Final author comments as response to reviewers on the original manuscript

Beyond wind-sector management: Optimal wind farm operational planning for balancing fatigue damage and extending lifetime

Niklas Requate, Lukas Vollmer, René Hofmann

May 2025

We kindly thank the reviewers for their detailed and valuable feedback. Below, we address each comment and describe the corresponding changes made to the manuscript.

- Reviewer comment in blue
- Our answer in black
- *Suggested revisions in Italic*

General comments

RV1: This paper presents an optimization approach that considers derating to extend turbine operational life and increase overall energy production over the farm's lifespan. The algorithm uses an engineering wake model with a surrogate model for damage. Results are presented for a small wind farm based on binned steady state conditions derived from ERA5 data, both on- and offshore. The selected operation plans were shown in simulation to improve net present value with moderate reductions in annual energy production, showing the promise of the approach.

RV2: This manuscript describes an approach for strategically using derating of individual turbines to achieve lifetime extension and NPV uplift for wind farms. The approach herein uses a time-monolithic nonlinear program-based optimization over the sector average-driven estimates of the AEP performance and lifetime load accumulation using both global and local (in the turbine/farm sense) strategies. The finding is that NPV increases are available by extending the operational lifetime of turbines while decreasing AEP in all present studies except for when limiting blade edgewise loads. Overall this paper is a strong contribution to the literature, and begins to elucidate the relationship between derating strategies and overall economic performance, though many outstanding questions remain, as the authors note, about how to act on the findings.

We thank both reviewers for their positive assessment and constructive feedback. We address all points below.

Surrogate models and optimization approach

RV1: Please discuss the creation and validation of the surrogate models. How well do they represent real world damage observations? What is the current state-of-the-art w.r.t. component failure predictions?

RV2: I would like to see a little more explication of the modeling approaches for 1) the surrogate modeling as well as 2) the development of the per-turbine optimization approach. The latter is likely to include some impact from ordering to capture non-linear effects, such as in Algo. 1, L4, where optimal choices will propagate to neighboring turbines.

We have deliberately kept the discussion of surrogate models and per-turbine optimization brief in this paper, as much of this topic has already been described and discussed in detail in our previous paper (Requate and Meyer, 2023). We now make this reference more explicit and add additional detail at the relevant locations.

Surrogate models: We add a clarifying sentence in Section 2.2 specifying the model type (polynomial regression fits through DELs from aeroelastic simulations). In the Discussion, we expand the paragraph on model uncertainties to explain the nature of errors coming from the surrogate modelling approach for the fatigue damage.

Revised text in **Section 2.2** (added sentence after the opening):

In this work, the surrogate models are polynomial regression fits through damage equivalent loads (DELs) obtained from aeroelastic simulations over a range of wind speeds, turbulence intensities, and derating setpoints. Their construction and the underlying simulation database are described in detail in Requate and Meyer (2023).”

Revised text in the **Discussion section** (replacing the existing paragraph on model uncertainties):

In addition to the uncertainties in the long-term evaluation of the economic benefits, the results also depend on the underlying models which are used on several levels, i.e. the configuration of the turbine controller, the surrogate models for the fatigue damage calculation, and the wake models. Controller-specific choices influence how loads and the wake respond to derating. The surrogate models used here are polynomial fits through DELs from aeroelastic time-domain simulations. These fits interpolate between discrete simulation points and thus introduce additional approximation errors. However, when the underlying simulations are sufficiently accurate and the turbine-controller interaction is tuned appropriately, these interpolation errors are small compared to the uncertainties inherent in the fatigue damage estimation itself. The surrogate approach based on DELs and linear damage accumulation is tractable and standard for design and control-oriented analysis of wind turbines, but it does not capture sequence effects or material *nonlinearities that can be relevant for composites. A detailed discussion of these modeling uncertainties, including the sensitivity to controller tuning, is provided in Requate and Meyer (2023) and applies equally to the farm-level extension presented here. In the wind farm context, the wake modelling approach also influences results. Firstly, dynamic effects cannot be represented with steady-state models. Secondly, the parameters of the steady-state model influence the effect of the turbines on the wake.*

Per-turbine optimization approach and loop ordering: We believe this aspect mainly needs some clarifcation: The local wind conditions for all turbines are computed *simultaneously* by the wake model (which internally resolves the upstream-to-downstream propagation for each wind direction). The subsequent per-turbine optimizations are therefore independent of each other and insensitive to the processing order, since they all operate on the same fixed local conditions. For the details on this, we refer to our previous paper again. Coupling between turbines only occurs *between* iterations, when updated derating setpoints modify the wake field. We add a clarifying passage in Section 2.3:

Within each iteration k, the local wind conditions for all turbines are computed simultaneously by the wake model (which internally resolves the upstream-to-downstream propagation order for each wind direction). The subsequent per-turbine optimizations (inner loop over i) are therefore independent of each other and insensitive to the order in which turbines are processed, since they all operate on the same fixed local conditions $\mathbf{X}^{(k)}$. Here, the exact same problem as described in Requate and Meyer (2023) is solved. Coupling between turbines only occurs between iterations, when updated derating setpoints change the wake field.

Integration of power prices

How feasible or valuable would it be to include power price in the optimizations, e.g., for farms with strong diurnal variations in power demand?

This is an excellent question and a very interesting topic. We have been working on including power price in our planning approach by using the wind-price correlation as an additional input dimension. We presented initial results at the Deepwind conference in January 2025 and are planning to publish a dedicated paper on this topic. The focus of the present paper was extending the method from turbine level to farm level. We add a short outlook in the Conclusion section:

A further step for practical usage is the integration of external influences beyond the environmental conditions. In particular, the integration of electricity prices into the planning process will be relevant for future wind farms operating in flexible markets, especially those with strong price variations. We have recently extended our single-turbine planning approach to account for the wind-price correlation as an additional input dimension and presented initial results at the Deepwind conference (Requate and Schmelter, [Presentation, 2025]). Extending this to the farm level is a natural next step.

In addition, it could also be investigated how the energy reserve from derating can be used for participation in the balancing market.

Code availability

Even if the code is not useful for general purposes, it should be shared anyway, as it fully describes the computational methods, whereas prose alone is insufficient. There is no need to make it user friendly or general purpose. If it is code that should be kept secret for commercial reasons, that is an acceptable reason to state as well.

We appreciate this comment. At the moment, the code can be shared with interested readers upon request. We are still in internal discussions regarding public availability. We have revised the code availability statement to be more transparent:

"The Python tool VIOLA was developed for an in-house workflow and is not yet publicly released. The code can be shared upon request to the corresponding author. The dataset of short-term DELs used to create the surrogate models is available at `https://doi.org/10.5281/zenodo.8385296`. For further inquiries about data or code, please contact the authors."

Notation

Overall, I would consider simplifying the notation a little bit and adding some more straightforward and immediate clarification for the definitions- it's taking me a few passes through a few sections to fully understand the material as it's presently presented.

We understand this concern. In the original manuscript, the notation was introduced primarily through two dense tables (one for variables with formal definitions and examples, one for indices), which required the reader to parse all symbols simultaneously before proceeding.

We have restructured the opening of Section 2. The first of the two tables is replaced by a step-by-step verbal explanation of the three key concepts—ambient conditions, local conditions, and control setpoints—each introduced as a named paragraph with descriptions, formal definitions, and concrete examples inline.The second table as a summary of dimensions and indices remains. We believe this structure is more accessible while retaining the mathematical precision needed for the subsequent subsections.

Revised opening of **Section 2** (replacing the original introductory paragraphs and Tables 1–2):

We formulate a nonlinear optimization problem to find an operational plan for arbitrary wind farms. An operational plan is a set of control setpoints to be applied under specific external (ambient) conditions. In this work, the external conditions are ambient wind conditions; the plan exploits the nonlinear relationship between energy yield and damage. In Requate and Meyer (2023), a four-step process for optimal operational planning of a *single turbine was introduced. Here, we extend the system boundary to the entire farm and develop a method that produces an operational plan for each turbine, accounting for wake interactions.*

The optimization problem is constructed to select an optimal control setpoint for each combination of discretized input conditions. The result is a lookup table that a wind farm control system can follow. We now introduce the key quantities and their relationships step by step. A summary of all indices and dimensions is provided in Table 1.

Ambient conditions and binning. *The wind farm inflow is described by an ambient input vector* $x^{\mathrm{amb}} = (x^{(\mathrm{amb},1)}, \ldots, x^{(\mathrm{amb},W)}) \in \mathbb{R}^W$ *with* W *ambient variables. These ambient conditions represent the undisturbed inflow to the farm. They are the same for all turbines, and serve as the external input to the optimization. The continuous range of ambient conditions is discretized into* B *bins, where* $B = \prod_{\ell=1}^{W} B_{x^{(\ell)}}$ *is the product of the number of bins per variable. Each bin* j *is represented by a central value* x_j^{amb}*, and all representatives are collected in the matrix* $X^{\mathrm{amb}} \in \mathbb{R}^{B\times W}$*. Example: With* $W = 3$ *ambient variables (wind direction* θ^{amb}*, wind speed* v^{amb}*, turbulence intensity* TI^{amb}*) and bins of* 10° *and* $3\,\mathrm{m/s}$*, one obtains* $B = 36 \cdot 6 = 216$ *bins with representatives such as* $x_1^{\mathrm{amb}} = (0^\circ,\, 7\,\mathrm{m/s},\, 0.08)$*.*

Local conditions. *In a wind farm, each turbine experiences different wind conditions due to wake effects. We denote the local (at-turbine) conditions by* $x = (x^{(1)}, \ldots, x^{(w)}) \in \mathbb{R}^w$*, representing the effective wind at the turbine rotor after accounting for wake deficits and added turbulence from upstream turbines. Local conditions depend on the ambient inflow, the farm layout, and the operational setpoints of all turbines. The collection of local conditions for all turbines and bins forms a tensor* $\mathbf{X} \in \mathbb{R}^{N_T\times B\times w}$*, with the per-turbine, per-bin vector* $x_{ij} = \mathbf{X}(i, j, :)$*. Example: With* $w = 2$ *local variables (wind speed* v*, turbulence intensity* TI*), a turbine operating in the wake of an upstream neighbor might see* $x_{0,0} = (7.5\,\mathrm{m/s},\, 0.12)$ *when the ambient conditions are* $x_0^{\mathrm{amb}} = (0^\circ,\, 8.0\,\mathrm{m/s},\, 0.08)$*.*

Control setpoints and operational plan. *For each turbine and each ambient bin, the controller applies a setpoint* $u \in \mathbb{R}^{N_u}$ *with* N_u *control channels. In this work, a single control channel is used: the derating factor* δ^P *that reduces the available power output by a given percentage (*$N_u = 1$*). The collection of all setpoints forms the operational plan* $U \in \mathbb{R}^{N_T\times B\times N_u}$*, or simply* $U \in \mathbb{R}^{N_T\times B}$ *for the derating-only case. Each entry* u_{ij} *specifies the derating setpoint for turbine* i *in ambient bin* j*. Example: The value* $u_{1,0} = 90\,\%$ *means that turbine 1 operates at 90 % of its available power in ambient bin 0, i.e. at* $x_0^{\mathrm{amb}} = (0^\circ,\, 8.0\,\mathrm{m/s},\, 0.08)$*. More generally, additional control channels such as yaw misalignment could be included, but all results in this paper use* $N_u = 1$*.*

Table 1: Dimensions and indices used in the problem formulation

Name	Notation
Number of ambient variables	$W \in \mathbb{N}$ (here: $W = 3$)
Number of local variables	$w \in \mathbb{N}$ (here: $w = 2$)
Number of bins for variable $x^{(\ell)}$	$B_{x^{(\ell)}} \in \mathbb{N}$
Total number of ambient bins	$B = \prod_{\ell=1}^{W} B_{x^{(\ell)}}$
Number of turbines	$N_T \in \mathbb{N}$
Number of control channels	$N_u \in \mathbb{N}$ (here: $N_u = 1$)
Turbine index	$i \in \{1, \ldots, N_T\}$
Ambient-bin index	$j \in \{1, \ldots, B\}$
Control-channel index	$m \in \{1, \ldots, N_u\}$

With these definitions, the following subsections describe how the wake model maps ambient conditions and control setpoints to local conditions (Sect. 2.1), how surrogate models translate local conditions into power and damage (Sect. 2.2), and how these components are combined into the optimization problem (Sect. 2.3).

In addition, regarding RV2's question on local conditions: the local conditions represent the effective incoming wind at each turbine's rotor, including wake deficits and added turbulence from upstream turbines. This is now clarified explicitly in the "Local conditions" paragraph above.

Interpretation of results

A key aspect I would like to see foregrounded a little more is what to do with the results here in terms of operational strategy.

If I were operating a farm in real time, I'm not necessarily compelled to work in this time-monolithic way, but rather projecting out the cost/value of a) what already happened to my farm/what I can detect about it and b) what I see, say 1 decade into operation, as the future circumstances: I might be in a totally different interest rate regime, or cost regime, or whatever, and I'm going to think of fatigue optimal choices w.r.t. this and my projections of the next 10-, 15-, 20-year period. There are likely future moments where repowering is the optimal choice given the trajectory I've taken to a given point, and futures in which it absolutely does not make sense to repower at all- and the perspective from which that question is relevant is not a time-monolithic one at all.

We fully agree with this concern. Circumstances can change significantly during 25–30 years of operation. On the one hand, operators have to consider such changes for the assessment of their investment and strategic decisions regardless of the underlying operational strategy—when operating a wind farm with its nominal strategy at all times, a time-monolithic operation is assumed implicitly. On the other hand, the optimization of operational strategies under consideration of fatigue introduces further degrees of freedom which make decisions more complex and sensitive to the uncertainties in the underlying methods.

Nevertheless, we consider our approach as a tool that can be applied not only once before commissioning but periodically updated throughout the farm's lifetime. In a previous study, we also investigated annually variable operational strategies and found surprisingly small benefits from year-to-year variability, because the nonlinear relationship between energy production and damage—which determines where lifetime can be saved efficiently—is largely independent of when during the lifetime the derating is applied.

In this paper, we concentrate on presenting the methods and comparing results on a global time-monolithic level. We expand the existing paragraph on replanning in the Conclusion/Outlook section to address how operational strategies can be used over time and in perspective of future operator decisions:

The evaluations in this study assumed that the optimized operational plan is applied consistently throughout the lifetime. In practice, however, economic and regulatory conditions change over decades of operation, and the turbine's condition evolve differently than predicted. The proposed framework can therefore be applied not only once before commissioning but periodically throughout the farm's lifetime, using updated site data, measured fatigue consumption, and revised economic assumptions. In this sense, the approach can be interpreted as a very slow model-predictive control scheme: at each replanning point, the remaining fatigue budget is assessed and the operational strategy for the remaining years is re-optimized. Nevertheless, the planning will always depened on the assumptions about future conditions and thus need to be carefully be compared to standard operation and potential drawbacks. n Requate (2024), we additionally investigated annually variable operational strategies where the derating plan is allowed to differ between years. Notably, the benefit of year-to-year variability was found to be small, because the nonlinear relationship between energy production and induced damage, which determines where lifetime can be saved most efficiently, is largely independent of when during the lifetime the derating is applied. However, the benefit also depends on external circumstances such as interest rates, price developments, and the farm's remuneration model, meaning that annually variable strategies may still be valuable under certain conditions

Minor comments

Reviewer 1

1. Line 7: The range 1–4 uses a hyphen instead of an en-dash.

Fixed.

2. The hyphen in the title should probably be a colon.

Fixed: title changed to *"Beyond wind-sector management: Optimal wind farm operational planning for balancing fatigue damage and extending lifetime"*.

3. Line 84: "he framework" should be "The framework".

Fixed.

4. Line 89: "tower bm" should be "tower BM".

We have consistently used "bm" and would prefer to keep this.

5. Figure 2: Italicization for sub- and superscripts does not match the text.

Fixed in Figures 1 and 2.

6. Table 3: The italicization here for subscripts does not match the text.

Fixed, e.g. from $d_{flap}(x, u)$ to $d_{\mathrm{flap}}(x, u)$. We hope that this is what the reviewer meant

7. Figure 3a: The farm layout appears to be measured in meters, not rotor diameters.

Fixed: caption of Figure 3a now reads *"Farm layout (distance in m; 1 rotor diameter = 164 m)"* to clarify the axis units.

8. Figure 4: If longer lifetime is preferable, perhaps green or blue would be a better color than red for the color bar.

We revised the colorbar to use a green color scheme where higher lifetime corresponds to a more positive color.

9. Line 476: Typo in the word "specifically".

Fixed.

We thank Reviewer 1 for the careful reading.

Reviewer 2

L17: "For energy, the objective is clear: ..." I'm not sure the objective is totally clear especially in the context of a fatigue life study rate of return, capacity factor/capacity utilization factor, PPA and debt service periods vs. post-contract pure-profit operations, etc. all should factor into what this objective actually might be highly variable depending on actor, regulatory climate, etc. L17, "For loads..." and subsequent L18 "In contrast, ..." lines are duplicative; suggest rewording

We merged the two sentences into one concise statement. At this point, we really mean pure energy production on a technical level. Economical considerations have to be considered an another level. We clarified be adding the word "pure"

"For pure energy production, the objective is clear: maximize farm-level annual energy production (AEP). For loads, however, the challenge is greater because energy production inevitably induces fatigue on multiple components of each turbine, and it is not straightforward to define a single optimization target."

L116: XambXamb isn't defined well enough? lead with Table 1?

We have expanded the introductory paragraph before Section 2.1 to provide verbal explanations of all key quantities before the formal notation (see response to "Notation" comment above).

local conditions (xx): is this upstream of the turbine w/ wake effects from upstream turbines but no induction effects? some clarification, please!

The local conditions represent the effective incoming wind at each turbine's rotor, including wake deficits and added turbulence from upstream turbines. This is now clarified in the expanded introductory text of Section 2.

L126: what are the surrogate models? are there details in here somewhere?

We added a sentence in Section 2.2 specifying the model type (polynomial regression fits through DELs from aeroelastic simulations) and referencing the previous paper for full details (see response to "Surrogate models" comment above).

L288, typo: your "NWP" I think should be "NPV"

Fixed: "NWP" corrected to "NPV".

Fig. 17: caption is unclear, appears to be incorrect or at least ambiguous

Revised to: *"Case 3, Evaluation of results: Offshore scenario with all failure modes, Comparison of AEP and NPV for the optimization approaches."*

on visualizations with "turbine index", since you have a 9-turbine farm, perhaps use "N", "NE", "E", ..., "NW", and "C" (for center) rather than indices, to help the reader interpret which row/column is impacted by derating choices

We added a brief explanation of the turbine indexing when the farm layout is first introduced in the Results section:

"Throughout this work, turbines are referenced by their index $i \in \{0, \ldots, 8\}$ *as shown in Fig. 3a, where indices 0-2 correspond to the bottom row (most most western), 3-5 to the middle row, and 6-8 to the top row (most eastern)."*

This provides the reader with the spatial context needed to interpret the per-turbine results.

Additional minor corrections

In addition to the reviewer comments, we corrected the following issues identified during revision:

- **Appendix A:** “NPC” corrected to “NPV”; clarified the sentence: *“While total energy production over the lifetime is higher than under nominal operation for all optimized results, NPV can decrease when the reduction in AEP is too large relative to the lifetime extension.”*
- **Code availability statement:** Revised for clarity and transparency.