state operating point (λ, θ) with pitch θ and rotor speed Ω held fixed. Neither pitch nor drivetrain

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

inertia can respond within a single 1P cycle. The collective pitch controller (CPC) we assume and use in our OpenFAST/-FAST.Farm calibration (Sect. 3.2) rate-limits pitch to 10°/s and low-pass filters its generator-speed feedback at 0.25 Hz. This places the loop bandwidth within a few percent of the NREL 5 MW's once-per-revolution frequency at rated speed ($\approx$0.20 Hz). Load fluctuations at or above 1P therefore receive little to no corrective re-pitching. Individual pitch control (IPC), which actively compensates periodic 1P loads and would attenuate the 1P-shear and 1P-TI terms accordingly, is not considered in this study as a simplification. We discuss a potential path for its inclusion in section 4.

Writing C_T as a function of tip-speed ratio and pitch, $C_T(\lambda,\theta)$, with $\lambda = \Omega R/v$ the tip-speed ratio, v the ambient hub-height wind speed, ρ the air density, A the rotor swept area, R the rotor radius, and $B = 3$ the number of blades, gives $\mathrm{d}\lambda/\mathrm{d}v\,|_\Omega = -\lambda/v$. Differentiating $F = \frac{1}{2}\rho A v^2 C_T(\lambda(v),\theta)$ at fixed θ,Ω gives

$$\left.\frac{\mathrm{d}F}{\mathrm{d}v}\right|_{\theta,\Omega} = \rho A v\, C_T(\lambda,\theta) - \tfrac{1}{2}\rho A v \lambda \left.\frac{\partial C_T}{\partial \lambda}\right|_{\theta}, \tag{8}$$

and the amplitude moment reference is

$$M_{\mathrm{blade,amp}} = \frac{v}{2}\left.\frac{\mathrm{d}F}{\mathrm{d}v}\right|_{\theta,\Omega}\frac{2R}{3B}. \tag{9}$$

$M_{\mathrm{blade,amp}}$ is evaluated at the rotor's real steady-state (λ,θ) operating point from the manufacturer's control schedule and used for all fast load-amplitude terms (Eqs. 12, 13, 20). When $\partial C_T/\partial\lambda = 0$, as holds under ideal closed-loop pitch control, Eq. (9) reduces to:

$$M_{\mathrm{blade}} = \frac{1}{2}\rho A v_{\mathrm{eff}}^2 C_T \frac{2R}{3B}, \tag{10}$$

following the annulus integration of Manwell et al. (2010). The factor $2R/3$ is the lever arm obtained by integrating a linearly-growing thrust distribution over the rotor disk, and the $1/B$ term splits the total rotor thrust evenly across the three blades. We use the ambient rather than the wake-reduced wind speed for v_{eff} here, so that wake exposure affects fatigue loading only through I_{eff} and not through a partially offsetting reduction in mean dynamic pressure. This follows the same convention as the effective-turbulence approach to wind-farm fatigue assessment described by Sørensen et al. (2008). Their design equations for a turbine in a wind farm integrate the same wind-speed distribution as for a turbine in free flow, with wake exposure entering only via an elevated turbulence standard deviation. The relative moment perturbation simplifies to:

$$M(v_{\mathrm{eff}} + \Delta v) \approx M(v_{\mathrm{eff}})\left(1 + \frac{2\Delta v}{v_{\mathrm{eff}}}\right) \implies \frac{\Delta M}{M} \approx \frac{2\Delta v}{v_{\mathrm{eff}}}, \tag{11}$$

which is also used for the yaw-induced moment range (Eq. 18).

2.2.2 Broadband turbulence and once-per-revolution turbulence-driven components

Using the relative wind-speed fluctuation from the effective turbulence intensity, $\Delta v/v_{\mathrm{eff}} = I_{\mathrm{eff}}$, and applying the linearization of Eq. (11) directly gives a turbulence-driven moment range which is proportional to $I_{\mathrm{eff}}\,M_{\mathrm{blade,amp}}$. Calibrating $M_{\mathrm{blade,amp}}$ against the OpenFAST simulations of Sect. 3.2 shows that this linear form leaves a systematic and predictable residual across

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

the tested turbulence range, so we instead raise the turbulence intensity to a power p and refit. This is an empirical closure rather than a derived result, and a power law is one of several possible closures with a single free shape parameter. We therefore tested the power law I_{eff}^p against two alternatives with the same number of free parameters, an additive offset $I_{\mathrm{eff}}+c$ and a saturating rational form $I_{\mathrm{eff}}/(1+cI_{\mathrm{eff}})$, refitting k_{asym} jointly with each candidate against the same two calibration datasets used below. Over the sampled turbulence range of 5 to 21 percent, the additive offset fits marginally better than the power law (2.7 percent against 4.6 percent maximum calibration residual on the turbulence sweep) and the saturating rational form fits worse (6.1 percent), while linear scaling performs far worse than any of the three (20.7 percent). We retain the power law because, unlike the additive offset, it returns exactly zero broadband and once-per-revolution turbulence-driven loading at zero turbulence intensity, consistent with the shear and gravity terms alone accounting for the load spectrum under an idealized turbulence-free inflow. A power-law exponent $p<1$ is also consistent with a saturating gust-amplitude response, since at high turbulence intensity the rotor disk samples an increasingly spatially averaged turbulent field and the effective load amplitude grows more slowly than the disk-average fluctuation intensity. To avoid overfitting on a single dataset, k_{asym} and p are fitted jointly across the single-turbine OpenFAST/TurbSim simulations (3 random seeds per condition) and the independent two-turbine FAST.Farm wake simulations of Sect. 3.3, with the gravity term of Eq. (19) active, giving $p\approx 0.66$. The broadband turbulence component adopts this estimate directly.

$$\Delta S_{\mathrm{turb}} = 2\,I_{\mathrm{eff}}^p\,M_{\mathrm{blade,amp}}. \tag{12}$$

The 1P-TI component scales the same quantity by an internal calibration factor k_{asym}, which relates the once-per-revolution azimuthal sampling asymmetry to the disk-averaged temporal fluctuation captured by Eq. (12):

$$\Delta S_{\mathrm{1P,TI}} = 2\,k_{\mathrm{asym}}\,I_{\mathrm{eff}}^p\,M_{\mathrm{blade,amp}}. \tag{13}$$

2.2.3 Wake dependence of k_{asym}

The coefficient k_{asym} is the effective transfer coefficient between bulk turbulence intensity and once-per-revolution azimuthal load asymmetry. It absorbs both the spatial decorrelation of turbulence across the rotor disk and the aerodynamic admittance of the blade, and is therefore calibrated rather than derived. Its wake dependence, however, follows from spatial coherence alone. Whether a turbulent fluctuation acts as a uniform gust, contributing to the broadband term of Eq. (12) without creating a 1P loading, or as an azimuthal gradient depends on the spatial coherence of the turbulence across the rotor disk. It is dependent on the ratio of the rotor diameter D to the turbulence integral length scale L. Using the standard exponential spatial autocorrelation that defines the integral length scale, $\rho(\Delta r)=\exp(-\Delta r/L)$, the decorrelated fraction of turbulent energy between two points separated by D is $1-\exp(-D/L)$. Since k_{asym} multiplies the turbulence intensity I_{eff} directly (Eq. 13), an amplitude quantity rather than an energy quantity, the azimuthal-asymmetry-available fraction must be expressed as an amplitude fraction as well, and is obtained by taking the square root of the decorrelated energy fraction:

$$k_{\mathrm{asym}}(L) \propto \sqrt{1-\exp(-D/L)}. \tag{14}$$

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

WIND ENERGY SCIENCE DISCUSSIONS

Thomsen and Sørensen (1999) measured, from the Vindeby wind farm, that wake-added turbulence reduces the local turbulence integral length scale by a factor of approximately two, relative to free-stream conditions for wind speeds below 12 m s^{-1}. Using the IEC 61400-1 Kaimal-model ambient length scale for a 100 m hub height ($\Lambda_1 = 42$ m, $L_0 \approx 8.1\,\Lambda_1 \approx 340$ m) and $D = 126$ m, Eq. (14) gives a decorrelated energy fraction of $1 - \exp(-D/L_0) \approx 0.31$ in free flow and $1 - \exp(-2D/L_0) \approx 0.52$ with the wake-halved length scale. This ratio defines $k_{\mathrm{asym,wake\_factor}}$:

$$k_{\mathrm{asym,wake\_factor}} = k_{\mathrm{asym,wake}}/k_{\mathrm{asym,base}} \approx 1.30. \tag{15}$$

A turbine generally experiences both ambient turbulence (length scale L_0) and superimposed wake turbulence (length scale $L_0/2$) at once. Since the effective turbulence intensity already combines the two contributions in quadrature, $I_{\mathrm{eff}}^2 = I_0^2 + I_{\mathrm{wake}}^2$ (Eq. 3), assuming that ambient and wake-generated turbulence contribute independently to the asymmetry-producing variance, the effective asymmetry-transfer coefficient can be approximated by a variance-weighted average of their respective transfer coefficients, $k_{\mathrm{asym,eff}} = k_{\mathrm{asym,base}}\,(I_0^2/I_{\mathrm{eff}}^2) + k_{\mathrm{asym,wake}}\,(I_{\mathrm{wake}}^2/I_{\mathrm{eff}}^2)$. Substituting $I_0^2 = I_{\mathrm{eff}}^2 - I_{\mathrm{wake}}^2$ gives

$$k_{\mathrm{asym,eff}} = k_{\mathrm{asym,base}} \left[1 + (k_{\mathrm{asym,wake\_factor}} - 1)\,\frac{I_{\mathrm{wake}}^2}{I_{\mathrm{eff}}^2} \right], \tag{16}$$

with $k_{\mathrm{asym,base}} = 1.44$ (the joint OpenFAST/FAST.Farm calibration of Sect. 3.2 and 3.3) and $k_{\mathrm{asym,wake\_factor}} \approx 1.30$ from Eq. (15). This recovers $k_{\mathrm{asym,eff}} = k_{\mathrm{asym,base}}$ for an unwaked turbine ($I_{\mathrm{wake}} = 0$) and smoothly increases toward $k_{\mathrm{asym,base}} \cdot k_{\mathrm{asym,wake\_factor}}$ as wake exposure comes to dominate the effective turbulence. We use $k_{\mathrm{asym,eff}}$ as k_{asym} in Eq. (13) throughout the remainder of this paper.

2.2.4 1P-yaw component

Under yaw misalignment Φ, a once-per-revolution flapwise load arises from the in-plane wind component $v \sin\Phi$, which normalised by the axial inflow $v\cos\Phi$ enters as $\tan\Phi$ in Eq. (17). This adds to the blade's rotational velocity Ωr on the advancing side of the rotor and subtracts from it on the retreating side, modulating the local relative velocity and angle of attack once per revolution. The steady-wind yawed OpenFAST simulations of Sect. 3.4 show that this relative-velocity modulation, with the perturbation ratio evaluated at the representative non-dimensional blade radius $r_{\mathrm{rep}} = 0.7$,

$$\left.\frac{\Delta v}{v}\right|_{\mathrm{yaw}} = k_{\mathrm{yaw}}\,\frac{v\,\tan\Phi}{\Omega\, r_{\mathrm{rep}} R}, \tag{17}$$

reproduces the measured yaw-induced 1P moment and its growth with Φ. Applying the linearization of Eq. (11) to this perturbation gives

$$\Delta S_{\mathrm{1P,yaw}} = 2\,M_{\mathrm{blade}}\,k_{\mathrm{yaw}}\,\frac{v\,\tan\Phi}{\Omega\, r_{\mathrm{rep}} R}, \tag{18}$$

where Ω is the rotor speed, R the rotor radius, and $k_{\mathrm{yaw}} = 1.3$ a single calibration constant fitted to the yawed OpenFAST simulations of Sect. 3.4. The term is zero at zero yaw and grows under wake steering. It uses M_{blade} rather than $M_{\mathrm{blade,amp}}$, since yaw misalignment is a sustained operating condition. The yaw and gravity terms enter Eq. (5) as magnitudes, $|\Delta S_{\mathrm{1P,yaw}}| + |\Delta S_{\mathrm{1P,grav}}|$, independent of the sign of the yaw angle. Since the relative phase of the two terms flips

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

with the yaw direction, this is the worst-case coherent combination, and the steering self-penalty is conservative under this convention.

2.2.5 1P-gravity and 1P-shear components

The blade's own weight produces an in-plane root moment $m_\mathrm{b}\,g\,r_\mathrm{cg}\cos\phi\cos\Theta$ that is maximal when the blade is horizontal. Stanley et al. (2022) use the same gravity-moment formulation in their reduced-order blade-fatigue model. They found gravity-driven loading to be a dominant contributor to azimuth-dependent fatigue damage. With the blade pitched by θ, this in-plane moment projects into the flapwise direction with $\sin\theta$, so a deterministic flapwise 1P oscillation appears whenever the rotor operates above rated wind speed and grows with the pitch angle. We take its peak-to-peak range over one rotor revolution:

$$\Delta S_\mathrm{1P,grav} = 2\,m_\mathrm{b}\,g\,r_\mathrm{cg}\,\cos\phi\,\cos\Theta\,\sin\theta(v), \tag{19}$$

where m_b is the blade mass, r_cg the radial position of its centre of gravity, ϕ the precone, and Θ the shaft tilt, and $\theta(v)$ is the steady-state pitch angle at wind speed v. For the NREL 5 MW blade we use the same values as Stanley et al. (2022), 17 537 kg, 20.65 m, 2.5°, and 5°. These data approximately correspond to the blade properties described by Resor (2013). The term contains no free parameter. Because the projected gravity moment peaks at the same horizontal blade positions as the yaw term, the two parts add coherently in Eq. (5). Below rated wind speed $\theta \approx 0$ and the term vanishes. Above rated it supplies the turbulence-independent load floor.

The vertical wind-speed gradient across the rotor disk contributes another deterministic 1P term. A blade tip at the top of its sweep sees a higher inflow speed than at the bottom, producing a once-per-revolution oscillation that is even present at zero ambient turbulence and zero yaw misalignment:

$$\Delta S_\mathrm{1P,shear} = 2\,c_\mathrm{shear}\,M_\mathrm{blade,amp}. \tag{20}$$

This oscillation occurs at exactly the 1P frequency that the pitch loop is designed to be insensitive to (see the controller-response discussion in Sect. 2.2), so it also uses $M_\mathrm{blade,amp}$ rather than M_blade. We set $c_\mathrm{shear} = \lambda_\mathrm{shear}\,(R/H_\mathrm{hub})$, with $\lambda_\mathrm{shear} = 0.2$ the IEC 61400-1 default power-law wind-shear exponent used directly in our TurbSim reference simulations, rather than an independently fitted constant.

$\Delta S_\mathrm{1P,TI}$ arises from random turbulent fluctuations that do not repeat in phase from one rotor revolution to the next and is therefore statistically independent of the deterministic terms. $\Delta S_\mathrm{1P,shear}$ peaks at the 12 o'clock blade position, where the wind-shear velocity excess is greatest, while $\Delta S_\mathrm{1P,yaw}$ and $\Delta S_\mathrm{1P,grav}$ both peak at 3 or 9 o'clock, where the advancing and retreating velocity modulation and the in-plane gravity moment are largest. The shear term is therefore 90° phase-offset from the coherent yaw and gravity contributions, giving the root-sum-square combination of Eq. (5).

Figure 2 illustrates the azimuthal phase relationship of the 1P load components. The deterministic shear, gravity, and yaw contributions are set by mean conditions alone and are identical regardless of wake exposure, while the stochastic broadband and 1P-TI contributions grow with the turbine's effective turbulence intensity. The combined 1P moment range ΔS_1P (Eq. 5) is cycled once per rotor revolution.

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

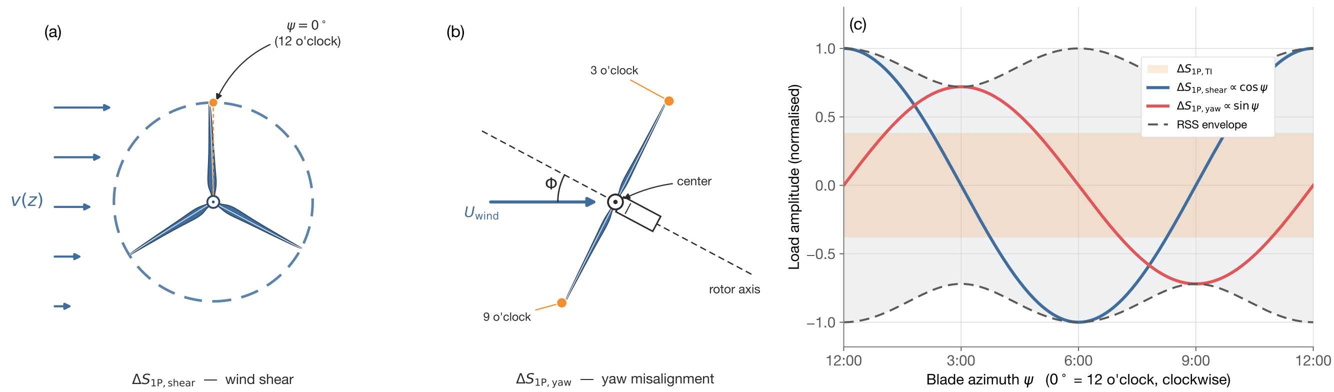

Figure 2. Physical origin and azimuthal phase relationship of the three 1P blade-root load components. (a) Wind-shear gradient causes a deterministic $\Delta S_{1\mathrm{P,shear}}$ peak at the 12 o'clock blade position ($\psi = 0°$). (b) Top-down view of the yawed rotor where the rotor axis is misaligned from the wind by the yaw angle Φ. The advancing and retreating relative-velocity causes $\Delta S_{1\mathrm{P,yaw}}$ to peak at the 3 and 9 o'clock blade positions. The pitch-projected gravity moment $\Delta S_{1\mathrm{P,grav}}$ peaks at the same positions and adds coherently. (c) Normalised 1P load amplitudes vs. blade azimuth ψ. The orange band marks the stochastic $\Delta S_{1\mathrm{P,TI}}$ (random phase), and the dashed curve is the root-sum-square envelope of the deterministic components.

2.2.6 Cycle counting and lifetime damage

The broadband cycling rate f_{bb} is derived from the IEC 61400-1 Kaimal longitudinal velocity spectrum (International Electrotechnical Commission, 2019; Chai et al., 2025) filtered by a first-order rotor spatial admittance. The Kaimal spectrum, normalized by the turbulence variance $\sigma_u^2 = I_{\mathrm{eff}}^2\, v^2$, is

$$\hat{S}_u(f) = \frac{4\,(L_k/v)}{(1+6\,f\,L_k/v)^{5/3}}, \tag{21}$$

where $L_k = 8.1\,\Lambda_1$ is the longitudinal integral length scale and $\Lambda_1 = 42\,\mathrm{m}$ for hub heights above 60 m. A first-order rotor spatial admittance,

$$|\chi(f)|^2 = \frac{1}{1+(f\,D/v)^2}, \tag{22}$$

accounts for spatial averaging of turbulence across the rotor disk such that eddies whose wavelength is shorter than D are attenuated. The expected broadband cycling rate follows a Rice-type spectral-moment estimate (Rice, 1944), applied to the admittance-filtered spectrum:

$$f_{\mathrm{bb}} = \frac{1}{\pi}\sqrt{\frac{m_2}{m_0}}, \qquad m_k = \int_0^\infty f^k\,\hat{S}_u(f)\,|\chi(f)|^2\,\mathrm{d}f. \tag{23}$$

Because $\hat{S}_u$ and $|\chi|^2$ both depend on f only through the reduced frequency fD/v and the fixed ratio L_k/D, the spectral moment ratio is $m_2/m_0 \propto (v/D)^2$, so it follows that $f_{\mathrm{bb}} \propto v/D$. The constant prefactor of Eq. (23) is not identified independently but is absorbed into the joint calibration of p and k_{asym} against OpenFAST. Equation (23) is evaluated numerically on a pre-computed wind-speed grid and looked up in the optimization loop at zero additional cost. For the NREL 5 MW rotor ($D = 126\,\mathrm{m}$), this yields $f_{\mathrm{bb}} \approx 0.007$ to $0.03\,\mathrm{Hz}$ across the operational wind speed range. Rotor speed below rated is estimated from a representative tip-speed ratio $\lambda = 7.5$, capped at the rated rotor speed $\omega_{\max}$. The cycle counts over the duration t_s of a wind-condition bin are:

$$n_{1\mathrm{P}} = f_{1\mathrm{P}}\, t_s, \qquad f_{1\mathrm{P}} = \frac{\omega}{2\pi}, \qquad \omega = \min\left(\frac{\lambda v}{R}, \omega_{\max}\right), \tag{24}$$

$$n_{\mathrm{turb}} = f_{\mathrm{bb}}\, t_s, \tag{25}$$

where $\omega_{\max}$ is the rated rotor speed. Together with $\Delta S_{1\mathrm{P}}$ and ΔS_{turb}, these cycle counts are exactly the quantities entering the accumulator and damage equations already stated as Eqs. (6) and (7) at the start of this section.

Equation (6) gives the cycle-count sum for a single wind direction and speed bin. The annual accumulator $\sum_j n_j\, \Delta S_j^m$ used in Eq. (7) is obtained by repeating this evaluation for every one of the 36 directions of the 10° grid and the Weibull speed bins introduced in Sect. 2.1, at the effective TI specific to that bin and turbine position, and summing the per-bin contributions weighted by the bin's annual occurrence probability and scaled to a full year of operation. AEP and fatigue therefore share one and the same directional and speed discretization.

2.3 Lifetime-aware LCOE

In the standard formulation, the park LCOE is:

$$\mathrm{LCOE} = \frac{\mathrm{CAPEX} + (\mathrm{OPEX}/T_d)\,\mathrm{rbf}(r,T_d)}{\mathrm{rbf}(r,T_d)\cdot \mathrm{AEP}}, \tag{26}$$

where $\mathrm{rbf}(r,T_d) = [(1{+}r)^{T_d} - 1]/[(1{+}r)^{T_d}\, r]$ is the annuity factor for real weighted average cost of capital (WACC) $r = 5\%$ and fixed design life $T_d = 30$ yr. Here OPEX denotes the total undiscounted operating expenditure accumulated over the design life. Consequently, OPEX/T_d is the annual operating cost and $(\mathrm{OPEX}/T_d)\,\mathrm{rbf}$ its discounted sum, which is consistent with the flat annual rate given in Sect. 2.5. In the proposed lifetime-aware formulation, each turbine i carries an individual remaining life determined by the weakest-link rule across all active degradation mechanisms:

$$T_i = \min_k \frac{1 - D_{\mathrm{acc},k,i}}{D_{\mathrm{ann},k,i}}, \tag{27}$$

where k ranges over blade fatigue and any additional damage mechanisms that may be considered in future extensions of this model, and $D_{\mathrm{acc},k,i}$ is the accumulated damage fraction at the start of the assessment period. The park LCOE is then the AEP-weighted mean of per-turbine values:

$$\mathrm{LCOE}_{\mathrm{park}} = \frac{\sum_i \mathrm{LCOE}_i(T_i)\, \mathrm{AEP}_i}{\sum_i \mathrm{AEP}_i}, \tag{28}$$

where each value $\mathrm{LCOE}_i(T_i)$ is computed with the annuity factor $\mathrm{rbf}(r,T_i)$ and AEP_i is the annual energy production of turbine i. Turbines for which $T_i < T_d$ carry a shortened annuity horizon, effectively penalizing layouts that concentrate

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

blade-root fatigue damage on downstream turbines. We deliberately use the energy-weighted mean of per-turbine cost ratios rather than the park-total ratio of summed discounted costs over summed discounted energy. Each turbine is treated as an asset whose capital must be recovered within its own remaining life, so a heavily loaded turbine cannot be cross-subsidized by the discounted energy of its neighbors, which is exactly the signal the optimizer should respond to. The two formulations coincide when all turbines reach the full design life.

Real turbines are not designed to fail at exactly T_d. Certification load cases, partial safety factors, and margins in the material S-N curve mean that a turbine reaching $T_i < T_d$ under this model has consumed its assessed fatigue budget rather than having a predicted failure date. Since the demonstration reference DEL of Sect. 2.2 already collapses these margins into a single reference value, T_i is best read as an economic surrogate signal that penalizes layouts concentrating fatigue on a subset of turbines, rather than as a realistic prediction of remaining service life.

We only consider blade-root flapwise fatigue in the weakest-link rule here. However, edgewise blade fatigue, tower-base, main-bearing, and yaw-system loads are outside the model. Also, we assume that T_i terminates the entire turbine's contribution to revenue in Eq. (28). However, in practice, a consumed blade fatigue budget would rather trigger a blade replacement or a life-extension assessment. Ending the revenue stream at T_i is the most conservative economic reading.

2.4 Optimization algorithm

The layout optimization uses a genetic algorithm implemented via the DEAP framework (Fortin et al., 2012). Turbine positions are encoded as continuous (x, y) coordinates with a population of 80 candidate layouts evolving over 150 generations. Each generation applies Deb tournament selection with tournament size 3. A blend crossover with $\alpha = 0.5$ is applied with probability 0.7, and Gaussian mutation with standard deviation $\sigma = 50$ m is applied to each coordinate with probability 0.2, for an overall mutation rate of 0.2. The number of turbines is fixed at $N = 9$, which was chosen to match the capacity density of the site described in Sect. 2.5. Two constraints are enforced as soft penalties on the fitness: a minimum inter-turbine spacing of $3D$ and containment within the site boundary. Each penalty is the summed violation distance in meters multiplied by a weight of 10^7. This penalty exceeds the objective scale by many orders of magnitude and renders any violating candidate effectively infeasible. After optimization, we validated that all final layouts reported in this paper satisfy the constraints strictly, with minimum inter-turbine separations of $4.5D$ (Scenarios A and B) and $4.8D$ (Scenario C).

2.5 Optimization scenarios

Three scenarios are defined to isolate the contributions of the objective function and the degradation model. Scenario A maximizes AEP without any degradation modeling. Scenario B minimizes LCOE using the standard fixed-lifetime formula (26). Scenario C minimizes the lifetime-aware park LCOE (28) by coupling the optimizer to the full fatigue pipeline at each function evaluation. Each scenario is optimized from multiple independent random seeds and the best layout is retained, to reduce the genetic algorithm's sensitivity to initialization.

All scenarios use the same 5 MW reference turbine ($D = 126$ m, hub height $h = 100$ m, rated tip speed $u_{\mathrm{tip}} = 80\ \mathrm{m\,s^{-1}}$) placed within a hexagonal site boundary of approximately $2500\,\mathrm{m} \times 1500\,\mathrm{m}$. The wind resource is synthesized from eight

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

Open Access WIND ENERGY SCIENCE DISCUSSIONS

directional modes representative of a North Sea coastal site. The mode-mean speeds range from 7.0 to 10.5 $\mathrm{m\,s^{-1}}$ and a dominant westerly component accounts for 22% of hours. Within each mode, the wind speeds follow a Weibull distribution with shape parameter $k = 2$. The optimizer evaluates this resource on the 10° bins introduced in Sect. 2.1. The sensitivity of the optimized layouts to a finer directional discretization is checked in Sect. 3.5. The mean wind speeds in our example site are typical for Northwestern European onshore and near-shore sites (Dörenkämper et al., 2020). The ambient TI is set to $I_0 = 5\%$, which is consistent with low-roughness onshore conditions. CAPEX is estimated using the cost regression of Lüers et al. (2024), which yields approximately 1470 EUR/kW for the reference turbine. OPEX is set to a simplified flat rate of 40 EUR/kW/yr.

3 Results

The next sections present the results in three steps. First, we evaluate the fatigue model using an example case to show the size of the degradation effects. Second, we compare the optimized layouts for the three scenarios. Third, we compare the resulting LCOE values using both the fixed-lifetime and lifetime-aware methods.

3.1 Blade fatigue and DEL distribution

Figure 3 shows the per-turbine blade-root DEL and annual Miner damage accumulation for a 3×3 grid layout at 5D uniform spacing under the site wind resource. The downwind column of turbines accumulates roughly sixteen times the annual Miner damage of the front column (18.3 to 18.6% $\mathrm{yr^{-1}}$ against 0.7 to 1.2% $\mathrm{yr^{-1}}$). This is driven by elevated effective TI amplified through the Wöhler exponent $m = 10$. The 1P-TI, 1P-shear, 1P-gravity, and broadband turbulence components together account for the entire DEL in this illustrative case, since the 1P-yaw component is inactive due to no consideration of wake steering in this example. By translating these damages to remaining lifetimes via Eq. (27), the downwind turbines at this deliberately tight spacing exhaust their design fatigue life in approximately 5 years and the middle column in approximately 13 years. This is substantially lower than the nominal 30 years.

3.2 Model calibration against OpenFAST

To assess how well the analytical DEL model of Sect. 2.2 reproduces realistic blade-root loading, we compare it with OpenFAST aeroelastic simulations of the NREL 5 MW reference turbine driven by TurbSim full-field turbulent wind. Five ambient turbulence levels ($I_{\mathrm{eff}} = 5\%, 9\%, 13\%, 17\%, 21\%$) are each simulated at four wind speeds (8.0, 11.4, 14.0, and 18.0 $\mathrm{m\,s^{-1}}$, spanning partial load to above rated), with three independent random seeds per condition (60 ten-minute time series in total) to reduce sensitivity to individual turbulence realizations. The blade-root flapwise moment (`RootMyb1`) from each run is rainflow-counted. Cycles from the three seeds are pooled into a single Miner accumulator per condition using the same Wöhler exponent and reference cycle count as Eq. (7). They are then combined across wind speeds for each turbulence level exactly as the annual DEL pipeline combines wind-resource bins.

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

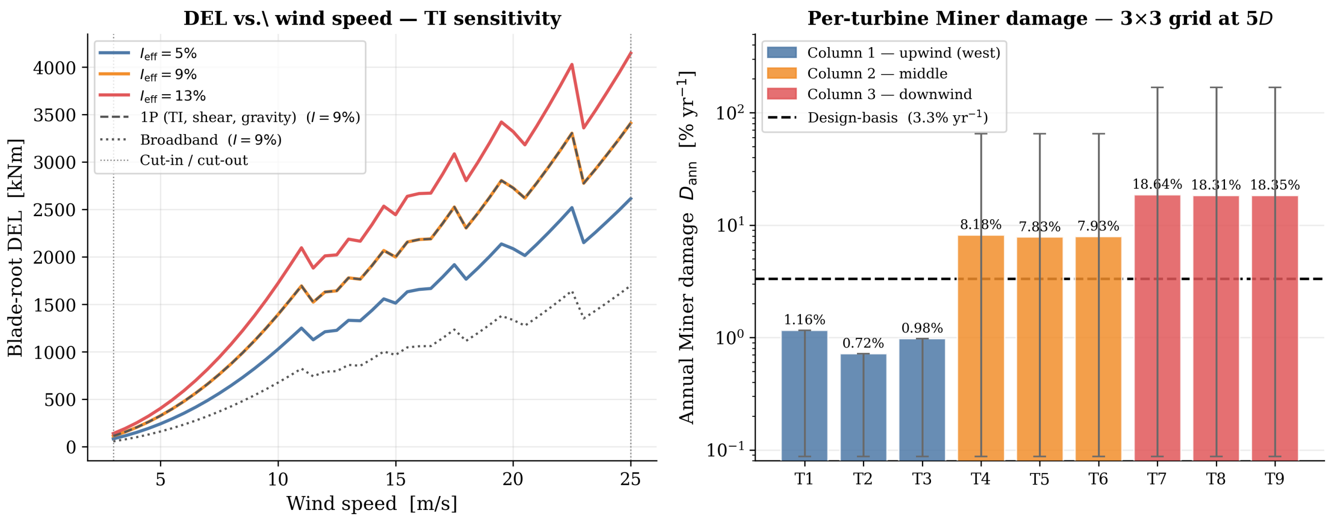

Figure 3. (a) Blade-root DEL as a function of wind speed for three effective turbulence intensity levels ($I_{\text{eff}} = 5\%$, 9%, 13%), with the combined 1P (TI-, shear-, and gravity-driven) and broadband components shown for $I_{\text{eff}} = 9\%$. (b) Per-turbine annual Miner damage for a 3×3 uniform grid with 5D spacing under the site wind resource. Turbines T1 to T3 are the upwind (west) column, T4 to T6 the middle column, and T7 to T9 the downwind (east) column. Error bars span the 10th to 90th percentile of the per-direction annual damage rate, weighted by sector frequency. The design-basis line at 3.3% yr^{-1} marks $1/T_d = 1/30$ yr^{-1}, the uniform damage rate that consumes the fatigue budget in exactly the nominal design life.

Figure 4 compares the resulting DEL against the analytical model evaluated at the same conditions. The model systematically underpredicts, with the combined-across-speed DEL agreeing within 15.7% across all five turbulence levels. Figure 5 decomposes this comparison by individual wind speed and turbulence level: all 20 individual conditions have a negative error (-3.4% to -32.5%, mean magnitude 14.4%), most strongly at 8 m s^{-1}, while the combined metric of Figure 4 is dominated by the highest-amplitude 18 m s^{-1} conditions under the $m = 10$ Wöhler weighting, where the per-condition underprediction is smallest. We assume that potential causes for these underprediction are the inflow-based cycle-count formulation of Eq. (23), the fixed-pitch C_T linearization at partial load, and the power-law turbulence closure. The consequences for the overall results of our model are, however, limited since $\mathrm{DEL}_{\text{ref}}$ is anchored to the model's own Scenario-A maximum DEL (Sect. 2.2). A bias common to all turbines cancels in the $\mathrm{DEL}/\mathrm{DEL}_{\text{ref}}$ ratio of Eq. (7), and only the condition-dependent residual can distort the comparison between layouts. This is achieved with the deterministic 1P terms of Eqs. (19) and (20). Without a deterministic 1P contribution from gravity and the rotor-disk wind-speed gradient, a model built only from turbulence- and yaw-driven terms underpredicts DEL substantially at low ambient turbulence and above rated wind speed, since every remaining term in Eq. (5) and Eq. (12) scales with I^{p}_{eff}. It would therefore vanish as $I_{\text{eff}} \to 0$, while a real rotor experiences 1P loading from shear and from the weight of its pitched blades alone.

As a targeted check of Eq. (15), we ran eighteen single-turbine OpenFAST/TurbSim cases at rated wind speed, covering three turbulence intensities, two coherence settings, and three seed-matched replicates. In each ambient/wake pair, the IEC

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

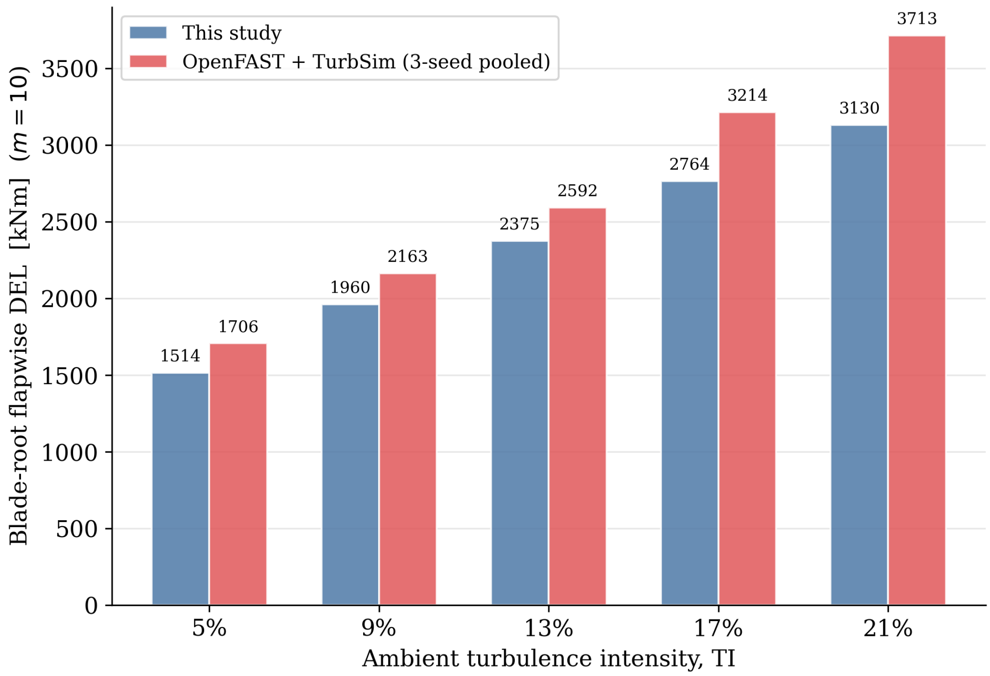

Figure 4. Blade-root flapwise DEL from the analytical model (this study) versus rainflow-counted, 3-seed-pooled OpenFAST/TurbSim simulations of the NREL 5 MW reference turbine, across five ambient turbulence levels.

Kaimal coherence length was set explicitly to L_0 (ambient) or $L_0/2$ (wake-equivalent), while TI, wind speed, random seed, and lateral/vertical coherence were held fixed. This isolates the streamwise coherence-length dependence predicted by Eq. (14) from every other source of run-to-run variability. It thereby isolates the specific length-scale-shortening mechanism measured by Thomsen and Sørensen. It is not a full emulation of wake turbulence, which also differs from ambient turbulence in spectral shape, intermittency, and mean shear within the wake's velocity-deficit profile. Across the nine seed-matched pairs, the wake-coherence case showed a higher DEL than its ambient-coherence pair in seven of nine cases, with a pooled DEL ratio of 1.05 ± 0.06 (mean $\pm$ s.d., standard error 0.02) and an implied $k_{\mathrm{asym,wake}}/k_{\mathrm{asym,base}} = 1.06 \pm 0.07$. This is below the 1.30 obtained from Eq. (15). The simulations in this paper use $k_{\mathrm{asym,wake\_factor}} = 1.30$, and the controlled OpenFAST check indicates that this remains an upper bound on the true wake-fatigue penalty from this mechanism.

3.3 Wake-fatigue calibration against FAST.Farm

The single-turbine simulations above cannot assert alone whether the ambient-speed convention of Eq. (10) is the right choice for a wake-exposed turbine. We therefore additionally simulate a two-turbine array with FAST.Farm (Shaler and Jonkman, 2021), using the same NREL 5 MW turbine and ROSCO controller (Abbas et al., 2022) as the single-turbine simulations. Six layout cases combine three downstream spacings (5D, 7D, 10D) with two lateral offsets (full-wake alignment and a 0.5D

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

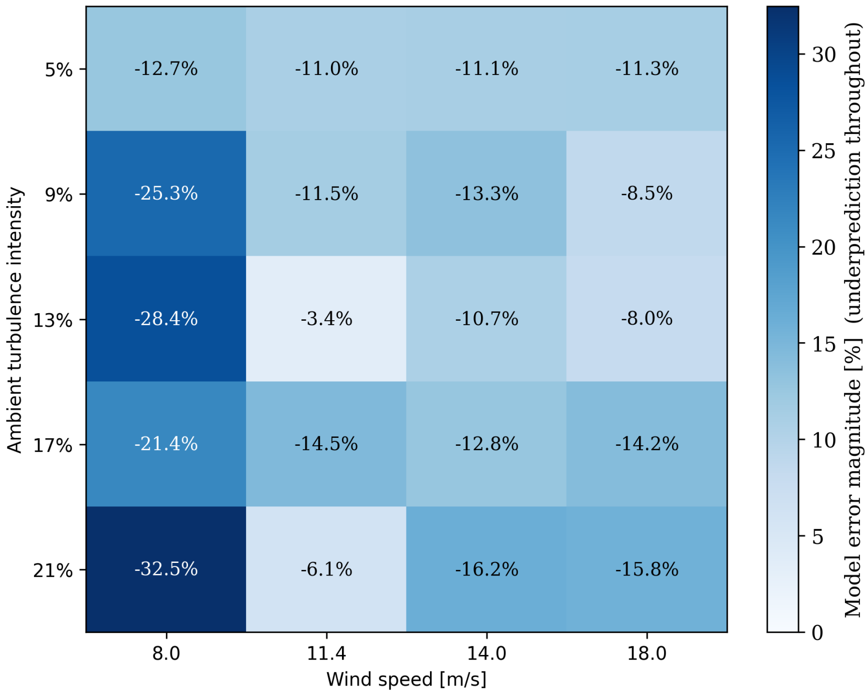

Figure 5. Model error against 3-seed-pooled OpenFAST (fixed-pitch C_T and TI-exponent model), decomposed by wind speed and turbulence intensity individually, showing the per-condition spread that the $m = 10$-weighted combination conceals.

lateral offset), shown in Figure 6, at a shared ambient condition of 11.4 $\mathrm{m\,s^{-1}}$ and 9% turbulence intensity. The upstream turbine (T1) is unwaked. The downstream turbine (T2) experiences the actual simulated wake deficit, meandering, and wake-added turbulence FAST.Farm produces, rather than an analytical wake model.

Blade-root DEL from FAST.Farm agrees with the analytical model, evaluated at each turbine's own simulated local wind speed and turbulence intensity, within 14.7% (mean 5.5%) across the 12 turbine cases (T1 and T2 across all six layout cases), with 2 of 12 exceeding 10% error and no consistent bias between the full-wake and half-wake-offset cases. The calibration constants k_{asym} and p used throughout this paper are fitted jointly across this wake simulation and the single-turbine simulations of Section 3.2, rather than to either dataset alone.

Because a rotor partially immersed in a wake samples the lateral velocity gradient once per revolution, a deterministic partial-wake 1P mechanism exists that the model does not explicitly represent, since wake exposure enters the load model only through I_{eff}. To test whether this omission matters where the gradient is steepest, we simulated four additional cases at 0.25D and 0.75D lateral offsets (at 5D and 10D spacing, same ambient condition) after the calibration above. The model reproduces these eight turbine-cases within 14.6% (mean 5.9%), with waked-turbine errors between $+0.9\%$ and -14.6% and no consistent dependence on lateral offset. Across all ten layout cases combined (20 turbine-cases), agreement is within 14.7% with a mean

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

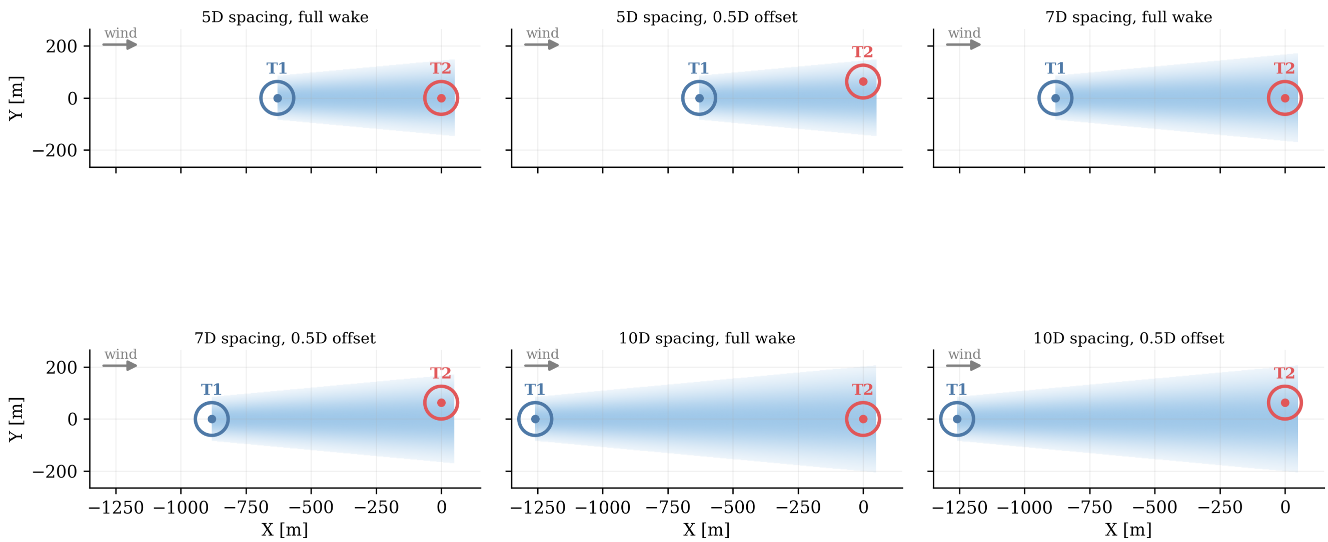

Figure 6. Geometry of the six FAST.Farm two-turbine wake-calibration cases (NREL 5 MW, $D = 125.88$ m). T1 (upstream) is unwaked. T2 (downstream) is exposed to the simulated wake.

of 5.6%. The omitted partial-wake 1P term therefore does not introduce a bias beyond the envelope of the aligned cases at this wind condition.

3.4 Yaw-fatigue validation

The wake-steering analysis of Sect. 3.9 rests on the model's yaw term, which neither the turbulence sweep nor the wake simulations above exercise, since every turbine in those runs operates aligned with the wind. We therefore add sixteen steady-wind OpenFAST simulations of the NREL 5 MW turbine with the nacelle locked at yaw angles of 0°, 10°, 20°, and 30° and uniform sheared inflow (power-law exponent 0.2) at 8.0, 11.4, 14.0, and 18.0 $\mathrm{m\,s^{-1}}$. Steady inflow removes all turbulence-driven loading, so the blade-root moment history contains only the deterministic 1P content. Fitting the first azimuthal harmonic to the flapwise root moment and vector-subtracting the zero-yaw baseline (shear and gravity) at each wind speed isolates the yaw-induced 1P amplitude. Figure 7 compares this measured amplitude with Eq. (18). The single constant $k_{\mathrm{yaw}} = 1.3$ reproduces the measured growth with yaw angle and wind speed, with errors below 18% at rated wind speed and a mean absolute error of 31.5% over the twelve yawed conditions. The scatter concentrates at small yaw angles and above rated, where the baseline-subtracted amplitudes are smallest and the pitch controller interacts with the yaw loading. The residual ≈30% uncertainty of the yaw term is carried forward as the dominant model uncertainty of the wake-steering results in Sect. 3.9.

To cross-check this result with actual steered wake, we repeat the 7D full-wake FAST.Farm case of Sect. 3.3 with the upstream turbine yawed by 20°. The model reproduces the absolute blade-root DEL of both turbines within 10% in the aligned and in the yawed configuration. It captures the direction of the downstream fatigue relief where deflecting the wake lowers the

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

WIND ENERGY SCIENCE DISCUSSIONS

downstream turbine's DEL by 4.6% in FAST.Farm and by 2.0% in the model. The upstream turbine's own steering penalty at this yaw angle is a few percent of its unwaked DEL in FAST.Farm (+3.3%).

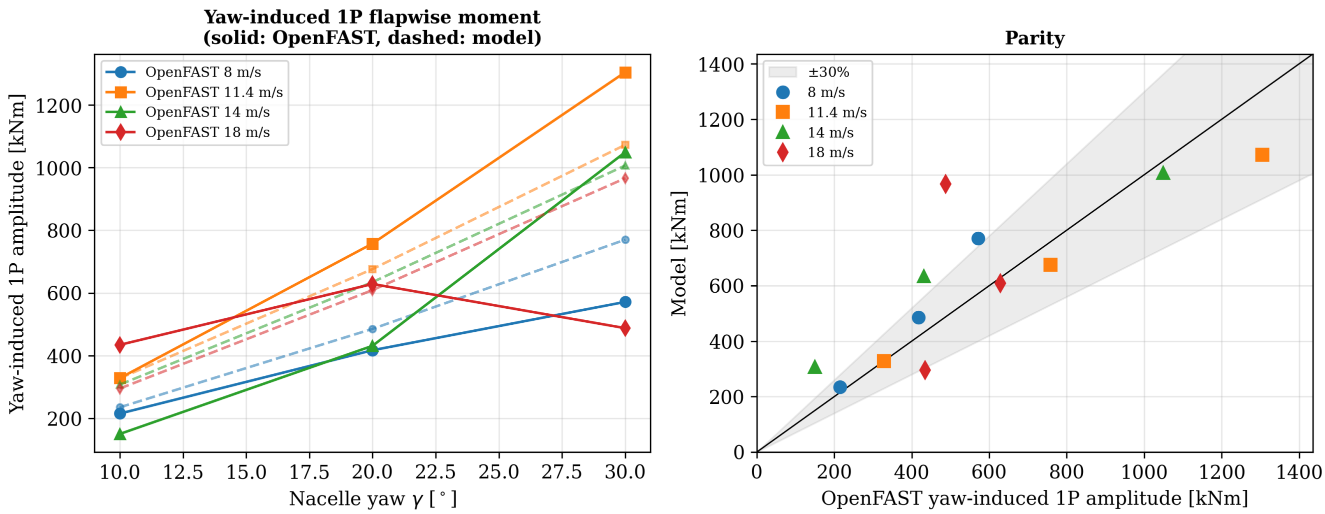

Figure 7. Yaw-induced 1P flapwise moment amplitude from steady-wind OpenFAST simulations of the NREL 5 MW turbine (nacelle locked, zero-yaw shear-and-gravity baseline vector-subtracted) versus the model term of Eq. (18) with $k_{\mathrm{yaw}} = 1.3$, across four wind speeds and three yaw angles. Left: amplitude versus yaw angle (solid: OpenFAST, dashed: model). Right: parity plot with the ±30% band.

3.5 Optimal layouts and remaining lifetimes

Figure 8 presents the optimal turbine positions and per-turbine remaining lifetimes for all three scenarios. Each scenario is optimized from eight independent random seeds and the best layout is retained. The AEP-optimal layout (Scenario A) achieves 203.9 GWh/yr but does not account for the fatigue penalty on downstream turbines. It carries a hidden lifetime-aware LCOE penalty of 2.1 EUR/MWh with the most wake-exposed turbine reaching a remaining fatigue life of 15.0 years (by the anchor convention of Sect. 2.2).

By construction, minimizing the standard fixed-lifetime LCOE (Scenario B) converges to the same layout as maximizing AEP (Scenario A) where neither objective contains a fatigue term, so the two scenarios share identical per-turbine TI, DEL, and remaining life. The lifetime-aware LCOE optimizer (Scenario C) finds a differently arranged layout that raises the worst-case remaining life from 15.0 to the full 30 years and closes the hidden penalty entirely. Across the eight random seeds optimized for Scenario A, the AEP ranges from 202.2 to 203.9 GWh/yr, so Scenario C's value sits just beyond the range of layouts the AEP objective itself finds under this optimizer. Section 4 examines this outcome further. Fatigue mitigation by layout choice is therefore penalty-free at this site. Wake exposure raises DEL only through I_{eff}, so layout choice has real leverage over which turbines carry the highest load. The maximum blade-root DEL across all nine turbines falls from 6442 kNm in Scenarios A and B to 5870 kNm in Scenario C, and the Miner damage on the most-exposed turbine drops correspondingly, as Table 1 shows.

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

eawe european academy of wind energy
Open Access
WIND ENERGY SCIENCE DISCUSSIONS

The Wöhler exponent $m = 10$ makes annual Miner damage highly sensitive to the turbulence-driven differences in DEL that layout choice can influence, which is why Scenario C closes the hidden penalty completely.

As a check of the directional discretization, both optimized layouts were re-evaluated on a finer grid with bins of 5° without re-optimization. The AEP decreases by about 1% for both layouts as narrower wake corridors are resolved, Scenario A's minimum remaining life shifts from 15.0 to 14.0 years, and Scenario C's from 30.0 to 19.7 years, while the lifetime-aware LCOE advantage of Scenario C over Scenario A reduces from 2.2 to 1.7 EUR/MWh. Part of Scenario C's nominal full-life result therefore reflects alignment with the 10° evaluation grid, but its ranking and the bulk of its economic advantage are robust to the finer discretization. In more complex terrains and under further site constraints, this result is expected to still hold in kind since the lifetime-aware LCOE captures a real trade-off and sends a genuine signal for layout choice.

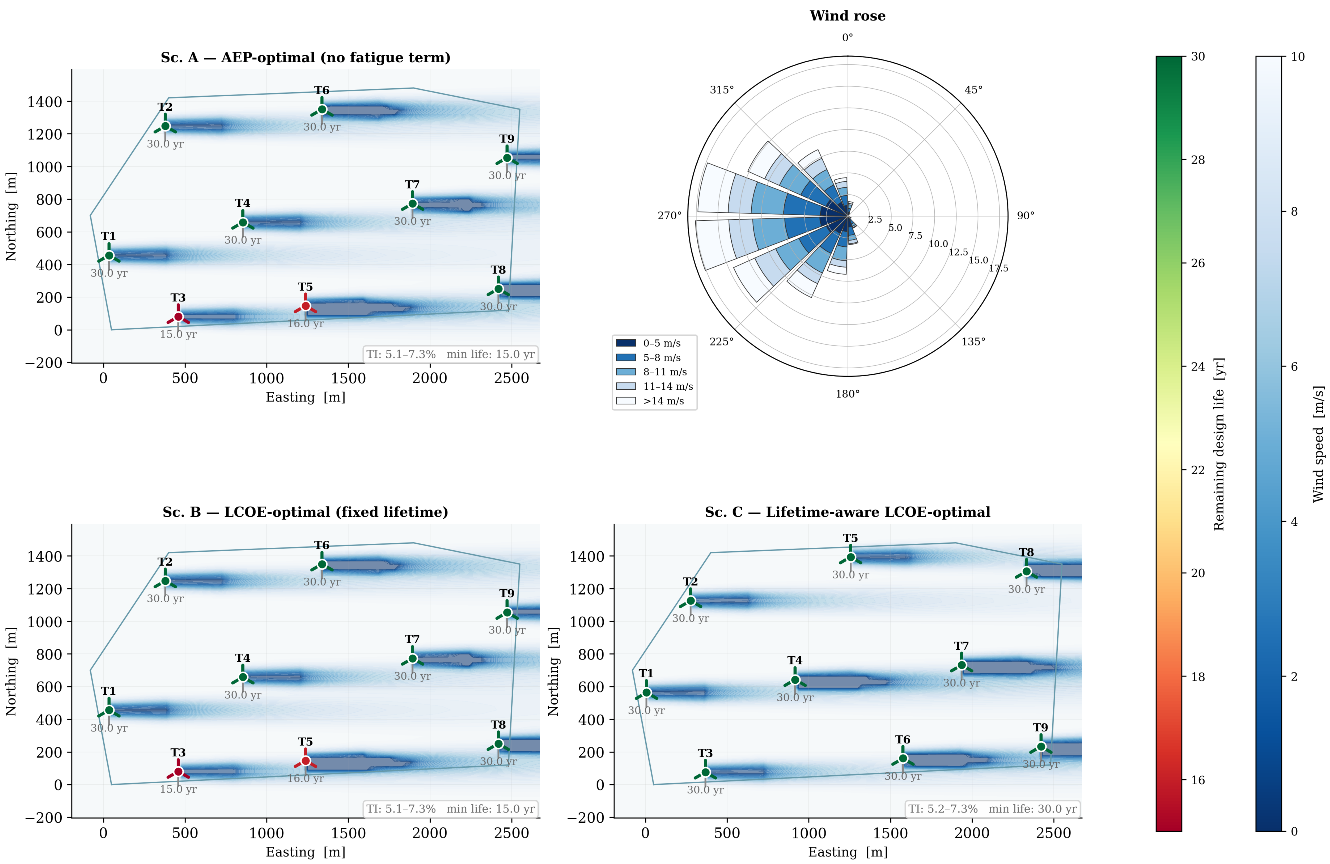

Figure 8. Optimal turbine positions for Scenarios A, B, and C, with the wind rose of the site wind resource (top right). Marker color indicates the per-turbine remaining design life (yr) under the lifetime-aware model. The prevailing wind direction is from the west.

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

3.6 Per-turbine breakdown

Table 1 lists the effective TI, blade-root DEL, and remaining life for each of the nine turbines under all three scenarios. DEL varies less than TI in relative terms across turbines within each scenario, since the wake-independent shear and gravity contributions (Eqs. 19 and 20) are the same magnitude for every turbine and combine with the turbulence-driven terms by root-sum-square. This limits how much of the turbulence-intensity spread carries through to the combined moment range. The Wöhler exponent $m = 10$ nonetheless converts this into a substantial life spread. In Scenarios A and B, T3 and T5 carry the highest DEL and are the only turbines below their full design life, at 15.0 and 16.0 years respectively. In Scenario C, the highest-DEL turbine (T3, 5870 kNm) remains at the full 30 years. Scenario C's TI spread (2.1 pp) is comparable to Scenarios A and B's (2.2 pp, Table 2). The lifetime-aware optimizer lowers the worst-case blade-root DEL rather than narrowing the overall TI spread.

Table 1. Per-turbine effective TI, blade-root DEL, and remaining life across the three scenarios. Turbines are numbered west to east within each scenario's own layout.

	Sc. A			Sc. B			Sc. C		
Turb.	**TI (%)**	**DEL (kNm)**	**Life (yr)**	**TI (%)**	**DEL (kNm)**	**Life (yr)**	**TI (%)**	**DEL (kNm)**	**Life (yr)**
T1	5.1	4428	30.0	5.1	4428	30.0	5.2	4450	30.0
T2	5.4	4542	30.0	5.4	4542	30.0	5.6	5187	30.0
T3	6.3	6442	15.0	6.3	6442	15.0	6.1	5870	30.0
T4	7.0	5891	30.0	7.0	5891	30.0	7.0	5634	30.0
T5	7.3	6401	16.0	7.3	6401	16.0	6.1	4812	30.0
T6	6.2	5105	30.0	6.2	5105	30.0	6.5	5464	30.0
T7	7.2	5651	30.0	7.2	5651	30.0	7.1	5654	30.0
T8	6.6	5300	30.0	6.6	5300	30.0	6.6	5359	30.0
T9	7.0	5757	30.0	7.0	5757	30.0	7.3	5646	30.0

Scenarios A and B yield identical values because both optimizations converge to the same layout.

3.7 LCOE comparison

Figure 9 and Table 2 compare the park-level LCOE across the three scenarios under both the standard fixed-lifetime formula and the lifetime-aware formula. Under the fixed-lifetime evaluation, Scenarios A and B are identical at 30.0 EUR/MWh, since both converge to the same layout. When the lifetime-aware formula is applied, this layout deteriorates by 2.1 EUR/MWh (7.0%), as it incurs wake-induced fatigue on downstream turbines, reducing the most-exposed turbine's life to 15.0 years. Scenario C achieves a lifetime-aware LCOE of 29.9 EUR/MWh, which eliminates the hidden penalty entirely at no AEP cost (204.5 vs. 203.9 GWh/yr). At this site, lifetime-aware optimization removes the hidden LCOE penalty completely on the 10° evaluation grid, and retains most of the advantage on the finer 5° grid (Sect. 3.5).

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

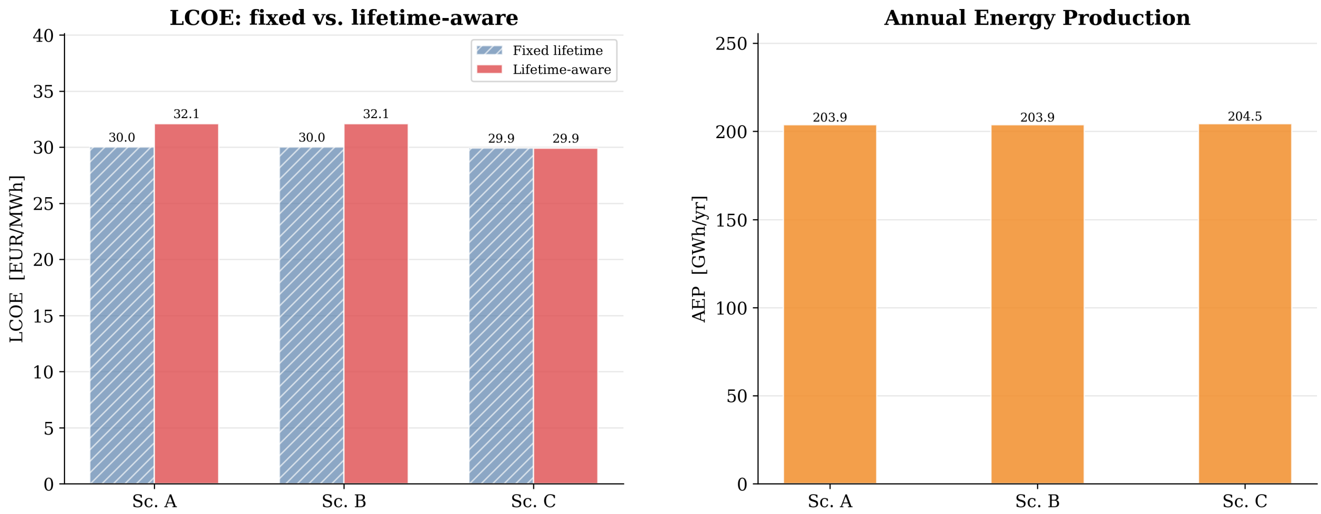

Figure 9. Park LCOE under fixed-lifetime (hatched bars) and lifetime-aware (solid bars) evaluation for the three optimization scenarios. The gap between the two bars for each scenario quantifies the hidden cost of degradation that the standard fixed-lifetime formula misses.

Table 2. Key performance metrics across optimization scenarios. AEP in GWh/yr, LCOE in EUR/MWh. Bold values indicate the metric for which each scenario's layout was optimized. Max. TI spread is the difference between the highest and lowest effective TI experienced across the nine turbines in the optimized layout, in percentage points (pp), and quantifies how evenly wake loading is distributed.

Metric	Sc. A	Sc. B	Sc. C
AEP (GWh/yr)	**203.9**	203.9	204.5
LCOE, fixed lifetime (EUR/MWh)	30.0	**30.0**	29.9
LCOE, lifetime-aware (EUR/MWh)	32.1	32.1	**29.9**
Fixed vs. lifetime-aware gap (EUR/MWh)	+2.1	+2.1	0.0
Min. remaining life (yr)	15.0	15.0	30.0
Max. TI spread (pp)	2.2	2.2	2.1

3.8 Sensitivity to the fatigue calibration parameter

Since $\mathrm{DEL_{ref}}$ is a demonstration value rather than a calibrated aeroelastic load, all three optimizations are repeated for eight $\mathrm{DEL_{ref}}$ values from 5317 to 6935 kNm (-12 to $+15\%$ around the anchor of Sect. 2.2). Figure 10 shows the resulting fixed vs. lifetime-aware LCOE gap for each scenario across this range. Scenarios A and B are identical at every tested value, since both converge to the same AEP-maximizing layout. Their hidden penalty decreases from 17.8 EUR/MWh at the lowest tested value to zero at the highest, since a larger $\mathrm{DEL_{ref}}$ allows turbines to tolerate higher loads before their predicted life is reduced. The lifetime-aware optimizer (Scenario C) keeps its gap small across the entire tested range, below 1.2 EUR/MWh from

5779 kNm upward and at most 3.2 EUR/MWh at the harshest anchor (5317 kNm, where the most exposed turbine of even a well-spread layout retains only 18 of 30 years). At the demonstration anchor $\mathrm{DEL}_{\mathrm{ref}} = 6010$ kNm, the gap is 2.1 EUR/MWh for Scenarios A and B and zero for Scenario C. The benefit of lifetime-aware optimization therefore holds robustly across the full tested range. It grows where fatigue binds harder and vanishes smoothly where it does not.

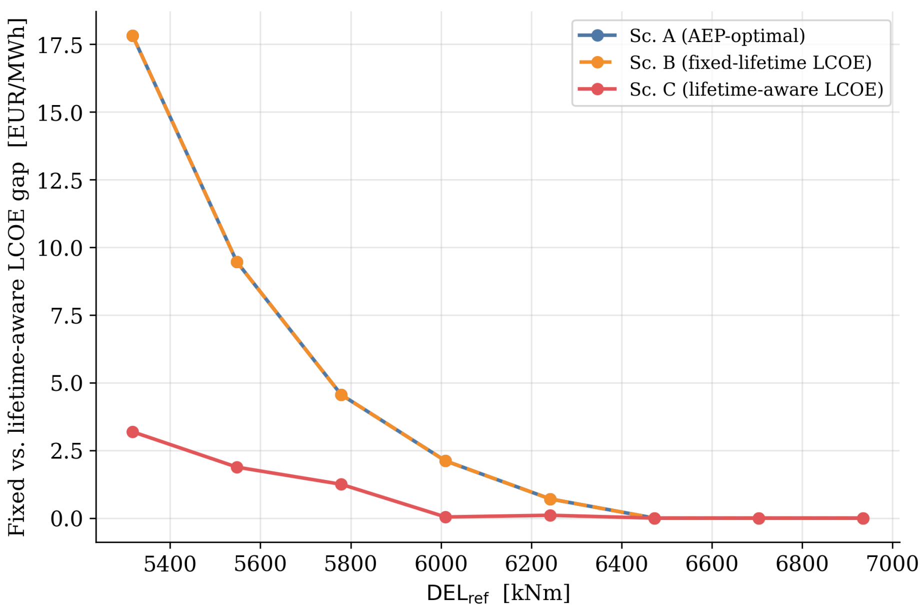

Figure 10. Fixed vs. lifetime-aware LCOE gap as a function of the fatigue calibration parameter $\mathrm{DEL}_{\mathrm{ref}}$, for all three scenarios. Scenarios A and B are identical at every tested value. The lifetime-aware optimizer (Scenario C) keeps its gap at or near zero across the full tested range.

3.9 Wake-steering trade-off

Wake steering yaws upstream turbines away from the wind to deflect their wakes laterally, reducing the velocity deficit seen by downstream turbines (Jiménez et al., 2010). Field trials report AEP gains of a few percent from this technique (Fleming et al., 2019), and the FLORIS framework established yaw co-optimization as a standard control strategy for wind farms (Gebraad et al., 2016). The same yaw misalignment also drives the 1P-yaw fatigue term in Eq. (18), so a purely power-based controller does not explicitly account for the resulting fatigue impact, whether positive or negative.

We test this by taking the three already-optimized layouts of Scenarios A, B, and C unchanged, and adding a single wake-steering pass on top of each. A coordinate-search yaw optimizer maximizes AEP per wind direction using the wake deflection model of Jiménez et al. (2010) and a cosine power loss model $\cos^{p_\mathrm{p}} \Phi$ with exponent $p_\mathrm{p} = 3$ for the yawed turbine itself. Figure 11 compares AEP and lifetime-aware LCOE before and after this yaw pass.

Wake steering raises AEP by 0.3 to 0.4% across the three scenarios. The yaw policy trades two fatigue effects against each other: deflecting a wake lowers the downstream turbine's turbulence exposure, while the yawed turbine itself incurs the 1P self-

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

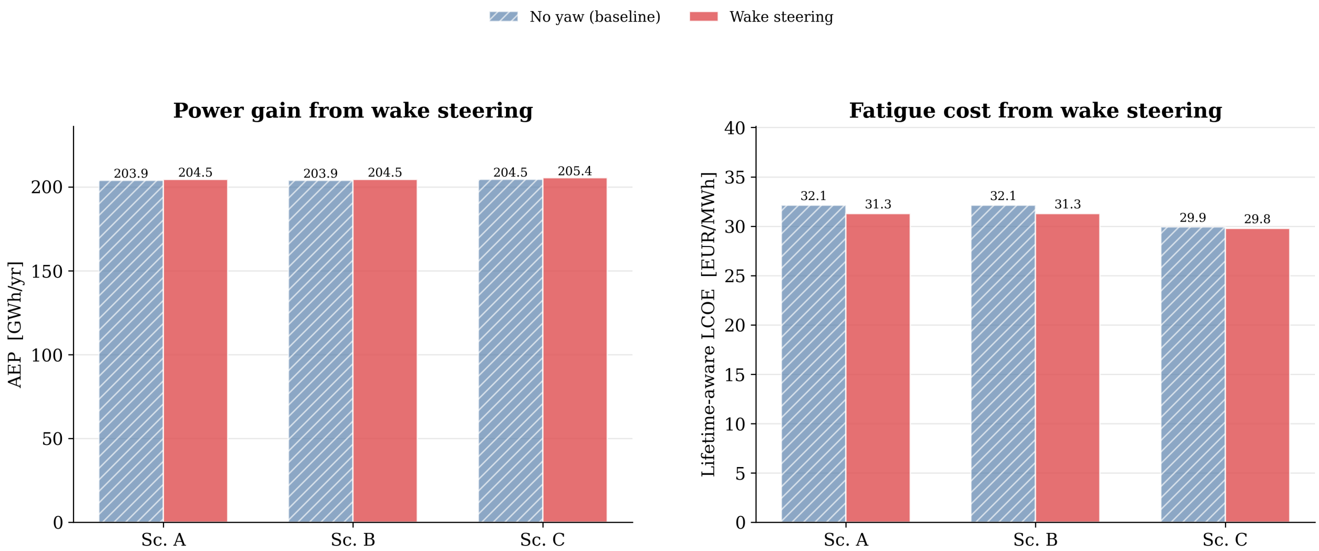

Figure 11. Wake-steering trade-off on the three optimized layouts. The left panel shows AEP gain from a power-maximizing yaw policy. The right panel shows the same yaw policy evaluated under the lifetime-aware LCOE, which accounts for the added 1P-yaw fatigue.

penalty of Eq. (18). At the moderate yaw angles the power-maximizing optimizer selects, the downstream relief dominates. The minimum remaining blade life rises from 15.0 to 16.2 years in Scenarios A and B and the lifetime-aware LCOE falls by 2.6% (32.1 to 31.3 EUR/MWh). In Scenario C, whose layout already avoids persistent wake exposure, every turbine remains at the full 30-year design life and the lifetime-aware LCOE falls by a further 0.4% to 29.8 EUR/MWh. This is the lowest value of any configuration in this study. A power-maximizing yaw controller with no explicit knowledge of fatigue therefore still improves the lifetime-aware outcome here. Wake steering and lifetime-aware layout design act as complementary rather than competing levers. However, the available steering gain is small once the layout itself is already well optimized for fatigue.

3.10 Co-optimized layout and wake steering

Scenario D evaluates a layout and yaw policy optimized jointly under the lifetime-aware objective, in contrast to the sequential, post-hoc wake steering of the previous subsection. Figure 12 shows the resulting layout, with effective TI ranging from 5.1% to 7.2% across the nine turbines and yaw angles of $-15°$ and $-10°$ on two turbines. With the yaw self-penalty in the model, joint co-optimization no longer removes the fatigue cost of steering but rather prices it. Scenario D attains the highest AEP of any configuration (205.7 GWh/yr) by accepting a minimum remaining life of 25.5 years, and lands at a lifetime-aware LCOE of 29.9 EUR/MWh, within 0.01 EUR/MWh of Scenario C and slightly above Scenario C with the sequential yaw pass (29.8, Sect. 3.9). Under this yaw-fatigue model, co-optimization thus uses yaw misalignment to buy additional energy at a fatigue price that the lifetime-aware objective keeps in check, rather than offering the strictly dominant solution that a model without the yaw self-penalty would predict.

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

eawe european academy of wind energy
Open Access WIND ENERGY SCIENCE DISCUSSIONS

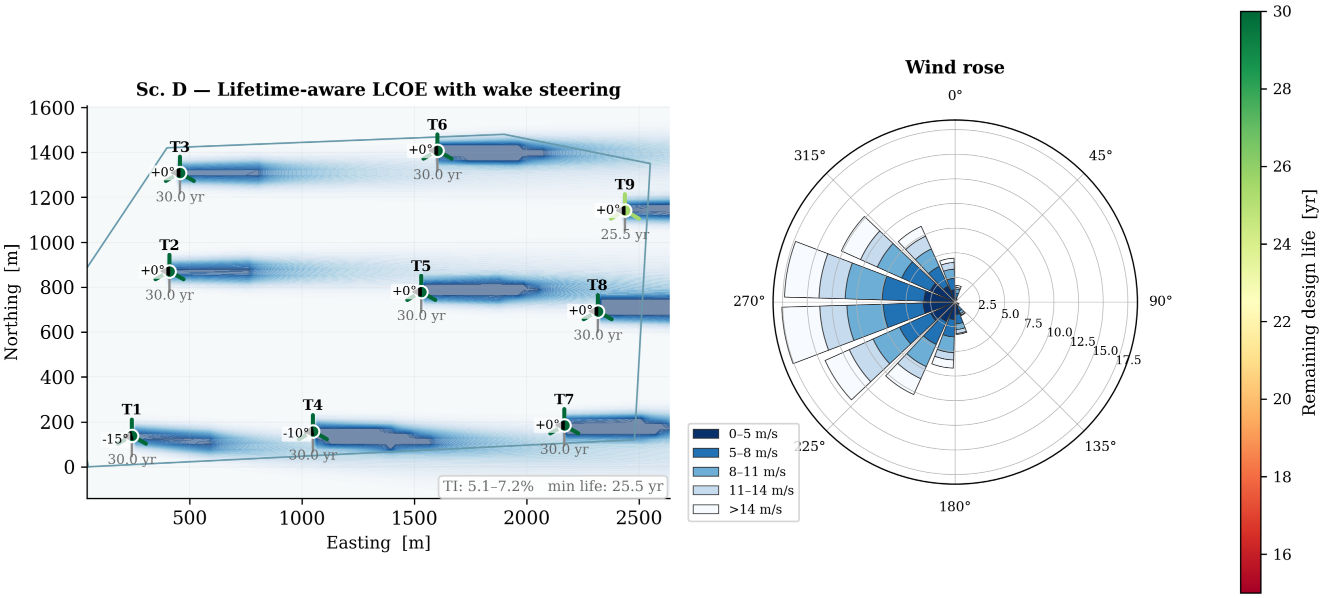

Figure 12. Optimal layout and yaw angles for Scenario D, where layout and wake steering are co-optimized jointly under the lifetime-aware LCOE objective. Marker color indicates per-turbine remaining design life. The prevailing wind direction is from the west.

4 Discussion

The results show that layouts optimized solely for AEP tend to underestimate long-term energy costs when the additional fatigue on downstream turbines is accounted for. Wake-added turbulence increases the DEL, and hence the annual Miner damage, in the same way that causes the layout differences seen in Figure 8.

Scenario C achieves a lifetime-aware LCOE of 29.9 EUR/MWh, compared to 32.1 EUR/MWh for Scenarios A and B, at no cost in energy yield for the simple scenario considered here (204.5 vs. 203.9 GWh/yr). The hidden penalties of 2.1, 2.1, and 0.0 EUR/MWh for Scenarios A, B, and C indicate that standard optimizers can select layouts that concentrate fatigue on the most exposed turbines, and that layout choice can avoid this at the studied site. The absolute size of the penalty scales directly with the demonstration anchor $\mathrm{DEL}_{\mathrm{ref}}$ (Sect. 3.8) and is expected to grow with larger rotor diameters, higher ambient turbulence, and higher WACC. Three model limitations bound the control-related conclusions. The yaw term carries an ≈30% uncertainty (Sect. 3.4) that propagates into the wake-steering gains. The finer-grid check of Sect. 3.5 shows that part of Scenario C's nominal full-life result reflects the 10° directional discretization. In addition, wake exposure enters the load model only through I_{eff}, so the deterministic 1P loading of a rotor partially immersed in a wake, which samples the lateral velocity gradient once per revolution, is not represented explicitly. The out-of-sample FAST.Farm check at four lateral offsets spanning the wake edge (Sect. 3.3) shows no consistent offset-dependent bias from this omission. The normalization of Miner damage against the design-basis DEL of the front-row turbine also introduces a turbine-type dependence that should be tested across a wider range of sites and wind climates.

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

Since Scenario C's AEP sits just beyond Scenario A's own range over various seeds, we evaluated whether the solver with AEP objective failed to find its global optimum. We tested this directly by taking Scenario C's final layout and re-optimizing annual energy production alone from that starting point, across eight further random seeds. Every one of the eight runs returned to exactly Scenario C's annual energy production value, without improving on it. Local search under the annual energy production objective therefore cannot do better than Scenario C's layout either. This indicates that Scenario C sits at or very near a genuine local optimum of the AEP landscape. The broader implication is that the AEP landscape at this site contains multiple, similarly good local optima, of which Scenario A's cold-start search reached one and the lifetime-aware objective reached another, better one, while the fatigue landscape across these near-optima is comparatively steep, varying the worst-case remaining life from 15.0 to 30.0 years between the two. A pure annual energy production objective is indifferent among these near-optimal layouts, so the lifetime-aware objective is doing additional work in selecting among them rather than merely trading energy for fatigue relief.

Individual pitch control (IPC) has been established as a 1P load-reduction strategy (Bossanyi, 2003). It is not modeled in the results above, since it targets exactly the deterministic once-per-revolution terms this model isolates, the shear, gravity, and yaw components of Eq. (5), while leaving the stochastic 1P-TI and broadband turbulence terms comparatively unaffected. Applying an illustrative attenuation factor of 0.81 to these deterministic terms only, corresponding to a reported 19 percent damage equivalent load reduction from individual pitch control (Lara et al., 2026), raises Scenario A's minimum remaining life from 15.0 to 17.8 years and reduces its hidden lifetime-aware LCOE penalty from 2.1 to 1.4 EUR/MWh. Scenario C is unaffected, since every turbine in that layout already reaches the full 30-year design life through layout choice alone, and individual pitch control has nothing left to correct. Layout optimization and individual pitch control therefore address the same fatigue penalty through different means, an AEP-optimal layout benefits from either, while a lifetime-aware layout has already captured the available benefit. The 0.81 factor is applied only to the deterministic terms here, whereas the cited reduction is a total DEL figure that may also reflect some attenuation of the broadband and 1P-TI content through the same feedback loop.

A turbine with a computed remaining lifetime of 15 years would not, in practice, be expected to fail due to fatigue exactly at that point, as engineering safety margins are incorporated into its design. However, we introduce our model as a more natural way of integrating fatigue minimization than using a combined objective with AEP, which requires selecting an explicit weighting factor for fatigue. Consequently, a consumed lifetime $T_i < T_d$ in our model should not be interpreted as indicating a likely complete failure at that time. Rather, it represents the consumption of the turbine's fatigue budget, beyond which continued operation may become economically unattractive due to increased maintenance frequency or the need for major overhauls.

5 Conclusions

This study presents an integrated wind farm layout optimization framework that combines wake-induced blade fatigue and lifetime-aware LCOE into one objective. The fatigue model uses the Palmgren-Miner cumulative damage rule with analytic blade-root bending moments, and is calibrated against single-turbine OpenFAST/TurbSim simulations (within 15.7% com-

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

bined across five turbulence levels, 3-seed pooled) and an independent two-turbine FAST.Farm wake simulation (within 14.7% across twelve turbine-cases, mean 5.5%). It is validated against yawed-rotor simulations in both tools and four out-of-sample partial-wake cases. Matching the OpenFAST data requires deterministic wind-shear and pitch-projected gravity contributions to the blade-root load spectrum, a fixed-pitch correction to the load-amplitude reference moment (Eq. 9) accounting for the pitch controller's by-design insensitivity to 1P-frequency excitations, and a sub-linear turbulence-intensity scaling, with loads assessed at each turbine's ambient wind speed. Matching the yawed-rotor data additionally requires the kinematic yaw term of Eq. (18). For turbines equipped with individual pitch control, the deterministic shear, gravity, and yaw terms should be attenuated, with the degree of attenuation depending on IPC bandwidth, sensor noise, and gain tuning. An illustrative attenuation reduces the AEP-optimal layout's hidden penalty by about a third while leaving the lifetime-aware layout unaffected, since the latter has already removed the fatigue constraint through layout choice alone. Under this model, the AEP-optimal layout carries a hidden lifetime-aware LCOE penalty invisible to a standard fixed-lifetime objective, while the lifetime-aware optimizer substantially reduces this penalty and extends the most exposed turbine's remaining life. It does so by selecting among near-degenerate AEP optima whose fatigue outcomes differ steeply, rather than by trading energy for fatigue relief. These figures are scenario-specific to the demonstration $\mathrm{DEL}_{\mathrm{ref}}$ used here and should not be read as general magnitudes. However, since the AEP of the fatigue-aware optimization remains largely unaffected, these results suggest that accounting for fatigue during layout optimization can improve the economic lifetime performance of a wind farm with little to no sacrifice in energy production.

Data availability. No additional data is required to produce the presented results. The source code is not published.

Author contributions. PDS developed the core concept, methodology, software implementation, and wrote the original draft. SW, TR, LS, and DF provided additional methodological input and contributed to the structure of the text. Supervision, manuscript revision, and final approval were provided by TR, SW, LS, DF, MK, and GJ. All authors have read and agreed to the published version of the paper.

Competing interests. None of the authors has any competing interests.

Acknowledgements. This open-access publication was funded by the RWTH Aachen University.

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

References

Abbas, N. J., Zalkind, D. S., Pao, L., and Wright, A.: A reference open-source controller for fixed and floating offshore wind turbines, Wind Energy Science, 7, 53–73, https://doi.org/10.5194/wes-7-53-2022, 2022.

Barthelmie, R. J., Hansen, K., Frandsen, S. T., Rathmann, O., Schepers, J., Schlez, W., Phillips, J., Rados, K., Zervos, A., Politis, E., et al.: Modelling and measuring flow and wind turbine wakes in large wind farms offshore, Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology, 12, 431–444, https://doi.org/10.1002/we.348, 2009.

Bastankhah, M. and Porté-Agel, F.: A new analytical model for wind-turbine wakes, Renewable energy, 70, 116–123, https://doi.org/10.1016/j.renene.2014.01.002, 2014.

Bossanyi, E. A.: Individual blade pitch control for load reduction, Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology, 6, 119–128, https://doi.org/10.1002/we.76, 2003.

Chai, Y., Wang, Y., Zheng, J., Zhang, J., and Ge, M.: Accuracy of Kaimal spectrum model for wind turbine load prediction in the Hexi Corridor: The impact of turbulence structure on wind turbine load, Results in Engineering, 26, 105 597, https://doi.org/10.1016/j.rineng.2025.105597, 2025.

Dörenkämper, M., Olsen, B. T., Witha, B., Hahmann, A. N., Davis, N. N., Barcons, J., Ezber, Y., García-Bustamante, E., González-Rouco, J. F., Navarro, J., et al.: The making of the new european wind atlas–part 2: Production and evaluation, Geoscientific model development, 13, 5079–5102, https://doi.org/10.5194/gmd-13-5079-2020, 2020.

Fingersh, L., Hand, M., and Laxson, A.: Wind turbine design cost and scaling model, Tech. rep., National Renewable Energy Laboratory (NREL), Golden, CO., https://doi.org/10.2172/897434, 2006.

Fleming, P., King, J., Dykes, K., Simley, E., Roadman, J., Scholbrock, A., Murphy, P., Lundquist, J. K., Moriarty, P., Fleming, K., et al.: Initial results from a field campaign of wake steering applied at a commercial wind farm–Part 1, Wind Energy Science, 4, 273–285, https://doi.org/10.5194/wes-4-273-2019, 2019.

Fortin, F.-A., De Rainville, F.-M., Gardner, M.-A., Parizeau, M., and Gagné, C.: DEAP: Evolutionary Algorithms Made Easy, Journal of Machine Learning Research, 13, 2171–2175, 2012.

Frandsen, S.: Turbulence and turbulence-generated structural loading in wind turbine clusters, 2007.

Frandsen, S., Barthelmie, R., Pryor, S., Rathmann, O., Larsen, S., Højstrup, J., and Thøgersen, M.: Analytical modelling of wind speed deficit in large offshore wind farms, Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology, 9, 39–53, https://doi.org/10.1002/we.189, 2006.

Gebraad, P. M., Teeuwisse, F. W., Van Wingerden, J., Fleming, P. A., Ruben, S. D., Marden, J. R., and Pao, L. Y.: Wind plant power optimization through yaw control using a parametric model for wake effects—a CFD simulation study, Wind Energy, 19, 95–114, https://doi.org/10.1002/we.1822, 2016.

Global Wind Energy Council: Global Wind Report 2025, Tech. rep., Global Wind Energy Council (GWEC), Lisbon, Portugal, https://www.gwec.net/reports/globalwindreport/2025, 2025.

Grady, S., Hussaini, M. Y., and Abdullah, M. M.: Placement of wind turbines using genetic algorithms, Renewable energy, 30, 259–270, https://doi.org/10.1016/j.renene.2004.05.007, 2005.

International Electrotechnical Commission: IEC 61400-1: Wind energy generation systems — Part 1: Design requirements, Tech. rep., International Electrotechnical Commission, Geneva, Switzerland, 2019.

https://doi.org/10.5194/wes-2026-130
Preprint. Discussion started: 4 August 2026

Jiménez, Á., Crespo, A., and Migoya, E.: Application of a LES technique to characterize the wake deflection of a wind turbine in yaw, Wind energy, 13, 559–572, https://doi.org/10.1002/we.380, 2010.

Lara, M., Ruz, M. L., Vázquez, F., Mulders, S. P., van Wingerden, J.-W., and Garrido, J.: Analysis of Collective and Enhanced Individual Pitch Control Strategies for Wind Turbines, in: 2026 IEEE 23rd Mediterranean Electrotechnical Conference (MELECON), pp. 1–6, IEEE, 2026.

Lüers, S., Wallasch, A.-K., and Heyken, M.: Kostensituation der Windenergie an Land – Stand 2024, Tech. Rep. SP24011A1, Deutsche WindGuard GmbH, erstellt im Auftrag des Bundesministeriums für Wirtschaft und Klimaschutz (BMWK), FKZ 03MAP434, 2024.

Manwell, J. F., McGowan, J. G., and Rogers, A. L.: Wind energy explained: theory, design and application, John Wiley & Sons, 2010.

Mendez Reyes, H., Kanev, S., Doekemeijer, B., and van Wingerden, J.-W.: Validation of a lookup-table approach to modeling turbine fatigue loads in wind farms under active wake control, Wind Energy Science, 4, 549–561, https://doi.org/10.5194/wes-4-549-2019, 2019.

Miner, M. A.: Cumulative damage in fatigue, https://doi.org/10.1115/1.4009458, 1945.

Mosetti, G., Poloni, C., and Diviacco, B.: Optimization of wind turbine positioning in large windfarms by means of a genetic algorithm, Journal of Wind Engineering and Industrial Aerodynamics, 51, 105–116, https://doi.org/10.1016/0167-6105(94)90080-9, 1994.

Mudafort, R. M., paulf81, Bay, C., misi9170, Hammond, R., ejsimley, nhamilto, Bachant, P., Doekemeijer, B., Fleming, K., jrannoni, Quon, E., Stanley, P., Ireland, P., Cia, P. B., Yalla, G., zerweck, scottryn, rctredgold, pduff code, mkrause strath, jfrederik nrel, hannahgielen, afarrell2, Martinez, T., Sortland, S., Simone, Riva, R., and Seim, K. S.: NREL/floris: v4.6, https://doi.org/10.5281/zenodo.18039100, 2025.

Palmgren, A.: Die Lebensdauer von Kugellagern, Zeitschrift des Vereines Deutscher Ingenieure, 68, 339–341, 1924.

Resor, B. R.: Definition of a 5MW/61.5m wind turbine blade reference model, Tech. Rep. SAND2013-2569, Sandia National Laboratories, Albuquerque, NM, https://doi.org/10.2172/1095962, 2013.

Rice, S. O.: Mathematical analysis of random noise, The Bell System Technical Journal, 23, 282–332, https://doi.org/10.1002/j.1538-7305.1944.tb00874.x, 1944.

Shaler, K. and Jonkman, J.: FAST.Farm development and validation of structural load prediction against large eddy simulations, Wind Energy, 24, 428–449, https://doi.org/https://doi.org/10.1002/we.2581, 2021.

Sørensen, J. D., Frandsen, S., and Tarp-Johansen, N.: Effective turbulence models and fatigue reliability in wind farms, Probabilistic Engineering Mechanics, 23, 531–538, https://doi.org/10.1016/j.probengmech.2008.01.009, 2008.

Stanley, A. P. J., King, J., Bay, C., and Ning, A.: A model to calculate fatigue damage caused by partial waking during wind farm optimization, Wind Energy Science, 7, 433–454, https://doi.org/10.5194/wes-7-433-2022, 2022.

Thomsen, K. and Sørensen, P.: Fatigue loads for wind turbines operating in wakes, Journal of wind engineering and industrial aerodynamics, 80, 121–136, https://doi.org/10.1016/S0167-6105(98)00194-9, 1999.